

Guidance and Epistemic Filtering¹

Maria Lasonen-Aarnio

This is a pre-published version; please cite the published version.

I. Guidance and Access

Here is a familiar thought: the normative is capable of *guiding* us. If a theory, whether in ethics or epistemology, cannot generate guidance-giving norms, that is a mark against the theory. For instance, classical utilitarianism has been thought by many to fail disastrously on the guidance front.² Epistemologists have likewise invoked guidance, most often to criticize various externalist norms and theories.³

My starting point will be a rather ambitious view on which the normative is capable of guiding us invariably, in every possible situation. Such a thought has been attractive to many who think that guidance is necessary for acting as one ought in a way that is *creditworthy*, and thereby not merely lucky or accidental. We should always be able to act as we ought in a creditworthy way, the thought goes, which in turn requires an ambitious kind of guidance.

Such guidance ambitions encounter a grave problem. For almost everyone agrees that guidance requires some form of *epistemic access* to a domain of facts: in order to be guided by a normative principle, one must have access to whether the conditions specified by the principle apply, and in order to be guided by a normative reason, one must have access to the reason itself.⁴ At first considerations having to do with guidance were seen by many as favouring more subjective theories, theories generating norms that make reference to a domain of facts within our ken.⁵ But more recently sophisticated work in epistemology has questioned whether appeal to the subjective is any solution at all, at least given the ambitious view on which normative principles must always be capable of guiding us. Internalizing or subjectivising the facts grounding normative truths only yields unflinching guidance if we invariably have access to these facts. But we are sometimes in no position to access our beliefs, desires, appearances, seemings, motivations, or the quality of our will – indeed, there is no domain of facts that we can invariably access.⁶ For this reason going internal or subjective does not guarantee guidance.⁷

This is the *Access Problem* for guidance. If there is no domain of facts we can invariably access, then it looks like there simply *are* no principles or norms that we are always in a position to be guided by in the desired way, though there might be a different, weaker kind of guidance that even more externalist or objectivist norms can provide. If the normative is essentially and invariably action-guiding, then we are forced to draw the grim conclusion that there is no such thing. If normativity is to be salvaged, we must, it seems, settle for less and tone down our guidance ambitions.

¹ This article heavily draws on Lasonen-Aarnio (2019), generalizing its main line of argument.

² Jackson (1991) is a classic example.

³ There is a vast amount of papers using guidance to argue for or against various epistemic norms. For a guidance-based argument against externalist norms, see e.g. Pollock (1987). For a guidance-based criticism of a truth norm on belief, see Glüer and Wikfors (2009, 2013). These are just some examples.

⁴ For some representative examples, see Jackson (1991), Raz (2011 :110), and Gibbons (2013, 132).

⁵ For instance, Hudson (1989) explicitly characterizes subjective theories in terms of guidance.

⁶ Williamson (2000, Ch 6) argues that there are no non-trivial *luminous* conditions, conditions such that whenever they obtain, we are in a position to know that they obtain. For a criticism, see Berker (2008). For a reply and clarification of the argument, see Srinivasan (2015b). For what I am calling the Access Problem, see Srinivasan (2015a) and Hughes (2018). See also e.g. Gibbons (2013: 130-131). Worries about epistemic access in connection with guidance have been expressed by numerous moral philosophers. See, for instance, Smith (1988, 2010, 2012). See also Lasonen-Aarnio (*Forthcoming C*).

⁷ See e.g. Srinivasan (2015a) and Hughes (2018).

My focus here will be on views, gaining in popularity, that seem to evade the Access Problem altogether, and that thereby seem to rehabilitate an ambitious kind of guidance. Let us concede that the project of identifying some domain of facts that we always have access to is doomed to fail. But here is a thought: perhaps, instead, we can use epistemic access as a *criterion* on facts that determine normative facts in the first place, thereby requiring that such facts pass an *epistemic filter*.⁸ Those who press the Access Problem can no longer complain that we don't always have access to the facts that ground normative truths, for access is now a requirement on being part of that domain! For instance, the most prominent defences of the Access Problem appeal to *anti-luminosity arguments*, which purport to show that there are no non-trivial *luminous* conditions, conditions such that whenever they obtain, we are in a position to know this. But if the facts that ground normative truths must pass a *knowledge filter*, then an anti-luminosity argument cannot be levelled: the challenge that we are sometimes be in no position to know these facts simply cannot be raised.

Epistemic filtering is normally implemented in a reasons-centered picture of normativity. It is first assumed that normative facts hold in virtue of facts about normative reasons: if I ought to ϕ (in a relevant sense of 'ought'), that is so in virtue of the fact that the overall balance of my reasons supports ϕ 'ing. But there is an epistemic filter on these reasons: a fact can only count as a normative reason for a subject if she has epistemic access to it.⁹ According to a different implementation of filtering, only normative reasons that one *has* or *possesses* determine the relevant normative facts, and there is an epistemic access condition on possession.¹⁰ On either view, the *potent* normative reasons—reasons that enter into determining normative facts—are subject to an access-constraint.¹¹

There are many ways of thinking about epistemic access, and correspondingly different kinds of epistemic filters. What makes for access to a fact? On one view, it is mere true belief. On another, it is justified or rational true belief. Perhaps the most popular (and to my mind plausible) view is that access is a matter of knowing. Indeed, knowledge has been the most commonly invoked epistemic filter: potent normative reasons must be known, or at least one must be in a position to know them.¹² But there are yet more options. Perhaps access requires a kind of apparentness, of something *seeming* to be the case: we have access to those facts that seem to us to be the case.

I will argue that, irrespective of the kind of filter deployed, epistemic filtering does not restore an ambitious kind of guidance. The reason for this is that a central good that guidance has been invoked to deliver is the availability of acting (or believing) as one ought in a way that is not merely lucky or accidental. But epistemic filtering does not guarantee the availability of such non-lucky ought-doing. In short, the problem is that given epistemic filtering, normative facts come to depend on *facts about what one has access to*, or *facts about what passes the epistemic filter*. But irrespective of the filter deployed, we are sometimes in no position to track these access-facts themselves. As a result, sometimes we can only act as we ought by luck.

⁸ This response goes back at least to H. A. Pritchard's discussion (see Pritchard 1932). Dancy (2000: 57) mentions the possibility of an "agent-relative epistemic filter". See also Raz's (2011: 110) discussion of epistemic filters.

⁹ I take Markovits's (2010) view to be that only known facts can constitute right-making reasons. Gibbons (2013) defends a similar view of normative reasons. Similarly, according to Kiesewetter (2017), only propositions that are part of one's evidence can be reasons to begin with, and being known is at least sufficient for being part of one's evidence (e.g. 162, 200). See also Raz (2011, Ch 6).

¹⁰ According to Lord (2018, Ch 3), in order to possess r as a reason (to ϕ), one must know r . This, Lord thinks, is not sufficient, for one might not 'see' the connection between r and ϕ 'ing – I might, for instance, know that the fish contains salmonella but not see that this is a reason not to eat the fish if I falsely believe salmonella to be a harmless bacterium.

¹¹ I borrow the term "potent normative reason" from Littlejohn (2018).

¹² For two recent books defending versions of it, see Kiesewetter (2017) and Lord (2018).

2. Guidance and Luck

Those who work within a reasons-centered picture of normativity tend to think of normative guidance as a matter of *responsiveness* to the normative reasons in virtue of which one ought to do this or that.¹³ This responsiveness, in turn, is a matter of being motivated to act by those reasons. Thus, Way and Whiting (2017: 364) simply express the idea that the normative must be guiding as the familiar thought that “normative reasons can be motivating reasons”.¹⁴ We can always be guided by our reasons, then, just in case the following is true:

Guidance

If one ought to ϕ , then one can ϕ for the normative reasons in virtue of which one ought to ϕ .¹⁵

In what follows, my focus will be on views that endorse *Guidance* together with some form of epistemic filtering.

The discussion below will be focused on a central good that guidance has repeatedly been invoked to deliver. It is often assumed that guidance is necessary and sufficient for acting as one ought in a way that is not a mere accident, fluke, coincidence, or matter of good luck: being guided by a principle is often opposed with mere accidental conformity to it.¹⁶ Indeed, a common complaint regarding objectivist theories and principles is that if these theories are true, and we lack access to the relevant domain of objective facts (as we often do), then sometimes we can only do what we ought in a haphazard, lucky or accidental kind of way. And many have found the thought that our normative lives are thus subject to fortune very difficult to live with.

But why is such non-lucky ought-doing important in the first place? Most parties agree that success that is not due to luck is closely connected with credit or praise. Doing what one ought not just by luck, or conforming to norms in a way that is not lucky or accidental, is necessary for conforming to them in a way that is *praiseworthy* or *creditworthy*. When one’s action has *moral worth*, for instance, one is praiseworthy for doing the right thing.¹⁷ A core motivation for guidance, then, is a need, via the notion of luck, to tie together the *deontic* and *hypological*.¹⁸ The thought is that it

¹³ This view of guidance is in contrast to one on which guidance is a matter of using a normative principle in deliberation about what to do. For discussions of the deliberative view, see Smith (2012) and Hughes (2018); see also Pollock’s (1987: 64) discussion of what he terms the “intellectualist model”.

¹⁴ See also Gibbons (2013: 135), Kiesewetter (2017: 11), Lord (2015, 2018).

¹⁵ There are many ways to express similar ideas. For instance, Lord (2015) says that “the facts that obligate us must be potentially action-guiding”, and that we must at least have the *ability* to act for the reasons determining what one is obligated to do. See Way and Whiting (2017) for a critical discussion of what it means to have the “ability” to act for the right reasons. See also Väyrynen (2006). It should be clear, however, that the ‘can’ here is not just that of metaphysical possibility.

¹⁶ See, for instance, Kant’s famous remark in *Groundwork* that “in the case of what is to be morally good it is not enough that it *conform* with the moral law but it must also be done *for the sake of the law*; without this, that conformity is only very contingent and precarious” (G 4:390). For more recent examples, Wedgwood (2002) says that being guided by a rule is incompatible with its being “purely a *fluke*” that one conforms to it. Smith (2012: 370) remarks that if one is not guided by a principle, one’s conformity to it is purely “coincidental”. Gibbons (2013: 128), discussing objectivist norms, complains that “if we did manage to comply with the objectivist’s norm, that would just be an accident”. Väyrynen (2006) discusses the idea that without guidance, a successful moral life would only be available “by luck or happenstance”. This is just a sample: it is difficult to find any discussion of guidance that does *not* draw on the idea that without guidance, one’s conformity to a principle or norm is merely accidental.

¹⁷ Doing the morally right thing not by mere luck or accident plays a pivotal role in discussions of morally worthy action. Most authors talk about morally praiseworthy action (e.g. Markovits 2010, Arpaly 2003, Sliwa 2015). Johnson King (2020) distinguishes between different types of praiseworthiness, arguing that when performing a morally worthy action, one is praiseworthy *for acting rightly*. I agree with Johnson King that one might be praiseworthy for doing something even if the action is morally right only in an accidental way.

¹⁸ For this terminology, see Zimmerman (2002: 554); see also Srinivasan (2015a). For instance, Lord (2015) thinks that non-accidentality is required for creditworthiness, and that it would be a big cost to admit the possibility of cases in which one cannot act in a creditworthy way.

should always be possible to act as one ought in a *creditworthy* or *praiseworthy* way—and hence, in a way that is not merely lucky. *Guidance* is supposed to guarantee this, in so far as non-lucky, creditworthy ought-doing just is a matter of ϕ 'ing for the potent normative reasons in virtue of which one ought to ϕ .

In what follows, I will assume that an adequacy condition on views of guidance is that being guided in the desired way entails that one acts as one ought in a way that is not merely lucky. I will argue that a responsiveness view of guidance, even when coupled with epistemic filtering, does not guarantee the availability of non-lucky ought-doing: either one cannot always be motivated to act by one's potent normative reasons, or acting for the potent normative reasons in virtue of which one ought to act does not rule out its being lucky or accidental that one acted as one ought.

My argument will draw on the idea that a certain kind of modal robustness or invariance of a success is *necessary* for its not being merely lucky. As I see it, such a modal condition is rather weak, and compatible with various full accounts of luck. Here are the bare bones of the modal condition I have in mind. Assume that a subject ϕ 's—chooses a particular course of action, performs that action, forms a belief, etc. Assume that her ϕ 'ing is a normative success: she chooses, acts, or believes as she ought, as a relevant norm tell her to. In order for her success to not be lucky, her way of ϕ 'ing must issue in normative success in a sufficiently invariant manner across a range of *relevant* counterfactual cases in which she chooses, acts, or forms a belief in that way. (The reader should not assume that such modal robustness is to be understood in terms of modally *close* cases.)¹⁹ ϕ 'ing in a way that is robustly successful in this sense is necessary for one's success to not be merely lucky. Think of the relevant cases, very roughly, as ones in which we allow the idiosyncratic features of the actual case to vary in acceptable, somewhat normal ways. The idea is that if one's success depends on the obtaining of some idiosyncratic feature of one's situation, then it is merely lucky.

The epistemology literature is replete with proposals for how to understand evaluations sensitive to, in particular, 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*. Ultimately, I prefer to think of these ways in terms of the dispositions that manifest as one's ϕ 'ing, though the main points I make below are compatible with alternative views.²⁰ The general proposal, then, is that in order for one's ϕ 'ing to be an instance of non-lucky ought-doing, ϕ 'ing in that way must issue in normative success in a sufficiently invariant way across a relevant range of counterfactual cases. One's way of ϕ 'ing must track what one ought to do across the relevant cases.²¹ Given a dispositional understanding of ways, non-lucky ought-doing requires manifesting dispositions that don't, across a range of normal cases, manifest as acting in ways that one ought not to act.²² If I am in no position to manifest dispositions that *discriminate*

¹⁹ Cases that are modally close, involving events that could easily have occurred, are often relevant. Ultimately, however, I don't think relevance is a matter of any kind of similarity relation (for more details, see Lasonen-Aarnio *forthcoming B*, *forthcoming C*). Hence, I am *not* assuming that praiseworthiness or creditworthiness requires that a success be *safe*.

²⁰ My reasons for adopting the dispositional view have to do with its theoretical utility. Understanding ways in terms of dispositions allows giving a unified account of non-lucky successes—whether ones involving conscious deliberation, automatic actions, doxastic revision, or athletic performances. Elsewhere I argue that it enables making sense of the verdict that victims of massive deceit can be forming their beliefs in ways that are just as good as those of ordinary subjects—thereby meriting a kind of praise—and that subjects who obstinately stick to their beliefs in putative cases of higher-order defeat are criticizable for managing their beliefs in bad ways. See Lasonen-Aarnio (*forthcoming A*, *forthcoming B*).

²¹ Tracking in this sense is not trivial even if I ought to perform the same action, or form the same belief, across all of the relevant cases. Assume, for instance, that we want to evaluate whether my way of forming a belief about the result of adding 126 to 296 tracks the truth of the matter. Even if the answer is the same across all cases, tracking the correct answer is not trivial: randomly guessing or adding up numbers using a defective method don't track the relevant mathematical fact across relevant cases.

²² Note that on my view it does not matter if the fact that one manifests the disposition in question is modally fragile. It might, for instance, be a very local and circumscribed disposition, even one that is only acquired in specific circumstances (see Lasonen-Aarnio *forthcoming B*). Thanks to an anonymous referee for bringing this up.

between relevant cases in which I ought to ϕ , and those in which I ought not to ϕ , then I am only in a position to act as I ought by a kind of luck. This sort of *dispositional discrimination* will play an important role in the discussion below.²³

It is helpful to see how the modal condition diagnoses why one is sometimes in no position to conform to objectivist norms or principles save by luck. Consider a well-known case from Jackson (1991: 463) involving *Dr Jill*, a physician who must choose which of three drugs to prescribe for her patient with a skin condition. Jill knows that drug A is likely to relieve the patient's condition without completely curing it. Further, Jill knows that one of drugs B and C will completely cure the condition, while the other will kill the patient. However, she has no way of telling which is which. Appeal to such examples has become the standard argument against objectivist theories and norms, such as a norm telling one to do what is best. In this case, this norm would require prescribing whichever of B and C cures the condition. But Jill is in no position to prescribe the best drug except by good luck! The modal account provides the following diagnosis of why this is so: no way of choosing between B and C is available to Jill that would track what is best across a range of counterfactual cases. Jill could, for instance, choose randomly, but randomly choosing would lead her to prescribe the wrong drug across a range of cases.

At first sight, it looks like a view endorsing epistemic filtering, filtering of potent normative reasons by some type of epistemic access, is not susceptible to the challenge that one can sometimes only act as one ought by luck. In what follows, I will argue that epistemic filtering does not solve the problem of luck, for it makes normative facts depend on facts about just what propositions we have epistemic access to, *and we are sometimes in no position to track these facts*. I first discuss knowledge filtering. I then generalize my arguments to other kinds of epistemic filters.

3. Knowledge Filtering: A Case Study²⁴

Knowledge is the most commonly invoked epistemic filter: potent normative reasons must be known, or at least one must be in a position to know them.²⁵ Recall that according to the view of guidance under focus, the *responsiveness* view, potent normative reasons must be capable of serving as our motivating reasons. There is an impressive array of arguments for the conclusion that something can be a motivating reason only if it is known.²⁶ Further, guidance is often taken to entail acting as one ought in a way that is not merely accidental or lucky, and several authors have argued that no epistemic standing short of knowledge can guarantee such non-accidentality.²⁷

Consider, then, the following view:

Knowledge Filtering

p is a potent normative reason for s just in case s knows p .²⁸

I will argue that knowledge filtering does not solve the guidance problem, for it does not guarantee the availability of acting as one ought in a way that is not lucky, a central good that guidance has

²³ For more on this notion of discrimination, and how it differs from an *epistemic* kind of discrimination, see Lasonen-Aarnio (*forthcoming B*, *forthcoming C*).

²⁴ This section draws very heavily on Lasonen-Aarnio (2019).

²⁵ For two recent books defending versions of it, see Kiesewetter (2017) and Lord (2018); see also the references in note viii.

²⁶ The view goes back at least to Unger (1975), and has been more recently defended by John Hyman, Timothy Williamson, and Jennifer Hornsby. For references, see Alvarez (2016).

²⁷ See, for instance, Sliwa (2015).

²⁸ Instead of being known, one might think that it suffices that r be in a position to know p . The discussion below applies equally to such views. Lord thinks knowledge is not sufficient (see note ix). However, his reasons don't have to do with the problems I raise below. As a result, as far as I can see, the extra condition he proposes on one's possessing r as a reason to ϕ doesn't solve these problems.

been invoked by many to deliver. Interestingly, the reasons for this derive from some of the very same structural features of knowledge that the Access Problem rests on.

Knowledge Filtering guarantees that ought-making facts are never beyond one's epistemic ken. But one of its immediate consequences is that normative facts now depend on epistemic facts, facts about what one knows, or is in a position to know. Indeed, this is one of the main selling points of views that appeal to epistemic filtering, often defended by considering Jackson-type cases.²⁹ In the original *Dr Jill* case, it is false that Jill ought to prescribe drug B. But now assume that she learns that drug B is in fact the cure, and C the killer. Given her new epistemic position, she ought to prescribe drug B. The only relevant change is her coming to know which drug cures the skin condition. True, the normative reason Jill acquires is *that drug B is the cure*, and not *that I know that drug B is the cure*: we can distinguish between what is in the set of potent normative reasons and criteria for getting there. But there is at least a modal, counterfactual kind of dependence between facts about one's epistemic position and normative facts. This point is worth emphasizing, for one of its consequences, as I will argue, is that sometimes non-lucky ought-doing will require being able to track the relevant facts about one's epistemic position in a modally robust way. And this, given the nature of knowledge, is not something we can always do.

I will now describe two kinds of problem cases, given in Lasonen-Aarnio (2019), in which the relevant subject is in no position to track those facts about what she knows that a relevant normative fact depends on and hence, *assuming knowledge filtering*, in no position to act as she ought in a way that is not lucky. These cases follow two kinds of templates: In the first, the normative fact that one ought to ϕ essentially depends on one's knowing (or being in a position to know) a relevant proposition p , which is a reason to ϕ : Given that this proposition is among one's normative reasons, one ought to ϕ . But all else being equal, if this knowledge condition was not met and hence, if p was not among one's normative reasons, it would be false that one ought to ϕ . The second kind of template is one in which the normative fact that one ought to ϕ essentially depends on one's *not* knowing.

Margins for error and precarious knowing

Assuming that knowledge is subject to a margin-for error principle, knowing can sometimes be precarious in the sense that one could easily have been in no position to know a relevant proposition.

Li works Saturday mornings at a climbing gym, when the place is flooded by young children and their parents taking over the auto belay devices. Li has to frequently estimate the weight of a child just by looking, for the scale at the gym is highly unreliable, and it is important that children too light are not allowed to use the auto belay devices. Li has become very good at estimating how much a child weighs just by looking: he can normally tell by a margin of 1 kg. Now consider the following example.

Auto Belay

Li knows that a child weighing exactly 15 kg or more is heavy enough to be lowered down by one of the auto belay devices. It would be a disaster, however, to let a child lighter than 15 kg climb: if they eventually got up, they would be left dangling, probably screaming, on the wall 16 meters above the ground, and it would be difficult to get them down – not to mention a *huge* PR disaster for the gym. Li sees a certain high-profile parent involved in local politics walk in with their daughter Ada. There is a lot of pressure to let the child climb, for the whole future of the gym depends on a favour from

²⁹ E.g. Kiesewetter (2015, 2017) and Lord (2015, 2018).

the parent. In fact, Ada weighs 16,05 kg. Li comes to know, just by looking, that she weighs over 15 kg.

I will assume a margin for error principle on knowledge. In particular, given the accuracy of Li's estimates, had Ada been just 100 grams lighter, Li would have been in no position to know that she weighs at least 15 kg. As a result, Li *just barely* knows that Ada weighs at least 15 kg. In this sense, Li's knowledge is precarious: he could easily not have known. Moreover, given the absence of a sufficiently reliable scale (or parent), he could easily have been in no position to know.

I will make the following assumptions about the case: *Ada weighs 15 kg or more* is an excellent reason to let her climb. Hence, given *Knowledge Filtering*, as long as Li knows this, he ought to let her climb. Moreover, it is difficult to see why the proposition that Ada weighs at least 15 kg couldn't be his motivating reason for letting her climb. It is, however, essential that this proposition be among his reasons, given that it would be very bad to let a child weighing less than 15 kg climb. Hence, propositions like *it is probable that Ada weighs 15 kg or more*, or *it seems that Ada weighs 15kg or more* just wouldn't cut it. (I revisit this below.) As a consequence, had she been 100 grams lighter, the normative facts would have been different due to Li's being in no position to know that the child weighs at least 15 kg.

Li's knowledge is precarious. He could very easily have not known, and no way of acting or choosing is available to him that discriminates between his case and the case in which Ada is just 100 grams lighter. Moreover, this case is relevant for evaluating whether Li's normative success is lucky: being in no position to track the normative facts given very slight changes to the child's weight means that Li is in no position to act as he ought in a manner that is not subject to luck.³⁰ Similarly, consider a counterfactual case in which the child is a bit lighter, weighting 15,95 kg. In that case Li ought not to let the child climb, since *Ada weighs 15 kg or more* is not among his potent reasons. But again, Li is in no position to act as he ought in a way that is not lucky, since he could easily have known, in which case the normative facts would have been different.

One might, of course, quibble about the details of the case. But if knowing can be precarious in the way assumed, and if a relevant normative fact can essentially depend on one's having such precarious knowledge, then examples with the structure described are bound to arise.

Those who resist anti-luminosity arguments might argue that the kind of precarious knowing I have described does not exist. But at this point it is worth reminding the reader of the dialectic. My focus has been on appeal to epistemic filtering as an answer to the Access Problem for guidance. In a nutshell, the problem is that since there is no domain of luminous facts, guidance ambitions are doomed to fail. Epistemic filtering, it seemed, offered a way to sidestep the problem without contesting anti-luminosity reasoning. I have argued, in effect, that the very structural facts about knowledge that anti-luminosity arguments rely on create a problem of luck even assuming *Knowledge Filtering*.

The second example does not rely on the assumption that knowledge is subject to a margin for error principle.

Unlucky not to know

Again, Jill must decide which of three drugs to prescribe for her patient with a skin condition. Several years ago Jill carefully researched the matter. Jill rates drugs using a scale from 1 to 10, 1 being maximally harmful (killing a patient), 10 being maximally helpful (a complete cure), and 5 being neutral. She remembers rating A as a 7. She remembers assigning either B or C a rating of 10, while assigning the other a 1, but she cannot remember which is which! Thankfully, she wrote the

³⁰ Note that I am *not* claiming that, if Li lets the child climb, it is accidental that the child gets down safely – arguably, a margin for error principle on knowledge guarantee that *this* success is not accidental. But that is not the question: the question is whether, assuming *Knowledge Filtering*, the success of *acting as he ought* is subject to luck.

conclusions of her research in notebook #158, which is in her office. There is just enough time to consult the notebook. Now consider:

Notebook trouble

Jill rushes to her office, and pulls out notebook #158. She opens the page on which she compiled her main findings on drug C, and written in the bottom of the page is the number 1: C is the killer drug, and therefore, Jill concludes, B is the cure! Alas, though this is true, Jill does not know it. About a week ago Jill had to leave her children alone in her office for 5 minutes. Four-year old Fanny, looking for something to apply her new eraser to, happened to pull out notebook #158 from the bookshelf and started randomly erasing numbers. She soon ended up on the page describing drug C, carefully erasing the '1' Jill had written on the bottom of the page. At this point eight-year old Eartha realized what was happening. Eartha asked what had been written on the page. The number 10, Fanny confidently (but falsely) told her. Not wanting to get into trouble, Eartha began carefully writing '10' on the bottom of the page. Very soon after she began, before she had time to write '0', she heard her mother's footsteps in the hallway, quickly closed the notebook, and put it back in its place.

Dr Jill holds a justified, true belief that drug B will cure the skin condition, but lacks knowledge, for she is in a Gettier case: she could very easily have falsely believed that C is the cure, and B the killer, which she would have done had Eartha had time to complete her task. In fact, she may currently not even be in a position to know which drug is the killer. Hence, assuming knowledge filtering, the proposition *drug B will cure the skin condition* is not among Dr Jill's potent normative reasons. Jill may know that it is likely on her evidence that drug B is the cure and that drug C the killer, but assuming the consequences of prescribing the wrong drug to be harmful enough, this does not suffice to make it the case that she ought to prescribe B. Similarly, she might know that it appears to her that drug B is the cure. But a mere appearance is compatible with B being harmful, and does not weight heavily enough in favour of prescribing B. In the absence of the proposition *drug B is the cure*, Jill's potent normative reasons just don't weight heavily enough in favour of prescribing drug B. (Below I discuss a reply insisting that they do.)

The totality of potent normative reasons that Jill has in *Notebook trouble* is not relevantly different from those in our original *Dr Jill* case. Since she doesn't know that drug B is the cure, what Dr Jill ought to do is prescribe medicine A (that is the standard verdict in the original case).³¹ The important question is: can she act as she ought in a way that is not merely lucky or accidental? I think not. Relatedly, if Jill ends up prescribing drug A, acting as she ought, her normative success is not creditworthy or praiseworthy. This verdict can be supported by appeal to the modal condition on luck. Jill cannot discriminate her case from counterfactual cases in which she is not in a Gettier-case, and does know that B is the cure on the basis of consulting her notebook. Indeed, had everything been normal, she would have known. Any way of making a choice that results in prescribing medicine A will lead her astray in cases in which she knows that B is the cure.

Compare a subject in a Gettier-case and a subject who forms a true belief on the basis of a random guess. Neither subject knows. However, as several epistemologists have pointed out, only the subject in a Gettier-case is *unlucky not to know*.³² Jill is unlucky not to acquire knowledge by consulting her notebook. The modal condition on luck outlined above explains this, for the way in which she forms her belief results in coming to know across a wide range of relevant, somewhat normal cases. As a result, her failure to know is merely unlucky.

One way to try to block my argument is to concede the claims I have made about knowledge, but to contest their implications for normativity in the context of *Knowledge Filtering*. In particular,

³¹ Note that nothing I have said explains how (or whether) proponents of *Knowledge Filtering* get the verdict they want in our original case of *Dr Jill*, but I am setting this problem aside.

³² See Sutton (2005) for a similar point about Gettier-cases. Sutton thinks this is true, more generally, of subjects who hold justified beliefs that fail to constitute knowledge. Hirvelä (MS) argues that subjects in Gettier-cases could *easily* have known.

one could contest the assumed dependence between facts about what Li and Jill know and facts about what they ought to do. Even had Li not known that Ada weighs at least 15 kg, its *seeming* to him that way would have sufficed to make it the case that he ought to let her climb. And in *Notebook trouble* Jill ought to prescribe medicine B after all, for even though she does not know that it is the cure, it *seems* or *appears* to her to be the cure, and she knows which number is written in her notebook, and these reasons are sufficiently strong on their own.

Instead of debating the details of particular examples, we should ask what must be assumed for such a reply to be available across the board. One would have to subscribe to the general claim that p can't be a weightier reason than an appearance or seeming that p : if p would (if known) be a potent normative reason for s to ϕ , then *its seeming to one that p* (or the seeming itself) would be a potent normative reason to ϕ with the same weight.³³ Moreover, in the overall balance of reasons, *its seeming that p* as it were screens off whatever weight p itself has: if, for instance, it first merely seems to me that p , and I then come to (be in a position to) know p , acquiring the new potent normative reason that p can make no difference for what I ought to do. Among our potent normative reasons, propositions concerning the world outside our appearances turn out to be normatively inert! I think this is extremely implausible. For instance, if p (obviously) entails q , then p is a maximally strong reason to believe q . But its seeming to me that p is not a maximally strong reason to believe that q , for it is compatible with q being false. One would think that the strength of one's reasons for believing p had something to do with probabilities. But unless one's prior probabilities *rule out* the possibility of misleading seemings or appearances, the probability of q on p will be higher than the probability of q on its seeming to one that p .³⁴ I deliberately chose the examples so that the consequences of acting if the relevant propositions are false are disastrous, and in such cases it matters whether one's reason is that it seems that p , or that p is the case.

But ultimately, even the backup view does not solve the problem. Let me once again remind the reader of the dialectic here. Anti-luminosity arguments, it seemed, simply have no bearing on views that have an epistemic filter on the potent normative reasons. But whether appeal to appearances or seemings, together with the backup view discussed, solves the problem I have raised depends on what it takes for facts about seemings to become potent normative reasons. If a proposition must be known in order to be among one's potent reasons, then the proposed view is susceptible to anti-luminosity -type reasoning: there are cases in which it seems to one that p , but one is in no position to know this.

It is also instructive to see why piling on more access conditions won't help. Consider a more stringent epistemic filter: in order for a proposition p to pass the filter, one must know that one knows p . Normative facts now depend on facts about whether one has such iterated knowledge. The problem is that there will be cases in which p is an essential normative reason weighting in favour of ϕ 'ing, but one is in no position to track the relevant epistemic facts, facts regarding whether one knows that one knows p . That is because knowledge that one knows p is susceptible to exactly the same phenomena as knowledge that p : one's knowledge that one knows can be precarious—one could easily have merely known—and one can have a Gettier-belief that one knows.

At this point those who appeal to seemings might simply reject knowledge filtering. True, they might concede, seemings are not immune to anti-luminosity reasoning. But we should simply reject knowledge as the relevant kind of epistemic access. Guidance requires access, but the access in question itself is a kind of apparentness: we have access to just those facts that *seem* to us to be the case. Alternatively, perhaps seeming-states *themselves* figure among our potent normative

³³ Kiesewetter (2017, Chapter 7) endorses a version of the backup view in response to an objection to his view of rationality. By contrast, Lord (2018: 193) concedes that what he calls the 'atomic' weight of a reason like *drug B is the cure* can be greater than that of *it appears that drug B is the cure*.

³⁴ See also Littlejohn's (2018) discussion of the backup view.

reasons, and do so *as long as we undergo them*.³⁵ How is anti-luminosity -type reasoning relevant against such views in the first place? In the next section my aim will be to generalize the problem to this alternative picture of epistemic access.

4. Seemings filtering: A Second Case Study

Some might worry that structural analogues of the anti-luminosity problem kept cropping up because of a knowledge-centric view of epistemic access. So it is worth discussing a wholly different kind of epistemic filtering.

Here is a thought. Only some facts are made apparent to us. Being made apparent in the relevant sense is not a matter of being known. It is a matter of these facts *seeming* to us to be the case, where a seeming is a *sui generis* kind of conscious mental state with propositional content and a distinctive phenomenology.³⁶ To be among our potent normative reasons, facts have to pass a *seemings filter*: in order for the fact that p to be among my potent normative reasons (or among the facts grounding the relevant normative truths), it must seem to me that p .

Here is another, slightly different thought. Ultimately, (potent) normative reasons are not facts or propositions, but mental states. And they are, in particular, *seeming-states*. Whether or not it is true that p , when it seems to me that p , the seeming-state itself is a potent normative reason, the totality of my potent reasons consisting in the totality of such seeming-states.³⁷ All it takes for a seeming-state to be among my potent normative reasons is for me to *undergo* that state.

These views differ in important respects. Assume that it seems to me that the pond is frozen, but it is in fact not frozen. Since it is not a fact that the pond is frozen, according to the first view there is no work for a seemings filter to do: at least the proposition that the pond is frozen is not among my potent normative reasons. According to the second view, whether or not the pond is frozen makes no difference, for either way, the seeming-state itself, with the content that the pond is frozen, is among my potent normative reasons. But despite their differences, according to both views, just what my potent normative reasons are – and hence, what the relevant normative facts are – depends on what seeming-states I undergo. That is, normative facts come to depend on facts about my seemings.

I argued above, drawing on previous work, that we are sometimes in no position to dispositionally discriminate between cases in which we know and cases in which we don't: sometimes we just cannot track the relevant knowledge facts and hence, the normative facts that depend on them. As a result, sometimes we are in no position to act as we ought in a way that is not lucky. Might seemings filtering immunize one from this sort of worry? A case in which one knows p , and a case in which one does not know p , can as it were look just the same from the inside – indeed, those who deploy seemings-talk will be inclined to say that things can *seem* just the same, for knowledge depends on facts about the world and one's relation to it. This was clearest in the second kind of case discussed above, which was a Gettier-case. By contrast, if it seems to a subject that p , a case in which it does not so seem cannot, trivially, be one in which things seem exactly the same!

But essentially the same problem, I want to argue, arises given seemings-filtering. Indeed, it was not even assumed in the precarious knowing case that Li's phenomenology is exactly the same

³⁵ I take this to be Kieseewetter's (2018) view, though his views is not that *only* such seeming-states constitute potent normative reasons.

³⁶ This phenomenology has been described as 'forceful' (Huemer: 2001: 77-9), 'assertive' (Tucker 2010: 530) as 'recommending' its content, and as having 'the feel of truth' (Tollhurst 1998: 298). As Siegel pointed out to me, the idea goes back at least to Price's 1932 book *Perception*. See further references in note 13 of Siegel (2017).

³⁷ According to a hybrid view, both mental states such as seemings, and propositions (presumably passing some epistemic filter) can be among one's potent normative reasons. I won't separately discuss such hybrid views, for the points I make equally apply to them.

in a case in which he knows and a case in which he does not. Even if seemings with different contents differ in some way regarding their phenomenology, substantive argument would be needed to establish that the possibilities regarding their contents always pattern with the limits of dispositional discrimination for a subject – that the limits of dispositional discrimination place principled limits on phenomenology. One would have to argue that the following kind of situation is impossible: while both a seeming that p_1 and a seeming that p_2 are in a subject's repertoire of possible experiences, she cannot discriminate between them. (Such dispositional discrimination would consist in reliably responding to these seemings in different ways across a range of relevant counterfactual cases – for instance, by coming to believe p_1 – not p_2 – when it seems to her that p_1 , and by coming to believe that p_2 – not p_1 – when it seems to her that p_2).

But why think that phenomenology is necessarily thus constrained by one's abilities of discrimination? Consider, for instance, a seeming that an object is red₂₇, and a series of cases by which a subject very gradually comes to experience a seeming that it is some quite different shade of red, say red₆₄ – perhaps, for instance, the light shining on the object very gradually and smoothly changes colour. Why couldn't her colour seemings change so gradually that she cannot dispositionally discriminate differences in how things seem from one moment to another, being unable to reliably track differences between very similar seemings across a range of relevant counterfactual cases? Once again precisely the kinds of considerations that anti-luminosity arguments rely on raise their head. Those who press the anti-luminosity argument will argue that seemings are not luminous: one is not always in a position to know how things seem. That, of course, is not what is at issue here. But underlying the plausibility of anti-luminosity reasoning is the assumption that it is just not feasible to be disposed to discriminate between two cases that are phenomenologically very alike – indeed, I think that lack of such dispositional discrimination *explains* lack of epistemic discrimination, explaining why one is not always in a position to know exactly how things seem.

If I am right, then we can press cases against seemings filtering that are structurally similar to the margin for error cases discussed above in connection with knowledge filtering. Consider, for instance, the first kind of view outlined above: one's potent normative reasons consist in those facts that seem to one to be the case. We can now build cases according to the following recipe. First, some proposition p_1 is true, and in fact seems to one to be true. (p_1 could, for instance, be the proposition that a child weighs more than 15 kg, or the proposition that a certain object is shade red₂₇.) Hence, p_1 is among one's potent normative reasons. Second, p_1 is not only a reason for one to ϕ , but an essential one, in the following sense: if p_1 were removed from one's stock of potent normative reasons, it would no longer be the case that one ought to ϕ . Finally, one cannot dispositionally discriminate between a seeming that p_1 , and various other seemings that one undergoes in other relevant counterfactual cases, such as a seeming that p_2 (that the child weighs just 15 kg, and no more; that the object is red₂₈ instead of red₂₇, etc.), where the contents of these other seemings, even if true, are *not* reasons to ϕ . That is, it is not feasible for one to be disposed to respond differently to these different seemings. As a result, one cannot discriminate one's actual case, in which one ought to ϕ , from other relevant cases in which one ought not to ϕ . But then, if one ought to ϕ and does so, it seems that one is lucky to have acted as one ought.

The worry is reminiscent of some discussions of the so-called *problem of the speckled hen*: one's phenomenology can have a determinacy that outstrips one's discriminative abilities.³⁸ Some have appealed to a distinction between *seemings* and *experiences* in order to solve the problem of the speckled hen.³⁹ But a principled case would have to be made that while the contents of experiences can outstrip our abilities of discrimination, this could not happen for seemings. And as

³⁸ See, for instance, Pace's (2010) discussion, who recaps the core of the problem as follows: "there are some properties presented in experience at a level of determinacy that outstrips the subject's recognitional capacities" (Pace 2010: 404).

³⁹ See, for instance, Tucker (2010), Brogaard (2013).

far as I can see, no such case has been made.⁴⁰ It is also worth noting that various strategies for dealing with such discrimination and access worries have no bite in the current dialectical context. For instance, in response to anti-luminosity worries applied to evidence – which he thinks consists in non-factive mental states – Declan Smithies (2012) invokes the distinction between propositional and doxastic justification: whenever I am in a given kind of mental state, I have propositional justification to believe that I am in that state, even if, due to my limited doxastic capacities, I cannot exploit that justification to come to justifiably believe that I am that state. Note that this strategy consists in conceding my points about lack of dispositional discrimination, while insisting that even if I cannot dispositionally discriminate a seeming that the object is red₂₇ and a seeming that it is red₂₈, I nevertheless have propositional justification to believe truths about how things seem. However, in conceding the point about dispositional discrimination, this strategy does nothing to help with the luck-related worry: when the fact that I ought to ϕ depends on its seeming to me that the object is red₂₇ (and not that it is red₂₈), my lack of discrimination prevents me from acting as I ought in a way that is not lucky.

5. Conclusions: Trouble for Epistemic Filtering

Normative guidance requires having access to facts in virtue of which normative facts hold. Think of these ought-makers as going into a box, the contents of which determine the relevant normative facts. The Access Problem arises because given a wide range of views about such ought making facts, we *don't* always have epistemic access to them, not even if they belong to an internal domain concerning our beliefs, experiences, or motivations. Epistemic filtering would seem to bypass the problem, for it *guarantees* epistemic access.

The new problem, however, is that epistemic filtering of any kind makes normative facts dependent on facts about what we bear the relevant epistemic relation to. But for any candidate kind of epistemic access, we are sometimes in no position to discriminate just which propositions we can access. As a result, we sometimes cannot act as we ought save by luck—and hence, we cannot act as we ought in a creditworthy way.⁴¹

References

Alvarez, Maria

2016 “Reasons for Action: Justification, Motivation, Explanation”, in *Stanford Encyclopedia of Philosophy*.

Arpaly, Nomy

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

Berker, Selim

2008 “Luminosity Regained”, *Philosophers' Imprint* 8.2: 1–22.

⁴⁰ These issues were discussed in an earlier, longer version of Hawthorne and Lasonen-Aarnio (*forthcoming*).

⁴¹ Many thanks to Daniel Drucker, Giada Frantantonio, Jaakko Hirvelä, Antti Kauppinen, Lisa Miracchi, and Niall Paterson for discussions and comments on an earlier draft, and to comments from two anonymous reviewers. This research has been funded by a grant from the *European Research Council*.

- Brogaard, Berit
2013 “Phenomenal Seemings and Sensible Dogmatism”, in Tucker (ed.) *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Oxford: Oxford University Press, pp. 270-289.
- Dancy, Jonathan
2000 *Practical Reality*, Oxford: Oxford University Press.
- Gibbons John
2013 *The Norm of Belief*, Oxford: Oxford University Press.
- Glüer, Kathrin and Wikfors, Åsa
2009 “Against Content Normativity”, *Mind* 118.469: 31–70.
2013 “Against Belief Normativity”, in Chan (ed.), *The Aim of Belief*, Oxford University Press, pp. 80–99.
- Johnson King, Zoë
2020 “Accidentally Doing the Right Thing”, *Philosophy and Phenomenological Research* 100.1:186–206.
- Hawthorne, John and Lasonen-Aarnio, Maria
Fortcoming “Not So Phenomenal”, *The Philosophical Review*.
- Hirvelä, Jaakko
MS “Justification and the Knowledge Connection”
- Hudson, James L.
1989 “Subjectivization in Ethics”, *American Philosophical Quarterly* 26.3: 221–229
- Huemer, Michael
2001 *Skepticism and the Veil of Perception*, Lanham: Rowman & Littlefield Publishers, Inc.
- Hughes, Nick
2018 “Luminosity Failure, Normative Guidance and the Principle ‘Ought-Implies-Can’”, *Utilitas*, 30.4: 439–457.
- Jackson, Frank
1991 “Decision-Theoretic Consequentialism and the Nearest and Dearest Objection”, *Ethics* 101.3: 461–482
- Kagan, Shelly
2018 “The Paradox of Methods”, *Politics, Philosophy & Economics*, 17.2: 148–168.
- Kiesewetter, Benjamin
2011 “‘Ought’ and the Perspective or the Agent”, *Journal of Ethics & Social Philosophy* 5.3: 1–24.
2017 *The Normativity of Rationality*, Oxford: Oxford University Press.

Lasonen-Aarnio, Maria

2010 “Unreasonable Knowledge”, *Philosophical Perspectives* 24(1): 1–21.

Forthcoming A “Competent Failure and Victims of Deceit”, in Dutant and Dorsch (ed.), *The New Evil Demon*, Oxford: Oxford University Press.

Forthcoming B “Dispositional Evaluations and Defeat”, in J. Brown and M. Simion (eds.), *Reasons, Justification and Defeat*, Oxford University Press.

Forthcoming C “Perspectives and Good Dispositions”, *Philosophy and Phenomenological Research*.

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.

Markovits, Julia

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

Pace, Michael

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

Parfit, Derek

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

Pollock, John

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

Pritchard. H. A.

1932 “Duty and Ignorance of Fact”, reprinted in *Moral Writings*, McAdam, Jim (ed.), Clarendon Press, Oxford, 2002, pp. 84–101.

Raz, Joseph

2011 *From Normativity to Responsibility*, Oxford: Oxford University Press.

Siegel, Susanna

2017 *The Rationality of Perception*, Oxford: Oxford University Press.

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* 22: 369–86.

Smithies, Declan

2012 “Mentalism and Epistemic Transparency”, *Australasian Journal of Philosophy*, 90:4, pp. 723–741.

Srinivasan, Amia

2015a “Normativity Without Cartesian Privilege”, *Philosophical Issues* 25-1: 273–299.

2015b “Are We Luminous?”, *Philosophy and Phenomenological Research* 90.2: 294–319.

Sutton, Jonathan

2005 “Stick To What You Know”, *Nous* 39.3: 359–396.

Tollhurst, William

1998 "Seemings", *American Philosophical Quarterly*, Vol. 35, No. 3, pp. 293-302

Tucker, Chris

2010 "Why Open-Minded People Should Endorse Dogmatism", *Philosophical Perspectives* 24: 529–545.

Unger, Peter

1975 *Ignorance: A Case for Skepticism*, Oxford: Oxford University Press.

Way, Jonathan and Whiting, Daniel

2017 "Perspectivism and the Argument from Guidance", *Ethical Theory and Moral Practice* 20: 361–374.

Wedgwood, Ralph

2002 "The Aim of Belief", *Philosophical Perspectives* 16: 267–297.

Williamson, Timothy

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