

Perspectives and Good Dispositions

Maria Lasonen-Aarnio

*Penultimate draft, forthcoming in Philosophy and Phenomenological Research;
please cite the published version*

I. Miners, the Access Diagnosis, and Perspectivism

I.1 Miners

My aim will be to identify an evaluative lacuna, and to discuss different approaches for filling it. To identify the lacuna, let me begin with well-known type of case, *Miners*. It has often been used in discussions of the semantics of ought-statements. As will become clear, my project here is not semantic. Rather, I want to use the case to highlight the need for making a certain kind of evaluation. The well-known case is as follows:

Miners. 10 miners are trapped in one of two shafts (shaft 1 or shaft 2), and floodwaters are rising. You must decide which shaft to block before finding out where the miners are. They are no more likely, given your evidence, to be in 1 or 2. You are able to block the water from reaching one of the shafts, but you don't have enough sandbags to block both. If you manage to completely block the shaft where the miners are, they are all saved; if you block the other shaft completely, they all drown. If you do nothing, letting both of the shafts will halfway with water, one miner will drown in any case.¹

What ought you to do?² Such cases are puzzling, many think, for we seem to be pulled in two different directions. It is clear what the *best* option is in this case: blocking whatever shaft the miners are in, thereby saving all of them. On some views, this gives rise to an objective sense of 'ought': if all the miners are in shaft 1, then you objectively ought to block shaft 1. However, that cannot be the end of the story. For everyone seems to agree that you would merit some kind of positive evaluation for letting both shafts flood halfway, and at the very least, you would have an excuse for doing so, or be blameless for doing so. And were you to choose to block shaft 1, perhaps you could be blamed.³ Further, many think that appeal to excuses and blame is not enough: it is true to say that you ought to do nothing, letting both shafts flood halfway.⁴ Setting aside ought-statements, I agree that appeal to mere blamelessness or excuses does not do justice to our inclination to positively evaluate the subject who blocks neither shaft. These categories are just too diffuse, and as a result, the tool too blunt. For instance, a subject who panics under pressure, randomly blocking one of the shafts, which happens to be the wrong (e.g. empty) one, may be blameless or have an excuse. But they are not positively evaluable.

¹ E.g. Regan (1980: 265), Parfit (1988). There are many structurally similar cases, such as Jackson's (1991: 462-463) case of *Dr. Jill*.

² The question is standardly understood to be what you overall *morally* ought to do, and this is how I will understand it. Those who think there are more subjective and objective senses of 'ought' typically think that such senses arise for different kinds of oughts (e.g. Wedgwood 2016).

³ See Olsen (2017, §4) for a discussion of the limits of the objectivist view appealing to notions like blamelessness.

⁴ For instance, Kolodny and MacFarlane (2010: 115) say that it is obvious – largely independently of one's moral views – that one ought to block neither shaft.

Though not everyone agrees, there is a long tradition of distinguishing between objective and subjective senses of ‘ought’, or objective and subjective notions of moral rightness.⁵ Similarly, many distinguish between the ought of advisability and the ought of rationality.⁶ According to a view running through much of the epistemology literature, a parallel distinction is needed when considering what we ought to believe. Perhaps it is objectively best to believe a proposition p just in case p is true. But consider a simple case like the following:

Coin toss. All you know is that a fair coin was just tossed in the next room. In fact, the coin landed heads. Ought you believe it landed heads, believe it landed tails, or suspend judgment about the matter?

Gibbard (2005), for instance, thinks that in the objective sense, you ought to believe that the coin landed heads, since that is true, but what doxastic states you ought subjectively to have depends on your evidence. Some talk about objective and subjective senses of ‘justification’ instead.⁷

As I already noted, my project here won’t be to provide a truth-conditional semantics for natural language ought-statements. What I want to focus on, rather, is how we should think of the kind of more “subjective” evaluative perspective that cases like *Miners* call for. Why are norms like *Choose the best!* and *Believe p just in case p is true!* inadequate?⁸ I will begin with the observation, often made in the literature, that even if you manage to block the shaft the miners are all in, you do what is best merely by *luck*, *accident*, or *fluke*: there is something haphazard about your conformity to the objectivist norm. As a result, the objectivist norm leaves open an evaluative lacuna. My first aim in what follows will be to criticize a broad *perspectivist* framework for diagnosing and solving the problem. My second aim will be to put forth an alternative diagnosis and solution.

I am here using *Miners* and the norm *Choose the best!* as a dialectical starting point, but a very wide range of norms, I will argue, face essentially the same problem, leaving open a similar kind of evaluative lacuna. Note at the outset that the problem does not essentially assume a consequentialist way of thinking.⁹ And of course, knowledge-firsters will claim that you shouldn’t believe that the coin landed heads in *Coin toss*, for you are in no position to *know* this. Still, there will be cases in which we want to positively evaluate, even praise, beliefs that fail to constitute knowledge. This is true, for instance, of diligent subjects with massively misleading evidence regarding a question, who fail to know because they hold a false belief (such thoughts give rise to the so-called New Evil Demon problem). The same seems true of a subject in a Gettier-case.¹⁰ A knowledge-norm on its own provides us with impoverished evaluative resources.

The queasiness many feel when considering situations in which a subject can only conform to a relevant norm by luck is often voiced in tandem with a range of other thoughts. A common complaint is that at least sometimes one cannot be *guided* by objectivist norms. Likewise, in *Miners* the success of conforming to the objectivist norm cannot be *attributable* to the agent, even if she manages to block the correct shaft, thereby saving all of the miners: she is not *creditable* for this

⁵ This was a rather standard view among moral philosophers working in the first half of the twentieth-century. For some representative examples from the past decades, see Jackson (1986), Gibbard (1990, 2005: 340), Oddie and Menzies (1992), Olsen (2017), and H. M. Smith (2010). Another kind of case often given in support of more subjective oughts is one in which a subject has every reason to think that she is doing what is objectively best, but in which her choice results in a tragic outcome (see e.g. Holly Smith’s (2010: 64) *Twin Towers I* example).

⁶ For a recent example, see Schroeder (2018).

⁷ Feldman (1988) opens his paper with the sentence “A view widely held view by epistemologists is that there is a distinction between subjective and objective epistemic justification, analogous to the commonly drawn distinction subjective and objective justification (or rightness or obligation) in ethics”.

⁸ I will talk of norms, but for now leave it open whether these are normative in some more heavyweight sense, or whether they are better thought of as standards of evaluation. (However, I don’t claim to fully understand, nor do I here want to subscribe to, the distinction often drawn between the normative and the evaluative in e.g. Kolodny 2005).

⁹ Cf. H. M. Smith (2010).

¹⁰ I have also argued (Lasonen-Aarnio 2010, 2021) that there is such a thing as “unreasonable knowledge”: there are also cases in which a subject is negatively evaluable, even to be blamed for believing a proposition p , even if she knows p .

success and, as a result, cannot be *praised* for saving the miners. Indeed, a very familiar meta-normative picture – one that I should say at the outset I don’t subscribe to – ties all of these diagnoses of what is problematic about objectivist norms together. On this picture, various evaluations are understood in terms of the relations subjects bear to norms (or rules). In particular, a subject can (a) conform to (or fail to conform to) a norm, and they can (b) be guided by or follow (or fail to be guided by/follow) a norm. The problem, on this picture, with many of the norms discussed is that at least in some situations one cannot be *guided* by them. It is no surprise that luck, guidance, and credit are often mentioned in one breath, for guidance is assumed to yield the kind of non-lucky conformity to a norm required for being creditworthy for choosing, doing, or believing as one ought.¹¹ Views exemplifying the *perspectivist* framework I will discuss are often explicitly formulated in terms of the norm-centered picture just outlined.

I will begin (in §1.2) by outlining a prominent diagnosis of just why we need more subjective norms in cases such as *Miners*, and an accompanying view of what those norms should look like, which I call **perspectivist**. Perspectivism arises from a rather natural and familiar diagnosis of why norms like *Choose the best!* or *Believe the truth!* are inadequate: they make reference to a domain of facts that sometimes (often) lies outside our perspectives. As will become clear, I don’t think this diagnosis is correct. In fact, perspectivism leads us to a bleak picture on which the problem is inescapable, for perspectivist norms face essentially the same problem as objectivist ones. I will outline an alternative diagnosis, and an alternative framework for how to think of a more subjective evaluative dimension, one that appeals not to perspectives but *feasible dispositions*. But first, let me say a bit more about perspectivism and its different forms.

1.2 Perspectivism

The problem with more objective norms, it is often argued, is that they make reference to facts that may not, in Jackson’s (1991: 467) words, be “present to the agent’s mind”. A natural understanding of being present to an agent’s mind is epistemic: the agent has some sort of epistemic access to the fact. According to the *access diagnosis*, the heart of the problem is that we often lack access to the facts that the recommendations of more objectivist norms depend on: you don’t know which shaft the miners are trapped in!¹² (Indeed, sometimes the subjective ought is referred to as an *information-relative* one.¹³) As noted above, here the notion of guidance is often brought in: objective norms don’t provide us with adequate guidance about what to do.¹⁴ And the reason for this is precisely that we often lack epistemic access to whether their application conditions obtain.¹⁵ As pointed out above, it is often assumed that guidance is necessary and sufficient for acting as one ought in a way that is not a mere lucky fluke or coincidence. Indeed, such non-lucky normative success is one of the central goods that guidance has been looked upon to deliver.¹⁶ My view is that we do best to keep guidance and non-lucky normative success apart. Though I will have much more to say about luck, I will largely set aside the issue of guidance.

¹¹ This is also assumed by many who think of normative guidance as a matter of responsiveness to the normative reasons in virtue of which one ought to do this or that, which in turn is a matter of being motivated to act by those reasons (see e.g. Way and Whiting (2017), Gibbons (2013: 135), Kiesewetter (2017: 11), Lord (2015, 2018)). On an alternative view, guidance is a matter of using a normative principle in deliberation about what to do (see H. M. Smith (2012) and Hughes (2018); see also Pollock’s (1987: 64) discussion of “the intellectualist model”).

¹² See, for instance, Gibbard (1990: 343): “The basic normative precepts that ground a subjective *ought* are subjectively applicable – applicable in light of information the agent has”. H. M. Smith (2010: 84) characterizes one of the attractions of the notion of subjective rightness as follows: “it could be used to identify a type of duty to which the agent has infallible access in his decision-making”.

¹³ E.g. Wedgwood (2016).

¹⁴ See e.g. Gibbard (1990: 344), Jackson (1991: 466-7), and Sepielli (2012).

¹⁵ For a discussion of guidance and epistemic access, see Hughes (2018).

¹⁶ For more discussion, see Lasonen-Aarnio (2019a).

I will characterize a very broad set of views – indeed, what I take to be a broad evaluative paradigm – in terms of the notion of a *perspective*.¹⁷ An agent’s perspective is whatever is, in the required sense, present to her mind. It is important to note at the outset that the central contrast of this paper between perspectivist views, on the one hand, and the kind of view I will defend appealing to feasible dispositions, on the other, is not tantamount to that between internalism and externalism. Though it is sometimes assumed that internally alike subjects have identical perspectives, the perspectivist framework as such is not an internalist one.

According to many views, a perspective is a set of propositions. It may consist of the totality of propositions constituting a subject’s *evidence*.¹⁸ Or, it may consist of the totality of subjective reasons one has.¹⁹ Or – and such a view is not incompatible with either of the previous ones – it may consist of the totality of propositions a subject *knows*, or the totality of propositions she is in a *position to know*. (Indeed, it is difficult to find a paper on the subject which doesn’t point out that in *Miners* one doesn’t *know* which shaft the miners are in.) Or, perhaps a subject’s perspective consists of the totality of propositions she justifiably believes, or has justification to believe. On one extreme are views on which a subject’s perspective consists simply of the totality of propositions she believes.²⁰

On some views, a subject’s perspective consist not of propositions, but of a subset of her mental states.²¹ For instance, some think that there is a *sui generis* kind of mental state called a *seeming*. According to a phenomenal conservative variant of perspectivism, perspectives consist of the totality of one’s seeming-states.²² Or, it might consist of the totality of mental states that are “phenomenally conscious”.²³ But even views that think of perspectives more propositionally typically take them to be fixed by a class of mental states. For instance, according to a dominant form of internalism, one’s perspective (e.g. one’s evidence) supervenes on one’s non-factive mental states.²⁴ Externalists might think instead that one’s perspective supervenes on the mental state of knowing, or on mental states that include factive perceptual states.²⁵

We don’t always represent a candidate fact as either obtaining or not obtaining, even in cases in which we have some view of the matter. Contrast a subject who lacks any opinion about the weather in London right now with someone who thinks it 50% likely to be raining, or yet another subject who thinks it is 70% likely to be raining. One might think that perspectives either are, or at least give rise to, probability functions. On one possible view, a subject’s perspective is her subjective credence function.²⁶ Or, perhaps we should look to credences one ought to have²⁷, to those a reasonable subject in one’s position would have, or to one’s *evidential probability* function.²⁸

¹⁷ For a similar notion of a perspective, see Alston’s (1986) discussion of what he calls “perspectival internalism”. See also Kieseewetter (2011), Way and Whiting (2017), and Gibbons (2010: 335).

¹⁸ For such evidentialist views, see Conee and Feldman (2004)

¹⁹ One prominent way of understanding the distinction between subjective and objective oughts is in terms of subjective and objective reasons (e.g. Schroeder 2018). See Lord (2018) and Kieseewetter (2017) for views of subjective reasons formulated either in terms of knowledge, or in terms of what a subject is in a position to know.

²⁰ Though such a view leads to norms that appear *too* subjective: Consider a subject in a case like *Miners* who happens to believe, as a result of some irrational process not in any way sensitive to where the miners in fact are, that all of the miners are in shaft 1. Cf. Gibbard (2005: 346).

²¹ Pollock & Cruz (1999) think some (foundational) reasons are mental states, while others are propositional.

²² Numerous questions arise, of course: the totality of one’s seemings at a time *t*, or in one’s history, or of a select subset of the totality of seemings in one’s history?

²³ Cf Sepielli, who (2014: 524) also uses the term ‘perspective’.

²⁴ E.g. Conee and Feldman

²⁵ Williamson (2000), Pritchard (2012).

²⁶ Again, such a view appears too subjective (consider a subject who is irrationally almost certain that the miners are all in shaft 1). Cf. Jackson (1991).

²⁷ Parfit (1988) characterizes subjective rightness in terms of what one has reason to believe, and Gibbard (2005: 346) in terms of what one ought to believe.

²⁸ See e.g. Williamson (2000) on evidential probability.

Though there are significant differences between different conceptions of perspectives, they all have important features in common. Perspectives are closely connected with some class of *mental states*: indeed, it is precisely because a class of mental states gives rise to a perspective that perspectives capture what is present to an agent's mind. Second, perspectives are *representational*. According to many views, a perspective is a set of true propositions corresponding to a subset of the facts – namely, those facts that a subject has access to (e.g. those constituting her evidence, her reasons, or her knowledge). And mental states like seemings have propositional, representational contents. Further, even a probability function can be seen as a degreed representation of the world. For instance, it is precisely because it is not likely enough on the relevant probability function in *Miners* that the miners are all in shaft 1 that the correct subjective norm doesn't tell one to block shaft 1.

However one thinks of perspectives, they can be **incomplete**: not all facts are present to one's perspective. And however one thinks of perspectives, they can be **misleading**. Even if one's perspective consists exclusively of true or known propositions, a body of truths can support (make likely) falsehoods. For instance, if many normally reliable testifiers I have no reason to distrust tell me that the miners are all in shaft 2, then true (and known) propositions regarding what these testifiers tell me can make likely the falsehood that the miners are all in shaft 2.

Just about all views that endorse subjective oughts in response to cases like *Miners* appeal to perspectivist norms. According to a prevalent way of thinking, what we subjectively ought to do is maximize some quantity, where that quantity is a function of

- (i) a probability function, and
- (ii) a value function assigning objective values to the available options.

According to a prominent class of views, one subjectively ought to choose the action that maximizes *expected value* (or an action among those with highest expected value) by the lights of the relevant probability function, a function that is fixed by (or identical to) one's perspective.²⁹ For instance, according to Jackson's (1991) consequentialism, one ought to *maximize expected moral utility*. According to Parfit (1988), the subjectively right act is one that *maximizes expected goodness*. This is, of course, not the only possible view of subjective oughts, but it is the most popular.³⁰ I will try to avoid in-house debates as far as possible. Instead, I will sketch what I take to be fundamental problems for any perspectivist norms. While I will focus on the norm *Maximize expected value!*, it will hopefully be clear that the problems raised are not specific to this norm.

My aim in this paper will be to outline an alternative diagnosis of why, it would seem, you ought not to block shaft 1 in *Miners* – or, setting oughts aside, why you would at least be negatively evaluable for doing so, and positively evaluable for letting both shafts flood halfway. Begin with the following observation: in the situation described it is just not feasible for you to choose what to do in a way that is sensitive to the location of the miners. You can only choose the best by luck, hostage to the fortunes of various contingencies of your situation. I will suggest that we think about feasible ways of choosing (acting, and believing) in terms of feasible dispositions. The core of the problem, then, is that any feasible disposition that manifests in the situation described as choosing to fully block one of the shafts issues in tragic choices across a range of cases in which we allow the seemingly idiosyncratic features of the actual case to vary. Hence, choosing to fully block one of the shafts cannot be a manifestation of a good (enough) disposition. To understand the more subjective

²⁹ What exactly is the quantity to be maximized? Such debates connect with what the right decision theory is: evidential or causal decision theory; risk-weighted value, etc.? I won't take a stance on these in-house debates.

³⁰ Alternatively, one might propose, for instance, that one ought to choose what is subjectively most likely to maximize actual value. Such a proposal has come under many attacks (e.g. Parfit 1988, Jackson 1991). For a good overview and criticism of various proposals, see H. M. Smith (2010).

dimension of evaluation we are after we must focus not on perspectives, but on the best feasible ways of choosing, acting, and believing – the best feasible dispositions. As will become clear, the resulting picture rejects the perspectivist diagnosis, but also the meta-normative picture described on which various evaluations are to be understood in terms of one of two relations a subject can bear to a norm (following/being guided by a norm, or merely conforming to it).

One feature of my view worth flagging at the outset is that it unifies the practical and theoretical/epistemic domains. The dispositional perspective can be straightforwardly applied to evaluate beliefs: a belief is to be positively evaluated just in case it is the manifestation of the best feasible dispositions, the goodness or success-conduciveness of a disposition being now determined in terms of epistemic values like truth or knowledge. The view also straightforwardly applies to doxastic states other than belief. Just like believing, suspending judgment, for instance, can be a manifestation of the best feasible dispositions: sometimes the best feasible dispositions, from the perspective of values like truth and knowledge, manifest as suspending judgment (see Lasonen-Aarnio 2021).

Facts about perspectives and those about feasible dispositions are intertwined in many ways – indeed, deceptively so. Facts about epistemic access often explain the feasibility, or lack thereof, of manifesting certain kinds of dispositions. For instance, the fact that you don't know where the miners are explains why it is not feasible to make a choice in a way that tracks their whereabouts across a portion of modal space – and hence, why managing to save all of them would be merely lucky. Indeed, the best feasible dispositions are often sensitive to our perspectives. For instance, often the best we can do is to proportion our beliefs to our evidence. We often manage to at least roughly track or be sensitive to how things stand with respect to some domain of facts by being sensitive to our own representations of those facts. However, I will argue that the perspectivist diagnosis of why objectivist norms are inadequate doesn't get to the heart of the matter. In fact, perspectivism forces us to draw the conclusion that the problem we started out with is unsolvable: just as it is sometimes not feasible to track facts about the whereabouts of objects in space, it is sometimes not feasible to track facts about one's perspective. As a result, in some situations one can only conform to perspectivist norms by luck or accident. The problem of luck starts to look *unsolvable*.

I will first argue that perspectivist norms are themselves susceptible to miner-type cases (§2).³¹ I then outline my alternative, dispositional evaluative perspective (§3), showing how it applies to our original case *Miners*. In §4 I outline a way of thinking about the kind of luck or accidentality at issue, providing an alternative diagnosis, one having to do with feasible dispositions, for why norms like *Choose the best!* are inadequate. In §5 I argue that there are principled reasons why a dispositional norm is not susceptible to the problem of luck. In §6 I conclude with some broader reflections on ways in which perspectivist thinking may have impeded progress in various debates.

2. Revenge of the *Miners*

I outlined a broad framework for thinking about the more “subjective” kinds of evaluations that cases like *Miners* seem to call for: what a subject ought to do is fixed by her *perspective*. Perspectives capture how things are by the lights of a subject. However, even though it is often motivated precisely by appeal to situations such as *Miners* involving structural uncertainty, perspectivism cannot escape the challenge raised by such cases, for essentially the same problem appears at a new level:

³¹ I by no means think this is the *only* problem with perspectivism; another will emerge in §6. There are also further problems, like the possibility of good dispositions that simply bypass one's perspective. I cannot get into these issues here, but there is mounting empirical evidence of complex unconscious mental phenomena guiding complex social behaviour and goal pursuit (for a good overview, see Hassin, et al 2004).

in some situations one can only conform to perspectivist norms by luck. We need to shift our diagnosis of the original problem, and the framework used for solving it.

A perspective is part of the world. Just as there are facts regarding where miners are located, about what the effects of administering a given medication are, and about what the weather in London is like, there are facts about what one's perspective is in the first place, and what one ought to do by the lights of one's perspective. For instance, there are facts about what we believe, what we know, and what evidence we have. And one's perspective can take a perspective on such facts, one that can be incomplete or even misleading. One can lack epistemic access to these facts, just as one can lack epistemic access to facts about which shaft the miners are in.³² As should become clear, I don't think the access diagnosis gets to the heart of our original problem – and hence, I don't think the ultimate reason why perspectivism fails is that we sometimes lack access to our perspectives. (Below I also consider the objection that perspectivism as such is not wedded to the access diagnosis.) But let me first say why, *given the access diagnosis*, we should expect perspectivism to fall prey to the same problem. In fact, the access diagnosis forces us to draw the grim conclusion that the problem is unsolvable.

To make the problem vivid, consider the following case, in which one's access to mathematical facts about expectations is incomplete:

Another Mining Disaster

You often find yourself in situations involving mining disasters. To prepare, you spend your evenings analyzing particular scenarios, and calculating the expected values of various actions. You now find out there has been another accident. Luckily, just last night you calculated the expected values of the available actions in the very situation you now face. But alas, you have forgotten the exact results of those calculations. There is no time for new calculations – if you don't act quickly, all miners will die with certainty. You face a choice between actions A, B, and C. You remember the following facts: one of A or B has the highest expected value (10), while the other one has the lowest expected value (0). You know, and are certain, that C has quite a good expected value (9), though not as high as that of the best of the three options (A or B). What should you do?³³

What we have looks like a version of the miners problem: intuitively, you would be positively evaluable for choosing option C, even though you know that doing so does not maximize expected value. There is a parallel problem of luck and accidentality here: though you could make a choice that in fact maximizes expected value, doing so would be a matter of luck, your conformity to the norm *Maximize expected value!* a matter of fluke. It looks like you cannot genuinely follow or be guided by this norm; and even were you to conform to it, you could not be credited or praised for this normative success. Exactly the same considerations that motivate introducing a more subjective ought in the original *Miners* case appear to show that norms like *Maximize expected value!* are not subjective enough.

Calculating expected values is itself not a trivial matter; non-ideal subjects don't always have access to such mathematical facts. But this is but one instance of the more general problem. Consider, for instance, the fact that relations of evidential support are often highly non-trivial. We can construct structurally similar cases in which instead of lacking access to mathematical facts about expectations, you lack access to facts about evidential support. Assume that you went through some painstakingly complicated reasoning earlier on to figure out how likely a salient proposition p is on

³² One might also lack access to what choices or actions are genuinely available to one. One might lack access to facts about values, and to facts about what the correct normative theories are in the first place. The discussion below is closely related to discussions of normative uncertainty, but for the purposes of this essay, I am setting the problem of normative uncertainty aside.

³³ Adapted from Kagan (2018: 155). Spencer and Wells (2019) discuss a case with a similar structure, *The Fire*. See also H. M. Smith's (1988: 98-99) much earlier discussion of what she calls the 'Problem of Doubt'. And there is, of course, a large literature on related problems having to do with normative uncertainty.

your evidence (it is in fact very likely). Unfortunately, you have forgotten the result of that reasoning, and there is no time for thinking things through again. All you remember is that p is either very likely or very unlikely. If it is very likely, A has the highest expected value (10), B the lowest (0), and C has an expected value that is close to but not as high as A (9). If p is unlikely, B has the highest expected value (10), A the lowest (0), and C again has a rather high, though not maximal, expected value (9). In such a case, many would say that you ought hedge your bets by choosing option C, even though you know that it doesn't maximize expected value.

One reply to such cases is to appeal to some notion of an idealized agent – agents who have unlimited computational powers and access to all necessary, *a priori* facts about evidential support relations (assuming such facets to be *a priori*) – and to insist that the relevant theory of subjective oughts only applies to such agents. But a theory of a more subjective kind of evaluation that doesn't even apply to us non-ideal agents would be spectacularly disappointing. More importantly, idealizing in these ways doesn't solve the problem. For perspectives don't always accurately represent facts about themselves – that is, the perspectives that perspectives take on themselves can be incomplete and even misleading. Let a perspective be *transparent* to itself if it has full access to facts about what it consists in. In so far as a perspective is a set of propositions, transparency requires that it is always certain on one's perspective exactly what this set of propositions is. It requires both a claim of positive and negative access. If p is part of one's perspective, then it is part of one's perspective that p is part of one's perspective. If p is not part of one's perspective, then it is part of one's perspective that p is not part of one's perspective. Such access conditions are extremely strong. But I doubt that even the positive and negative access conditions suffice to make perspectives transparent to themselves. To see this, consider the following example.

A constellation of Greek gods (and only Greek gods) is trapped in an escape room together, and you are to figure out whether, given their known powers, they will be able to escape. You know that *most* arbitrary subsets of Greek gods could escape, given their reasonable level of intelligence and resourcefulness. For each god trapped in the room, you know of that god that they are trapped there (cf. *positive access*). For each god who is not trapped, you know of that god that they are not trapped (cf. *negative access*). As a matter of fact, it is Dionysios and Ares who are trapped. So you know that Dionysus is there, and that Ares is there. And, for instance, you know that Hera, Aphrodite and Zeus are not. Doesn't this give you knowledge of exactly what gods are in the room – don't you then know that Dionysus and Ares are in the room, *and that no-one else is*? It does not. For assume that your knowledge of Greek gods is rusty: you do your best to write down a list of each and every god, but for all you know, you have omitted someone! As a result, even though you have gone through your list, crossing out or underlining names, you cannot be sure that the two remaining underlined names constitute a full list of the contents of the room. You know that Dionysus and Ares would not be able to escape: they would end up getting drunk and fighting. But given an addition of almost *anyone* else, they could probably sort things out. So unfortunately, you do not know whether the trapped gods would be able to escape.

The point of the above example was to demonstrate just how difficult it can be to achieve full transparency, whether about the exact contents of a room, or about the exact contents of one's perspective (for instance, of one's evidence). Just as one's perspective can be incomplete when it comes to the world, one's perspective can be incomplete when it comes to one's perspective, by failing to correctly represent true facts about what is and what isn't part of one's perspective. Not only might it be incomplete, but it might even be misleading when it comes to such facts. It is blatantly clear that transparency fails on more externalist theories of perspectives – for instance, views on which one's perspective consist of all and only the propositions one knows. It is more or less uncontroversial that I can fail to know p , without being in any position to know this. But internalism as such provides no guarantee against transparency failures. The currently most popular versions of internalism are formulated in terms of the notion of a *non-factive mental state*. Knowing p is not a non-factive mental state, but believing p is. Our own beliefs, however, aren't transparent to us.

Empirical evidence confirms that we are very far from even having positive access to our own beliefs.³⁴ The non-transparency of our beliefs creates trouble for any view on which a proposition p is part of one's perspective *only if* one believes p . And the non-transparency of our mental states more generally creates trouble for views on which our perspectives consist of our mental states.

I am not the first person to discuss a revenge-type worry for perspectivist views arising from the possibility of limited epistemic access to our perspectives.³⁵ But proposed solutions normally assume some form of perspectivism. I take perspectivists to have three main lines of response. The first is to accept what I have said, and to concede that it shows that we need a plethora of different subjective oughts of different orders, or of different kinds of subjective evaluations.³⁶ The second is to deny that Miner-type situations could arise for perspectivist norms. This strategy could take one of two forms. The first denies that perspectives can be incomplete or misleading when it comes to facts about perspectives, or facts about what particular perspectivist norms recommend in one's situation. The second denies the intuition regarding *Another Mining Disaster*, or discounts its significance. The third is to try to defend some theory of subjective oughts that takes into account what perspective a perspective takes on itself, on various mathematical facts, facts about evidential support, etc. According to such a view, for instance, it may be that in the original *Miners* case one subjectively ought to maximize expected value, but that in iterated cases such as *Another Mining Disaster*, one ought instead to do something else entirely (one might suggest it is to maximize expected expected value, but unfortunately that would be of no help³⁷). Such views might either appeal to a single rule that takes into account all of these facts about one's perspectives, or to some sort of hierarchy of rules.³⁸

I want to motivate a completely different diagnosis of the original problem and correspondingly, a different account of the kinds of "subjective" evaluations needed to solve it. I take what has been said above to be enough to show that perspectivist views of subjective oughts face a serious challenge – ultimately, I think perspectivism makes the problem of luck unsolvable.

Here is where we are. In some situations one can only conform to objectivist norms like *Choose the best!* by luck, accident, or fluke. These are paradigmatically situations like *Miners* in which one has limited access to relevant facts about the world. According to the access diagnosis, the fundamental source of the problem is that we sometimes lack epistemic access to the facts that the recommendations of norms like *Choose the best!* depend on. Given this diagnosis, however, we should not expect perspectivist norms to solve the problem, for we sometimes have limited access to our perspectives. In fact, if the access diagnosis is right, and there are very few facts we can

³⁴ See Srinivasan (2015) for references. H. M. Smith (2010) also discusses problems for theories of subjective oughts (or "subjective rightness") raised by the fact that we don't always have access to our beliefs.

³⁵ For a seminal discussion of luminosity, see Williamson (2000); see also Srinivasan's (2015) on what she calls Anti-Cartesianism. For discussions of problems in this ballpark for theories of subjective oughts in particular, see H. M. Smith (1988, 2010), Sepielli (2014), Kagan (2018), and Spencer and Wells (2019).

³⁶ Though I haven't discussed normative uncertainty, Sepielli (2014) argues that there are incommensurable 'orders' of rationality that might disagree about what one ought to do when faced with different kinds of uncertainty.

³⁷ Expectations, as well as any higher-order expectations, depend on a relevant probability function, be it evidential probabilities or subjective credences. It was not assumed that in *Another Mining Disaster* one lacks access this probability function itself; rather, it's the relevant mathematical facts that one has forgotten. In fact, the relevant probability function Pr might be perfectly luminous to itself: for any proposition p and value r , if $\text{Pr}(p) = r$, then $\text{Pr}(\text{Pr}(p) = r) = 1$. If this is so, higher-order expectations just equal lower-order ones. If the expected value of action A is highest, then so is the expected expected value of A. That's why appeal to expected expected values won't give us the result that one ought to choose C. Thanks to Tim Williamson for discussion.

³⁸ See Lasonen-Aarnio (2014), for a critical discussion of so-called Über-rules, Kagan (2018, footnote 5) for a brief discussion of a "multilayered subjectivised principle" and a worry that no such principle would solve the problem, H. M. Smith's (1988) critical discussion of "rules of thumb", and Spencer and Wells (2019) and Spencer (*manuscript*) for an attempt to work out how such a view would work. Smith (2010) discusses problems raised by cases in which a subject lacks access to her own beliefs. She sees these as problematic because of reasons having to do with guidance. As a remedy, she (2010, 1988) suggests a hierarchy of principles, and the idea that what is subjectively right is given by the most highly ranked principle a subject is able to use.

invariably access, the situation might look desperate, and our original problem unsolvable! But perhaps the situation is not that desperate after all.

Before turning to dispositionalism, I want to offer a promissory note concerning a perspectivist objection to the arguments offered in this section. Someone might worry that what I have shown is that we should not construe what it is for something to be “present to one’s mind” – and hence, part of one’s perspective – in terms of epistemic access. Rather, we should think of one’s perspective as consisting of some class of (non-factive) mental states, states such as experiences or seemings. (Indeed, my characterization of perspectivism did not essentially involve epistemic access.) The applicability of perspectivist norms depends on which mental states one is in, and non-lucky conformity to those norms only requires responding in an appropriate way to being in those mental states. One need not engage in any kind of explicit deliberation involving beliefs about what mental states one is in, and about what the correct norms recommend in one’s current situation – one need not take an epistemic perspective on one’s mental states.³⁹

In spelling out the revenge problem I have so far gone along with the access diagnosis, arguing that if it is right, we should expect perspectivism to face the very same problem. But as I indicated at the outset, I don’t think the access diagnosis is correct. At this point it is time to kick away the ladder. The ultimate problem, I will argue, has to do with limits on *feasible alternative dispositions*. With this diagnosis on the table, the problem for perspectivism will emerge in its full generality.

3. Dispositional evaluations

I will here sketch the rough contours of my big evaluative picture, showing how it applies to *Miners*.⁴⁰ On the view I will propose the more subjective dimension of evaluation we seem to need, across different normative domains, is focused on the goodness or success-conduciveness of the dispositions manifested by agents. Positive evaluability is a matter of manifesting good dispositions, dispositions that compare favourably with the best *feasible* ones. As I see things, this captures an important sense in which an agent might be doing the best she can, given her situation.

We need to evaluate doxastic states, choices, and actions in a manner that is sensitive to the way in which those doxastic states are formed and retained, the choices are made, and the actions performed. The epistemology literature is replete with proposals for how to understand evaluations sensitive to ways of forming beliefs: a belief is formed in a good way if it is properly based on sufficiently good, undefeated *reasons* or *evidence*, if it is the output of a *reliable process*, or if it is formed by a *reliable method*. In this connection epistemologists often talk of *doxastically* (as opposed to merely *propositionally*) justified beliefs. But we need to be able to make similar kinds of evaluations in the practical realm: one must choose or act in the right way for one’s choices and actions to be *ex post*, and not merely *ex ante*, justified.

My proposal is to let dispositions that manifest themselves as one’s ϕ ’ing (coming to believe something, retaining a belief, making a choice, performing an action) do the work of identifying these ways.⁴¹ I won’t here take on questions regarding the metaphysics of dispositions. I will simply assume that there are dispositions of a wide variety of different kinds, and that inanimate objects and rational subjects alike are constantly manifesting them. I will assume that we can draw a distinction between abundant and sparse dispositions at every level of reality, not just the microphysical level.⁴² For instance, the kinds of dispositions of interest when evaluating beliefs and choices will be at the

³⁹ Many thanks to David Christensen for pressing me on this.

⁴⁰ For more details, see Lasonen-Aarnio (2020) and a book still under progress.

⁴¹ Note though that the basic structure I will outline is compatible with different accounts of ways: in principle, the theoretical role played by dispositions in my account could be played by something else, such as *methods*, *rules* or *strategies*, assuming one could give a satisfactory account of their feasibility (suffice it to say I am skeptical). Thanks to Bart Streumer for discussion.

⁴² See Schaffer (2004).

psychological level, broadly understood. My starting point will be that when an action or mental act (such as a doxastic transition) is attributable to an agent, then even if it is atypical for her in the sense that it may not reflect her character, it is the manifestation of some of her dispositions.⁴³ If, for instance, my mind comes to be controlled by some external force causing me to be in a neural state that makes for believing p , then coming to believe p is not a manifestation of my dispositions, not something properly attributable to me.

On my picture the goodness of a disposition – and hence, of one’s ways of believing, choosing, and acting – depends on the *values* of its manifestations across a range of relevant counterfactual cases. Hence, dispositional evaluations are always relative to a standard of success and a related value function. For instance, whether believing that one’s lottery ticket will lose just based on the odds is positively evaluable depends on whether epistemic success is a matter of believing truly or of knowing. Moral choices will be evaluated very differently depending on whether acting successfully is a matter of doing things with good consequences, of conforming to deontological rules, or of something else entirely. Hence, the basic framework can be implemented in very different ways. Moreover, in some cases all that matters is whether a manifestation is a success or failure: coming close to succeeding doesn’t get assigned any extra value (e.g. penalty kicks in football). In other cases success itself comes in degrees. In cases like *Miners*, it is plausible that the value assigned to choice is a function of the number of lives saved.

The goodness of a disposition, then, depends on the values of its manifestations across a relevant portion of modal space – or, in so far as relevance comes in degrees, across cases weighted according to relevance.⁴⁴ How should we think of such relevance? Dispositional evaluations are malleable and context-dependent: what counts as relevant can easily shift depending on our focus and context as evaluators. Nevertheless, there are general structural points to be made. Relevance should not be understood in terms of a relation that any situation trivially bears to itself (to a maximal degree). Hence, relevance is not a matter of *relevant similarity*. The situation under evaluation might be deviant or abnormal and hence, it – and ones very much like it – may have little or no relevance when evaluating a subject’s dispositions.⁴⁵ Indeed, one potentially promising account of relevance makes use of the notion of *normalcy*. Even lacking a precise account, we have some initial grasp on what sorts of cases count as normal.⁴⁶ Various circumstances involving highly misleading evidence, for instance, are abnormal: it is abnormal to encounter a real-looking fake tree among real ones, or to experience an intricately crafted perceptual illusion with no hints that one is being deceived.⁴⁷ The dispositions manifested by a subject who comes to form a belief on the basis of such an illusion might be good, for across somewhat normal cases they manifest as knowledge-constituting (or otherwise epistemically successful) beliefs.

Here is a toy model that at least comes close to how I think about relevance. As evaluators, we consider some features of a case to be idiosyncratic, while holding others fixed. Consider, for instance, *Miners*. Some features, largely given by the description of the case, we hold fixed: that there are several miners trapped in one of two shafts, that floodwaters are rising, that all the miners in the blocked shaft will die and all in the flooded shaft will be saved, that if both shafts are flooded one miner will die in any case, that the subject must make a choice without delay, that they cannot

⁴³ This should not be conflated with the false claim that actions attributable to an agent always arise out of *habitus*, or that we only ever do things we are generally disposed to do. For more discussion, see Lasonen-Aarnio (2021).

⁴⁴ More precisely: the goodness of a disposition is given by a *score*, which is fixed by the weighting of situations according to relevance, and a value function. The value function assigns values (which may be positive or negative) to manifestations of the profile. I am assuming the score of a disposition to be a *weighted average* of the values of its manifestations in relevant counterfactual cases, the weightings being by relevance. I think this is a natural view, though it is not the only possible one.

⁴⁵ Cf Lasonen-Aarnio (2010), where I talk about dispositions to know and avoid false belief “across a wide range of normal cases”.

⁴⁶ Appeal to normalcy is gaining purchase in epistemology. See M. Smith (2010, 2016), Goodman (2013), Goodman and Salow (2018).

⁴⁷ See Goodman (2013) for a discussion of normalcy in connection with perceptual experience.

communicate with the miners, etc. Other features we consider idiosyncrasies of the case at hand: in particular, the precise location of the miners, as well as irrelevant features such as the colour of the sandbags. We consider, then, a portion of modal space in which the features we are holding fixed obtain – we could think of these features as fixing a contextually determined *type* of situation – while letting other features vary in somewhat normal ways.

Hence, how good a disposition is depends on the values of its manifestations across a relevant portion of modal space.⁴⁸ Ultimately, I will be interested in the comparative goodness of various dispositions, on how successful the manifestations of a particular disposition are compared with the manifestations of *feasible alternatives*. According to the view I see as having the most promise of capturing the more subjective kind of evaluation we are after, dispositional goodness is a matter of manifesting dispositions that are sufficiently good compared to the best ***feasible alternatives***.⁴⁹

There are various constraints, of different strengths, on what dispositions we can have and manifest. Some are logical, some nomological (arising, for instance, from the nature of causation), some metaphysical, and some flow from features of our cognitive design. Though there is nothing metaphysically impossible about such a disposition, we cannot be disposed to only use safety belts on car rides that end in crashes, or to buy insurance only for trips that end up involving some sort of calamity. Climbers cannot be disposed to only check their knots when they have mis-tied them. Or, consider the fact that our total evidence is sometimes misleading: we sometimes have overwhelming evidence in favour of a false proposition *p*. We couldn't be disposed to only take into account total evidence making a proposition *p* very likely when that evidence is not misleading regarding whether *p*. In general, we cannot be disposed to be sensitive to our perspectives only when those perspectives are neither incomplete nor misleading regarding a given subject matter.

Given the context-sensitivity of dispositional evaluations, feasibility in different contexts might involve different senses of “could”. As a general rule, feasibility is constrained by our broad cognitive architecture, allowing for some idealizing away from distorting influences.⁵⁰ If I reason poorly because of fatigue, then I am under the influence of a distorting condition: I may be excused, but I am not manifesting good (enough) dispositions. The same is true when I am under the influence of a bias or prejudice, or when I have been drugged. Human cognitive architecture leaves open a very wide range of more specific capacities, abilities, and dispositions, and here context – and ultimately, perhaps, our expectations as evaluators – will again play an important role.

A *feasible alternative* disposition is one that would manifest itself as a relevant doxastic state, choice, or action in one's situation. For instance, relevant alternative dispositions in *Miners* must manifest in the situation at hand as making a choice among the available options. Assuming dispositions to have stimulus conditions, a disposition can only manifest if its stimulus conditions obtain: hence, it must be feasible not only to *have* the relevant disposition, but also to *manifest* it.

⁴⁸ Though I will here just talk about feasible alternative dispositions, my official view compares feasible sets of dispositions constituting *dispositional profiles*. The kind of positive evaluability I am interested in is a matter of manifesting the best dispositions, *given the type of situation one is in*. But the actual disposition one manifests might not even be available to be manifested across some other instances of that type, for the simple reason that its stimulus conditions don't obtain. A dispositional profile is a way of being disposed as far as certain types of situations are concerned; we can model it as a function from instances of a type of situation to dispositions manifested in those instances. In order to manifest a profile, a subject must have the profile by having the dispositions constituting the profile, and they must manifest the right disposition, the one that the function described outputs in the subject's situation. Below I say more about feasibility, but note for now that for a profile to be feasible, it is not enough that each of the dispositions involved are feasible on their own; it must be feasible to have all of them.

⁴⁹ If infinitely many dispositions are feasible, then there may be no upper bound to their goodness. I am setting such cases aside; if they do arise, we need some threshold for what counts as good enough. Also, for the purposes of stating my view, I will assume for simplicity that in a given situation a single disposition is manifested. It would, however, be more realistic to assume that one's Φ 'ing is the joint manifestation of multiple dispositions.

⁵⁰ Thanks to discussions with Rachel Fraser and Michael Fiorica. Foley (1993: 160), for instance, appeals to a distinction between abilities or capacities, on the one hand, and distorting conditions, on the other, in his discussion of what kinds of idealizations are legitimate, and what kinds are not.

Moreover, its manifestation must be the making of a choice among the available options. It is feasible, for instance, for the subject in *Miners* to manifest a disposition to scratch her head, but this disposition is not an alternative in the relevant sense.

On my view, we often deploy something very much like the dispositional evaluative perspective.⁵¹ At this point the reader might worry this cannot be so, for we have little or no access to the precise dispositions manifested by others. However, considerations of feasibility often provide us with an important shortcut: we often have a good grasp of the kinds of dispositions an agent could feasibly manifest in a given situation. If none of the best (or close enough) feasible dispositions would manifest as choosing (believing, acting) as the subject does, this is all we need to know in order to correctly negatively evaluate her choice (belief, action).

Consider again a human subject in *Miners* who chooses to block shaft 1. This is objectively the best course of action, for all of the miners are saved. But assuming that we are dealing with a human subject, could such a choice be a manifestation of good dispositions? We know that the subject doesn't know where the miners are, and her evidence favours neither the hypothesis that they are in shaft 1 nor the hypothesis that they are in shaft 2. It is not feasible in the situation described to manifest a disposition that is sensitive to the location of the miners, one that would manifest as different choices depending on their location. We could also talk about a non-epistemic sense of *discrimination* here: there are no feasible dispositions the subject could manifest that discriminate between cases in which the miners are in one shaft and those in which they are in the other shaft. Whatever dispositions a human subject manifests by choosing to block shaft 1, those dispositions would indiscriminately manifest across a range of relevant cases as tragically choosing to let the shaft the miners are all in flood.

As we saw, the goodness of a disposition is always relative to some value function. In *Miners*, the objective value of a choice is plausibly a linear function of the number of human lives saved (indeed, this is also assumed by perspectivist views on which one should let both shafts flood halfway because this maximizes expected value). The best feasible dispositions are ones that score the best, across relevant counterfactual cases, given such a value function. I have proposed thinking of the relevant cases as instances of a given type of situation, weighted by how normal they are. The type here is largely given to us in the description of the case. And it is very natural to assume that there is nothing abnormal about the miners being trapped in one shaft rather than the other.

To sum up: the problem with choosing to block shaft 1 is that any feasible way of choosing – any feasible alternative disposition – that manifests as fully blocking one shaft will indiscriminately manifest across a wide range of relevant cases as blocking the shaft containing all 10 miners. By contrast, it is feasible to make a choice that saves 9 lives across all cases. Uniformly saving 9 lives is better than haphazardly saving all 10 in roughly half of the relevant cases, but taking action that results in all 10 miners being killed in the other half. Choosing to fully block one shaft cannot be a manifestation of a good enough disposition: the best dispositions manifest as choosing to let both shafts flood halfway.

The dispositional view often mimics perspectivism because our perspectives constrain what dispositions it is feasible for us to manifest, and the best feasible dispositions are often sensitive to our perspectives – to how we, or our evidence, represent the world as being. Indeed, in paradigm cases (and certainly ones epistemologists have mainly focused on) our only means of tracking how things are in the world is by tracking facts about our perspectives. Nevertheless, it should be clear by now that the underlying normative stories are different. Below I discuss cases in which they yield different verdicts.

⁵¹ I have argued that it explains a large class of “data” that has been used to support internalism in epistemology, like intuitions about victims of deceit. For how my view solves the New Evil Demon Problem, see Lasonen-Aarnio (*forthcoming A*). For defeat, see Lasonen-Aarnio (2010, 2021).

I have offered a solution to one of the problems we started out with, which is providing a fuller evaluative story of *Miners*. But to fully state my alternative diagnosis of why norms like *Choose the best!* are inadequate – in particular, why in *Miners* you can only choose the best course of action by luck – we need a better understanding of what such luck is in the first place.

4. Luck and attributability

I have argued that a very wide range of norms – including perspectivist ones – can sometimes only be conformed to in a way that we want to describe as merely “accidental”, “lucky”, a “coincidence”, or a “fluke”. But how should we think of such luck?

I will assume that whatever we say here should fall out of a general account of what it takes for a normative success – or any success of an agent, for that matter – to not be merely accidental or lucky in the relevant sense. I take this to rule out a view, often gestured at in the literature on moral luck, that cashes out luck in terms of control, notions like intentionality sometimes being brought in. Setting aside other problems, it is difficult to see how such control-based accounts could apply to belief. It is commonly assumed that the main alternative is a modal account. An idea familiar in the epistemology literature is that for a success to not be lucky, failure must be avoided in every case (or most cases) within a certain portion of modal space. The best-known view takes the relevant portion of modal space to consist of cases that are modally close to the case under evaluation, modal closeness being understood in terms of a similarity relation.⁵²

I will also assume that the kind of luck we are interested in is intricately connected with whether a success is *attributable* or *creditable* to an agent. Consider a 5-year old attempting to bake a cake. If the child’s father corrects everything she does, taking ingredients out of the bowl and adding new ones in, then there is perfectly good sense in which it is no fluke or accident that the resulting cake is edible. Yet, the success in question is not suitably connected with the actions of the child: it is not suitably attributable to her. This is one reason to be skeptical of the simple modal account: if the child’s father is very diligent, then we should wonder why the cake isn’t edible in all modally close cases, even holding fixed the child’s haphazard method. A second, related problem has been persuasively spelled out by Ernest Sosa over the years: the fact that something could very easily have interfered with one’s success does not make the success any less creditable to an agent. For instance, the fact that a gust of wind could very easily have blown the arrow off its course does not make the archer’s success less creditable to her.⁵³ These points are two sides of the same coin: the modal robustness of a success can be gained or lost in ways that have little to do with attributability.

However, a control account and simple modal account are not our only options. I will outline the rough contours of an alternative way of thinking about the kind of non-lucky (normative) success at issue.

4.1 Flukes and explanations

As a starting point, consider the contrast between being a fluke or coincidence and having an explanation.⁵⁴ For instance: it’s no fluke that very few people in a given district fell ill with Covid, for a very large proportion were vaccinated, and being vaccinated explains not contracting the virus.

⁵² Pritchard (2005, 2006), Lackey (2008), and Coffman (2015) all assume the relevant cases to be modally close. See Pritchard (2005, 2006) for a defence of a modal account. Lackey (2008) also contrasts the modal and control accounts, arguing against both, but she only considers the kind of modal account Pritchard defends. See also Coffman (2015), whose main reason for rejecting the modal account is the kind of case (*Buried Treasure*) given by Lackey against the modal account.

⁵³ See e.g. Sosa (2007).

⁵⁴ I will assume, in line with much of the recent literature on explanation, that there is a continuity between scientific and non-scientific explanations (e.g. Woodward 2003: 19). For an account of coincidence in mathematics (for instance, is it a coincidence that two equations have exactly the same solutions?) appealing to explanation, see Lange (2010)). Faraci (2019) draws on Lange’s account to develop an explanationist account of knowledge.

The rough idea will be that an agent's successful ϕ 'ing is not lucky or a mere fluke in the relevant sense connected with attributability when the success *depends* in a specific way on her way of ϕ 'ing. I suggest cashing out the relevant kind of dependence here in terms of the notion of an explanatory generalization: there is a sufficiently modally robust, invariant connection between the dispositions and the success, a connection that explains the success on this occasion. The dependence itself is modal, but such an explanationist account is distinct from the simple modal view.

I will draw on what is perhaps the most popular approach among philosophers of science, the contrastive-counterfactual theory of explanation defended by Woodward (2000, 2003). He intends it to be a theory of causal explanation (causal in a rather broad sense), but the basic approach has been extended apply to non-causal explanations.⁵⁵ Begin with the rather common thought that good explanations make reference to explanatory *generalizations*, generalizations that are not merely true by fluke or accident. We can think of a generalization as stating connections between *variables*, the aim being to uncover systematic and objective relations of dependence between these. Woodward's thought is that explanatory generalizations need not be lawful or exceptionless, but they must be sufficiently *invariant*. Invariance involves a kind of modal robustness: the relations of dependence at issue must hold across a range of counterfactual cases corresponding to different *interventions* and changes in background conditions.⁵⁶

An intervention is a causal processes changing the value of some variable *X* under interest in an appropriately exogenous way.⁵⁷ Consider, for instance, a generalization connecting being vaccinated and exposed to a virus with not falling ill. Let *V* represent whether a subject receives the vaccine, *E* represent whether the subject is exposed to the virus, and *C* represent whether the subject contracts the virus. All of these variables take binary values 0 or 1. An intervention would involve changing the value of *V* from 0 to 1, or from 1 to 0. The intuitive idea here is that for the connection between being vaccinated and not falling ill to be invariant, it must hold over a range of interventions and changes in background conditions (for instance, different beliefs held by a subjects concerning the effectiveness of the vaccine, other medications she is taking, lifestyle choices, etc.). In the case of indeterministic or probabilistic causes, what is needed is a invariant connection between one variable and the probability distribution involving another.

When does a generalization successfully explain a singular event or fact? Generalizations track patterns of dependence between variables. Whether a generalization can explain a singular event or fact depends on the structure of the system of actual relationships of dependence at issue. Consider, for instance, a generalization linking treatment with penicillin and recovery from strep throat. The generalization holds invariantly in virtue of an underlying mechanism by which the antibiotic acts in the body. Assume that I take penicillin, and recover from strep. However, for whatever reason the penicillin lacks its usual causal efficacy in my case. Nevertheless, its bad taste causes me to reach for another glass in front of me, a glass that happens to contain a more potent antibiotic that then ends up killing the strep. In the situation just described what is causally relevant about the penicillin is its taste, not its antibiotic powers. Further, the causal route to recovery involves another antibiotic: the penicillin does not act in my body *via* its usual mechanism. To accurately model the system we need to build in a variable for whether or not I take the other antibiotics. In this case, the actual system of dependencies does not instantiate the pattern in virtue

⁵⁵ See e.g. Saatsi & Pexton (2013).

⁵⁶ Woodward (2000) thinks of invariance in terms of families of "active" counterfactuals the antecedents of which describe interventions. He (2000: 236-239) is clear that holding across *nearby* or *close* cases is neither sufficient nor necessary for invariance.

⁵⁷ Consider a generalization claiming a connection between a given value of a variable *X* and a value of a variable *Y*. Exogenous causal processes change the value of *X* in a way that is entirely due to the intervention, and in such a way that "if a change in *Y* occurs, occurs only in virtue of the change in *X* and not as a result of some other set of causal factors" (Woodward 2000: 199-200).

of which the generalization holds, a pattern that involves penicillin acting through its usual mechanism.

A paradigm way of successfully explaining token events or facts is by pointing to a systematic dependence of a certain kind, a dependence that holds in the case at issue, explaining the token event. Our interest here, of course, is in the non-coincidental nature of human success, and it should be immediately clear that not just any explanation will do. Did the 5-year old bake an edible cake merely by fluke? An explanation in terms of the state of the world at a given time and deterministic laws of microphysics is clearly not what we are after. First, we are interested in an explanation at a different level, one having to do with human agency. And second, recall the connection between the specific kind of luck, accidentality, or flukiness we are interested in and attributability. To make sense of such attributability, the successes in question must be suitably connected to the agent: we need an explanation of the successes that makes reference to something pertaining to the 5-year old. My suggestion is that what we need, in particular, is a suitable dependence between the child's way of baking the cake and its edibility.

Talk of ways here is just a placeholder for something playing a given theoretical role. To play this role, these ways themselves must bear a suitable connection to the agent. If, for instance, the child randomly pours ingredients into a bowl, then her way cannot be characterized in terms of the precise quantities she poured in (e.g. "putting in the quantities indicated by a cake recipe"). In my view, the dispositions manifested – the manifestations of which constitute the child's attempt at baking the cake – are suited to play this role. They are dispositions *of the agent*, properties she has. Since they manifest themselves, they are the properties at play on the particular occasion. Their manifestations are or constitute mental events and actions with causal powers. And we can ask how successful the manifestations of these dispositions are in various counterfactual circumstances. One and the same disposition can, in different circumstances, manifest as different choices, actions, or doxastic states. For instance, the dispositions manifested by the 5-year old can on different occasions manifest as pouring in different, rather random quantities of ingredients.

The suggestion, then, is that a given success (of a subject) is not a mere fluke in the sense we are after just in case there is a sufficiently invariant, modally robust generalization tracing a relation of dependence between the dispositions manifested by the subject and the success in question, a generalization that explains the token success on this occasion. In such a case it is no fluke or coincidence that the dispositions in question manifested as successful choices, actions, and beliefs. For a generalization to explain a token success, it must accurately reflect the structure of the actual relationships of dependence at issue. That is, there must be an actual dependence between the subject's way *W* of acting (choosing, believing) and the token success *S* in question, a dependence of the kind in virtue of which the generalization holds. Consider again the 5-year old baking a cake. Why (and how) did she manage to produce an edible cake, rather than no cake at all, or at best an inedible mess? Because of the way in which she went about the cake-baking: she followed a simple recipe, measuring the correct ingredients sufficiently accurately, mixing them, and asking an adult to cover the part involving an oven. Baking a cake in this way generally leads to success, and it did so on this occasion, in the usual manner. The relevant qualities of the resulting cake (its edibility) depended, in the right kind of manner, on the child's way of baking the cake. This is a paradigmatic, good explanation of why a human activity succeeds.

Even when it comes to exemplary attempts at cake-baking we can, of course, easily imagine cases of failure. Various things can interfere with one's success, and in any case some cooperation from the environment is required (for instance, the oven not breaking). But this is also true of a wide range of generalizations deployed in scientific explanations.⁵⁸ Relatedly, on a given occasion a

⁵⁸ Microeconomists deploy as explanatory generalizations the principles of rational choice theory, which fail to hold under a wide class of circumstances. Social scientists standardly deploy generalizations with very limited domains of invariance (limited, for instance, to specific spatio-temporal intervals). Even generalizations deployed in the natural sciences are not fully invariant:

generalization can be explanatory even if it does not model everything that an explanandum depends on, so long as the values of the other variables are being held fixed.⁵⁹ And it is important to bear in mind that it is not only degree, but also kind of invariance that makes for good explanations: certain changes are often more important given a specific subject matter or domain, and it is invariance under those changes that counts.⁶⁰

On the account sketched, the kind of attributable success we are interested in requires the success to depend on the dispositions manifested by the agent, in accordance with a sufficiently invariant generalization.⁶¹ It follows that the menu of feasible alternative dispositions constrains whether a subject is even in a position to succeed in a way that is not merely lucky or accidental. As pointed out above, there are limits, first, to which dispositions a subject can *have* in the first place, and second, to which ones are *alternatives* in the sense that they can *manifest* as an available choice (action, doxastic state) in the situation at hand. One might have some perfectly good dispositions that cannot be manifested because their stimulus conditions don't obtain. For instance, if I have a clear perception as of a cat on the street, I can form a belief by manifesting dispositions that are not available if I have no perceptual experience as of a cat, or if the cat is surrounded by dense fog, its form barely visible. If I know where the miners are I can make a choice by manifesting dispositions that are simply not alternatives if I am uncertain about their whereabouts. In *Miners* there is no feasible way of choosing that successfully conforming to the norm *Choose the best!* invariantly depends on. More generally, non-lucky normative success requires the availability of ways or dispositions that track what one ought (according to the norm) to do at least across some portion of modal space – otherwise success could not depend on the dispositions one manifests *via* a robustly true generalization.

Return to the thought that the norm *Choose the best!* provides us with inadequate evaluative resources, for sometimes (perhaps often) one can only conform to it by luck. On the account I have proposed, the reason why the problem of luck arises is that in some situations there is no feasible way of making a choice that success robustly depends on – hence, one is in no position to choose in a way that explains the success. Such a connection between a way *W* of choosing and making the best choice would have to hold in a sufficiently modally robust way. Across cases in which one chooses in way *W* that involve different background conditions (for instance, the miners being located in different shafts), one still chooses the best course of action. And across cases involving an intervention on one's way of choosing – cases in which one doesn't choose in way *W* but in some other way, or fails to make any choice at all – one fails to make the best choice. In *Miners* there is no feasible way of choosing that is connected with making the best choice in a sufficiently invariant manner: this is true of guessing, going by hunches, lucky numbers, or maximizing of expected value. By contrast, consider a case in which you know where the miners are, for you are able to communicate with them using a phone. You make a choice regarding which shaft to block based on this knowledge. You can then manifest dispositions with the right explanatory connection to success. This is so even if the dispositions in question are taken to be intrinsic⁶², and their stimulus conditions are assumed to be internal states, states such as having a normal-seeming experience as of talking with one of the miners. Manifesting such dispositions is, after all, connected with successfully saving all of the miners in a relatively invariant way. Whether such a generalization can explain your token success depends on the actual relations of dependence at issue (e.g. between

for instance, the field equations of general relativity break down when quantum gravitational effects become important (see Woodward 2000, 2003).

⁵⁹ See Woodward's (2000: 216-217) discussion of explaining the speed of a car by appeal to the relationship between its speed and the depression of the gas pedal.

⁶⁰ Cf. Woodward (2000: 219-220).

⁶¹ But recall that this is not sufficient: the explanatory generalization must also explain the token success. A generalization can only explain if the actual relations of dependence instantiate the pattern of dependence in virtue of which it holds.

⁶² Indeed, everything I have said is compatible with the orthodox view that dispositions are intrinsic properties of their bearers.

your mental states and the location of the miners). But here the fact that you actually *know* where the miners are is relevant: it is precisely because the usual relations of dependence hold that the generalization in question can explain your token success.

At this point I want to take a step back and return to an objection to my discussion (in §2) of the revenge problem faced by perspectivists.

4.2 Back to Perspectivism: Access and mental states

The discussion in §2 largely assumed that we should think of perspectives in terms of propositions one has epistemic access to. But perhaps the problem disappears if we think of perspectives in terms of mental states (such as experiences or seemings).

First, note that even those who think of perspectives in terms of mental states need some way of getting from states such as experiences or seemings to choices. A dominant perspectivist norm is to maximize expected value. But given any story of how to get from mental states such as experiences to the required probabilities view, cases like *Another Mining Disaster* will prove to be a problem: sometimes it will be a matter of luck that a subject manages to maximize expected value, for there is no feasible disposition they could manifest that is robustly connected with successfully maximizing expected value.

More generally, I have argued that the problem or luck arises from limits on what dispositions are feasible alternatives in a given situation. And as far as I can see, *there are no necessary entailments of the sort the perspectivist would need between perspectives and feasible dispositions*. In particular, mental differences between cases can outstrip our abilities of dispositional discrimination – mental differences do not always pattern in the right way with the feasibility of dispositions.⁶³ Consider, for instance, a series of cases over which one gradually transitions from one experiential state to another. Perhaps, for instance, in the first case c_1 one has a clear perception as of a cat some distance away on the road, and in c_n one has no such experience, but an experience as of a road covered with fog. As we move along the series of cases, the fog gets denser, so that by the end the cat is no longer in sight. At some point in the series, one no longer has an experience as of a cat. As an example, consider a candidate perspectivist norm on which one ought to believe there is a cat just in case one has an experience as of a cat. Is there any reason to think that, in a situation in which one *barely* has this experience, it will always be feasible to manifest a disposition robustly connected with holding a belief that reflects one's experience? One of the core points I argued above was that features of one's situation constrain the menu of feasible alternative dispositions: if I know where the miners are I can make a choice by manifesting dispositions that I cannot manifest if I am uncertain about their whereabouts. Similarly, whether a disposition that suitably tracks features of my perceptual experience is available for me to manifest depends on the nature of my experience. In a borderline case of barely having an experience as of a cat, the cat being almost completely immersed in the fog, I very much doubt that it is feasible to manifest a disposition that tracks whether or not one has an experience as of a cat in the first place. I see no reason to think that the contents of our own minds cannot thus outstrip our dispositional discriminative abilities. Indeed, the perspectivist would need to establish that facts about what dispositions are feasible *necessarily* constrain the repertoire of mental states we can be in, or experiences we can have.

For any norm, what the norm tells one to do on a particular occasion will depend on how things stand with respect to some *domain of facts* (or multiple such domains), whether facts about

⁶³ The reader might here be reminded of the speckled hen problem, though note that my points are not restricted to contents of experience. See Chisholm (1942) for the speckled hen problem. It was initially pressed against views on which we have a kind of infallible access to certain mental items, starting with Gilbert Ryle's objection to Ayer's views on sense data. More recently especially Pace (2010, 2017) has pressed the problem for foundationalist views that don't require such access. For more discussions of the problem, see also Sosa (2003), Sosa and Bonjour (2003), Markie (2009), Schellenberg (2016), and Smithies (2019).

what actions have the best consequences, about what is true, about what one knows, or about one's experiences or seemings. Conforming to a norm in a way that is attributable to the agent requires manifesting dispositions that track how things stand with respect to the relevant domain across at least some portion of modal space. The problem is that sometimes it is not feasible to manifest such dispositions, even if the domain in question concerns one's own perspective.

Perspectivism does not close the evaluative lacuna created by the problem of luck. In §5 I will argue that dispositionalism does. But before doing so, I want to bring the different parts of the evaluative framework I have outlined together.

4.3 Success, non-accidental success, and good dispositions

Just as manifesting the best feasible dispositions does not entail succeeding, it does not, of course, entail succeeding in a way that is attributable to the agent. In some unfortunate circumstances, even the best feasible dispositions may fail, and none bear an invariant enough connection to success. For instance, if I face a rescue operation in a storm at sea, the very best thing to do might be to randomly toss a life raft out of the ship, hoping for the best. Similarly, a doctor in a war zone who has access to nothing but one kind of antibiotic might be manifesting the best feasible disposition when she gives this antibiotic to a patient with a fever and an undiagnosed condition. The dispositional perspective correctly evaluates these actions positively.⁶⁴

However, we may also care about whether a given success is appropriately attributable to an agent. On the account I have sketched, this requires manifesting dispositions that are connected with success *via* a sufficiently invariant generalization that explains the token success. Such attributability seems relevant, for instance, for whether, and the degree to which, a subject is to be *praised* for succeeding.⁶⁵

I have, in effect, outlined a framework that distinguishes between three different kinds of evaluations. First, we can evaluate a choice (action, belief) by asking how *successful* it is. It is natural to think of success in terms of value. For instance, human lives are valuable, and in cases such as *Miners* a choice is more successful the more lives it saves. Maximal success is a matter of doing what is best: of saving all 10 lives. Second, we can ask how good the disposition(s) one manifests are: is it the best feasible one given the situation at hand, or at least close enough to the best? The goodness of dispositions here is relative to what we take the relevant success to be. Third, if one does succeed, we can ask whether and how attributable the success is – whether the dispositions one manifested are connected with success in an invariant manner that explains the token success. (Similarly, if one fails, we can ask whether one failed merely by fluke.)

One of the often touted merits of virtue epistemology is its ability to offer a more general axiological framework applicable to beliefs and actions alike. Ernest Sosa famously draws a threefold distinction between accuracy, adroitness, and aptness by means of which all sorts of performances can be evaluated. At first sight the notion of adroitness might look promising in capturing a more subjective kind of evaluation of the sort we have been looking for. Adroit performances manifest competence, where Sosa thinks of competences as dispositions to succeed (to believe what is true, to hit one's targets in archery, etc). A false belief can manifest competence, and so can a well-aimed shot that misses its target because of a gust of wind. However, it is far from clear how the performance-normative framework can be applied to cases like *Miners*. If competences are dispositions to achieve optimum states (with high enough probability), how can a choice that

⁶⁴ And indeed, in our original case *Miners* manifesting the best feasible dispositions is incompatible with the most successful course of action, saving all 10 lives. An analogue in the case of belief is that sometimes the best feasible dispositions will manifest as suspending judgment, but suspending judgment cannot be successful by constituting true belief or knowledge (see Lasonen-Aarnio 2020).

⁶⁵ We should be sensitive to the difference between praising a subject for doing something (e.g. making a given medical decision), and praising her for succeeding (e.g. curing a patient). See Johnson King (2020).

inevitably and knowably results in a sub-optimal outcome manifest competence? Indeed, it is difficult to see how a wide range of real-world actions and choices can be evaluated from the confines of the performance-normative framework. The goodness of many choices and actions is a matter of doing the best one can, not a matter of manifesting a disposition to succeed. My account, which brings in the idea of the best feasible available disposition, seems to offer an even more general evaluative framework.

At this point the reader may be wondering what reason there is to think that dispositionalism escapes the kind of revenge problem I have outlined for perspectivism. If the problem ultimately had to do with access, then given the fact that there is no domain of facts we can invariably access, the problem would indeed be inescapable. But the ultimate problem, on my view, is not one of access. Similarly, if the right account of attributable, non-lucky success was the simple modal one, the problem would again threaten to be inescapable. But I have offered a different, explanationist account of attributable success.

5. Dispositionalism to the rescue

According to the perspectivist diagnosis, the reason why objectivist norms are susceptible to the problem of luck is that they make references to facts outside our perspectives. I have argued that this cannot be right, for perspectivism is susceptible to the same problem. The ultimate problem is not limits of perspectives, but limits of feasibility: limits on what dispositions – ways of choosing, acting, and believing – are feasible alternatives in one's situation. Of the ways of making a choice available to one, none might bear an invariant enough connection to the relevant normative success, whether one involving conformity to an objectivist or perspectivist norm.

But at this point the reader might wonder about the dispositional evaluative norm *Manifest the best feasible disposition!*. If it, too, is susceptible to the problem of luck, then it is not clear whether the recurrence problem has much bite against perspectivism. Indeed, one might think that what I have identified is an unfortunate, pervasive feature of the normative, an evaluative lacuna that cannot be closed. But we shouldn't be too quick to draw such grim conclusions. If one manifests D^* , and D^* is in fact the best feasible disposition, how could it be merely lucky, a mere fluke, that one manifests the best feasible disposition? It cannot, for D^* is in fact the best feasible disposition, and it is not a mere fluke that D^* is the best feasible disposition. Let me address two objections.

First, one might still worry that it might be merely by fluke that one manifests D^* , as opposed to some other, worse disposition. I have argued that when it comes to the kind of luck tied with attributability (or lack thereof), what is at issue is whether the success depends, *via* an explanatory generalization, on the dispositions in fact manifested by the agent. The agent's ability to manifest these dispositions (and not others) may depend on highly contingent features of her situation – perhaps their stimulus conditions could very easily have not obtained – but this does not detract from the attributability of her success.⁶⁶

Second, one might worry that it might, after all, be merely by fluke that D^* is a good enough disposition: even holding fixed the set of feasible alternative dispositions, the worry is that the comparative ordering of these might not be at all modally robust. Here is a line of reasoning that might lead one to this conclusion: “Your account of the relative goodness of dispositions deployed the notion of normalcy: dispositions are assigned scores which determine their relative goodness across a (contextually determined) set of relevant cases, weighted by normalcy. Assume that D_1 and

⁶⁶ Though I don't endorse a view of knowledge as attributable true belief, one might here be reminded of discussions of the kind of epistemic luck incompatible with knowing. For instance, the fact that I could easily have used a different method that would have led me to form a false belief does not make my belief true by luck (see e.g. Pritchard 2005). See also Sosa's (2010: 469) discussion of the fragility of an agent's competence.

D_2 are feasible alternative dispositions, and that D_1 is in fact better. However, the ordering of cases by normalcy depends on contingent features of the world, and is not itself robust.”⁶⁷

The worry is that due to the contingency of normalcy, it might be a mere fluke that by manifesting D^* one manifests the best feasible disposition, for it might be a mere fluke that D^* is the best (or even close enough to the best) of the feasible alternative dispositions. I won't here take a stance on whether normalcy orderings are contingent, for given my account of attributable normative success, the question is irrelevant. Seeing why will help understand the difference between my explanationist account and the simple modal account.

The account I have sketched contrasts being a fluke and systematically depending on something. For instance, it is no fluke that so few people in the district contracted Covid because their not contracting the virus depends, in accordance with an explanatory generalization, on their having been vaccinated. However, such systematic dependence is not tantamount to modal robustness. Indeed, though explanatory generalizations must have a kind of modal robustness, it does not follow that the *explananda* themselves must be modally robust.⁶⁸ An explanation consists of an invariant generalization (linking the of a variable X – or n -tuple of variables – with value of a variables Y), together with a statement of initial conditions (involving the value of variable X). If these initial conditions are not modally robust, then neither is the explanandum. So, for instance, if members of the district could very easily not have been vaccinated, then the low Covid rate among them may be highly modally fragile. But this does not make it a mere fluke.

Consider also the following example. The meanings of our words depend on how they are used. But then, as facts about use shift from one possible world (or time) to another, so do meanings. In fact, many of our words could very easily have had slightly different meanings: perhaps, for instance, the extension of 'table' could very easily have been just slightly different. Hence, facts about the (exact) meanings of words are highly fragile, and not at all modally robust. But does it follow that it is a fluke that our words mean what they do? It is no fluke, for meanings depends on systematic ways on how words are used; they are not created by gods flipping coins. There are systematic, invariant “laws” connecting values of use-variables with values of meaning-variables.

Are facts about the comparative goodness of various dispositions a matter of fluke? Much the same could be said here as in the meaning case: facts about the comparative goodness of various dispositions depend in systematic, albeit context-dependent ways, on features of one's situation and of the world one inhabits. Because, at least given a context, there is a systematic dependence between a range of facts and the relative goodness of various dispositions, it is not a fluke that D^* is a good, success-conducive disposition. And so it is not a fluke that, by manifesting D^* , one manifests the best (or close enough to the best) feasible disposition. Consider again *Miners* and *Another Mining Disaster*, and the relevant objectivist and perspectivist norms. Even if the subject conforms to these norms, it is merely by fluke or accident: any available way of making a choice only haphazardly issues in a norm-conforming choice. By contrast, if D^* is the best feasible disposition, it is not by mere fluke or coincidence that by manifesting D^* one manifests the best feasible disposition.

Before concluding, I want to consider possible cases in which one in fact manifests the best (or close enough) feasible disposition, but one's perspective is misleading regarding the goodness of the disposition manifested. These provide a test case for my view: if I am right, subjects in such cases

⁶⁷ While many discussions assume that facts about comparative normalcy are necessary, several authors have argued that they are contingent (see e.g. M. Smith 2007 and S. Carter 2019).

⁶⁸ Daniel Greco (2016) seems to think that on interventionist views the explananda of good explanations must be modally robust; Pavese (2018) also draws on this idea. Greco's central example (involving a marble that stops in the bottom of a bowl, irrespective of where it is dropped) involves an equilibrium explanation. But not all good explanations are equilibrium explanations (see Kuorikoski 2007).

are not negatively evaluable. Rather than providing trouble for my view, I will argue that what emerges is a further problem for perspectivism.

Consider, then, cases in which a choice I make, or belief I hold, is the manifestation of good dispositions, but my perspective is misleading regarding the matter:

A subject s φ 's, and

(a) s 's φ 'ing is a manifestation of the best (or close enough) feasible dispositions, but

(b) by the lights of s 's perspective, it is likely that her φ 'ing is not the manifestation of the best feasible dispositions.

If one thinks of perspectives in terms of a subject's evidence, then we can cash out (b) as follows: it is likely on s 's evidence that her φ 'ing is not the manifestation of the best feasible dispositions. My view predicts that, since they are in fact manifesting the best feasible dispositions, subjects in such cases shouldn't be evaluated negatively.

On some ways of filling in the background details, one might think that many examples of (putative) knowledge defeat by *higher-order evidence* have the above structure: a subject's belief in p constitutes knowledge, and is in fact the manifestation of the best feasible dispositions, but she subsequently acquires seemingly strong albeit misleading evidence that her belief in p was the output of a flawed cognitive process. The following case follows a somewhat standard template of such putative defeat:

Resident. Rezi is a resident working through some medical cases for practice. After carefully reflecting on a hypothetical patient's symptoms, labs, and other relevant information, she becomes confident, and comes to believe, that the appropriate treatment is a 10mg dose of *Wellstrol* (proposition *Well*). Rezi knows that due to her constantly sleep-deprived state, rarely an isolated cognitive *blip* will occur: an error in her reasoning that results in her arriving at a random conclusion by a perfectly cogent-seeming process. When blips occur, Rezi cannot detect them herself. As it happens, Rezi's performance is being monitored by a team of neuroscientists. The neuroscientists now tell her that a blip occurred: her diagnosis is the output of a process no better than a random guess at tracking the correct treatment. Though the neuroscientists are all but infallible, this time they are mistaken: Rezi's original reasoning was impeccable, and a 10mg dose of *Wellstrol* is appropriate in the case at hand.⁶⁹

Assume, however, that Rezi fails to adjust her belief in any way: she continues to believe *Well*, despite the testimony of the neuroscientists. If conditions (a) and (b) above were satisfied, this would be a problem, for the majority of epistemologists think that Rezi is to be negatively evaluated for continuing to believe *Well* – indeed, many think that she now *ought* to suspend judgment.⁷⁰

However, I have argued in length elsewhere (Lasonen-Aarnio 2010, 2021) that subjects who retain their beliefs in such standard examples of putative defeat *don't* satisfy condition (a). It is not enough that the dispositions manifested when the subject came to originally form the belief were good; in addition, those manifested in its *retention* must be good. And given natural assumptions about feasibility, subjects who retain beliefs despite acquiring seemingly strong evidence about their own cognitive failure are typically not manifesting good dispositions. When subjects retain beliefs despite acquiring seemingly strong evidence that they are manifesting bad dispositions, an important question to ask is whether this retention manifests dispositions that indiscriminately manifest as ignoring such evidence *whether or not it is misleading*. Is it feasible to manifest dispositions that discriminate between cases in which the evidence about one's own cognitive failure is misleading, and relevant cases in which it isn't? If not, then the dispositions manifested are problematic: those

⁶⁹ From Lasonen-Aarnio (2019b).

⁷⁰ For some representative literature, see Christensen (2007, 2010, 2011), Feldman (2005), Horowitz (2014).

very dispositions manifest across a wide range of relevant cases as retaining botched beliefs, beliefs formed through flawed processes. If anything, it is more normal for expert testimony of the sort that Rezi receives *not* to be misleading – the neuroscientists in *Resident* are stipulated to be all but infallible. Suspension of judgment does better across relevant, somewhat normal cases. This is so whether we think of good dispositions as being conducive to true belief or knowledge.⁷¹ Hence, given facts about feasibility, it is not that easy to (continue to) manifest the best feasible dispositions upon acquiring evidence that one is manifesting problematic ones. Persisting in one's belief despite such evidence often manifests a problematic kind of obstinacy.

Let me now consider cases in which both conditions (a) and (b) above are genuinely satisfied.

I will make a bold claim, and offer tentative support for it: cases in which both (a) and (b) are true are what I call *Huck Finn* -type cases. In them a subject cannot help but manifest good dispositions, and is positively evaluable for doing so, despite being trapped in a perspective that misleads them into thinking that those dispositions are bad. Not surprisingly, I think Mark Twain's Huckleberry Finn is a good candidate, at least given a certain interpretation of the text: Huck manifests good dispositions in helping Jim, a fugitive slave, escape, while believing that what he is doing is morally wrong. From his perspective, it looks like he is manifesting bad dispositions, but he is in fact manifesting good ones.⁷² Indeed, the case of Huckleberry Finn is often used in the literature on moral worth as an example of a subject whose actions are praiseworthy, even though he believes himself to be doing something that is morally wrong.⁷³ My account offers a straightforward diagnosis of why Huck Finn is praiseworthy: he manifests good dispositions, dispositions that robustly track what is morally right. It is no fluke or accident that his actions are morally right.

We can describe epistemic analogues. I think it is possible to have misleading evidence about logic. Assume that Logan comes to believe some proposition p based on an impeccable deductive inference. She then acquires seemingly strong but misleading testimonial evidence that the rule she just deployed is not truth-preserving. However, she cannot but follow correct, truth-preserving rules: the right logic is deeply ingrained in the ways she reasons, and so she continues to believe p . From her perspective, it looks like she is manifesting bad dispositions, but she is in fact manifesting good ones. There is an important dispositional difference between Logan and Rezi. It is natural to think, I argued, that if Rezi retains her belief despite the testimony of the neuroscientists, she manifests a problematic, indiscriminate kind of obstinacy. However, there is no reason to think that Logan manifests such obstinacy: there is a particular domain of facts regarding which her evidence is systematically misleading, but her correct reasoning dispositions as it were override this evidence, leading her to reason in ways dissonant with it.⁷⁴

As Amia Srinivasan points out, not all cases that seem to fit the defeat template invite intuitive verdicts of defeat.⁷⁵ Indeed, what I think emerges here is another deep problem for perspectivism, one related to but distinct from the revenge problem discussed above: perspectivism is unable to pull apart cases that are structurally similar if we only focus on facts about a subject's perspective. Consider again the cases of Rezi and Logan. Initially both of their perspectives supported believing a relevant proposition (in Rezi's case *Well*, a proposition concerning the appropriate treatment; in Logan's, a proposition deductively entailed by her evidence). They then acquire misleading

⁷¹ See Lasonen-Aarnio (2021).

⁷² It is somewhat controversial, of course, whether he genuine has *evidence* that what he is doing is morally wrong. Elizabeth Harman (2011: 460-462) claims that false moral beliefs that arise not from ignorance of non-moral facts, but from ignorance of moral facts, are not epistemically justified. I disagree: setting Huck Finn aside, I think our evidence regarding normative claims can be misleading.

⁷³ Discussions of the case go back at least to Bennett (1974). Several authors have appealed to the case in support of their views (e.g. Markovits (2010), Arpaly (2003)). Authors who dissent typically discuss the case as an objection to their views (e.g. Sliwa 2015, Johnson King 2020).

⁷⁴ The fact that the case involves a deductive inference is not essential to the point I want to make.

⁷⁵ See Srinivasan's (2020) discussion of the *Classist College* case. I have also seen Elise Woodard discuss such cases. I am grateful to Elise for conversations we had in Ann Arbor in 2017.

testimonial evidence casting doubt on whether these beliefs were ever appropriate by the lights of their perspectives after all. In order to generate a standard verdict of defeat, perspectivists better be able to say that believing *Well* is no longer rational or justified by the lights of Rezi's perspective. (Telling such a story even in the case of Rezi is, I think, a big challenge that I won't get into here.) But the worry is that any such story will have to apply to Logan as well: if our only normative resources are perspectivist ones, we are forced to say that it is just as bad for Logan to retain belief as it is for Rezi. But in so far as one thinks that Logan can even be praiseworthy for tracking the correct logical relations (just as Huck Finn is praiseworthy for helping Jim), this cannot be right. In some cases pressing on despite having evidence about one's own epistemic or moral failure seems problematic; in others it doesn't. Dispositionalism is well-equipped to make sense of and explain this difference: in some cases it is feasible to manifest good dispositions despite such evidence, in other cases it isn't.

6. Conclusions

I started out with a well-known kind of case, *Miners*, which has been used by many to motivate the need for a more subjective kind of ought governed by perspectivist norms. I introduced the case as a starting point for uncovering why a wide range of norms, on their own, provide us with impoverished evaluative resources. In *Miners* we want to positively evaluate choices that fail to conform to the norm *Choose the best!* and to negatively evaluate ones that do conform. It seems that one can only choose the best (save all miners) by luck or coincidence, by happening to block the shaft all of the miners are in, but choosing to fully block one of the shafts strikes many of us as deeply problematic. Perspectivism arises naturally from a diagnosis of why norms like *Choose the best!* are inadequate: we need norms that only make reference to facts about one's perspective.

I have argued that perspectivism cannot deliver on its promise. The ultimate problem has to do *not* with epistemic access or with what is present to our perspectives, but with limits on what dispositions are feasible alternatives in one's situation. In general, the applicability of a norm depends on how things stand with respect to some domain (or multiple such domains) of facts, and for a very wide range of such domains, sometimes no feasible dispositions track how things stand with respect to that domain across a relevant portion of modal space. While the feasibility of tracking how things stand with respect to a given domain of facts across modal space is constrained in systematic ways by facts about what we have access to – for instance, in *Miners* the fact that you don't have access to where the miners are located explains why you can only choose what is objectively best by luck – it is important to distinguish between what I have called *epistemic* and *dispositional* discrimination. There are no necessary connections, no pre-established harmony of the sort the perspectivist would need, between perspectives, on the one hand, and feasible modal profiles, on the other.

A broader theme that emerges is that a perspectivist focus on issues of epistemic access, or on what is present to an agent's mind, may prevent us from seeing the full range of options available: too often both sides of a dispute have been locked in what is essentially a perspectivist framework. Externalists have pointed out that we sometimes lack access to our perspectives, no matter how deeply buried within our minds they are (indeed, especially if they are so buried). On this basis they have painted what many see as a bleak picture on which we are inevitably hostages to normative fortune: sometimes only a bout of good luck will enable us to believe and act as we ought.⁷⁶ No matter how hard we try, sometimes the most that can be said in our favour is that our normative violations are blameless. Indeed, according to this outlook the conclusion to draw from my discussion of perspectivism is that the problem with seeking more subjectivist norms that it is always possible to conform to in a non-accidental manner is that *there are no such norms*. Perspectivism

⁷⁶ See e.g. Srinivasan (2015).

doesn't solve our original problem, not because its diagnosis of the problem is wrong, but because its diagnosis is right, and the problem is unsolvable: there is no transparent domain of facts. The most we can hope for is to find norms for which the problem of luck is less pervasive.

To some such bleak conclusions are intolerable: acting as we ought in a creditworthy way, not merely by luck or accident, *must* always be possible.⁷⁷ Internalists unable to accept the externalist's outlook have tried to shun access worries, sometimes resorting to internal states stipulated to be immune to anti-luminosity -style arguments.⁷⁸ Recently several authors have tried to evade access worries altogether by resorting to views that deploy *epistemic filtering*, insisting that the normative reasons that determine a relevant range of normative facts about what one ought to do or believe in the first place must be known (or otherwise epistemically accessed). But such views are, of course, just a form of perspectivism, perspectives being now understood in terms of knowledge. And as such, we should expect them to face the kinds of problems I have spelled out above.⁷⁹

All parties to these disputes are focused on issues of access: whether we have access to the conditions that the recommendations of a norm depend on determines whether we are in a position to conform to the norm in a non-accidental manner. But ultimately, I have argued, what matters is the feasibility of certain kinds of dispositional, modal profiles. With this shift in focus there is hope, I have argued, of closing the evaluative lacuna we started out with. In place of evaluative norms focused on a subject's perspective, I have proposed norms focused on feasible ways of choosing, acting, and believing – and ultimately, on feasible dispositions. Here we find evaluative norms that are not too demanding (because we sometimes cannot conform to them in the first place) or alienating (because we can sometimes only conform to them by luck). My view is that *everyone* will need such evaluative resources. To arrive at an adequate evaluative framework, perspectivist norms – whether internalist or externalist; whether ones appealing to knowledge, evidence, experiences, or seemings – will have to be supplanted with dispositional ones.⁸⁰

References

Alston, William

1986 "Internalism and Externalism in Epistemology", *Philosophical Topics* XIV, 179-80.

Arpaly, Nomy

2003 *Unprincipled Virtue: An Inquiry Into Moral Agency*, Oxford University Press.

Bacon, Andrew

⁷⁷ For a recent example, see Lord (2015, 2018: 236-237).

⁷⁸ A recent example of this is appeal to *seemings* as a way to answer the so-called speckled hen problem as it arises for contents of experience. See, for instance, Tucker (2010).

⁷⁹ I have argued elsewhere that epistemic filtering does not solve the problem of luck, for it makes normative facts depend on facts about our knowledge (or more generally, on facts about what passes the epistemic filter) and it is sometimes not feasible to dispositionally track facts about what we know (see Lasonen-Aarnio 2019a, 2020).

⁸⁰ I am very grateful for discussions with David Christensen, Daniel Drucker, Branden Fitelson, Giada Frantantonio, Dmitri Gallow, Jeremy Goodman, John Hawthorne, Jaakko Hirvelä, Sade Hormio, Antti Kauppinen, Jaakko Kuorikoski, Igal Kwart, Max Lewis, Matt McGrath, Lisa Miracchi, Ram Neta, Daniel Nolan, Niall Paterson, Jim Pryor, Samuli Reijula, Keshav Singh, Jack Spencer, Martin Smith, Teru Thomas, Teemu Toppinen, Ralph Wedgwood, Tim Williamson, Elise Woodard, and audiences at the University of Groningen (February 2019), the *Evidence and Knowledge* workshop in Glasgow (March 2019), the weekly *Competence and Success* seminar in Helsinki (April 2019), the *Spring Epistemology Workshop* of my ERC project in Helsinki (April 2019), the *Rutgers Epistemology Conference* (May 2019), the *Southampton-Humboldt Normativity Conference* (June 2019), Research Seminars at Uppsala University and Stockholm University (October 2019), the *23rd Annual Oxford Philosophy Graduate Conference* (November 2019), a graduate seminar at USC (November 2019), a colloquium at UT Austin in February 2021, and a graduate seminar run by Matt McGrath and Ernest Sosa at Rutgers in March 2021. Special thanks to detailed written comments from Jaakko Hirvelä, Nick Hughes, Carolina Flores, Carl Montan and Tim Williamson.

2020 “Inductive Knowledge”, *Nous* 54.2: 354-388.

Beddor, Bob and Pavese, Carlotta

2020 “Modal Virtue Epistemology”, *Philosophy and Phenomenological Research* 101.1:61-79.

Bennett, Jonathan

1974 “The Conscience of Huckleberry Finn”, *Philosophy* 49. 188: 123– 134.

Carter, Sam

2019 “Higher-order ignorance inside the margins”. *Philosophical Studies* 176: 1789–1806.

Chisholm, Roderick M.

1942 “The Problem of the Speckled Hen”, *Mind* 51: 368–73.

Christensen, David

2007 “Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals”

2010 “Higher-Order Evidence”, *Philosophy and Phenomenological Research*, 81(1): 185– 215.

2011 “Disagreement, Question-Begging and Epistemic Self-Criticism”, *Philosophers’ Imprint* 11.6.

Coffman, E.J.

2015 *Luck: Its Nature and Significance for Human Knowledge and Agency*, Palgrave Macmillan.

Conee, Earl & Feldman, Richard

2004 *Evidentialism*, Oxford: Clarendon Press.

Feldman, Richard

1988 “Subjective and Objective Justification in Ethics and Epistemology”, *The Monist* 71.3: 405-419.

2005 “Respecting the evidence”, *Philosophical Perspectives*, 19.1: 95–119.

Feldman, Richard & Conee, Earl

2004 “Evidentialism”, reprinted in E. Conee & R. Feldman *Evidentialism* (Oxford: Clarendon Press) with Afterword, pp. 83-107.

Foley, Richard

1993 *Working Without a Net: A Study of Egocentric Epistemology*, Oxford University Press.

Gibbard, Allan

1990 *Utilitarianism and Coordination*, New York: Garland Publishing.

2005 “Truth and Correct Belief”, *Philosophical Issues*, Vol 15, pp. 338-350.

Gibbons, John

2010 “Things That Make Things Reasonable”, *Philosophy and Phenomenological Research* 81.2: 335-61.

2013 *The Norm of Belief*, Oxford: Oxford University Press.

Goodman, Jeremy

2013 “Inexact Knowledge without Improbable Knowing” *Inquiry*, 2013 56.1: 30–53.

Goodman, Jeremy and Salow, Bernhard

2018 “Taking a Chance on KK”, *Philosophical Studies*, 175.1: 183–196.

Greco, Daniel

2016 “Safety, Explanation, Iteration”, *Philosophical Issues*, 26, pp. 187–208.

Greco, John

2003 “Knowledge as Credit for True Belief”, in DePaul & Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press, Ch 5.

Hassin, R. R., J.S. Uleman, & J.A. Bargh (eds),
2006 *The New Unconscious*, New York: Oxford University Press.

Harman, Elizabeth
2011 “Does Moral Ignorance Exculpate?”, *Ratio*, XXIV.4: 443-468.

Horowitz, Sophie
2014 “Epistemic Akrasia”, *Noûs* 48.4: 718-744.

Hughes, Nick
2018 “Luminosity Failure, Normative Guidance, And The ‘Ought-Implies-Can’ Principle”, *Utilitas* 30.4: 439–457.

Jackson, Frank
1986 “A Probabilistic Approach to Moral Responsibility”.
1991 “Decision-Theoretic Consequentialism and the Nearest and Dearest Objection”, *Ethics*, 101.3: 461-482.

Johnson King, Zoe
2020 “Accidentally Doing the Right Thing”, *Philosophy and Phenomenological Research* 100.1: 186-206.

Kagan, Shelly
2018 “The Paradox of Methods”, *Politics, Philosophy & Economics*, 17.2: 148-168.

Kiesewetter, Benjamin
2011 “‘Ought’ and the Perspective of the Agent”, *Journal of Ethics & Social Philosophy* 5.3: 1-24.
2017 *The Normativity of Rationality*, Oxford: Oxford University Press.

Kolodny, Niko
2005 “Why Be Rational?”, *Mind* 114.455: 509–563.

Kolodny, Niko, and MacFarlane, John
2010 “Ifs and Oughts”, *Journal of Philosophy* 107: 115–143.

Kuorikoski, Jaakko
2007 “Explaining with Equilibria”, in J. Persson and P. Ylikoski (eds.), *Rethinking explanation* (Boston studies in the philosophy of science; Vol. 252), Springer, pp. 149-162.

Lackey, Jennifer
2007 “Why we don’t deserve credit for everything we know,” *Synthese*, 158: 345–361.
2008 “What Luck is Not”, *Australasian Journal of Philosophy* 86.2: 255-267.
2009 “Knowledge and credit,” *Philosophical Studies*, 142: 27–42.

Lasonen-Aarnio, Maria
2010 “Unreasonable Knowledge”, *Philosophical Perspectives* 24(1): 1–21.
2019a “Guidance, Epistemic Filtering, and Non-Accidental Ought-Doing”, *Philosophical Issues* 29.1: 172-183.
2019b “Higher-Order Defeat and Evincibility”, in M. Skipper and A. Steglich-Petersen (eds), *Higher-Order Evidence: New Essays*, Oxford University Press.
2021 “Dispositional Evaluations and Defeat”, in Brown, Jessica and Simion, Mona (eds.), *Reasons, Justification and Defeat*, Oxford University Press, pp. 93-115, 2021.
Forthcoming A “Competent Failure and Victims of Deceit”, in Dutant and Dorsch (ed.), *The New Evil Demon*, Oxford: Oxford University Press.

Lord, Errol

2015 “Acting for the Right Reasons, Abilities, and Obligation”, Shafer-Landau R (ed.) *Oxford Studies in Metaethics*, Vol 10, Oxford University Press, Oxford, pp 26–52.

2018 *The Importance of Being Rational*, Oxford: Oxford University Press.

Markie, Peter

2009 “Classical Foundationalism and Speckled Hens”, *Philosophy and Phenomenological Research* 79.1: 190–206.

Markovits, Julia

2010 “Acting for the Right Reasons”, *Philosophical Review*, 119. 2: 201 – 242.

Moore, G. E.

1912 *Ethics*, New York: Oxford University Press, 1965.

Nagel, Thomas

1979, *Mortal Questions*, New York: Cambridge University Press.

Oddie, Graham, and Menzies, Peter

1992 “An Objectivist’s Guide to Subjective Value”, *Ethics* 102: 512–533.

Olsen, Kristian

2017 “A Defense of the Objective/Subjective Moral Ought Distinction”, *The Journal of Ethics*, 21.4:351–373.

Pace, Michael

2010 “Foundationally Justified Perceptual Beliefs and the Problem of the Speckled Hen”, *Pacific Philosophical Quarterly* 91: 401–441

2017 “Experiences, Seemings, and Perceptual Justification”, *Australasian Journal of Philosophy* 95:2, 226-241.

Parfit, Derek

1988 “What We Together Do”, *Unpublished manuscript*.

2011 *On What Matters*. Vol. 1., Samuel Scheffler (ed.), Oxford: Oxford University Press.

Pavese, Carlotta

2018 “Know How, Action, and Luck”, *Synthese* 198: 1595-1617.

Pollock, John L.

1987 “Epistemic Norms”, *Synthese* 71.1: 61–95.

Pollock, John L. & Cruz, Joseph

1999 *Contemporary Theories of Knowledge*, 2nd edition, Lanham, MD: Rowman & Littlefield.

Pritchard, Duncan

2005 *Epistemic Luck*, Oxford: Clarendon Press.

2006 “Moral and Epistemic Luck”, *Metaphilosophy* 37.1: 1-25.

2012 *Epistemological Disjunctivism*, Oxford: Oxford University Press.

Regan

1980 *Utilitarianism and Cooperation*, Clarendon Press.

Saatsi, J. & Pexton, M.

2013 “Reassessing Woodward's account of explanation: regularities, counterfactuals, and non-causal explanations”, *Philosophy of Science* 80.5: 613-624.

Schaffer, Jonathan

2004 “Two Conceptions of Sparse Properties”, *Pacific Philosophical Quarterly* 85 (2004): 92–102.

Schellenberg, Susanna

- 2016 “Phenomenal Evidence and Factive Evidence”, *Philosophical Studies* 173.4: 875-896.
- Schoenfield, Miriam
2015 “A Dilemma for Calibrationism”, *Philosophy and Phenomenological Research* XCI.2: 425–455.
- Schroeder, Mark
2018 “Getting Perspective on Objective Reasons”, *Ethics* 128.2: 289-319.
- Sepielli, Andrew
2012 “Subjective Normativity and Action Guidance”, *Oxford Studies in Normative Ethics: Volume 2*, 45-??
- Smith, Holly M.
1988 “Making Moral Decisions”, *Noûs*, 22.1: 89-108.
2010 “Subjective Rightness”, *Social Philosophy and Policy* 27.2: 64-110.
2012 “Using Moral Principles to Guide Decisions”, *Philosophical Issues*, Vol. 22, *Action Theory*, pp. 369-86.
- Smith, Martin
2007 “Ceteris paribus conditionals and comparative normalcy”, *Journal of Philosophical Logic* 36(1), pp. 97–121
2010 “What else justification could be”, *Nous*, 44, 10–31.
2016 *Between probability and certainty: What justifies belief*, Oxford: Oxford University Press.
- Smithies, Declan
2019 *The Epistemic Role of Consciousness*, Oxford: Oxford University Press.
- Sosa, Ernest
1991 *Knowledge in Perspective*, Cambridge: Cambridge University Press.
2003 “Privileged Access”, in Q. Smith and A. Jokvic (eds), *Consciousness: New Philosophical Essays*, Oxford: Oxford University Press.
2007 *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. 1, Oxford: Oxford University Press.
2010 “How Competence Matters”, *Philosophical Perspectives* 24: 465-475.
- Sosa, Ernest and Bonjour, Laurence
2003 *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, Oxford: Blackwell Publishers.
- Spencer, Jack, and Wells, Ian
2019 “Why Take Both Boxes?”, *Philosophy and Phenomenological Research* 99.1: 27-48.
- Srinivasan, Amia
2015 “Normativity Without Cartesian Privilege”, *Philosophical Issues* 25-1: 273–299.
2020 “Radical Externalism”, *The Philosophical Review* 129.3: 395–431.
- Sutton, Jonathan
2005 “Stick To What You Know”, *Nous* 39.3: 359–396.
- Tucker, Chris
2010 ‘Why Open-Minded People Should Endorse Dogmatism’, *Philosophical Perspectives* 24, pp. 529-545.
- Way, Jonathan and Whiting, Daniel
2017 ‘Perspectivism and the Argument from Guidance’, *Ethical Theory and Moral Practice* Vol 20.2: 361-374.
- Wedgwood, Ralph
2016 “Objective and Subjective ‘Ought’”, *Deontic Modality*, Nate Charlow and Matthew Chrisman (eds.), Oxford: Oxford University Press, pp. 143-168.
- Woodward, James
2000 “Explanation and Invariance in the Special Sciences”, *Brit. J. Phil. Sci.*, 51.2: 197-254.

2003 *Making things happen: A theory of causal explanation*, Oxford University Press.

Williamson

2000 *Knowledge and its Limits*, Oxford: Oxford University Press.

Zagzebski, Linda

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