GOOD THINKING

by

Timothy R. Kearl

_____

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

-2022-

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by: Timothy R. Kearl

titled: Good Thinking

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

Juan Comesana
_____   Date: May 2, 2022
Juan Comesana

_____   Date: May 2, 2022
Stewart Cohen

Michael McKenna
_____   Date: May 2, 2022
Michael McKenna

_____   Date: May 2, 2022
Carolina Sartorio

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Juan Comesana
_____   Date: May 2, 2022
Juan Comesana
Philosophy

Contents

Abstract

Good Thinking is a collection of papers about abilities, skills, and know-how and the distinctive but often overlooked—or explained away—role that these phenomena play in various foundational issues in epistemology and action theory.

Each chapter, taken on its own, represents a fairly specific intervention into debates in (i) epistemic responsibility, (ii) the nature of inferential justification, and (iii) connections between inference and action. But taken collectively, these chapters constitute fragments of a larger mosaic of commitments about the explanatory priority of abilities in normative theories.

One distinctive argumentative strategy employed throughout *Good Thinking* is its placing special emphasis on what might be called "bad thinking": defective judgments borne out of cognitive short-circuiting, incoherence or self-doubt, depression, or anxiety. The underlying motivation for this is that much of what we can learn about good thinking is only revealed at the margins, where thinking has in some respects gone bad without being entirely spoiled.

An Introduction to *Good Thinking*

This dissertation is dedicated to my mother, Mari, and to my late father and very first editor, Steve. Without their love and support, I would have given up on philosophy long ago.

Good Thinking is a collection of papers about the nature of epistemic and practical agency. In very rough outline, it deals with topics such as: what is it for one to *think* or to *act*, rather than to merely behave? What is it for a thinker to be *responsible* for her thoughts, and how does one square that with accounts of what it is for a practical agent (an "acter") to be responsible for her actions? And so on. When the dust settles, it is my hope that the reader sees the topics of *epistemic agency*, *practical agency*, and *the relationship between them* in a new and kinder light.

Each paper (or chapter) is meant to stand on its own merits, and I summarize them below. But before that, I will humor myself by providing some personal context and history.

In February 2009, when I was 17, my father, Steve, suffered a debilitating stroke. He had lived, before the stroke, as an incredibly capable, witty, and dependable person. He was such that if you knew him, you tended to admire him. Stories of his youth, his career, and his various adventures made him seem larger than life, ones which, if told by anyone else would probably strike the listener as untrue. But that was really Steve.

After the stroke, he was left disabled in various ways and faced a number of health-related challenges, some of which he never overcame, even after years of recovery. A few days after my 30th birthday, almost 13 years after the stroke, Steve died peacefully and in the arms of someone he loved, my mother, Mari. When we last spoke, he and my mother sang me Happy Birthday, and he said he'd be watching *The Last of the Mohicans* to commemorate the occasion; it was, as we often said to one another, "one of the better all-time greats." As I write this, I am so proud to say that he edited the

papers collected in this volume; we were able to talk philosophy as equals, even as he insisted that he couldn't make heads or tails of some of the finer points, or of the various professionalisms that can permeate contemporary academic writing. To borrow a phrase from Robert Wallace, my father and I became "friends in thinking."

I spent years struggling to process Steve's stroke—we all did. I was fortunate enough to be able to escape the day-to-day strain at Dartmouth College, where I began studying mathematics before eventually deciding to study philosophy too. In a series of events that were no more than sheer dumb luck, I took a course with Adina Roskies and Christine Thomas in which we read about theories of personal identity, or accounts of the conditions under which someone is *the same person* across time. Although it was not so obvious to me then, it is obvious to me now that philosophy gave me a therapeutic outlet to grapple outwardly with questions that I could not overcome inwardly. In a perversion of Wittgenstein, philosophy began for me as therapy, and then I kicked the ladder away.

Over time, I found mentors in David Plunkett, Timothy Rosenkoetter, and John Kulvicki. They were incredibly smart and empathetic people who, knowing a bit about my personal situation, helped me come to enjoy and appreciate philosophy in and of itself. David, in particular, convinced me to take the idea of graduate school seriously when I was otherwise aimless. I don't know how I'll ever adequately thank any of them for the role they've played in kickstarting my philosophical career.

At Arizona, I was able to situate myself within an incredible group of philosophers, all of whom I consider dear friends: Juan Comesaña, Carolina Sartorio, Michael McKenna, and Stew Cohen. Each of them has, in their own way, helped shape me as an author, a critic, and a teacher, but I will spare them the embarrassment of too much praise. There are, of course, many more people than that at Arizona who have had a lasting and positive impact on my philosophical upbringing: Connie Rosati, Jason Turner, Houston Smit, Terry Horgan, Robert Wallace, Chris Howard, Santiago Sanchez-

Borboa, Sarah Raskoff, Jacob Barrett, Hannah Tierney, Tyler Milhouse, Bryan Chambliss, Bjorn Wastvedt, Sean Whitton, Rhys Borchert, Eyal Tal, Will Schumacher, and many more besides. I am also indebted to my mother, Mari MacMillan Kearl, my partner, Tia Lewis, and to many of my friends: Michael Harmon, Kevin Kennedy, Cecil Qiu, Roland Mansilla, Luke Peters, and countless others, each of whom could not only tolerate the occasional philosophical rant but also listen, wax philosophical, and teach me something I was too preoccupied to see. To anyone reading this thinking that they belong on this list, you're probably right, and the list should be emended.

Besides the requisite acknowledgments, why say any of the rest? Why the personal history? In my mind, setting aside the philosophical name-dropping in pursuit of situating these ideas conveniently but somewhat artificially in the recent history of thought, there is a thread running through *Good Thinking* that is borne out of dealing very directly and viscerally with what might be called "bad thinking": my own bad thinking through grief and depression, my father's limited and sometimes short-circuited thinking after his stroke, and the many more unsexy errors of judgment and action that get idealized away in theory-crafting. Any philosophical theory of rational thought (and action) worth its salt must be alive to the many ways in which thinking (and acting) goes bad, sometimes without being entirely spoiled, and not all of which fall neatly into traditional philosophical categories or paradigms. (As Steve once said, "Tim, all of us are a little fucked up.")

And so the dissertation begins with *A plea for exemptions*, a paper concerning the inadequacy of extant accounts of epistemic responsibility. Theories of epistemic responsibility tell you, among other things, the conditions under which an agent is "on the hook" for her failure to follow her evidence, or her failure to have coherent beliefs, or… According to a commonly accepted picture of epistemic responsibility, there are *rational judgments* and *irrational judgments*, and among the latter there are the *excusable* and *inexcusable* ones. I argue that this three-way classification cannot capture what is

distinctive about a certain class of failures of rationality: some failures of rationality are borne out of varying degrees of incompetence or lack of skill, rather than being the result of poor epistemic performance. These "competence errors" are, on the commonly accepted picture, just more irrationality, sometimes excusable, sometimes not. This picture is one that I go to great lengths to refute, arguing for the cogency of a *sui generis* category of epistemic exemption. (In a companion piece not included here, *The quality of thought*, I focus on the poverty of extant accounts of epistemic excuses *per se*, rather than on the poverty of the landscape within which theories of epistemic responsibility are housed.)

The second chapter, *Evidentialism and the problem of basic competence*, argues against certain forms of evidentialism about inferential justification. According to (one version of) the view I criticize, for an agent A to be justified in inferring hypothesis H from evidence E, A must possess not only *E* as evidence but also the conditional *if E, then H* as evidence. In a slogan, these conditionals "link" one's evidence to various hypotheses, allowing rational inference to H to flow from E by way of the link, itself another piece of evidence. Against this picture, I argue for a form of foundationalism about inferential justification; what the agent needs to "link" her evidence to various hypotheses is not more evidence, but instead a certain kind of basic knowledge-how. Thus, the epistemology of inference resembles the epistemologies of perception and intuition much more than has been appreciated, at least when the latter two are captured in foundationalist frameworks.

The third and final chapter, *What we know when we act*, ostensibly reconciles causalist views of action with Anscombian views of action, but just beneath the surface it is an extension of the arguments developed in chapters one and two. Intentional action, like rational inference, requires the agent have a special kind of knowledge-how, what I call "ability-constituting knowledge". Only by appealing to an agent's ability-constituting knowledge *of inference* can one explain what is sometimes rational about

certain judgments borne out of incoherence or self-doubt. (These cases are sometimes lumped under the heading of "inadvertent epistemic virtue", resembling cases of "inadvertent moral worth".) Likewise, only by appealing to an agent's ability-constituting knowledge *of action* can one explain what distinguishes what might be called "inadvertent" intentional actions from unintentional ones, non-intentional ones, and mere behavior. The resulting picture of practical agency is causalist to the extent that, necessarily, ability-constituting knowledge is among (or among the grounds of) the proximal and sustaining causes of an agent's behavior in cases of intentional action. It is Anscombian to the extent that it gives shape to the suggestive—if widely contested—idea that, necessarily, in acting intentionally, one knows what one is doing as one does it.

As I said above, despite the etiological undertow connecting these chapters, they are not meant to rise or fall together; they were written so as to be modular but mutually supportive. I cannot guarantee that any of the ideas defended in *Good Thinking*, singly or collectively, will not be overturned in short time by arguments I failed to consider, connections I was too thick to see, by the changing winds of philosophical sentiment, or whatever. For anyone wishing I had done more groundwork in this introduction to soften you up to the philosophical ideas expressed in the following chapters: read them and weep.

<div style="text-align: right">

Timothy Rion Kearl

The University of Arizona

As of September 26, 2021

</div>

Chapter 1:

A Plea for Exemptions

Currently popular theories of epistemic responsibility rest on the (perhaps implicit) assumption that *justification* and *excuse* exhaust the relevant normative categories. One gets the sense that, once we've laid down the conditions for justified belief, and once we've laid down the conditions of excusably unjustified belief, the work is done; all that's left is to clock out.

Against this backdrop, one is naturally led to think that if an agent's doxastic state—her various beliefs and belief-like attitudes, or a subset thereof—fails to be justified, it is thereby unjustified, perhaps excusably so. The aim of this paper is to argue that that natural thought is mistaken; some agents are epistemically incompetent, and in virtue of their incompetence, their doxastic states are *neither* justified nor unjustified (even excusably). Instead, the doxastic states of such agents are *exempt from epistemic evaluation altogether.* I argue that what underlies this point about exemptions is that epistemic competences or abilities play an important and typically overlooked role in epistemology, especially in theories of epistemic responsibility. Here, I am interested in uncovering that role and explaining what it is, and also in explaining how one could accommodate it within various epistemological frameworks.

Here is the rough motivating idea: call any belief-like state that fails to be justified an "epistemic error". There are, speaking very generally, two ways in which epistemic errors arise. The first, "epistemic performance errors", are the result of failing to properly manifest some underlying epistemic competence. The second, "epistemic competence errors", are the result of lacking sufficient underlying competence. Extant theories of epistemic responsibility mistakenly lump competence errors with various

performance errors. My own "epistemic abilities view" of epistemic responsibility has the notable advantage of avoiding this mistake.

1. Excuses and the problem of epistemic competence

It is by now common to distinguish between merely unjustified beliefs and excusably unjustified beliefs, the latter exhibiting some epistemic good not found in the former. There is no consensus as to the conditions under which an unjustified doxastic state is excusable, but the candidate theories can be helpfully sorted into three varieties.

First, there are "norm-based conceptions" of epistemic excuse. As an example of a theory of this sort, consider Williamson (forthcoming). There, he distinguishes between the *primary* norm of belief, and various *derivative* norms which hold only in virtue of this primary norm. Assuming that the primary norm of belief is the knowledge norm, various derivative norms will require that agents be variously *disposed* to satisfy the knowledge norm, by, for instance, being disposed to believe as a knower would; when an agent fails to satisfy the knowledge norm but nevertheless satisfies a relevant derivative norm, their irrationality may be excusable on those grounds.

Next, there are "virtue-based conceptions" of epistemic excuse.[1] Lasonen-Aarnio (forthcoming) offers a view of this sort that starts by distinguishing between "epistemic success", or satisfying the knowledge norm, on the one hand, and "success-conducive dispositions", or manifesting epistemic virtues, on the other. The epistemic virtues, as Lasonen-Aarnio conceives of them, are dispositions to form beliefs that constitute knowledge in an environment like ours. She argues that subjects of a so-called "new evil demon scenario"[2] *virtuously* fail to satisfy the knowledge norm, insofar as the agent displays a kind of good-making epistemic feature, which she calls "reasonableness",

---

[1] See, e.g., Lasonen-Aarnio (forthcoming); Boult (2016)
[2] This is originally due to Cohen (1984)

and which can obtain despite failing to know.[3] Radically deceived subjects' failures to know are excusable because reasonable, and reasonable because virtuously formed.[4]

Finally, many authors have entertained a view of epistemic excuses that draws heavily on the control-based theories of responsibility now popular in the moral responsibility literature, according to which an agent is responsible for their actions only if those actions are performed with some *level of control*, in a sense of 'control' to be specified by the theory in question. Call these "control-based conceptions" of epistemic excuse.[5] These theories are relatively heterogeneous, differing in the degree of control over one's own doxastic state that is attributable to normal agents, the degree of control over one's own doxastic state that one must lack in order to count as excused, and the sorts of phenomena that count as control-diminishing. What is common to various

---

[3] Perhaps some virtue-based theories of epistemic excuse collapse into norm-based theories. If, for instance, one accepted *both* that knowledge is the norm of thought, *and* that epistemic virtuous are comprised simply of dispositions to know or to believe as a knower would, there might appear to be little by way of interesting theoretical space between norm- and virtue-based conceptions of epistemic excuse. If, on the other hand, one had a conception of "epistemic virtue" that involved certain traits like *thoughtfulness* or *inquisitiveness*, virtue-based views would more clearly come apart from norm-based views.

[4] Littlejohn (2012; forthcoming) and Sylvan (2015) treat justification and rationality as distinct standards of epistemic evaluation. (See Cohen (2016) for a criticism of this methodology.) Both accept that a belief must satisfy the knowledge norm in order to be justified, and both deny that a belief must satisfy the knowledge norm in order to be rational. According to Littlejohn, a belief is justified just in case it conforms to the epistemic norms, while a belief is rational just in case it is the product of an "excellent use of a subject's rational capacities" (See his 2012, section 3.1). Like Lasonen-Aarnio, Littlejohn maintains that the display of epistemic excellence can serve as an excuse for one's failure to so conform.

In contrast, Sylvan argues that meta-epistemology needs the meta-ethical distinction between objective and apparent reasons. According to him, beliefs are *justified* only if they fit the *objective* reasons the agent possesses, while beliefs are *rational* only if they fit the *apparent* reasons the agent possesses. I take it that 'apparent' is only functioning to capture a non-factive use of the term 'reason', according to which one can have merely apparent —but not "objective"—reason to believe a false proposition. Like Lasonen-Aarnio and Littlejohn, Sylvan suggests that responding to apparent reasons "excuse[s] you in the best way, by showing that your rationality was in working order" (9-10).

[5] See, for instance, Owens (2000), Levy (2007), and McHugh (2013).

control-based conceptions of epistemic excuse, and what distinguishes them from norm- and virtue-based conceptions, is that excusably unjustified beliefs need not involve *any success* on the part of the agent. Rather, the agent's failure must be *outside of her control,* in the relevant sense of 'control'.[6] Paradigmatically, though views may differ on the details, an agent is epistemically excused on control-based theories when their judgments are influenced by fear, depression, high-stakes, distraction, and the like.[7]

One might sensibly wonder which of these uses of 'excuse', if any, is correct. We can set that issue aside for now, however, since for our purposes the important point is that every epistemologist already admits of epistemic excuses, in light of cases where either (i) the agent satisfies some derivative norm of belief while failing to satisfy the primary norm of belief; (ii) the mechanism by which one forms epistemically good beliefs is intact, but its operation results in an epistemically bad belief; or (iii) the mechanism by which one forms epistemically good beliefs is, in some sense, intact, but its proper operation is interfered with in some or another way outside of the agent's control.

Crucially, none of these views can adequately accommodate and explain failures to form justified beliefs that are the result of epistemic *incompetence*, where the mechanisms by which one forms epistemically good beliefs are not intact. Below, I

---

[6] Since control-based views of epistemic excuses do not imply that agents with excusably unjustified beliefs agents have in any way *succeeded*, these are particularly suited to non-factive theories of epistemic justification. While non-factive theories of justification do not need to appeal to excuses to explain the ways in which irrational agents are nevertheless *successful*, as factive theories of justification do, they do not thereby avoid the need to appeal to excuses entirely. Thus, non-factive theories of justification may avoid what Gerken (2011) has called an "*ad hoc* excuses maneuver", but this is consistent with insisting that one must appeal to excuses *somewhere*, and this is enough to see the problems for control-based views of excuse.

[7] Various examples of control-based views range from Gardner's (1998) "Kantian" conception of legal excuses (this is Gardner's term for a view he finds wanting), where excuses are tied to an agent's specific abilities in a context; and forms of direct doxastic voluntarism, where agents can exercise direct control over at least some of their beliefs (see Ginet 2001 and Weatherson 2008 for discussion).

present a pair of cases that are meant to make salient the contrast between unjustified belief formed out epistemic incompetence and unjustified belief formed out of epistemic poor performance. The motivating thought behind these cases is that there is *some interesting normative difference* between the doxastic states of each agent, one of whom forms unjustified beliefs out of incompetence, the other of whom forms unjustified beliefs while retaining his competence but manifests it poorly or not at all. And the problem for each of the views of epistemic excuse falling under (i)-(iii) is that it fails to respect this normative difference.

Consider the following pair of vignettes:

1. Smith, a diehard Democrat, volunteers to help gather political party information for the upcoming US census. He flies to northern Florida to conduct door-to-door polling. Smith polls 100 households and all 100 identify as Republican. Naively optimistic and prone to wishful thinking about the presence of an enclave of Democrats in north Florida, each Republican household strengthens Smith's resolve that the next household will identify as Democrat. By the 101st household, Smith fully believes that the next household will identify as Democrat.

2. Jones happens to volunteer with the same organization as Smith. He flies to Mississippi to conduct door-to-door polling for the upcoming US census. Jones polls 100 households and all 100 identify as Republican. Some of the locals resent these political out-of-towners and conspire with the resident nefarious epistemologist, Black, to surreptitiously drug Jones on his way to the 101st household. The drug is administered as an odorless gas, the effect of which is to permanently remove one's ability to believe on the basis of one's evidence.[8] In this state, Jones has a kind of normative blindness; he

---

[8] We can make Jones's predicament more vivid if we imagine Jones's phenomenology under the influence of the drug is of merely representing things as being the case, but not for this or that *reason*; rather, Jones represents things as being the case because (causally speaking) it suits him

cannot recognize anything as a reason to believe anything else; whatever beliefs Jones forms, he will form them merely because (causally speaking) it strikes his fancy to do so. As it so happens, it strikes Jones's fancy that the next household will identify as Democrat, and he comes to fully believe that the next household will identify as Democrat.

In each case, the agent's belief fails to be justified or rational.[9] After all, each agent has overwhelming inductive evidence in support of the hypothesis that the 101st household will identify as Republican. It seems that the only rationally permissible attitude each agent can take towards that hypothesis is one of belief (as opposed to disbelief or suspension of judgment).

Because both Smith's and Jones's beliefs fail to be rational, I'll say that Smith and Jones are each committing "epistemic errors". Nevertheless, their epistemic errors are

---

at the moment. Jones's phenomenology regarding epistemic reasons resembles a patient suffering from a kind of sensory neglect, as when someone experiences serious brain trauma, sometimes the result of a stroke, which causes her to neglect information from parts of her visual field, behaving as though she is not receiving that visual information at all. For instance, she may fail to eat the food on the left half of her plate despite complaining of hunger. Such a subject has lost the ability to attend to, and so respond to, certain visual information. Jones, by way of analogy, has lost the ability to attend to, and so respond to, evidential considerations *tout court*.

[9] Admittedly, some authors will deny that Jones has 'cognitive states' at all, properly so-called, and so deny that Jones counts as *believing* anything at all after the gas has taken effect. On views like Ichikawa and Jarvis (2013), a state only counts as a cognitive state insofar as it is governed by the "rules of thought", and Jones, by stipulation, has lost whatever capacities ensure this governance. Whatever the merits of Ichikawa and Jarvis's view, it seems like denying the cogency of Jones's position is not a plausible way to proceed, at least if that denial rests merely on the claim that a representation can rise to the level of genuine thought only when that representation is governed by the rules for thought (if 'governance' is understood constitutively, as opposed to normatively), since, in this context, that would just beg the question. But even setting that aside, denying that Jones is *genuinely* thinking does not fully address the problem; surely his representations, despite their disconnectedness from the rules of thought, insofar as he affirms them, and insofar as they guide his actions, may nevertheless rise to the level of "quasi"-belief, and the problem of what to say about the normative status of *those things* re-emerges.

importantly different. Smith's epistemic error is a paradigm case of irrationality. This seems true whether we accept a norm-, virtue-, or control-based account of epistemic excuse. Whether and the extent to which Smith's doxastic state is *excusably* irrational will be determined by some further conditions specific to theories of each stripe, and I'll comment on that below.

First, let me offer what I take to be a natural explanation of these cases. Smith's error is due to poor epistemic performance and not because of any lack of epistemic competence.[10] To wit, Smith can reason inductively just fine; the issue is rather that Smith's diehard commitment to the Democratic Party and naïve optimism about the prevalence of an enclave of Democrats in northern Florida interferes with the proper exercise of his underlying competence in fitting his beliefs to his evidence; he instead fits them to things that he is not warranted in taking for granted.

By contrast, Jones's error is not the result of his failure to manifest his underlying competence; Black's odorless gas removed it. An agent like Jones who cannot, in general, treat things as reasons to believe, will not be epistemically competent, in any interesting sense of 'competent'. He lacks the ability[11] to defeasibly base his beliefs on perceptual evidence, or to perform enumerative induction, and so he will not be able to learn any empirical propositions. Moreover, Jones cannot *revise* his beliefs, since engaging in belief-revision is matter of recognizing and responding to reasons which may not support one's present beliefs.

---

[10] To forestall any misunderstanding on this point, to say that Smith retains his epistemic competence is not to make any claim about his having a higher-order belief about what his evidence supports while having a conflicting, first-order belief; in that case, Smith would be epistemically akratic, both believing that P, and believing that it is irrational to believe that P. If Smith doesn't have a higher-order belief about what his evidence supports, but nevertheless fails to properly react to it because his naïve optimism and party affiliation interferes with his epistemic abilities, he is irrational.

[11] This is ambiguous between *local* and *global* ability; While the differences between local and global abilities need not concern us for the exposition of the cases, I address this ambiguity in section 2.2 in the course of defending an abilities-based view of epistemic responsibility.

This is, at any rate, a sketchy but plausible explanation of the difference between Smith and Jones, one which will be sharpened as we go along. But I suspect that a careful reader might already know where this is going: namely, norm-, virtue-, and control-based theories of epistemic responsibility are each forced to deny that Smith's and Jones's cases present any interesting normative difference at all. Let's take those views in turn.

Let's consider norm-based views of epistemic responsibility, using Williamson as a stalking horse. According to that sort of position, an agent is excusably irrational for believing that P only when they both fail to know that P and satisfy some derivative norm of belief in so believing. Williamson's own suggestion is that, for any norm N, there is a derivative, secondary norm DN to have a general disposition to comply with N, and a tertiary norm to do what someone who complied with DN would do in the situation at issue.[12] The knowledge norm of belief gives rise to secondary norms requiring agents have the general disposition(s) to know, and tertiary norms requiring that agents believe what an agent generally disposed to know would believe. Clearly, both Smith and Jones fail to satisfy the primary norm of belief, and their beliefs are thereby irrational. As I've told the story, Smith and Jones fail to satisfy the secondary norm of belief too; we can imagine that Smith's judgments about political affiliation in the Deep South are systematically influenced by his wishful thinking, so that he lacks any interesting disposition to know in the case in question. Likewise, Jones lacks any disposition to know by stipulation. And finally, it is plausible that both Smith and Jones fail to satisfy the tertiary norm of belief. Neither agent believes as a knower would believe; knowers would believe that the 101st household is Republican. As such, neither Smith nor Jones has any claim to excuse for their irrational belief; both agents' doxastic

---

[12] Williamson (forthcoming; pp 7-10)

states are *merely* irrational.[13] This does not seem to leave room for any interesting normative difference between Smith's and Jones's doxastic states.

Next, consider virtue-based theories of epistemic responsibility, treating Lasonen-Aarnio as an exemplar. According to such views, an agent is excusably irrational only when she both fails to know and manifests the disposition of a knower. But whatever competences underlie being a knower, at the very least these include the competence to respect one's evidence. And in this regard, neither Smith nor Jones manifests the competence of a knower, so neither Smith nor Jones has any claim to excusable irrationality. Both agents have *merely* irrational beliefs. Again, this leaves no theoretical space within which to articulate the interesting normative difference between Smith's and Jones's doxastic states.

---

[13] Williamson (*ibid*) discusses certain "brain scrambler" cases that affect an agent's reasoning capacities, making fallacious inferences seem obviously correct. He claims that an agent is excusably irrational in such cases, on the grounds that the brain scrambled agent believes the same proposition that someone with the general disposition to believe only what they know would believe in their circumstances (for the simple reason that the circumstances include having your brain scrambled in that way). This is an interesting claim, and I will devote part of 2.2 to exploring it. For now, however, we can note that the problem drawn out by contrasting Smith's case with Jones' is not addressed by this observation. If this sort of condition—believing as one who followed one's cognitive dispositions would believe—is sufficient for Smith to count as satisfying a tertiary norm of belief, it is also sufficient for Jones to count as satisfying a tertiary norm of belief. Consequently, the view would predict that *both agents are excusably irrational*, and this is just another way to deny the interesting normative difference between their doxastic states.

What's more, even if one cannot bring oneself to acknowledge *further,* sui generis *normative categories*, as I will suggest in section 2.1, one should take the contrast between Smith and Jones to support the claim that *there are importantly different kinds of epistemic excuses*. On one way of understanding this claim, it may be no more than a terminological variant on my own suggestion, although I would recommend caution: one might think that there are importantly different kinds of epistemic excuses, for reasons entirely unrelated to those given here in defense of the need for epistemic exemptions. See *The Quality of Thought* for elaboration on this claim.

Finally, consider control-based views. These views might appear to have the resources to explain the normative difference between Smith's and Jones's doxastic states, namely that Smith's doxastic states are under Smith's control in a way that Jones's aren't under Jones's control. Although the details will turn crucially on an account of doxastic control, it is not totally unreasonable to think that naïve optimism and ideological commitment of the sort that plagued Smith are features of one's cognitive life over which one has some relevant kind of control. And in virtue of this difference, one might insist that Smith's belief is merely irrational (his irrationality was under his control), while Jones's belief is *excusably* irrational (his irrationality was not under his control). One might think that control-based views of epistemic responsibility have a clear and decisive advantage over norm- and virtue-based views, at least in light of the performance/competence distinction.

Be that as it may, we can easily imagine variants on Smith's case that are problematic for control-based views too. Smith*, we can suppose, is someone very much like Smith who is tasked with polling the same area, but who is not led to believe that the 101st household will identify as Democrat by naïve optimism and ideological commitment. Instead, Smith* may be led to this irrational belief because the mere thought of finding no Democrats in northern Florida is so triggering as to always be accompanied by acute panic and feelings of dissociation, and he simply cannot bear it. Whatever control-based views say about the difference between Smith and Jones, it seems that they are forced to count Smith*'s doxastic state as being on a par with Jones's doxastic state; both agents will be excusably irrational for their lack of control. But surely there is a residual difference between Smith*'s severely panicky reasoning and Jones's diminished underlying ability to reason.

Here is what I hope you take from the foregoing example: even if control-based views have the resources to adequately distinguish between *Smith* and Jones, perhaps by offering an account of epistemic control according to which Smith is merely

irrational because sufficiently in control, while Jones is excusably irrational because insufficiently in control, we can bring the very same problem into relief by varying details of Smith's case. By taking *Smith*'s* error further and further outside of his control *without removing his underlying epistemic competence* entirely, we can force control-based views to predict that Smith* is *excusably* irrational. Since control-based views should already say that Jones's error is excusably irrational (because of a lack of control), one has to say deny that there is any interesting normative difference between the doxastic states of Smith* and Jones.[14]

In sum, views of epistemic excuse that fall under (i)-(iii) all have the same problem: they incorrectly imply that there is no interesting normative difference between Smith's (or Smith*'s) doxastic state and Jones's doxastic state. Norm-based views treat their doxastic states on a par because neither agent satisfies a derivative norm of belief; virtue-based views treat their doxastic states on a par because neither agent manifests an epistemic virtue; and control-based views treat their doxastic states on a par because neither agent is sufficiently in control of their beliefs.

Earlier, I offered a natural explanation of what differs between Smith and Jones: Smith commits a performance error, and Jones commits a competence error, and this difference seems to ground a difference in the normative status of their resulting doxastic states. Because none of the views of epistemic excuses canvassed above are consistent with this natural, epistemic competence-based explanation, call this "The Problem of Epistemic Competence".

---

[14] One might insist that this result is not problematic on the grounds that Smith* has claim to only a "partial" excuse, while Jones has claim to a "full" excuse. But, as was emphasized in footnote 13, even if one insists on a framework that appeals only to justification and excuse, the cases presented above support the thought that *there are irreducibly different kinds of epistemic excuses*; this already introduces more structure into a theory of excuses than can be accommodated by someone who claims that Smith* and Jones differ from one another only in degree.

In the next section, I outline an "epistemic abilities" view that offers a principled solution to The Problem of Epistemic Competence. An abilities view links the normative status of one's doxastic state to one's underlying epistemic abilities, thereby allowing room to say something interesting and distinct about the normative status of doxastic states which are formed as the result of epistemic incompetence, or lack of epistemic ability. In particular, an abilities view can make room for the normative category of *exemption*, which applies to the doxastic states of agents like Jones, whose error is the result of sheer epistemic incompetence, but not to the doxastic states of agents like Smith, whose error is much more clearly the result of failing to properly manifest some underlying competence.

2. Epistemic abilities

The distinction between normative competence and performance resembles the Chomskian distinction between linguistic competence and performance, which is roughly that a speaker's linguistic competence consists in their mastery of certain linguistic rules, and a speaker's linguistic performance consists in their employment of those rules in speaking and listening.[15] Clearly, not all linguistic error is the result of failing to manifest an underlying linguistic competence; when a freshman Greek student (who cannot yet speak or understand a bit of Greek) utters a string of Greek words that fail to form a grammatical sentence—perhaps by failing to recognize the

---

[15] 1965, pp.3-15. The task of the linguist, as Chomsky conceived of it, "is to determine from the data of performance the underlying system of rules that has been mastered by the speaker-hearer and that he puts to use in actual performance" (4). I have little interest in defending a particular reading of a particular phase of Chomsky's thinking on this distinction. Instead, I hope to employ a bit of conceptual machinery that enables us to distinguish between types of errors, and to simply note that this bit of conceptual machinery bears a familial resemblance to Chomsky's distinction between linguistic performance and competence.

difference between the dative and accusative cases—she is exhibiting a kind of linguistic error due largely or entirely to linguistic *in*competence.

Similarly, not all *normative* error is the result of failing to manifest an underlying normative competence. For instance, sometimes an agent violates some normative standards as a result of lacking the relevant normative competence. And, at least in some cases, normative performance errors and normative competence errors have different normative standing. This is an operative assumption of the moral responsibility literature, according to which at least *some* agents are not appropriate candidates for having moral obligations at all.

More specifically, there is remarkably broad consensus within the moral responsibility literature that an agent acts morally responsibly only if she has an ability to recognize and react to certain reasons, namely *moral reasons*, or reasons given by *the good*.[16] Call these an agent's *moral abilities*. In these terms, there is broad consensus that a necessary condition on an agent's being morally responsible for her actions is that she possess (sufficiently robust) moral abilities.[17] An agent exercises her moral abilities when, say she recognizes that the fact that X-ing causes needless suffering to an innocent person (or breaks a promise, or is unfair) is decisive (or very strong) grounds for refusing to X, when she treats this reason as decisive (or very strong) in deliberating, and when she tends to hold herself and others accountable to acting in accordance with this reason by blaming X-ers and perhaps praising non-X-ers. But an agent lacking in many or all of those abilities would not be the right kind of thing to be held accountable

---

[16] See, for instance, Wolf (1990), Wallace (1994), Fischer and Ravizza (1998), Nelkin (2011), McKenna (2012; 2013), Vihvelin (2004; 2013). There are detractors, and below I will point out where their detraction is relevant to my arguments. See, for instance, Arpaly (2002), Weatherson (2019), and Mason (2016).

[17] Vihvelin (2013, p.188) argues that this is both a necessary and sufficient condition on being morally responsible.

to moral standards; they are not, so to speak, in the space of moral reasons, rendering them exempt from certain forms of moral responsibility.[18]

This summary observation can help us formulate the epistemic abilities view, and, at least as a first pass, to motivate it. Let us start with the following claim, which I'll call the "Epistemic Abilities Constraint" (or "**EAC**"):

EAC: An agent S is epistemically rational or irrational for being in doxastic state X if, and only if, S has sufficiently robust epistemic abilities. Otherwise, the agent's doxastic state is epistemically exempt.

There are two pressing questions about **EAC**. First, what sort of *abilities* does **EAC** appeal to? Second, why should we say that doxastic states of agents lacking in sufficiently robust epistemic abilities are *exempt*? This and the next section address these questions, respectively.

Let's start with an answer to the first question. The phrase "sufficiently robust epistemic abilities", as it appears in **EAC**, is purposefully vague, for two reasons.

The first respect in which "sufficiently robust epistemic abilities" is vague is that it is meant to be compatible with a range of views about the nature of epistemic abilities themselves. There are those like Sosa (2007) and Lasonen-Aarnio (forthcoming) who understand epistemic abilities as extrinsic dispositions to know, "anchored" or indexed to certain environments. Comesaña (forthcoming), in contrast, understands epistemic abilities as intrinsic, unanchored dispositions to form rational beliefs. Whether the

---

[18] Even free will skeptics, notably van Inwagen (1983; 2015) and Pereboom (2001; 2014), according to whom moral agents are not morally responsible for their actions, accept that the actions of moral agents can be subject to distinctly moral evaluation; In this thinner sense of 'moral responsibility' that serves skeptics, we can still distinguish moral errors that are the result of normative incompetence from violations of obligations which are the result of poor normative performance.

abilities Jones loses are best understood as extrinsic, anchored dispositions to know, or intrinsic, unanchored dispositions to form rational beliefs, he loses them. Thus, **EAC** and the cases that motivate it have the advantage of remaining neutral on any specific account of the nature of epistemic abilities.

The second is that there is significant debate within the metaphysics of abilities, especially (but not exclusively) as those abilities relate to free will, about the relative explanatory priority of *local* and *global* abilities. So far, I have remained silent on how this difference may matter to epistemic responsibility.[19] On the global reading of abilities, an agent's abilities are determined solely by how the agent is constituted (psychologically, physically, etc.). For instance, Roger Federer has the (global) ability to play tennis, even when he is without his racket snorkeling in the Bahamas, given the facts about how he is intrinsically constituted at that time (which include his training, skills, and so on). On the local reading of abilities, an agent's abilities are determined by how the agent is constituted, plus features of the agent's environment which determine the agent's options. Roger Federer lacks the (local) ability to play tennis when he is without his racket snorkeling in the Bahamas, because facts about his environment preclude his exercising his tennis-playing ability as an option.[20] Federer regains this (local) ability once he gets back on court with a racket.

Recall that the motivating cases for **EAC** were ones in which one agent, Smith, retained his global epistemic abilities, but circumstances conspired to prevent the

---

[19] This terminology is due to Whittle (2010). Vihvelin (2013) uses the terms "narrow ability" and "wide ability", and Mele (2002) uses the terms "general ability" and "specific ability" to pick out a similar distinction to that of "global" and "local" abilities, respectively.
[20] As Lewis (1976) notes, ability attributions often employ the modal verb 'can', as in "Roger Federer can play tennis", and to say that Roger Federer can play tennis is to say his playing tennis is compossible with certain facts, namely facts about Federer's physiological makeup and training. But what is compossible with one set of facts may not be compossible with another, such as the set of facts that include Roger's currently snorkeling without his racket in the Bahamas.

proper exercise of his local epistemic abilities, and another agent, Jones, lost both their local and global epistemic abilities. It is natural to suggest on these grounds that "sufficiently robust epistemic abilities" means something like "sufficiently robust *global* epistemic abilities", but the relationship between local and global abilities is contentious[21], and I intend **EAC** to be neutral with respect to that relationship, the details of which should be settled in the metaphysics of abilities *per se*, not here. Still, without attempting to settle the question of whether local or global epistemic abilities have a kind of explanatory priority over the other, it is nevertheless extremely plausible that an agent like Jones, who lacks *both* global and local epistemic abilities, fails to meet conditions necessary for his doxastic states being rational or irrational.

This may feel like I'm dodging an important question, but the local/global ability distinction is much less important than the masked/finked ability distinction, at least for purposes of explicating and motivating **EAC**. Allow me to explain.

For exposition and simplicity, let's speak in terms of "dispositions" rather than "abilities" *per se*. In those terms, masks prevent the manifestation of a disposition in the presence of its stimulus conditions without removing the underlying bases for that disposition. As a paradigmatic example: a fragile glass is disposed to break when struck, but if a fragile glass is packed in bubble-wrap, its disposition to break when struck is masked, not lost entirely. After all, the underlying structural and molecular features of glass, the ones that account for its fragility, do not change in the presence of bubble-wrap, and removing the bubble-wrap is all one needs to do to "reconnect" the stimulus conditions of fragility (dropping the vase, say) with the manifestation conditions of fragility (the vase shattering).

Finks are a different beast. Finks, unlike masks, prevent the manifestation of a disposition *by removing the bases for that disposition* in the presence of its stimulus

---

[21] See, e.g. Maier (2013)

conditions. For instance, Martin (1994) introduced the famous case of an "electro-fink", which, when attached to a dead wire, makes the wire live at the instant it is touched by a conductor. A dead wire is disposed *not* to send electricity (its manifestation condition) through the conductor when touched (its stimulus condition), but a dead wire hooked up to an electro-fink, when touched, loses the bases for this disposition and electricity flows from wire to conductor.

Now, it is a small step from talk of "dispositions" to talk of "abilities", but it is a step, nonetheless. I'm assuming here that abilities can be helpfully explained in terms of—if not also reduced to—various dispositions to act and think, and to the extent that that's correct, the lessons of masking and finking carry over straightforwardly. (While the reductive claim is controversial, the idea that there is an important explanatory connection between abilities and dispositions is not.) Cases of *masked epistemic abilities*, as far as **EAC** is concerned, are cases in which the agent retains her abilities and so remains the appropriate target of epistemic evaluation; masks do not operate by removing the underlying bases of the ability. In contrast, cases of *finked epistemic abilities* are cases in which the agent's abilities are lost, and so **EAC** implies that the agent is not the appropriate target of epistemic evaluation, at least on matters concerning her finked abilities. After all, finks work precisely by removing the underlying bases of an ability. Questions about the relative priority of local and global abilities are simply orthogonal to the point about masking and finking.

With those points of clarification in mind, it may help to see how **EAC** interacts with recent thought experiments presented in Schaffer (2010) and Williamson (forthcoming). Schaffer's "de-basing demon" operates by "doom[ing] the basing stage [of belief] by forcing you to guess but disguises the belief stage so that your belief seems properly based" (234). Such a demon systemically and improperly rewires the bases for an agent's beliefs, while making it seem to the agent as though their beliefs are properly based. And Williamson's "brain scrambler" is a device that "emits waves of some sort

with a selective scrambling effect on brains. The waves inflict no permanent damage, and do not even change what 'programme' the brain is running, but they occasionally alter the contents of unconscious short-term working memory, so that some computations produce incorrect results." Both authors suggest that victims of such manipulation are irrational.

Williamson's brain scrambler is most plausibly a mask of the agent's local epistemic abilities; again, masks prevent the manifestation of a disposition in the presence of its stimulus conditions without removing the underlying bases for that disposition. The brain scrambler case, then, is just a vivid way to arrive at the same verdict as we did about Smith (case 1, section 1); the brain scrambled agent retains her epistemic abilities, but the proper manifestation of those abilities is somehow impeded, resulting in (perhaps excusable[22]) irrationality. **EAC** thus agrees with the Williamsonians' line on brain-scrambling, for whatever that's worth.

Schaffer's demon, on the other hand, might be thought of as either a mask or a fink of the agents' epistemic abilities. Evil demon scenarios have their dialectical force only on the assumption that agents in evil demon scenarios are in some sense counterparts to agents in "normal" scenarios; at the very least, whatever epistemic abilities normal agents have, agents in evil demon scenarios have too. On this

---

[22] This is Williamson's (forthcoming) suggestion, on the grounds that the brain scrambled agent satisfies a tertiary norm of belief, that of believing as someone disposed to know would believe in her situation. I accept Williamson's characterization of the case as one of excusable irrationality, but I believe that Williamson's argument for this conclusion commits the conditional fallacy (Shope 1978). The idea is that one might be suspicious about whether these sorts of brain scrambled agents are *believing as someone disposed to know would believe*. For instance, it might be that brain scrambled agents are no longer disposed to know, so that the antecedent of the subjective conditional never obtains. In that case, grounding the brain scrambled agent's excusable irrationality by way of the truth of that subjective conditional looks unmotivated. At the very least, it is an open question whether having one's brain scrambled could leave one's knowledge-conducive dispositions intact, one which I cannot hope to address here. However that shakes out, we needn't agree with Williamson *for Williamson's reasons*.

assumption, Schaffer presents yet another vivid case of epistemic irrationality, wherein the demon acts to prevent the agents from properly manifesting their epistemic abilities (in particular, the ability to base their beliefs on their evidence) without stripping them of those abilities.

On the other hand, if we jettison the assumption that the victim of this evil demon has *the very same epistemic abilities* as her normal counterpart, it seems like we should say that part of the demon's demonizing consists or results in removing certain of the victim's epistemic abilities. In this way, we might view Schaffer's demon as a fink, in particular of the victim's abilities to base her beliefs on her evidence and to revise her beliefs in light of new evidence. Whenever the conditions under which the victim would normally exercise those epistemic abilities obtain (surely these conditions involve *acquiring new evidence*), the demon goes to work crossing wires, stripping away the agent's epistemic abilities in the conditions that prompt their manifestation. If this is the right way to understand Schaffer's case, then the doxastic states of agents deceived by a finkish de-basing demon would, then, be in error only because of her unfortunate epistemic incompetence. *Pace* Schaffer, **EAC** says that agents deceived by this sort of finkish demon are not epistemically irrational (not even blamelessly). Instead, because the demon robs them of certain epistemic abilities, they are epistemically exempt.

3. Why exemptions?

Why should we say that doxastic states of agents lacking in sufficiently robust epistemic abilities are *exempt*, as I have been suggesting? Why not say, to put a spin on Aquinas, that *excusabilitas est multiplex*?[23]

---

[23] This is probably correct, but one should not be tempted to think, by implication, that allowing more complexity into one's theory of excuse will obviate the need for exemptions. In fact, I have argued elsewhere that it is both true that excuses are more complicated than extant theories allow, and that one needs exemptions anyway. See *The Quality of Thought* for further discussion.

For an answer, we can look for cues from accounts of moral responsibility. Some authors, like Wallace (1994), argue that morally blaming or excusing someone presupposes that that agent possesses the normative competences to recognize and respond to moral reasons. It is only because the agent has certain global moral abilities that we can *fairly* demand of them that their actions fit the moral reasons, even if circumstances conspire to prevent the agent from manifesting that ability, so that we are inclined to excuse them from blame (perhaps the agent is severely depressed, and this impedes her ability to respond to the moral reasons she recognizes). That is, the possession of certain global moral abilities is a necessary condition on being held morally responsible, on pain of morality being unfair. On this account, for instance, it is precisely because a small child cannot yet recognize and react to moral reasons, or perhaps cannot do so very well, that it would be unfair to demand of them that their actions fit those reasons. As the child goes on to develop their moral abilities, they become a full-fledged moral agent, at which point the fair demands on their conduct change.

This argument seems to apply no less to epistemic reasons than to moral ones; it is only because the agent has certain global epistemic abilities that we can fairly demand of them that their doxastic states fit their epistemic reasons, even if circumstances conspire to impede their local epistemic abilities. Accordingly, the possession of sufficiently robust global epistemic abilities is a precondition on an agent's doxastic states being rational or irrational, on pain of epistemology being unfair, and the doxastic states of agents who lack such abilities are exempt from rational evaluation.[24]

---

[24] Boult (2019) offers a position that interacts fruitfully with what I claim here; he argues that small children, for instance, count as "prospective epistemic agents", which we hold epistemically responsible "as a kind of heuristic or method aimed at turning them into adult human epistemic agents" (156). This position is entirely consistent with my own view; one could accept **EAC**, which implies that the doxastic states of insufficiently competent agents are

Of course, accepting this argument would involve accepting a substantive position on the relationship between global and local epistemic abilities, one according to which global abilities are explanatorily prior to local ones. Moreover, 'fairness' is a moral notion, so an argument from premises concerning fairness to the conclusion that the doxastic states of epistemically incompetent agents are epistemically exempt is a moral argument.[25] While these features of the argument are not in and of themselves objectionable, I wish to remain neutral on both counts. In particular, many epistemologists will be suspicious of moral encroachment, or whether the moral features of a situation can affect whether an agent's doxastic states are (ir)rational. Moreover, the arguments from fairness might be thought to rely on the intuitive idea that it is unfair to *morally blame* someone who is morally incompetent, where this involves distinctively *moral emotions* like resentment or indignation. But it is less intuitive—at least without further argument—that there are epistemic analogs of blame which might play the corresponding roles in the argument above.[26]

---

not—strictly speaking—irrational, while also insisting that there is value—perhaps practical value—to holding such agents epistemically responsible (or at least the subset of them that have the potential to acquire more robust epistemic abilities), with the aim of making them sufficiently competent.

[25] Whether fairness can itself be understood in terms of more general, non-moral normative considerations is a question I leave open. For instance, fairness might be grounded in some form of reasonableness or reasons-responsiveness. If it can be, then these arguments are not (at least straightforwardly) moral encroachment arguments, and all the better for me.

[26] It is worth noting, as an anonymous referee points out, that these arguments are not probative in all corners of the moral responsibility literature, and this puts pressure on my analogical reasoning. There are those, like Arpaly (2002) and Weatherson (2019), whose "right reasons" positions make no mention of capacities at all; what various normative domains require is that agents respond to the right reasons (in each domain), or be sensitive to the right-making features (in each domain). These are powerful and attractive theories, but they do not rid themselves of the problem of explaining what to make of agents who *cannot track the right-making features* in question. This is just a more general version of The Problem of Epistemic Competence. In section 3, I suggest that these "right reasons" views may have the right tools to think about excuses *per se*, rather than about responsibility writ large.

And there are those like Mason (2016) who advocate for positions that admit of an

We can appeal to other, non-moral considerations of when it is intelligible to hold an agent accountable to certain standards to arrive at a similar conclusion. For instance, Smith (2012), building off a view in Scanlon (2008), argues that an agent is morally responsible for X-ing only if X-ing "bears a rational connection to the agent's evaluative judgments… [such that] the agent is open, in principle, to demands for justification regarding [X]" (577). In order for an agent's being open in principle to demands for justification, the agent's behavior must be connected to their evaluative judgments in a way that renders such demands "intelligible" (578). For an agent lacking in sufficient moral ability, in the first place it would not be *unfair* but *unintelligible* to hold the agent answerable to moral standards.[27]

This non-moralized criterion for holding agents morally accountable applies no less as a criterion for holding agents epistemically accountable; for an agent lacking in sufficient epistemic ability, it would likewise be unintelligible to hold the agent answerable to epistemic evaluation. Agents lacking sufficient epistemic abilities are such that holding their doxastic states answerable to epistemic standards—in part, by judging that they *ought* to have adopted some other doxastic state, or that they are

---

abilities condition, but deny that it plays this role of delimiting the exempt from the non-exempt. Mason's view is thoughtful and articulate, but I do not see it, ultimately, as conflicting with EAC (or, more precisely, with the picture according to which one's abilities determine whether one is exempt from responsibility). Her claim is that moral exemption tracks a special sort of moral ignorance, but on some popular meta-epistemological views around the nature of competence, competence consists in some or another propositional knowledge state, perhaps under a special mode of presentation. These so-called "intellectualist" views of competence, which I think are independently plausible, can accept her point about moral ignorance, while also maintaining that this ignorance constitutes a lack of competence. For articulation and defense of an intellectualist view of this kind, see Stanley (2011).

[27] Watson (2011) suggests a similar criterion of "unreachability" for demarcating exemption. Even on these intelligibility or reachability standards, incompetent agents may be *evaluated* without being held responsible in any way (even by being excused); morally incompetent agents may do morally bad things (cause needless harm, say), and epistemically incompetent agents may do epistemically bad things (make baseless judgments, say), but neither has done something for which they are responsible.

irrational for failing to fit their beliefs to the states these standards recommend—is unintelligible. *A fortiori*, the idea that such agents might be *excusably irrational* is unintelligible. Moreover, we needn't endorse any particular first-order theory of justification to recognize our collective need to appeal to epistemic exemptions; finally, something that we can all agree on!

Of course, that's too glib. There is, to be sure, a great deal to say that must be left unsaid, at least in this paper. I have not, for instance, explored the fruitful connections between my discussion of epistemic exemptions in terms of epistemic abilities, on the one hand, and theories of cognitive disability, on the other, and because of this sin of omission some readers may worry that my motivating examples are somewhat *unhinged*; for instance, Smith's and Jones's stories (section 1) are highly stylized and idealized examples, and Jones's borders on science fiction.

But this should not give readers the impression my epistemology is tethered to intuitions about recherche cases. Instead, the idealization (if not the stylization too) is meant to abstract away from the range of events or circumstances that would put an agent like us into a position relevantly similar to Smith's or Jones's. There are, for instance, a slew of everyday examples that the epistemic abilities view predicts easily by emphasizing certain dimensions of similarity to Jones's predicament: those born with certain severe cognitive limitations, or who are abused or systematically misled as their capacities to reason develop, or who suffer certain traumatic brain injuries later in life, are not irrational—not even *blamelessly*—for certain of their epistemic errors. Agents constitutively unable to recognize their own wishful thinking as such—as might happen in cases of severe brain damage—are exempt from rational evaluation in those judgments that, unbeknownst to the agent herself, manifest wishful thinking. I hope that these predictions about exemptions in extreme cases are not especially contentious, but I recognize that saying only this much constitutes a promissory note for a fuller theory.

4. A final upshot

My aim in this paper has been, in the first place, to motivate the epistemic abilities view and to situate it within a number of contemporary debates in epistemology. More generally, I hope to have made the case that epistemologists must attend to a largely neglected normative category, that of epistemic exemption. Once we admit of epistemic exemptions in terms of epistemic abilities, we can not only explain the interesting normative differences between the doxastic states of (rather unfortunate agents like) Smith and Jones, but also see epistemic responsibility as part of a broader normative framework that offers explanations of obligation, excuse, and exemption in terms of normative abilities.

I will conclude by pointing to an important upshot of the epistemic abilities view for certain foundational meta-epistemological debates. To fully explore and justify this implication would require another paper altogether, so I mention it here as a point for further research.

It is tempting to think that there are some epistemic principles *so basic* that whenever an agent's doxastic states violate these principles, the agent is thereby irrational. Some epistemologists have suggested certain logical (or quasi-logical) principles as candidates for basic epistemic principles. For instance, Hartry Field (2009) suggests that if one's evidence E "obviously entails" a proposition H, one shouldn't believe E without believing H. Call this the "obvious entailment principle". Agents who violate the obvious entailment principle are at best *excusably* irrational. On one reading of Field's view, what counts as an obvious logical entailment will depend on what is in fact the correct logic.[28] But this view—or this reading of it—cannot distinguish between

---

[28] Field says, "I've been deferring a worry: what counts as obvious? In my view, there is no general answer to this, it depends on both who is being assessed and who is doing the assessing; but this is not obviously a problem for using the notion in describing normative

violations of the obvious entailment principle due to epistemic incompetence, and those due to poor epistemic performance; both are deemed irrational in virtue of violating the obvious entailment principle. This fails to acknowledge the role that an agent's epistemic abilities play in determining whether the agent is an appropriate or intelligible object of epistemic evaluations in the first place—that is, this fails to answer The Problem of Epistemic Competence.

Bayesian epistemologists are united in accepting that an agent's degrees of belief must be probabilistically coherent, on pain of being irrational. Call this "Probabilism". Probabilists similarly cannot distinguish between probabilistically incoherent degrees of belief that are due to epistemic incompetence, and those that are due to poor epistemic performance.[29]

Not everyone agrees that principles of logic or probability are candidate basic epistemic principles,[30] but even less contentious candidates for basic epistemic

---

requirements, for normative requirements are relative in both these ways" (259). One way of understanding 'obvious entailment', which I take issue with above, appeals to *the correct logic*, encapsulated in his principle "D*(alt)" (261-2). Another way of understanding 'obvious entailment' appeals only to the agent and that agent's norms of thought (262). For all I have argued, this second way of understanding Field's view may be consistent with the thesis advanced in this paper, and I have provided reason to reject the first way of understanding Field's view.

[29] The epistemic abilities view of course predicts this, and it is important to note a point of empirical support. In Tversky and Kahneman's (1983) famous work on the conjunction fallacy, they suggest that competence with the conjunction rule consists at least in the subject recognizing that violating the conjunction rule is a decisive reason to overrule the representativeness heuristic. Naïve subjects' naïveté consists, perhaps among other things, in their failure to "fully understand" certain principles of probability and the conditions under which those principles "prevail" over conflicting ones (300). Violating the conjunction rule is fallacious precisely because, in assigning probabilities, the axioms of probability prevail over conflicting, representativeness heuristics. But if a naïve respondent's violation of the conjunction rule is analogous to a child's proconservation reasoning about volume, their failure is surely not *irrational*; we generally do not think that children are irrational for their probabilistic ignorance, nor should we think that probabilistically naïve adult subjects are irrational for their flouting a probabilistic rule.

[30] Harman (1986)

principles will suffer the same problem. For instance, it is tempting to think that among the basic epistemic principles, if there are any, is the following:

> **Induction principle**: If one observes sufficiently many Fs that are G, then one has a (prima facie) reason to believe that the next observed F will be G.

Even the induction principle is non-basic, in the sense that not all agents whose doxastic states violate the induction principle will thereby be irrational. Instead, what is properly called "basic" about the induction principle (or Probabilism, or the obvious entailment principle) is that, of agents who are sufficiently epistemically competent, if their doxastic states violate the induction principle, those states will thereby be irrational, perhaps excusably so.[31]

Chapter 2:

Evidentialism and the problem of basic competence

According to evidentialists[32] about inferential justification, an agent's evidence—and only her evidence—determines which inferences she would be justified in making, whether or not she in fact makes them. But there seem to be cases in which two agents would be justified in making different inferences from a shared body of evidence, merely in virtue of the different competences those agents possess. These sorts of cases suggest that evidence does not have the pride of place afforded to it by evidentialists; competence seems to play at least as important a role as evidence in explaining which inferences an agent would be justified in making.[33]

In this paper, I consider how two versions of evidentialism about inferential justification might try to account for the role of competence in inference, and I present problems specific to each version. I end by sketching and briefly defending an alternative to these evidentialist views, "inferential dogmatism".[34] While dogmatic views have gotten some attention in debates around non-inferential justification, they

---

[32] "Evidentialism" is, of course, an umbrella term for a somewhat heterogeneous collection of views. In the next section, I spend some time clarifying exactly what sort of evidentialism is at issue.

[33] As I understand him, Ned Hall (1994) uses the terms "analyst expertise" and "database expertise" (resp.) to track largely the same distinction. What's more, something like this distinction is implicit in the so-called "anti-intellectualist" tradition of thinking that knowledge-how is not fully reducible to knowledge-that. See, e.g., Ryle (1949), and Stanley (2011) for an intellectualist rebuttal (although the latter position is plausibly consistent with a non-reductive account of analyst expertise).

[34] One of the most forceful challenges to basic justification views is the so-called "easy knowledge problem" (Cohen (2002; 2005; 2010); Vogel (2000); White (2006); Weisberg (2012)); I respond to this challenge in section 4, drawing on work by Alston (1980; 1983), Cohen (2010), Pryor (2000), and Wedgewood (2013).

have largely been ignored in debates around inferential justification and are to that extent novel.

1. Evidentialism and competence

Evidentialists form a fairly heterogeneous group, and I cannot discuss every view falling under that heading. For our purposes, we can simplify the terrain a bit and think of all evidentialists as committed to (at least) the following principle:

> **E:** An agent A's evidence at a time, t, determines which inferences are justified for A at t.

As stated, this principle requires two points of clarification. First, **E** is meant to be a principle of *propositional* (or *ex ante*) justification, rather than a principle of *doxastic* (*ex post*) justification.

According to a standard way of understanding that distinction, in order for A to be doxastically justified in adopting doxastic attitude D, A must (i) be propositionally justified in adopting D, and A must (ii) adopt D on the basis of that which propositionally justifies it. Notice that condition (ii) of doxastic justification requires that an agent *properly base* her doxastic attitudes on her evidence, which already requires more of her than *merely having* that evidence. I will simply proceed as if the standard way of drawing the propositional-doxastic distinction is correct, in which case **E** is much more attractive and plausible as a principle of propositional — rather than doxastic — justification. In less technical terms, **E** concerns which inferences an agent would be justified in making, whether or not she makes them.

Second, this evidentialist principle is fairly restrictive in one sense: **E** is meant to apply only to cases of *inferential* justification, not to cases of non-inferential justification. An evidentialist principle more general than **E** might concern the relationship between

one's evidence and one's *total doxastic state*, but this more general principle would expose itself to problems specific to debates around the nature of non-inferential justification.[35] (Some of the attitudes in one's total doxastic state, perhaps one's simple perceptual beliefs, will be non-inferentially justified.) It is enough for our purposes to note that **E**, restricted as it is to cases of inferential propositional justification, has some initial plausibility.

With those qualifications laid out, the target of this paper is someone who accepts **E**, to whom I'll refer as an "inferential evidentialist". The general complaint I want to lodge against inferential evidentialists is that they are forced to accept an implausible account of an agent's inferential competences. The idea, laid out more fully below, is fairly straightforward: if there are cases in which possessed inferential competence makes a difference to which inferences an agent would be justified in making, that competence is just more evidence. That follows from a more general commitment of inferential evidentialism: *anything* that makes a difference to which inferences an agent would be justified in making is just more evidence, since evidence is the only thing that makes such a difference.

Before I launch into a criticism of this view, I want to spend some time fleshing out the kind of explanation it provides. To that end, I'll discuss an example that is meant only to be illustrative, not critical.

---

[35] For instance, many debates around evidentialism in the philosophy of perception turn on whether one understands *experiences themselves* as evidence, or as *providers of* evidence. According to a view of the first kind, variously held by Conee and Feldman (1985), Pollock (1971; 1987), and Pryor (2000), perceptual beliefs are quasi-inferential, with experiences playing the role of premises. According to another view, attributable to philosophers like Williamson (2000) and Comesaña (2020), evidentialism is false precisely because experiences are a non-evidential basis for perceptual knowledge.

Birdwatching

Ed is an expert ornithologist; among other things, he knows how to recognize the species of a bird by the sounds of its songs. Neil is a novice birdwatcher; while he is not totally incompetent, Ed is his superior across most of the relevant dimensions. Suppose that Ed and Neil each hear a certain birdsong out in the brush, one which both men recognize as sounding like the song of a predatory bird. While Neil is quick to infer that the two hear a predatory bird, Ed is not. Instead, Ed's makes a more modest inference: that the two either hear a predatory bird or hear an avian mimic.

Birdwatching is a scenario in which an expert and a novice share a body of evidence but arrive at different conclusions on its basis. Moreover, the expert and novice are each rational to arrive at different conclusions.

Just in case that characterization of the cases raised any eyebrows, here's a fairly intuitive and (I hope) innocent way to think of what's going on. If Neil is simply unaware of the possibility of avian mimics in the area, or if he had never considered such things as avian mimics, it would not make sense for him to infer, on the basis of hearing such a birdsong, that it might have come from an avian mimic. Since it would make no sense for him to consider the possibility of avian mimics, it makes no sense to fault him for ignoring it; most of us reject a conception of rationality that requires agents be sensitive to *the evidence that there is*, rather than to *the evidence that they have*.[36] In contrast, part of what makes Ed the expert of the two is that he has acquired, through years of training, a certain competence in distinguishing bird species by their song, one which is fairly sensitive to the possibility of avian mimics. It's natural to think of what's

---

[36] Moreover, this does not seem to be a case in which Neil *should have had* certain background evidence that he in fact lacks.

going on here as a situation in which Ed's competence makes a difference to which inferences he's justified in making. In particular, Ed's competence justifies him in making a more cautious inference than Neil, in virtue of making him sensitive to a possibility about which Neil is unaware.[37]

It's worth emphasizing that **Birdwatching** is not a recherche thought experiment meant for the only the most hardened epistemological intuitions; instead, it is meant to be a stylization of a totally ordinary and familiar situation, one that *any* plausible epistemology of inference should be able to explain.[38] And, fortunately, inferential evidentialists can easily explain it.

In rough outline, the inferential evidentialist explanation of **birdwatching** and related cases goes like this. Differences in competence can give rise to differences in propositional inferential justification, but two agents cannot differ in competence while having *exactly* the same evidence. Instead, the expert and novice share only *some* evidence; in the case of Neil and Ed, their shared evidence might be something like *that the birdsong sounds thus-and-so.* But the expert has a great deal more evidence, perhaps only in the background, that the novice lacks. Ed, for example, plausibly knows that *if a birdsong sounds thus-and-so, it comes from either a predatory bird or avian mimic*, while Neil does not. It is the expert's possessing certain background evidence that the novice lacks that explains why their inferences could rationally diverge.

Here is what I hope to draw out of the discussion so far: the story that inferential evidentialists must tell about competence is a reductive one. If there are cases in which

---

[37] I am supposing that neither Ed nor Neil has prior information about the preponderance of avian mimics in the area; this is not a case where Ed (but not Neil) already knows he is likely to run into some.

[38] Generally, authors do not deny that epistemic competence can ever make a normative difference, in light of the fact that it arguably does not make an evidential difference, plus one's commitment to **E**. As a purely sociological fact, evidentialists instead opt for the reductive account of criticized here—this is explicit in Conee and Feldman (1985). Others, like Sosa (2007; 2011), offer a virtue theoretic alternative to evidentialist accounts.

competence makes a difference to which inferences an agent would be justified in making, competence is just more evidence. That follows from a more general commitment of inferential evidentialism: *anything* that makes a difference to which inferences an agent would be justified in making is just more evidence, since evidence is the only thing that makes such a difference.

Given that inferential evidentialists have to accept a reductive account of competence in terms of evidence, one might reasonably wonder: if competence just is evidence, which evidence is it? Presumably, the most plausible thing to say on behalf of inferential evidentialists is that this competence-constituting evidence must in some way *link* an agent's first-order evidence to various hypotheses. Call this the "linking principle":

> **Linking principle**: For an agent to be justified in inferring H from first-order evidence E, given her competence, her competence-constituting evidence must link E to H.

The linking principle is meant to be somewhat of a precisification of, and somewhat of a substantive constraint on, the reductive view of competence. Of course, the language of "linking" leaves room for ambiguity. As I explain in the next two sections, how we think of the "linking" characteristic of competence-constituting evidence will depend on the version of evidentialism in question. But as a rough and ready pass, competence-constituting evidence could perform this linking function if we thought of it as consisting of conditionals, the antecedent of which is a body of first-order evidence, and the consequent of which is a hypothesis supported by that evidence. (We will consider other ways to flesh out "linking" besides this one.) And if we accept that competence consists in nothing more than having these conditionals as evidence, we could straightforwardly explain how a difference in competence between two agents could

result in divergent rational inferences from shared first-order evidence. For instance, in **Birdwatching**, Ed is justified in making a different, more cautious inference than Neil because Ed's competence-constituting evidence (including, perhaps among many other things, the conditional *if a birdsong sounds thus-and-so, it comes from either a predatory bird or avian mimic*) provides a link to a different space of hypotheses.

I want to briefly comment on potential objections from two directions. First, some readers might think that inferential competence, whatever it is, is only relevant to questions of doxastic justification. Presumably, the motivation for this position stems from the idea that *manifesting inferential competence* is a way of making properly based inferences, and that rings true in my ears. But we should not be tempted to slide from the idea that *manifesting* inferential competence is exclusively relevant to questions of doxastic justification to the idea that *possessing* inferential competence is exclusively relevant to questions of doxastic justification.[39] In fact, in the **Birdwatching** case above, we have already shown how possessing inferential competence is relevant to propositional justification: inferential competence is sometimes constituted by relevant background evidence. As such, the problem with this objection is that it goes too far; it blocks my criticisms at the cost of undermining the very plausible story that evidentialists already tell about cases like **Birdwatching**. A working assumption of this paper is that there is some initial plausibility to the idea that inferential competence performs a kind of theoretical double-duty: its manifestation is relevant to doxastic

---

[39] Authors like Turri (2010) have argued, on entirely different grounds, that in order for an agent to be justified in believing some proposition P, she must have some means available the employment of which would result in her justifiably believing P. On this sort of view, possessing certain inferential competences is obviously relevant to propositional justification, since manifesting such competence is very plausibly a means of coming to justifiably infer. This does, however, invert the orthodox picture of the relationship between propositional and doxastic justification, and I have not taken on that contentious commitment.

justification, while its possession is relevant to propositional justification. The interesting question, then, is not *whether* but *how* inferential competence is relevant.

Second, some readers might be inclined to explain the relationship between inferential competence and inferential justification by appeal solely features of *non-inferential* justification, not by way of the **linking principle**. Authors like Siegel (2010) and Chudnoff (2020), for instance, have defended positions according to which experts and novices sometimes differ in what they are *non-inferentially* justified in believing by differing in the contents of their experiences, and these differences in non-inferential justification "bubble up" to give rise to differences in inferential justification. Somewhat roughly put, experts sometimes "just see" things differently than non-experts, and this difference explains why experts and non-experts can differ in which inferences they would be justified in making. For instance, there may be cases where a radiologist "just sees" a compound fracture in the x-ray, whereas the patient might "just see" a broken bone. In such a situation, the radiologist would, having different non-inferentially justified perceptual beliefs, much more plausibly be in a position to settle the question of whether surgical intervention would be required.

But even if we accept these sorts of "cognitive penetration"-inspired views of non-inferentially justified perceptual belief and its connection to inferential justification, there is a residual question—the one that we are concerned with here—about what *further* differences might explain why sometimes experts and non-experts who really do share a body of evidence differ in what they could justifiably infer from that evidence. It strikes me as implausible that *all* interesting questions about inferential justification bottom out in a theory of non-inferential justification, and to deny this much seems to simply beg the question. And anyway, it is my hope that exploring these issues will supplement existing views about the relationship between expertise and non-inferential justification, rather than compete with them.

Having addressed some early concerns, let's briefly take stock. I have presented a core inferential evidentialist thesis, **E**, and I showed how a commitment to **E** brings with it a reductive view of inferential competence in terms of evidence. I then offered the **linking principle** as a plausible way to characterize *which* evidence constitutes one's competence, according to the reductive view. I have acknowledged a few reasons to doubt how I've set up the problem: perhaps one is skeptical about the role of competence in propositional inferential justification, or perhaps one thinks that all interesting questions about propositional inferential justification bottom out in appeals to non-inferential justification. Both of these concerns overreach.

In the next section, I'll look at cases that are difficult for the reductive view of competence to explain. The difficulty, I suggest, is due to the fact that the reductive view is only suited to explain *acquired* competences, and not all competences are acquired.

2. Propositionalism and difficult cases

There are two different ways of developing evidentialism: one can take a propositionalist or a non-propositionalist view of evidence. Both of these views, I will argue, have troubles adequately explaining the role of competence in inferential justification. In this section, I focus on propositionalist versions of evidentialism.

According to "propositionalism", the following claims are true (in addition to **E**):

- Evidence is, by its nature, propositional.
- Having evidence is a matter of standing in some special epistemic relation to the propositions which constitute one's evidence.

The first commitment distinguishes propositionalism from non-propositionalism, which I consider in the next section. The second commitment is meant to cast a wide net across

a few battle-lines in epistemology over the nature of the special epistemic relation; one might think that in order to have a proposition P as evidence, an agent must *know* that P, *rationally believe* that P, or *merely believe* that P. I wish to remain neutral on any of these particular characterizations of propositionalism.

Qualifications aside, here is what propositionalist versions of inferential evidentialism might look like: consider Fumerton's (1990) "Principle of Inferential Justification" ("PIJ"), according to which, for an agent to be propositionally justified in inferring H on the basis of E, the agent must (1) be justified in believing E, and (2) be justified in believing that H makes E probable.[40,41] A PIJ-style view satisfies the **linking principle** articulated in the last section; evidence of the form *if E, then H is probable* is well-suited to "link" E and H.

On a view like Fumerton's, inferential competence consists in nothing but justified beliefs in certain conditionals, the antecedent of which is a body of evidence, and the consequent of which is a hypothesis. This sort of view is equipped to explain **Birdwatching** and similar cases: Ed is justified in making a more cautious inference than Neil because Ed justifiably believes, perhaps only in the background, that *if a birdsong sounds thus-and-so, it comes from either a predatory bird or avian mimic*, while Neil does not.

I admit that very many cases of justified inference can be handled within this sort of framework; very many cases of expert inference can be explained by pointing out background commitments, justifiably held, that are characteristically present in experts but absent in novices. But, as I hinted at the end of section 1, this framework is only

---

[40] Also see Hasan (2013); Foley (1993), e.g., chapter 3, section 2.
[41] The proposition *if E, then H is probable*, at least as Fumerton discusses it, seems to be explicitly higher-order and fairly intellectualized; the agent needs to have sufficient mastery of the concept 'probable' to be justified in making *any* inference. But this requirement could be relaxed to require only a *de re* or sub-personal grasp of the linking evidence. Nothing in my arguments depends on settling this issue.

suited to explain *rationally acquired expertise* in inference, and not all expertise is rationally acquired.

To flesh out this claim, start with the idea that some competences are "basic"; they are deployed in the rational acquisition of other (non-basic) competences without themselves needing to be rationally acquired. Short of offering a general account of what distinguishes the basic from the non-basic competences, there are certain clear candidates for *bona fide* basic competence: competence with reasoning by enumerative induction, representativeness, and inference to the best explanation all have some claim to count as basic, if any inferential competence does. For instance, one does not *learn* how to use enumerative induction by employing other rational capacities; one's capacity for induction is a kind of starting point for rational learning.

Of course, if the reader disagrees with me on the particular candidates for basic competence, I encourage them to read the cases with their preferred candidates in mind. If, instead, the reader is suspicious about the existence of basic competences, I hope to show by the end of this section that the cost of denying their existence is rather steep.

To get to that point, it will help to consider a few more cases, ones which focus on (putatively) basic inferential competences. Each case presented below involves an agent getting evidence by observation that the first 99 balls pulled from an opaque urn known to contain 100 balls are all black. Call this evidence "O". Antecedently ignorant of the colors of any of the balls in the urn, each agent considers whether the color of the 100th ball is black. Call the proposition that the 100th ball is black "B". Here, then, are the cases:

> **Case 1**: Larry gets evidence O, and then, taking some time to reflect on his
> observations, justifiably believes that *if O, then probably B*. Being fully aware of
> the colors and quantity of the balls so far observed, Larry's confidence that B

increases; suppose he began inquiry by suspending judgment on B, and he now believes that B.

Case 2: Moe gets evidence O, but he has also been presented strong but misleading evidence by his most trusted colleagues that *it is not the case that if O, then probably B*, and he believes the reports of his colleagues.[42] Being fully aware of the colors and quantity of the balls so far observed, Moe's confidence that B nevertheless increases; suppose he began inquiry by suspending judgment on B, and he now believes that B.

In case 1, Larry is meant to be a paragon of rationality for propositionalists; he satisfies PIJ and so would be all-things-considered rational to infer B from O. Moreover, Larry seems to possess whatever background evidence constitutes competence with enumerative induction: in this case, his competence would consist in his having as evidence conditionals of the form *if O, then probably B*. Let's grant this for the sake of argument.[43]

In case 2, Moe fares very badly by the lights of propositionalism. Not only does Moe infer B from O *despite* his (misleading) linking evidence, it seems that he does not have the right kind of linking evidence to count as fully competent with enumerative induction in the first place. (We can imagine, in Moe's case, that the right kind of

---

[42] Moe's testimonial evidence is *not* that there is at least one non-black ball in the urn. That is, he does not get, as some additional first-order evidence, the proposition that not-B; rather, his testimonial evidence is "higher-order", supporting the negation of the conditional *if O, then probably B*.

[43] It is worth noting, however, that one could also read Case 1 as one in which Larry *lacks evidence altogether* about relationship between O and B, as opposed to having evidence—even implicitly or sub-personally—that O makes B probable. If we were inclined to read case 1 in this way, Larry would be a case of competence in the absence of evidence, which would itself be problematic for a PIJ-style view.

linking evidence has been defeated by his colleagues' testimony.) Moe is, at worst, *both* irrational for failing to respect his linking evidence *and* less-than-fully competent with enumerative induction, or, at best, neither rational nor irrational because totally incompetent. This strikes me as entirely the wrong result; not only is Moe *not* incompetent with enumerative induction, he does something remarkable.

To claim that Moe does something remarkable is *not* to claim that Moe is all-things-considered rational or fully epistemically justified to believe that B, at least if that is taken to mean that Moe is believing as an ideal epistemic agent would believe. Among other things, ideal epistemic agents are perfectly coherent, and Moe is plainly not. But here is something to say in favor of Moe's belief: his first-order evidence in fact supports it, and he comes to believe it by disregarding higher-order evidence that is in fact misleading. At the very least, Moe's inference does not wear its irrationality on its sleeve; it is not a paradigm case of all-things-considered irrationality.

To get a sense of how Moe's inference might be afforded some *positive* epistemic status, compare Moe's predicament to that of Huckleberry Finn: Huck is in a position of deciding whether to turn Jim in to slave catchers, and he feels deeply conflicted about doing so. On the one hand, Huck seems to believe that morality requires that he turn Jim in. After all, Jim is a fugitive slave, and Huck has been raised to host a range of negative beliefs about slaves in general, and fugitive slaves in particular. On the other hand, Huck cannot bring himself to do it. Jim is also Huck's friend. *Something* drives Huck to free Jim, but it is certainly not Huck's "better judgment" about what morality requires of him. On pain of classifying Huck as a moral monster, or a merely accidental do-gooder, many authors in the moral responsibility literature have attempted to accommodate "inadvertent moral virtue" into their accounts of morally worthy action. Perhaps, as Arpaly (2002) and Arpaly and Schroeder (2013) have suggested, the fact that Huck is moved by some deep-seated, intrinsic concern for the right-making features of action (in this case, Jim's humanity) explains why Huck is virtuous, albeit inadvertently.

Put in those terms, Moe's predicament looks very much like an instance of a related phenomenon, one that is sometimes called "inadvertent epistemic virtue".[44] Cases of inadvertent epistemic virtue are instances of rationally permissible *akrasia*, a particular kind of inner-conflict between one's first- and higher-order beliefs. Brian Weatherson (2019) discusses a case of inadvertent epistemic virtue in which a testimonial skeptic, Aki, comes to believe something on the basis of her friend's testimony. (Testimonial skepticism is the view that one cannot gain knowledge by testimony.) Aki, hearing her friend testify that the Tigers won the night before, comes to believe that the Tigers won the night before, and in so doing ignores her belief in testimonial skepticism. (This belief is firmly held in the seminar room, we can imagine.) Weatherson is inclined to treat Aki as a "paragon of rationality" (171). He defends that claim on the grounds that Aki's first-order evidence *evidentially screens-off* her judgment about what that evidence supports.[45] But we don't have to go as far as defending Aki as a *paragon* of rationality by appeal to evidential screening-off to recognize some positive epistemic status in Aki's belief, any more than we have to defend Huck as a *paragon* of morality to recognize some positive moral status in his freeing Jim.

Moe, like Aki, infers that B *despite* his "better judgment" –which in this case provides him with evidence linking his first-order observations to various hypotheses— about what counts as evidence for what. Of course, someone might rightly insist that Moe would be "epistemically improved" or "closer to a rational ideal" if he were not merely *inadvertently* epistemically virtuous. Granted, not all cases of *akrasia* are rationally permissible, and perhaps no case of *akrasia* is rationally ideal, but there is a lot

---

[44] See Weatherson (2019) for explicit discussion of inadvertent epistemic virtue, and Arpaly (2002) and Arpaly and Schroeder (2013) for a discussion of inadvertent moral virtue, from which Weatherson draws heavily.
[45] See Weatherson 2019, chapter 11.

of daylight between accepting that coherence between first- and higher-order evidence is a rational ideal, on the one hand, and accepting that less-than-fully-ideal rational agents can sometimes make rational inferences despite this kind of incoherence, on the other.

The point here is not to defend any particular account of inadvertent epistemic virtue, nor to defend the implication of any such account that there are cases of rationally permissible *akrasia*. I accept that implication, and others have defended it extensively.[46] The point here is much more modest: Moe's belief that B seems to have a positive epistemic status, perhaps short of ideal rationality, one that is shared with Aki's testimonial belief, and one that is analogous to the positive moral status of Huck's freeing Jim, an action which itself might fall short of ideal morality. Propositionalism seems to be insensitive to these finer details of epistemic assessment. If Moe is irrational to infer B *and* less-than-fully competent with enumerative induction, it's hard to see what positive epistemic status an inference to from O to B could have for him. There seems to be nothing about Moe that would even *prima facie* justify an inference from O to B.

What about the idea that Moe is totally incompetent? If so, this would surely undermine my claim that he is *prima facie* rational to infer B from O.

One consideration that brings out the implausibility of this idea is that, if we take seriously the idea that competence consists in having certain linking evidence, we will hold competence hostage to arbitrarily higher-order defeaters. This proposal has severe costs; to see why this is so, compare Moe's case to Curly's:

---

[46] See, e.g., Williamson (2011; 2014), Wedgwood (2011), Lasonen-Aarnio (2014), and Weatherson (2019)

Case 3: Curly gets evidence O, and he, like Larry, has come to justifiably believe that if O, then probably B. Curly has, however, been presented with strong but misleading evidence by his most trusted colleagues that it is not the case that *if (O and if O, then probably B), then probably B*, and he believes the reports of his colleagues. Being fully aware of the colors and quantity of the balls so far observed, Curly's confidence that B nevertheless increases; suppose he began inquiry by suspending judgment on B, and he now believes that B.

On one way of reading Curly's case, he has the same first- and second-order evidence as Larry, the paragon of rationality for propositionalists. In that case, Curly would, like Larry, satisfy a PIJ-style constraint, and his inference would be all-things-considered rational. But this would render Curly's third-order evidence doubly normatively irrelevant; that third-order evidence would neither affect which inferences would be rational for Curly to make, given his first-order evidence, nor would it constitute Curly's competence with enumerative induction.

More plausibly, I think, a propositionalist would insist that Curly's second-order linking evidence is defeated by his third-order linking evidence. What they would require in so insisting is a kind of "mesh" between higher-order linking evidence. Unless one's second-order linking evidence appropriately meshes with one's third-order (and so on) linking evidence, one is not genuinely competent in the first place, or so says the propositionalist. But if this is the line one takes, then it is not clear that one can stop at any *finite* level of higher-order evidence; competence would, in some sense, consist of an infinity of "higher-order" evidence linking one's "lower-order" evidence to some hypothesis. It seems plainly wrong that *any* justified inferences (for things like us) are justified by an *infinity* of further, higher-order beliefs about what our evidence supports.

Compare this position to Achilles' predicament in "What the Tortoise said to Achilles".[47] There, the Tortoise asks Achilles to say what must be added to someone who accepts some proposition A, and who accepts that A logically entails B, but who fails to recognize that they must thereby accept B (or reject a premise). Achilles goes through a process of appealing to ever-higher-order premises, linking A, A implies B, and B; the absurdity of this response is apparent. And the dialogue is often taken to show, at the very least, that to follow an inference rule (in this case *modus ponens*) is not to infer by way of an additional (perhaps suppressed) premise, on pain of regress. The moral of Carrol's dialogue generalizes; for *any* inference rule one can create a Carrol-style regress to motivate the thought that to infer competently is not merely to infer by way of an additional premise. The contrast between Larry, Moe, and Curly is merely a case in point.[48]

These sorts of regress considerations, I hope, present a bit of a defense of the cogency of basic competence, even to those who disagree with me about the particular candidates I suggested (enumerative induction, representativeness reasoning, and inference to the best explanation). Unless one appeals to basic inferential competence *somewhere*, one is left with a view according to which *putatively* basic competences consist of an infinitely large mesh of ever-higher-order linking evidence. Even if, as I've granted in section 1, there are plenty of everyday cases in which a particular inferential competence may consist in nothing more than the possession of certain linking evidence, this concession should not be taken to indicate that, in general, to be

---

[47] Carrol (1895)

[48] In some ways, I agree with the upshots of Fumerton (1990); if one accepts PIJ (or a plausible weakening of it as we've done here), then one faces either (i) an infinite justificatory regress, or (ii) foundationalism about inferential justification. Finding both options wanting, Fumerton suggests a kind of meta-epistemological skepticism about inferential justification more generally. In section 4, I will argue for option (ii), a version of foundationalism about inferential justification.

competent is to possess the right sort of evidence; it should certainly not lead us to think that, rather implausibly, that inferential competence consists of evidence "all the way down."

Despite its problems, the *initial* appeal of propositionalism might derive from the fact that following one's evidence is an ideal of rationality, and someone who ignores any of her evidence falls short of this ideal, whether or not the evidence she ignores turns out to have been misleading. Moe, insofar as he ignores some of his higher-order evidence, falls short of this ideal. Similar remarks apply to Aki, and perhaps also to Huckleberry Finn. But I have tried to stress that it is one thing to fall short of a rational ideal, and it is quite another to be incompetent with a rule of inference. These problems suggest that being *basically* inferentially competent does not consist having evidence that satisfies the **linking principle**, at least if evidence is construed propositionally.

In the next section, I argue that non-propositionalist versions of inferential evidentialism have the virtue of avoiding both problems for propositionalism. Non-propositionalism does, however, invite a different sort of problem, one which should push us to look for alternatives to inferential evidentialism.

3. Non-propositionalism and idle explanations

The arguments in section 2 attacked a propositionalist version of inferential evidentialism. Someone who espoused that sort of view would endorse three claims: **E**; that evidence is, by its nature, propositional; and that having evidence is a matter of standing in some special epistemic relation to the propositions that constitute one's evidence. In contrast, non-propositionalist versions of inferential evidentialism[49] can be characterized by their acceptance of **E**, plus the following:

---

[49] See, e.g., Conee and Feldman (1985; 2001) and Plantinga (1983)

- An agent's evidence at a time consists of her *total mental state* at that time.

Non-propositionalists have a strictly more expansive notion of 'evidence' than propositionalists. For any bit of evidence, P, a propositionalist would countenance an agent A as having, they would do so only because A stood in a special epistemic relation to P. Standing in such a relation to P is plausibly a mental property of A. Consequently, the non-propositionalist would countenance the agent's standing in such a relation to P as part of her evidence. So, whatever propositions a propositionalist would countenance as part of A's evidence, it seems a non-propositionalist would countenance a corresponding bit of evidence that is A's attitude towards those propositions. But the non-propositionalist would also countenance further, *non-propositional* mental states (mental dispositions, skills, bare sensations without propositional content, and so on) as part of an agent's evidence.

For non-propositionalist versions of inferential evidentialism, basic competence is still evidence, but it is not evidence that is determined entirely by the propositional mental states of the agent. For instance, competence with enumerative induction may consist, on such a view, in nothing over and above a cluster of mental dispositions to make certain inductively supported inferences. While the stimulus and manifestation conditions of such dispositions may implicate certain propositional mental states of the agent, the disposition itself is not a propositional mental state of the agent.

Non-propositionalist versions of inferential evidentialism avoid the two problems I raised for propositionalism because they are not forced to accept that inferential competence consists in *propositional* mental states, but they nevertheless retain the idea that inferential competence consists in having certain evidence (by occupying certain mental states). For one thing, there is no Carrol-style regress threatening a view that appeals to non-propositional evidence; non-propositional mental dispositions are not the kinds of things that could serve as *premises* in inference,

54

so there is no vicious regress of ever-higher-order premises looming. Moreover, non-propositionalists do not obviously face the problem of denying inadvertent epistemic virtue. Non-propositional linking evidence is not the agent's "better judgment" against which they can act in cases of inadvertent virtue, since judgments (explicit or not) are propositional mental states.

These considerations suggest that non-propositionalism is better suited to explain the relationship between competence and inferential justification than the view considered in the last section. But here I want to stress that the relative advantage non-propositionalism enjoys should not distract us from recognizing that that view has a very different sort of problem. In particular, non-propositional linking evidence does not perform what might be naturally thought of as *any* of the characteristic normative functions of evidence. In short, the sort of non-propositional linking evidence in question does not make the *right kind of normative difference* to count as evidence at all.

Of course, non-propositional linking evidence makes *some* normative difference on these views. For instance, if an agent lacks the relevant competence-constitutive evidence, then she is incompetent, and so she is not propositionally justified at all in making inferences that would be the manifestation of such a competence. Were Larry incompetent with the rule of enumerative induction, for instance, he would lack propositional justification for believing that B, given O. But competence with enumerative induction does not itself raise the probability of inductive inferences; rather, it is what explains why his first-order observations O raise (rather than lower) the probability of various inductively supported hypotheses. Competence looks rather like a mere enabling condition on justified inference, normatively on a par with other, more mundane enabling conditions, like there being oxygen flowing to the agent's brain—these are conditions that must be satisfied in order for an inference to be justified, but not because they bear directly on the *justification* of the inference.

Perhaps in and of itself, the claim that basic competence is an enabling condition on justified inference is not objectionable; it may be that the first-order evidence, all on its own, supports one inference over another, but it is only in the presence of basic competence that that inference is, in some sense, 'available' to the agent. Of course, if non-propositional linking evidence is a *normatively relevant* enabling condition, we are owed an account of what this normative relevance comes to, one that does not simply amount to special pleading.[50] Absent that account, it would appear that the non-propositionalist avoids the problems of propositionalism only at the cost of positing "idle" evidence.[51]

It is worth taking some time to emphasize what is objectionable about appealing to idle evidence. Importantly, appeals to idle evidence flout a plausible meta-epistemological view about the nature of evidence, one according to which evidence just is that which performs certain characteristic normative functions. It is widely accepted that evidence is that which makes a difference to the degree to which an agent is justified in adopting certain attitudes towards various propositions,[52] it is that which makes a difference to the resilience of one's degree of belief (as when one gets evidence supporting what one already knows), and it is that which could serve as a basis for justified belief.[53]

---

[50] One might, for instance, claim that competence is a normatively relevant enabling condition *because it constitutes a relevant ability*, but the problem is that the normative relevance of this ability must be understood, by the lights of non-propositionalism, in terms of *evidential* relevance. In the next section I argue that, while one can avoid the problem of idle evidence by appeal to sources of immediate justification, one does so at the cost of rejecting evidentialism.

[51] Thanks to Juan Comesaña for raising this possibility in conversation.

[52] See, e.g., Comesaña and Sartorio (2014). Thinking of evidence as a difference-maker in this way follows straightaway from commonly accepted probability-raising conceptions of evidence. According to conceptions of evidence as probability-raising: E is evidence for H just in case, for some antecedently specified probability function, $Pr(-)$, which models a subject's credences, $Pr(H|E) > Pr(H)$.

[53] There are, of course, cases of so-called "blind spot propositions" that an agent cannot use as a basis for justified belief, such as the propositions that *I am content to live in an ice hut and I doubt*

This meta-epistemological claim is meant to be as neutral as possible. It is neutral with respect to the nature of justification more generally, with respect to whether the explanatorily fundamental doxastic states are binary or admit of degrees, with respect to the relationship between binary and degreed doxastic states, whether evidentialism is correct, etc. It is, to put it plainly, just meant to encode a truism. That "idle evidence" fails to satisfy any of these difference-making conditions characteristic of evidence would seem to undermine its explanatory power *as evidence*. Again, idle evidence does not seem to make the right kind of difference to count as evidence at all. Absent independent motivation, one wonders what, beyond an unshakeable commitment to **E**, would make this option look attractive.

To sum up: the last two sections have motivated the thought that inferential evidentialists face serious problems in accounting for basic competence, whatever their favored account of evidence. Either evidence is propositional, or it is not. If it is, then inferential evidentialists must either deny the phenomenon of inadvertent epistemic virtue or accept a vicious regress; if it isn't, they draw normative distinctions that seem to mark no normative difference. Both routes lead to the idea that the normative contribution of basic competence is not captured by the **linking principle**. In light of the problems outlined in this and the last section, one should look for alternatives to inferential evidentialism.

Before I present my own view, however, it is instructive to compare the position criticized here to others that bear a family resemblance to it. First, take *mentalism*, the view according to which any difference in the epistemic status of a doxastic state between two agents supervenes on a mental difference between them.[54] (For all

---

it. But it would be rather extreme to deny that the normative function of evidence is to serve as a basis for justified belief in light of these exceptional cases, rather than pointing to peculiar features of, say, *de se* belief.

[54] E.g., see Conee and Feldman (2001).

mentalism says, these mental differences may or may not be evidential differences.)

Mentalism does not have a problem accommodating my claim that basic competence

makes a non-evidential difference to justification, since mentalism does not claim that

evidence is the only thing that makes a difference to justification; mental dispositions,

skills, and bare sensations without propositional content, perhaps among other things,

can be normatively relevant without being evidentially relevant. On these sorts of

views, it is possible for two agents to share a body of (first-order) evidence but differ in

non-evidential mental states, and as a result differ in inferential propositional

justification. This is how someone broadly sympathetic to mentalism would explain

how expertise could be normatively relevant to inferential justification in a distinctly

non-evidential way.

Next, take various forms of Bayesian confirmation theory. These views treat

evidential support as a three-place relation between a subject's evidence (typically a set

of propositions or sentences expressing them), a hypothesis, and that subject's

conditional credences. Conditional credences are not evidence—not even "higher-

order" evidence. Instead, they encode something like an agent's *evidential standards*,

what she takes to be evidence for what.[55] On these sorts of views, it is possible for two

agents to share a body of (first-order) evidence but differ in evidential standards, and as

result differ in inferential propositional justification. This is how someone broadly

sympathetic to Bayesian confirmation theory would explain how expertise could be

normatively relevant to inferential justification in a distinctly non-evidential way.

Both mentalism and Bayesian confirmation theory have the theoretical resources

to respect what seems to be a fundamental epistemological distinction between the

mental states that function so as to provide an agent with a basis for justified inference

---

[55] See, for instance, Titelbaum's "Fundamentals of Bayesian Epistemology", chapter 4 for an extended discussion of evidential standards.

and the mental states that function so as to guide how one bases one's inferences in a way that counts as justified. It is a serious mark against non-propositionalist version of inferential evidentialism that it runs these two things together.

In the next section, I lay out my own view, "inferential dogmatism", so-called because it is structurally very similar to various forms of dogmatism defended in the epistemology of perception. After presenting a sketch of inferential dogmatism, I try to differentiate it from standard forms of Bayesian confirmation theory while situating it among other forms of mentalism.

4. An alternative to inferential evidentialism

I have motivated the idea that being basically inferentially competent is not simply a matter of having certain evidence. We should reject the reductive account of basic competence in terms of evidence, and the more general reduction of being a competent reasoner in terms of having certain reasons.

This idea is captured by the following constraint on any plausible epistemology of inference:

BC: The normative contribution that *basic inferential competence* makes to what an agent is rationally permitted to infer is *not* an evidential contribution.

This constraint is purely negative; it says what the normative contribution of basic competence cannot be. It does not say what basic competence *is*, nor does it say what the normative contribution of basic competence is, if not an evidential one.[56]

---

[56] As an aside, **BC** is meant to be neutral on the metaphysical question of what basic competences are. Basic competences, on my view, are whatever our best account of competence in general, of which basic competence is a special case, says they are. And, roughly, the two most prominent accounts of competence in general are *intellectualist*, on the one hand, and *Rylean*, on the other. Intellectualist accounts of competence hold that all competence can

This section has two aims. The first is to sketch and defend a positive proposal, "inferential dogmatism". This will go some way towards explaining what the normative contribution of basic inferential competence comes to. The second is to situate inferential dogmatism within the range of views that plausibly satisfy **BC**. For instance, In the last section I indicated that mentalism and certain forms of Bayesian confirmation theory seem to satisfy **BC,** at least to the extent that those views permit non-evidential factors to make a difference to inferential justification. But, depending on how we think of the nature of the conditional credences to which Bayesian confirmation theorists appeal, inferential dogmatism will have a better explanation of how cases of inadvertent epistemic virtue are both possible and *prima facie* rational.

4.1 Inferential dogmatism sketched

To get a sense of what dogmatism about inferential justification looks like, let's consider more traditional forms of dogmatism concerning non-inferential, perceptual justification.

Dogmatic views in the epistemology of perception maintain that merely having certain perceptual experiences—as of a hand (say)—immediately but defeasibly justifies the perceiver in believing certain propositions (that she has a hand), whether or not she is also justified in believing that perception, on that occasion, worked. In short, justified perceptual beliefs are not justified in the first place because of prior and independent evidence that perception works. This is consistent with the claim that a perceiver could, on some particular occasion, also justifiably believe that perception works, and so be

---

ultimately be reduced to states of propositional knowledge, perhaps under some special mode of presentation. Rylean accounts of competence deny that all competence can ultimately be reduced to states of propositional knowledge; perhaps competence consists in *both* having states of propositional knowledge under special modes of presentation *and* in the agent possessing a certain set of dispositions. It is not my aim to wade into this debate, let alone to settle it. Instead, I just want to flag that **BC** is consistent with either account of competence in general.

justified in believing that she has a hand in some way mediated by this background evidence. But that would simply be a matter of rational overdetermination; it would not show that a competent perceiver's merely having certain perceptual experiences, all on its own, did not provide a source of immediate defeasible justification.[57]

If not prior and independent evidence that perception works, what could justify an agent in adopting certain perceptual beliefs? Different forms of dogmatism give different answers. It is possible to endorse dogmatism about perceptual justification on largely *a posteriori* grounds. For instance, one might hold that believing what perceptually seems true is, within certain limits, in fact a reliable means of forming true beliefs.[58] In other words, for those who treat reliability as a mark of justification, the outputs of reliable perceptual belief-formation mechanisms have claim to a kind of immediate but defeasible justification.

A perhaps more common way to endorse dogmatism about perceptual justification is on *a priori* grounds. In some cases, this is put in explicitly phenomenological terms: perceptual appearances are "assertive", or present-as-true (Huemer 2006, 2013; Pryor 2000). There is a tight connection between the *nature* of one's perceptual experience—its distinctive presentation of its contents as true—and what one is justified in believing (namely, its contents). On other views, it is not as much the phenomenology as the etiology of one's perceptual experience that accounts for perceptual justification. On a view like Markie's (2013), immediately justified perceptual beliefs involve the manifestation of knowledge-how; for example, when Gus the gold prospector, expert at identifying gold, looks at his nugget and "just sees" that it is gold, Gus is exercising his knowledge of how to visually identify gold nuggets. His resulting perceptual belief (that his nugget is in fact gold) has a kind of distinguished,

---

[57] Pryor (2000), p.535; Alston (1983), p.79; Pollock (1971)
[58] See Siegel and Silins (2015) for discussion.

justification-conferring etiology, that of being the manifestation of knowledge-how. This distinguished etiology is missing in someone who also believes he's come across a gold nugget, "just seeing it" out of desperation for a payday.

Let's move from a sketch of various forms of perceptual dogmatism to a sketch of inferential dogmatism. Recall Larry, Moe, and Curly. Each possesses excellent evidence (99 balls drawn from an opaque urn known to contain 100 balls are black) for a simple inductive inference (to the conclusion that the 100th ball is likely black). According to inferential dogmatism, each of Larry, Moe, and Curly, merely in virtue of having the inductive evidence and competence that they do, are *prima facie* justified in making the inference that they in fact make. Unlike Moe and Curly, Larry *does* have linking evidence at his disposal. But all this shows is that, for agents like Larry, an inference may be rationally overdetermined, both immediately and mediately justified. It would not show that a competent reasoner's merely having certain first-order evidence did not provide a source of immediate but defeasible inferential justification.

What, then, if not the possession of linking evidence, could even *prima facie* justify these agents in making the inferences that they do? Here, as in the case of perceptual dogmatism, there are range of options. One might appeal to the reliability of the mechanisms by which one performs basic inductive inferences, the phenomenology of simple inductive inferences (perhaps, given a set of premises, certain conclusions are presented as true in a way analogous to the way that perceptual or quasi-perceptual contents are said to be), or to the distinguished etiology of basic inductive inferences to ground the claim that such inferences are immediately justified.

The last option is the one I will explore in what follows, namely that *manifesting one's basic knowledge of how to perform simple inductive inferences* is a distinguished, justification-conferring etiology. I cannot, of course, hope to resolve disputes *within* dogmatism about how these various formulations fare against one another, whether one has claim to be more fundamental than others, for instance. Instead, I will do

something much more modest. I will defend the idea that a version of inferential dogmatism that appeals basic knowledge of how to infer is better equipped than both versions of inferentialist evidentialism to explain the relationship between competence in reasoning and inferential justification.[59]

4.2 Inferential dogmatism motivated

Recall that propositionalism faced two, related problems. The first was that the view could not accommodate inadvertent epistemic virtue; an agent like Moe who inferred that the 100[th] ball was likely black while harboring doubts about the connection between his evidence at that conclusion would, at best, lack *prima facie* justification for his inference, and, at worst, be incompetent with induction. I argued that there was a great deal more to say in favor of Moe's inference, even if he was not *all things considered* rational, and even granting that he fell short of rational ideals. The second problem was that, in order to avoid saying implausible things about agents like Moe and Curly, propositionalists had to accept a Carrol-style regress of premises to explain competent inference.

---

[59] My own suggestion aligns with some of Markie's work on dogmatism and knowledge-how (2013, 2015). According to Markie, to know how to X is to have a "special ability" to X, which itself is a matter of hosting a range of X-related dispositions. While I am sympathetic to the general spirit of Markie's work in these papers, I do not endorse his neo-Rylean view. I treat it as an open question whether the X-related dispositions constitutive of one's special ability are themselves explained in terms of further states of propositional knowledge under a distinctly "practical" mode of presentation. But if this reading of Markie is incorrect, and he is instead not committed to a form of Ryleanism, I am happy to think of my arguments as supporting his position for reasons different than the ones he offers. His aim, at least in his (2013), is to find a version of qualified dogmatism about non-inferential justification that is consistent with mentalism and foundationalism. My aim here has been, firstly, to criticize a family of views about the nature of propositional inferential justification, and secondly, to point to a more plausible alternative. This more plausible alternative is a form of qualified dogmatism about inferential justification that is consistent with mentalism and foundationalism. To the extent that my arguments converge with Markie's, I take it to indicate the plausibility of this particular brand of qualified dogmatism, one which appeals to knowledge-how.

Inferential dogmatism avoids these two problems. If basic inferential competences are a source of immediate justification, then there is no need to appeal to linking evidence as premises that would mediate the inference, and so no Carrol-style regress looms. And cases of inadvertent epistemic virtue are simply cases where a basically epistemically competent agent adopts a doxastic state that is in fact supported by their first-order evidence, but which goes against their "better judgment". But their better judgment is simply some other, familiar propositional mental state that they instantiate. In this way, inferential dogmatism shares certain advantages with non-propositionalism.

What about the objection leveled against non-propositionalism in section 3: basic competence cannot make a normative difference at all, given that it appears only to *enable* competent inference? More specifically, I criticized non-propositionalism for appealing to "idle" evidence, evidence which fails to perform what might be thought of as *any* of the characteristic functions that evidence performs. Not all readers will be impressed by the meta-epistemological principle about the nature of evidence to which I appealed. They might, for instance, insist that we mean different things by 'evidence'.

Now, the debate is not meant to be terminological; if the inferential evidentialist is committed to using the term 'evidence' in a way that flouts that meta-epistemological principle, so be it. The real issue is which way of employing normative language carves nature at the joints. My suggestion, terminology aside, is that an important normative difference is marked by the fact that a certain subset of one's mental economy can affect the degree to which one is justified in accepting various propositions (or can make a difference to the resilience of an attitude, or can serve as a basis for justified belief), and it seems that basic competence does not affect the normative status of an agent's doxastic state in any of these ways. Instead, basic competence is the sort of thing that makes a difference to what an agent is justified in inferring from a body of evidence by not only enabling but *guiding* how one makes inferences so as to count as justified. This

non-evidential *guiding role*, I submit, is what renders one's basic knowledge of how to infer more than a mere enabling condition on justified inference.

4.3 Inferential dogmatism, mentalism, and Bayesian confirmation theory

I've only given an initial sketch and motivation for inferential dogmatism. Still, I hope to have shown at least two things. First, that dogmatism about basic inferential justification inherits some plausibility by way of resemblance with dogmatism about basic perceptual justification. Second, that inferential dogmatism avoids the problems I raised for both forms of inferential evidentialism.

But there are surely some lingering concerns about how inferential dogmatism fits within the broader epistemological landscape. For instance, I earlier claimed that certain forms of mentalism and Bayesian confirmation theory are consistent with **BC**, the claim that basic competence makes a non-evidential difference to inferential justification. What reasons might there be to prefer my own view over these?

The first thing to note is that we do not have the choose between inferential dogmatism and mentalism; inferential dogmatism is just a form of mentalism. If we think—as seems plausible—that having certain basic competences is a mental property of an agent, then the claims of inferential dogmatism will be consistent with the somewhat broad supervenience claim that characterizes mentalism ("no difference in justification without a mental difference"). Thus, rather than seeing these as competing views, inferential dogmatism is a precisification of a rather broad but plausible stance on the conditions relevant to justification.

The relationship between inferential dogmatism and Bayesian confirmation theory is trickier. In certain respects, the two views are sympathetic; to the extent that Bayesians treat conditional credences as representing one's "evidential standards", or underlying dispositions to infer, Bayesians carve out a space in their theory of

inferential justification for something very much like basic inferential competence. Evidential standards are not further evidence, after all.

Despite these points of agreement, here is one point of divergence: Bayesian confirmation theory describes a certain sort of rationally ideal agent, one whose attitudes are probabilistically coherent, and one whose conditional and unconditional attitudes stand in a particular relationship. Specifically, if "C" is a credence function and "H" and "E" are propositions, a Bayesian agent's attitudes will satisfy the ratio formula:

RATIO: $C(H|E) = C(H\&E)/C(E)$

In other words, when a Bayesian agent has a certain conditional credence $C(H|E)$, she will also have certain unconditional credences $C(H\&E)$ and $C(E)$, the quotient of which equals $C(H|E)$. If we take **RATIO** at face-value, Bayesian confirmation theory treats conditional and unconditional commitments as a kind of package deal; having certain conditional commitments guarantees, all by itself, that one has certain unconditional commitments, and *vice versa*.

The important point is this: if an agent's attitudes are constrained by **RATIO**, certain cases of *prima facie* rationality will be excluded from one's theory of rationality. Think of Moe, who harbors doubts about the connection between his first-order evidence ("O") and the hypothesis that the 100[th] ball is black ("B"). $C(O)$ is high (perhaps he is certain of his evidence), but given his doubts about the connection between B and O, $C(B\&O)$ is exceeding low. This would force his conditional commitment, $C(B|O)$ to be exceedingly low. But then when Moe goes on to infer B from O, he does not rely on his exceedingly low conditional credence for B, given O. If the inference to H from E is even *prima facie* rational, its positive epistemic status is not

66

grounded in facts about how his conditional and unconditional commitments hang together, since those commitments would not even *prima facie* support that inference.

To the extent, then, that we are inclined to see Moe (or Curly) as manifesting inadvertent epistemic virtue, we cannot vindicate that idea purely by appealing to the non-evidential role that conditional credences play in inference, at least not if we accept that conditional and unconditional credences are mutually constrained by **RATIO**. Instead, what grounds the (perhaps merely) *prima facie* rationality of Moe's inference is that Moe has basic knowledge of how to perform enumerative induction; in well-ordered cases, but perhaps only in well-ordered cases, this more basic explanatory fact is captured by facts about his credences.

To be clear, I am not suggesting that here we've found the thread that, if we were to pull hard enough, would unravel Bayesian confirmation theory. To the contrary, I hope that the reader sees my position and the arguments for it as broadly in line with the considerations that would push someone to adopt a ternary view of evidential support, as Bayesians do.[60] Here, I mean only to suggest that there is a deeper explanation for what makes certain inferences justified than what Bayesian confirmation theory provides. That deeper explanation involves an agent's underlying basic knowledge of how to infer, not the surface-level features that, at least for agents that approximate certain Bayesian rational ideals, might be thought to represent or encode that underlying knowledge of how to infer. Bayesian confirmation theory may be entirely adequate to describe most cases of inference, and perhaps all cases of *ideally rational* inference. It is only by looking at non-ideal cases of inference that we can see where the Bayesian story and my own diverge. In such cases, an agent might be probabilistically incoherent to some degree, or may fail to respect a bit of their evidence

---

[60] Thanks to an anonymous referee at *Ergo* for pointing me to this parallel.

to some degree, and that agent may nevertheless make a (*prima facie*) rational inference by manifesting their basic knowledge of how to infer.

5. Conclusion

Inferential evidentialism is a popular and attractive view that explains what I've called "basic" epistemic competence, which provides "basic" inferential justification, in terms of derivative or mediate justification by way of background evidence linking one's first-order evidence to what it is evidence for. I raised a number of problems for two versions of this view—for propositionalist versions, that they faced a regress of premises, and that they failed to countenance inadvertent epistemic virtue; and for non-propositionalist versions, that they only avoided those problems at the cost of positing idle evidence.

These problems suggested that the normative contribution of basic competence is not an evidential one; to be competent with a rule of inference is not, in general, simply to possess certain linking evidence. I then sketched a dogmatist view of basic inferential justification according to which one's basic competences give immediate, defeasible justification for certain inferences. While some authors have been attracted to dogmatic views for cases of non-inferential justification, dogmatic views have been largely ignored in cases of inferential justification. This under-explored option, inferential dogmatism, avoids the objections I raised to inferential evidentialism in its various forms. And ultimately, this position vindicates the natural thought that being a competent reasoner is not simply a matter of having certain reasons to believe this or that, even if a good deal of (derivatively) competent reasoning can be adequately explained in those terms.[61]

Chapter 3:

What we know when we act

Two traditions in action theory offer different accounts of what distinguishes intentional action from mere behavior.

> **The causalist tradition:** intentional action, as opposed to mere behavior, has certain distinguished *causal antecedents*. In particular, when an agent acts intentionally (to X, say), her beliefs, desires, and intentions cause her X-ing (in the right way).[62]

> **The Anscombian tradition:** intentional action, as opposed to mere behavior, has certain distinguished *epistemological features*. In particular, when an agent X-s intentionally, she knows what she is doing as she does it.[63]

These traditions, as stated, may seem to place only necessary constraints on what it is to act intentionally, and, as such, one might wonder how they are in conflict at all. But that's because a lot of important theoretical work is hidden in parenthetical clause "in the right way". Causalists often speak as though *a non-deviant causal connection* between an agent's behavior and her antecedent beliefs, desires, and intentions is both necessary and sufficient for that behavior to count as an intentional action. This, at any rate, is the kind of causalist that I have in mind throughout the essay.

According to (typically causalist) critics, Anscombians have a hard time explaining a range of very ordinary—even if not *paradigm*—cases of intentional actions,

---

[62] E.g., see Davidson (1971; 1978), Paul (2009), Mele (1992b; 2001), Bratman (1984; 1987).

[63] E.g., see Anscombe (1957), Velleman (1989), and Setiya (2009). Typically, these authors tend to deny that having certain distinguished causal antecedents *suffices* for behavior to count as intentional action, rather than denying the relevance of those causal antecedents full stop. There are important differences between Anscombians that, while of independent interest, I will set aside here.

intentional actions performed in the absence of certain first-personal doxastic attitudes. For instance, when distracted drivers make it home from work, they do so intentionally, but it seems odd to describe them as "knowing what they are doing as they do it", if for no other reason than they are distracted. And when athletes cannot explain why or how to perform complicated physical sequences, even as they do them intentionally, they plainly lack whatever knowledge underwrites an ability to explain what they are doing. For instance, competent cyclists often describe how to turn a bike one way but then go about turning another way. Their firm convictions expressed in *describing* how to turn do not count as knowledge, since those convictions are false.

Taken at face-value, each of these problematic cases involves an agent performing an intentional action in the absence of knowledge of what they are doing as they do it. And if Anscombian views have problems accommodating these cases, it seems reasonable to conclude *so much the worse for an epistemic constraint on intentional action*. One might suspect that, even if central or well-ordered cases of intentional action are ones in which the agent stands in a special epistemic relation to what they are doing as they do it, standing in such a relation is not particularly important to action theorists. In fact, Sarah Paul has argued that practical agents only have a kind of contingent, *inferential* knowledge of what they are doing as they do it. When an agent knows what she is doing as she does it, on Paul's view, it is because the agent knows what she (previously) decided to do, and she can infer, perhaps by way of background beliefs concerning her own efficacy as a practical agent, that she is doing what she (previously) decided to do. Paul notes, rightly to my mind, that even if this sort of contingent, inferential knowledge of what one is doing is desirable and important for practical agents, it is not the mark of intentional action *per se*.

But cases of distraction or absent-mindedness, cases of inarticulability, and the like only threaten a fairly narrow band of Anscombian views. This fairly narrow band of views accept a highly intellectualized, discursive conception of the knowledge one

has of what one is doing as one does it. Characteristically, discursive knowledge is the kind of knowledge one can *state* or *articulate,* one can *bring to mind*, one can *offer reasons in favor of,* and so on. This sort of knowledge is plainly valuable to practical agents, but we should all agree with causalist critics of Anscombianism that *it* is not the mark of intentional action *per se*.

What sort of knowledge is a better candidate? Below, I'll argue for the following quasi-Anscombian position:

Ability-constituting knowledge of action ("AKA"): When an agent X-s intentionally, she manifests ability-constituting knowledge of action as she X-s.

The first two sections of this paper are devoted to explaining what ability-constituting knowledge of action is, and how it differs from other kinds of knowledge. The last two sections of this paper are devoted to explaining why action theorists should take ability-constituting knowledge of action seriously. First, appealing to such knowledge is specially suited to distinguish between cases of "inadvertent" [64] intentional action from merely unintentional ones. In this respect, my position fits into a unified account of inadvertent virtue across various normative domains, of which action theory is one.

Second, while the view I defend may naturally fit within the Anscombian tradition, it is consistent with views in the causalist tradition; in principle, there is nothing inconsistent about both accepting a causal theory of action and accepting AKA. In particular, one might think that an agent's beliefs, desires, and intentions cannot cause her behavior *in the right way* (or "non-deviantly") unless the way in question is

---

[64] I hope that the stipulated terminology is innocent enough; I mean to use "inadvertent" in the sense of "inadvertent virtue" or "inadvertent epistemic virtue" used by Arpaly (2002) and Weatherson (2019). It is natural enough, in ordinary contexts, to use "inadvertent" and "unintentional" as rough synonyms, and I break from that natural usage here.

itself a manifestation of ability-constituting knowledge of action. By appealing to an agent's ability-constituting knowledge of action, causalists can explain how one's antecedent and occurrent mental states causally sustain one's subsequent behavior so as to count as "guiding" or "controlling" it. If one accepts AKA, one can offer an illuminating account of what this guidance consists in, and why it is intuitively absent in unintentional actions and mere behavior, in terms of the possession and manifestation of ability-constituting knowledge. Thus, to the extent that it is possible, my view reconciles causalism and Anscombianism.

So, what is ability-constituting knowledge of action, and how does it differ from other, more familiar kinds of knowledge?

1. Ability-constituting knowledge and rational inference

In order to characterize ability-constituting knowledge of *action*, it helps to start with the nature and normative role of ability-constituting knowledge of *inference*. Ability-constituting knowledge of inference "connects" evidence to hypotheses, without itself serving as evidence.

Consider Carrol's What the Tortoise said to Achilles. Therein, the tortoise asks Achilles what to make of an agent who knows that P, and that P implies Q, but who simply fails to see that Q. Achilles, naively, suggests that the agent may simply be missing a premise: that if (P and P implies Q), then Q, but he is then faced with the question of what to make of an agent who knows that P, that P implies Q, and that if (P and P implies Q), then Q, but who simply fails to see that Q. Achilles, absurdly now, suggests that the agent may simply be missing a premise. And so on. While this particular Carrollism has received much attention since its publication, often giving rise to divergent sophisticated analyses of what has gone wrong with Achilles' suggestion, there is general consensus that Carroll has demonstrated something important about

competence with rules of inference. What it is, in this case, to be competent with deduction is not to deduce by way of additional, perhaps suppressed premises.

Still, when a rational agent competently deduces Q from her evidence *P* and *If P, then Q*, she manifests her knowledge that *If (P and if P, then Q), then Q.* We should just deny that this knowledge is the sort of thing that serves as *further evidence* for judging that Q. Instead, this knowledge *guides an agent's judgment* that Q from her evidence *P* and *if P, then Q*.[65]

Talk of "guidance" may prompt a variety of reactions, and I mean to use that term in a fairly deflationary sense.[66] First, "guidance", as I understand it, is merely *de re* sensitivity to certain propositions, not *de dicto* sensitivity. In general, it is not true that when one is guided by a proposition, one can articulate the guiding proposition, or "hold it before one's mind" under a particular description. Second and relatedly, "guidance" may operate merely sub-personally. We should not think that the fact that a particular proposition is guiding an agent's behavior is in principle accessible to the agent herself; she may be unable to hold the guiding proposition before her mind at all.[67]

One might wonder: what good is ability-constituting knowledge of inference? Why can't frugal epistemologists get by with all and only *evidence-constituting* knowledge? For epistemologists, ability-constituting knowledge of inference plays at least one crucial normative role: that of determining whether an agent has *properly based her beliefs on her evidence*.

---

[65] I further characterize ability-constituting knowledge of inference at the end of section 2.
[66] Some are suspicious of "guidance"-talk because they think Ryle's (1949) critique of intellectualism is more or less correct, and one of its main targets is a "guiding proposition" conception of intelligent behavior. I do not want to endorse the view that Ryle rightly criticized as overly intellectual.
[67] This, I take it, is no more controversial than the epistemological thesis that one's mental life is not "luminous". A mental state M is luminous (for agent A) just in case: when A is in M, A is in a position to know that she is in M. See, e.g., Williamson (2000).

Imagine an agent, Ham, who believed both P and if P, then Q, and, in the course of considering whether Q, the agent is hit on the head with a hammer, and the precise force and angle of impact scrambles his brain so as to make him believe that Q, while retaining his earlier evidence. Despite "fitting" the evidence, there is something normatively deficient about Ham's belief. Many authors would explain this deficiency as a failure of proper basing; the connection between Ham's evidence and his judgment is just too *lucky* to credit the agent with any sort of epistemic achievement. Despite a kind of "fit" between evidence and inference, Ham does not infer Q *because of its connection to his evidence*.[68] After all, being hit on the head at just the right angle so as to come to believe all and only the hypotheses supported by one's evidence is patently not to manifest any sort of knowledge of the connection between evidence and hypothesis.

In general, manifesting ability-constituting knowledge of inference rules out a form of *luck* in epistemology. This sort of luck diminishes the *agent's contribution* to her behavior. If an agent with evidence E comes to believe hypothesis H *but not by an exercise of her epistemic agency*, her belief is too lucky, in this sense, to count as rational.[69]

---

[68] Sosa (2007; 2010) has long defended a view of this kind. In his terms, we could say that, absent manifesting ability-constituting knowledge of inference, the inference would not be *apt*. Apt inferences are *accurate because adroit*. And being hit on the head is not an adroit or skillful way to form accurate inferences on one's evidence. Of course, one does not have to be a Sosa-style virtue epistemologist to accept these sorts of "ability constraints" on proper basing. See, for instance, Turri (2010).

[69] Of course, the predominant view about what it is for an agent to properly base her belief that P on her evidence E is that E must cause the belief that P (in the right way). Thus, a causal theory of basing faces the same deviance worries as a causal theory of action. (See, e.g., Turri (2010) for an extended defense and elaboration of this claim.) The most promising accounts of "non-deviant" basing consist not in rejecting a causal theory of basing, but in articulating constraints on which belief-formation mechanisms—understood causally—are rationality-grounding. A number of authors have suggested that these mechanisms must themselves be a kind of knowledge. See, e.g., Wedgwood (2013) and Silva (2017) for very different versions of this idea.

So far, I've suggested that ability-constituting knowledge of inference connects evidence to hypotheses, but on pain of a Carrol-style regress, this knowledge cannot be further evidence (not on a natural construal of 'evidence', anyway). Moreover, that an agent manifests this ability-constituting knowledge of inference (that if E, then H) in moving from E to H explains why the judgment that H is attributable to the agent as a rational inference, not, for instance, as something that merely happens to her.

Below, I suggest that ability-constituting knowledge of inference and ability-constituting knowledge of action are species of a common genus: Ability-constituting knowledge. Species of this genus play a special role in various first-order normative theories: that of "linking" reasons to reasons-based performances. The knowledge that links reasons to reason-based performances is not itself a reason, on pain of a Carrol-style regress, and this knowledge plays an ineliminable role in explaining why certain inferences, despite "fitting" one's evidence, are deviant.

2. Ability-constituting knowledge and intelligent action

Why think that action theory needs the same conceptual resources as epistemology? More specifically, why think that action theory needs to "connect" reasons for action[70] to action by way of manifesting some ability-constituting knowledge, as epistemology needs to connect premises to conclusion by way of ability-constituting knowledge?

The short answer is that we need to be able to explain why *some* cases of "inadvertently" successful performance are merely unintentional actions, while others

---

[70] I mean to be ecumenical about how 'reasons for action' is best understood. Someone who is suspicious of the broadly Davidsonian framework I'm using could make do with talk of 'intention' instead. For instance, on a Bratman-style planning view of intention, intentions are *sui generis* psychological states with distinctive motivational profiles and connections to diachronic agency, and they would be ineliminable in the causal explanation of an agent's behavior (alongside her reasons for action, explained in terms of her beliefs and desires). I hope that I can simply bracket this nuance for the sake of exposition.

are *bona fide* intentional actions. Here, "inadvertent" only serves to mark that the action in question is performed in the absence of a certain kind of knowledge, whatever knowledge would enable the acting agent to correctly judge or explain what they are doing as they do it. Consider, as an instance of cases of the first kind, the novice dart-thrower:

> **Darts:** Al meets up with some friends at a local bar, *Che's Lounge. Che's* happens to have a lively darts competition one night each week, and it is slated to begin as Al finishes his second beer. Never having played a game of darts in his life, but brimming with confidence, Al signs up to play, fully intending to win. On his first turn he looks at the bullseye, makes some arm movements that, for all he knows, resemble the arm movements of genuine dart players, and sends the dart on a wing and a prayer. Lo and Behold! Al hits the bullseye. As it turns out, Al goes on to whiff every subsequent throw and loses badly.

In **Darts**, Al hits the bullseye, but he does not *intentionally* hit the bullseye. This should be an acceptable verdict for causalists and Anscombians alike. After all, Al has never played a game of darts before in his life, and it's pretty obvious, after the game is over, that his first bullseye was beginner's luck. We might say, of Al's hitting the bullseye, that a stopped clock is right twice a day. (At any rate, this is the verdict often taken for granted in thinking about cases like **Darts**, and I will accept it for the sake of argument.[71])

---

[71] See, e.g., O'Shaughnessy (1980), Harman (1986). Mele (1992a) officially registers agnosticism but offers similar cases of luck to undermine attributions of intentional action. Others, like Kenny (1975), suggest that the novice darts player lacks a pertinent kind of *ability*. Still others, perhaps Mele (1992b), might suggest that the novice lacks a pertinent kind of *control*. But it is plausible to think that 'ability' and 'control' are causal notions; for instance, it seems like the particular causal route by which the novice hits the board is not *control-grounding*, and it thereby falls short of counting as an intentional action.

In claiming that Al's hitting the bullseye is not an intentional action, I do not mean to deny that Al's hitting the bullseye is an action *full stop*. For instance, there is a clear difference between Al's hitting the bullseye as described in **Darts**, and Al's hitting the bullseye as a result of being shoved in the direction of the dart board, or as a result of sneezing while noting the heft of a dart. If Al were to hit the bullseye as a result of these events, it would be odd to think that *Al* did anything to bring about the dart hitting the bullseye. In those sorts of cases, hitting the bullseye is just something that *happens to him*.

This simply brings out two ways in which an agent's behavior may fail to count as intentional action: that behavior might count as action that is not intentional, or it might fail to count as action at all. Why think that, in **Darts**, Al's hitting the bullseye counts as any kind of action, as opposed to something that merely happens to him? Short of answering this question with a complete action theory, one plausible thought is that Al's hitting the bullseye is an action at all in virtue of other actions[72] that Al performs intentionally: his throwing the dart at the board, perhaps his *trying* to hit the bullseye, and so on. In other words, the fact that Al does *something else* intentionally can explain why his hitting the bullseye counts as an action at all, albeit an unintentional one.

Supposing that Al's hitting the bullseye is an action, but one which falls short of intentional action, what ingredient is missing? For Anscombians, Al succeeds in hitting the bullseye in the absence of knowledge of what he is doing as he does it, and that accounts for our judgment that Al does not intentionally hit the bullseye. After all, he surely does not know that he is throwing his dart in a way that makes it likely to hit the

---

[72] If one prefers a Davidsonian way of talking, then Al's hitting the bullseye counts as an action because there is some *description of the event of his hitting the bullseye* under which Al brings it about intentionally.

bullseye, nor does he know anything about what he's doing that would enable him to explain to others how to do it.

For causalists, Al's success is somehow "causally deviant"; his mental states antecedent to throwing, which perhaps include his belief that the dart board is in front of him, his belief that dart players throw *thusly*, his desire to hit a bullseye, and his intention to hit a bullseye (by throwing *thusly*), constitute Al's reasons for action. Those reasons for action somehow or another cause Al's successfully hitting the bullseye, but however the causal story goes, it is not the right one to ground Al's *intentionally* hitting the bullseye.

Now, compare Al's case to Billy's:

Bike: "Consider the following counterintuitive fact about cycling: to turn right, you typically start by steering left. All competent cyclists are in a position to deploy that information for the purposes of making turns. But few are able to deploy it for the purposes of explaining how to ride a bike. In fact, when giving verbal explanations, most cyclists are disposed to report exactly the opposite" (Elga and Rayo, 15-6). Billy, a competent cyclist, believes and reports that to turn right, you start by steering right. He then hops on a bike to demonstrate a right turn, and he starts by steering left. This inconsistency is lost on Billy.

Clearly, Billy lacks knowledge of how to ride a bike that would enable him to correctly explain how to do so. He may have accumulated lots of misleading testimonial evidence that precludes him from being knowledgeable in this fairly specific way. But just as clearly, Billy nevertheless knows how to ride a bike; after all, he hops on and manifests this knowledge in intentionally performing a right turn.

What could explain why, in **Bike**, the agent's success is attributable to them as intentional action, but in **Darts** it is not? One answer suggests itself: Billy manifests a special kind of knowledge in behaving as he does. Compare Billy to his brother, Willy.

Willy, never having ridden a bike himself, has diligently documented and memorized all of Billy's (mistaken) advice. Willy might hop on Billy's bike and, having largely the same reasons as Billy did—certain beliefs concerning how one turns a bike, certain desires to avoid falling, certain intentions to turn—succeed in turning. But Willy's case is importantly different than Billy's; it seems clear that Willy's success in turning right is a merely unintentional action. But all that distinguishes Billy's turning right from Willy's is that Billy knows, perhaps *de re* or sub-personally, how to turn a bike, even if each has false *de dicto*, person-level beliefs about how to turn a bike.

Bike bears some structural similarities to cases of inadvertent epistemic virtue in epistemology, and to cases of inadvertent moral worth in theories of moral responsibility. For instance, Brian Weatherson (2019) imagines that Aki is convinced of testimonial skepticism by powerful philosophical arguments. (Testimonial skepticism is the view that one cannot get reasons to believe propositions on the basis of testimony.) Aki's friend, whom Aki has every reason to trust, tells her that *p*. Aki, despite her skeptical philosophical leanings, comes to believe that *p* on the basis of her friend's testimony. According to Weatherson, that is what precisely what Aki should think, given the circumstances and her evidence. Her case is peculiar in that "she forms the right belief, for the right reasons, *while thinking these are bad reasons*" (170, emphasis added).

Nomy Arpaly (2002) has argued convincingly that theories of moral worth must explain what it is about Huckleberry Finn's freeing Jim that distinguishes him from a merely "accidental good-doer", and which elevates him above those "with liberal principles who are still viscerally prejudiced against people of different races" (10). Ultimately, her own explanation, developed in more detail in later work (Arpaly and Schroeder (2014)) is that Huck acts from intrinsic desires that are *de re* sensitive to the right-making features of action. Abstracting away from the details of her positive

proposal, Huck's case is one in which he does the right thing, for the right reasons, while thinking these are bad reasons.

If I'm correct, these similarities between cases of inadvertent epistemic virtue, cases of inadvertent moral worth, and **Bike** should not be surprising; rational judgments, morally worthy actions, and intentional actions are *reasons-based performances attributable to the agent*.[73] Epistemic agents, morally responsible agents, and practical agents have, by virtue of the kind of things that they are, capacities to recognize and respond to reasons of certain sorts (epistemic, moral, and practical reasons, respectively). But it is not necessary, in any of these normative domains, for the agent responding aptly to reasons *to be able to explain what they are doing as they do it*, in order for us to attribute the reason-based performance to her. This rules out *one sort* of knowledge as necessary for the proper attribution of reason-based performances, but

---

[73] An anonymous referee points out that cases of inadvertent epistemic virtue and inadvertent moral worth involve agents who are mistaken about *what to believe* or *what to do, morally speaking*, whereas my cases of inadvertent intentional action seem to involve agents who are mistaken about *how to do* what they aim to do. This seems to put pressure on my claim that inadvertent successes (in rational judgment, morally worthy action, and intentional action) are all species of the same genus.

The extent to which this is a disanalogy, I think, turns on the relationship between knowledge-how and knowledge-that, which I comment on briefly at the end of this section. For instance, on an intellectualist account of knowledge-how, Billy's ability to ride a bike might be constituted by his knowing, of various ways w, w', …, that w is a way to pedal, w' is a way to initiate turns, etc. Thus, Billy might be mistaken in his de-dicto, person-level judgments about which ways are ways to initiate turns, while still having propositional knowledge about those matters in a de-re, sub-personal way.

Supposing, however, that one rejects intellectualism, it still seems plausible that Huck manifests certain moral ability-constituting knowledge—for instance, to recognize the weight of Jim's humanity and its relative importance to mores of the antebellum South. Huck manifests such knowledge despite being mistaken, at the level of his own deliberation, about how to weigh those considerations together.

Ultimately, however, I am open to the idea that there are residual differences between the conditions that render intentional actions inadvertent, on the one hand, and the conditions that render morally worthy actions inadvertent, on the other, so long as there is a family resemblance between cases of these kinds.

again, this is a fairly narrow and intellectualized conception of epistemic relation that agents stand in to their behavior.

3. Ability-constituting knowledge as knowledge-how

So far, I've indicated that ability-constituting knowledge (of action or inference) is *not* the kind of knowledge one manifests by correctly explaining what one is doing as one does it. Instead, it may be merely *de re* and sub-personal. Below, I'll argue that it is a form of knowledge-how.

Here is the general idea: it is Billy's manifesting his knowledge of how to ride a bike that renders his properly executed turn intentional. And recall Ham, who possesses P and if P, then Q as evidence, but only believes Q because he is hit on the head with a hammer; it is Ham's failure to manifest his knowledge of how to perform simple, one-step deductions that renders his belief (that Q) irrational. While appeals to know-how already have some purchase in action theory,[74] that phenomenon tends to be neglected (or at least under-appreciated) in epistemology, and so the parallels between action theory and epistemology to which I am pointing go unnoticed.

To clarify my proposal, I want to address two, related concerns. First, Kieran Setiya (2008) has argued that we sometimes X intentionally without knowing how to X. If that's correct, my account of ability-constituting knowledge of action is problematic. Second, my account of ability-constituting knowledge *of inference* makes the relevant knowledge appear propositional, while my account of ability-constituting knowledge *of action* makes the relevant knowledge appear non-propositional; it may seem that there are really two things going on where I say there is only one. My aim is not, in the course of addressing these concerns, to offer a complete metaphysics of ability-constituting knowledge; rather, I hope to situate my proposal within a number of related debates.

---

[74] See Mele & Moser (1994) and Setiya (2008; 2009).

In "Practical Knowledge", Kieran Setiya argues that we don't have to know how to do everything that we do intentionally. In other words, one can *intentionally X* without *knowing how to X*. He says:

Bomb

"There are cases of intentional action that are not accompanied by knowledge how. For instance: I am trying to defuse a bomb, staring with confusion at an array of coloured wires. Which one to cut? In desperation, not having a clue what the wires do, whether they will trigger the bomb or not, I disconnect the red wire—and the timer stops. Even though I did not know how to defuse the bomb, and managed to do so through dumb luck, I count as having defused the bomb intentionally. That is certainly what I meant to do, despite my uncertainty."
(2008; 404)

In **Bomb**, the protagonist, call him "KS", does not know how to defuse the bomb, in the sense that KS does not know which way of cutting the wires counts as a way of defusing the bomb. Setiya thinks that if KS goes about cutting wires, driven by something like his desire to avoid being blown up and his belief that cutting some wire or another might defuse the bomb, and thereby manages to succeed in defusing the bomb, there is some intuitive pressure to say that he intentionally defused it without knowing how.

Here is a simple way to make inroads on this rather dire case: either **Bomb** and **Darts** are relevantly similar, or they aren't. If **Bomb** and **Darts** are relevantly similar, we should simply deny that Setiya has produced a *bona fide* case of intentional action. Even if, let's suppose, KS's bomb-defusing is merely unintentional, there is surely *something* that KS does intentionally (*trying to defuse the bomb*). But the connection between what KS does intentionally and his defusing the bomb is just too lucky for the latter to count as intentional too. If so, KS's situation in **Bomb** resembles Al's situation in **Darts**, and

82

my view can explain why: KS and Al succeed merely unintentionally because each lacks knowledge of how to do that which they're trying to do.

If, however, **Bomb** and **Darts** aren't relevantly similar, it does not follow that **Bomb** is a case of intentional action, as Setiya suggests. One might treat KS's defusing the bomb as a "middling action", one that is neither intentional nor unintentional. Following Mele (2012; 1992) and Mele and Moser (1994), we could say that KS defuses the bomb "non-intentionally". The category of non-intentionality applies, perhaps among other things, to "side-effect actions", actions that, for all we know we might bring about without, strictly speaking, aiming to do so.[75] Manifesting one's knowledge of how to take certain relevant means (cutting a wire) to bring about a desired end (defusing the bomb), despite being largely ignorant of *which* of his available means will likely succeed in bringing about that desired end, may be enough to rise above the level of merely unintentional action without thereby earning the honorific "intentional". On my view, it matters quite a bit that KS's knowledge of the means to defuse the bomb is incomplete but not entirely absent, while Al's knowledge of the means to hit the bullseye is entirely absent.

We've just gone over two ways of denying Setiya's claim that **Bomb** is a case of intentional action, and the plausibility of these denials, together with the natural explanation my account provides, relieves some dialectical pressure. But let's consider, if only for the sake of argument, the possibility that **Bomb** is dissimilar to **Darts** because

---

[75] Consider Harman's (1976) example of the sniper who fires his gun trying to kill an enemy soldier, knowingly alerting the enemy to his presence. That he knowingly alerts the enemy to his presence in firing his gun does not entail that he intentionally alerts the enemy. But since his alerting the enemy is done knowingly and non-accidentally, there is surely pressure to resist labeling it a merely unintentional action. See Mele and Moser (1994, section 3) for extended discussion.

KS intentionally defuses the bomb. What, if not knowledge of how to defuse the bomb, could account for this verdict? For Setiya, it is knowledge of how to do something else:

> "When I do something intentionally that I do not know how to do, I must at least know how to take some relevant means. In the present case, I know how to cut the red wire, and I think it might defuse the bomb, even though I can't be sure." (*ibid*)

Setiya's suggestion is that X-ing intentionally requires knowing how to X, or else it requires knowing how to do something else (cutting the red wire, perhaps) that would count as way to X in the context. At this point, a friend of Setiya would need to say more about the relationship between knowing how to X, on the one hand, and the variety of more basic things one knows how to do that count as ways of X-ing in a context, on the other.

Without such an account on offer, one might reasonably wonder whether the right thing to say about **Bomb** is precisely what Setiya says about it: that KS defuses the bomb intentionally *without knowing how to defuse the bomb*. Hear me out. Consider what happens in a significantly lower-stakes situation, as when you use the old xerox machine in the office to make copies of a handout; you know that, in order to make copies on the old xerox, you've got to hit either the red or green button, but you can't remember which (the other button scans the document).[76] You try the red one, say, and it works – copies abound! Here, for whatever it's worth, I am not inclined to say that you've succeeded in intentionally making copies *without knowing how to make copies*, even if we can all agree that your knowledge of how to make copies is incomplete. If this verdict is at all plausible in low-stakes situations, as when making copies, it is plausible in high-stakes situations, as when defusing bombs.[77] The inclination to

---

[76] Thanks to Carolina Sartorio for suggesting this example.

[77] Granted, there are certain "pragmatic encroachment" views of knowledge-that according to which a difference in practical circumstances or stakes can constitute a difference in knowledge-

withhold the *attribution* of (more complex forms of) knowledge-how to KS, *if* one is so inclined, may simply reflect the stakes in his context, not of what KS in fact knows how to do.[78]

The second concern is more or less that, contrary to what I've suggested, ability-constituting knowledge of inference and ability-constituting knowledge of action are *not* species of the same genus. This concern might be motivated as follows: in section 1, I said that ability-constituting knowledge of inference, at least in cases of competent, single-premise deduction, was knowledge of the form *if (P and if P, then Q), then Q*. One might think that I am appealing to an especially important *proposition*, knowledge of which guides one's inferences (at least in simple, single-premise deduction) so as to render them rational.

And earlier in this section, I argued that certain cases of inadvertent intentional action are only distinguished from merely unintentional ones because the protagonist manifests ability-constituting knowledge of action, and I relied on certain classic examples of knowledge-how (like Billy riding a bike) to make this point. Many (but not all) accounts of knowledge-how treat it as essentially *dispositional* rather than

---

that (Stanley (2005); Fantl & McGrath (2009)). That is already a highly contentious thesis regarding knowledge-that, and I cannot here assess the plausibility of an analogous claim about knowledge-how, but it seems fairly risky to anchor one's arguments to this latter thesis.

[78] In a footnote (2017, ch5, fn3), Setiya recognizes that one might preserve (a version of) the link I'm defending between intentional action and knowledge-how by appealing to *partial* knowledge-how, rather than to his disjunctive view, although he opts not to explore the former option. If he is sympathetic to appeals to partial knowledge, he may be at least sympathetic to treating KS's bomb defusing as non-intentional, being the manifestation of *merely partial* or *incomplete* knowledge-how, if not also sympathetic to the claim that manifesting incomplete knowledge-how can satisfy the link in question. At any rate, my aim is not to prove that Setiya is barking up the wrong tree; it is just to point out the theoretical space that exists between our two suggestions and certain considerations that might attract a neutral reader towards my position. To the extent that Setiya would be happy to countenance partial knowledge-how, my arguments may simply exhibit a novel route to a somewhat familiar position, and that's fine by me.

propositional. Thus, one might think that I am appealing to an especially important *disposition*, the possession and manifestation of which guides one's actions so as to render them intentional. A careful reader might wonder how these two ideas square with one another.

The first thing to emphasize in response to this sort of concern is that, even though ability-constituting knowledge of inference (for simple, single-premise deduction) can be represented by pointing to knowledge of an indicative conditional *If (P and If P, then Q), then Q,* this does not necessarily render such knowledge propositional.[79] It might simply be the most perspicuous way of representing that an agent has *knowledge of how to infer*. On this reading, the surface-level 'propositionality' of ability-constituting knowledge of inference is a red herring.

The second thing to emphasize is that I have been discussing the relationship between ability-constituting knowledge and knowledge-how, trying to set aside the relationship between knowledge-how and knowledge-that. As I see it, questions surrounding the relative metaphysical priority of knowledge-how and knowledge-that are orthogonal to whether there is a special form of knowledge-how, ability-constituting knowledge, necessarily manifested in rational inference and intentional action (whatever knowledge-how turns out to be, "at bottom").

Still, some accounts of knowledge-how would, more or less by fiat, rule out ability-constituting knowledge of *inference* as a form of know-how, if only because inference is a "theoretical" concern, rather than a "practical" one.[80] I am, admittedly, not

---

[79] The so-called "Ramsey test", derived from Ramsey (1929), treats one's commitment to indicative conditionals as a kind of marker for how agents would change their attitudes in light of various suppositions. But on these sorts of views, indicative conditionals do not have truth conditions; "knowing" such conditionals, then, is a matter of encoding an agent's propensities to accept certain propositions on the supposition that others hold. For contemporary proponents, see Edgington (1991; 2009), Sturgeon (2020).

[80] Some authors have argued that knowledge-how is to intention as knowledge-that is to belief; *attempts* at knowing-how leave a *residue of intention*, just as *attempts* at knowing-that leave a

tempted to think that ability-constituting knowledge *of inference* leaves a residue of intention, residues of intention being a purportedly characteristic mark of knowledge-how.[81]

Perhaps this means only that extant accounts of knowledge-how are raised on a one-sided diet of cases: ones that are *action-* and *intention-centered*. If, however, we abstract away from the action- and intention-centeredness of the cases upon which accounts of knowledge-how are typically built, then we might see ability-constituting knowledge as something like the kernel of knowledge-how, which, rather than having certain characteristic outputs (e.g., action as opposed to judgment), has certain characteristic normative functions: in particular, that of "linking" one's reasons to one's reasons-based performances in various first-order normative domains.

Let's briefly take stock. Short of offering a metaphysics of ability-constituting knowledge, I hope to have accomplished two lesser tasks. First, I hope to have characterized ability-constituting knowledge in general, and ability-constituting knowledge of action in particular, in terms of its normative role. Ability-constituting knowledge in general plays the role of connecting reasons to reason-based performances, and ability-constituting knowledge of action plays the role of connecting

---

residue of belief. Others, influenced by Ryle, think that knowledge-how is essentially dispositional, while knowledge-that is essentially propositional. And still others think that the difference between knowledge-how and -that is broadly Fregean; it consists in a difference of mode of presentation, not in metaphysical status (Stanley 2011; Stanley and Williamson 2001). More recently, some have suggested that the difference between knowledge-how and -that is explained in terms of the aims to which certain information is indexed; sometimes information is indexed to giving explanations, sometimes to performing more distinctly practical tasks, sometimes to both (Elga and Rayo 2018).

[81] I am bracketing the view that "actions" and "intentions" might be understood in a thoroughly deflationary way so as to treat certain inferences or judgments as intentional actions without collapsing into a form of doxastic voluntarism. This view may, at bottom, be no more of a terminological variant on my own suggestion, but I am trying to remain as neutral as possible on the nature of intentions *per se*, so I will not discuss it at length. Thanks to XXXX for raising this possibility.

reasons for action to actions. Second, I hope to have shown that the thing playing this particular normative role is not metaphysically suspicious. If, as I have argued, ability-constituting knowledge (of both inference and action) is a special kind of knowledge-how, we could think of ability-constituting knowledge as propositional knowledge under a special mode of presentation, as propositional knowledge indexed to certain contexts or aims, or as something non-propositional, perhaps as a bundle of special dispositions (whichever turns out to be *the* correct view of the nature of knowledge-how and its relationship to knowledge-that, which, again, is a further question I have not attempted to address).

4. Whither inferential knowledge of action?

Recall that, when causalists argue against various epistemic constraints on intentional action, they often focus on cases in which, intuitively, an agent acts intentionally, but it is difficult to attribute a certain sort of knowledge to them. Consider two such cases:

> **Carbon copier:** The man writing heavily on the page with the intention of producing ten carbon copies does not know or believe he is succeeding. In fact, he positively believes he will fail. Nevertheless, he succeeds in producing ten carbon copies. (Davidson 1971)

> **Distracted driver:** The distracted driver, who is operating on 'autopilot', stops at the light, signals to turn, and arrives home safely without realizing that she is doing so. The driver, since she is operating on autopilot, has no opinion one way or the other as to what she is doing. (Paul 2009)

One might wonder what, according to causalists, distinguishes **Distracted driver** and **Carbon copier** cases from that of **Darts**. Why isn't it the case, for instance, that the

distracted driver is just "too lucky" to count as having intentionally turned right? Here is Sarah Paul's explanation:

> "To express the fact that his turning and signaling are purposive, have a means-end structure, and are performed in response to his reasons, it is natural to say that he turns and signals intentionally. If his driving is habitual enough, however, he may have no belief that he is doing these things and might only be able to discover that he is doing them by observation (e.g, noticing that the turn signal is on)." (p.5, 2009)

According to Paul and many prominent accounts of action, what grounds the distracted driver's intentionally turning right is *that her intentions bring about her turning right in the right way.*[82] It is plausible to think that, in **Darts**, Al's intention to hit the bullseye does not bring about his hitting the bullseye in the right way, one which is sufficiently "purposive, [has] a means-end structure, and [is] performed is response to his reasons". For those inclined to use the language of "control", it seems fairly uncontroversial that Al exercises an impoverished kind of control, if any at all, over his hitting the bullseye.

Like many of the cases we've been concerned with, it may be tempting, at least at first blush, to think that **Carbon copier** and **Distracted driver** are ones in which the agents act intentionally without knowing what they are doing as they do it; if Paul is correct, each agents' intentions do all the explanatory heavy lifting. From this tempting thought, it is a short step to conclude *so much the worse for an epistemic constraint on intentional action*.

But as I've argued, ability-constituting knowledge can come apart from the knowledge one would express by, perhaps, explaining what one is doing. We should grant that the carbon copier and distracted driver lack knowledge of the latter sort. (Of

---

[82] Bratman (1987); Davidson (1978); Grice (1971), although he uses the language of 'the will' instead of 'intentions'; Paul (2009); Mele (1992).

course, the carbon copier's lack of knowledge consists in his having a false *de dicto* belief about what he is doing as he does it, and the distracted driver's lack of knowledge consists in her having no opinion one way or the other, but these are just two ways of failing to have a kind of knowledge that is not particularly relevant to intentional action *per se*.)

Critics of Anscombian views have not taken these cases to indicate that the epistemic constraints on intentional action should be deflated, as I have. Instead, they have tried to explain away purported epistemic constraints as merely epiphenomenal. For instance, in the same article as cited above, Sarah Paul argues that sometimes agents who act intentionally have knowledge of what they are doing as they do it, but this knowledge, when an agent has it, is merely *inferentially justified*. The thought seems to be that agents have a kind of privileged but fallible access to the contents of their own decisions, from which they can infer what they are (now) doing. Suppose than an agent, A, at some time t, decides to F. Shortly thereafter, A begins to make particular motions with her body, motions that in fact count as a way of F-ing. Because A decided to F and can remember this decision, A is in a position to know by inference that she is F-ing. (Perhaps this inference needs further premises, like *that A typically follows through on her decisions*, but that is not objectionable in and of itself.) This, the causalist seems to think, is the closest thing to a true thesis in the vicinity of non-causalism:

> **(Merely) inferential knowledge of action ("IKA"):** When an agent acts intentionally, she is typically in a position to know what she is doing as she does it in virtue of the fact that she can remember what she decided to do, and this (perhaps through some background premises) provides her with evidence that she is doing what she decided to do.

Nothing about merely inferential knowledge as such distinguishes intentional actions from unintentional ones, or actions from non-actions. Other agents can have the very same inferential knowledge of what you are doing as you do it, justified by knowledge of a prior decision of yours, and the knowledge that you tend to stick to your decisions. Surely *my* knowing what *you* are doing is not an important part of the explanation of what makes it the case that you are doing it intentionally.

Nevertheless, we should not downplay the value and importance of (merely) inferential knowledge of action. This sort of knowledge is especially important as a way of, among other things, facilitating skill acquisition (both intra- and inter-personally), and in particular by allowing the agent to correct course. Here is an example: Mike is trying to learn how to flip eggs, but he's never before stepped foot in a kitchen (Mike always orders takeout). I explain to Mike, in very precise terms and making the corresponding gestures, that you flip eggs by making certain grabbish and wrist-flickish motions. Suppose that all the conditions are in place so that Mike thereby gains testimonial knowledge of how to flip eggs. Since he's seen my gestures accompanying the testimony, he is not totally in the dark as to how to flip eggs, but because he's never been in a kitchen, it seems safe to deny that he thereby gains ability-constituting knowledge of how to flip eggs. Now I present Mike with a pan of eggs ready to flip, and I ask him to give it a try. Mike struggles at first, but he is a quick learner; he is good about observing himself as he tries things, whereby he learns whether he is acting in a way that counts as successfully flipping eggs. To have that sort of knowledge would seemingly enable agents to "correct course", so to speak, by (dis)confirming that they are succeeding in bringing out what they aim to bring about. Were Mike to learn that that way of grabbing and flicking is not a way to successfully flip eggs, he would correct course in the next attempt by grabbing and flicking in another way, and after sufficiently many tries at this, he may have internalized enough of the relevant information (either by acquiring enough of the relevant dispositions, or acquiring

91

enough of the relevant propositional knowledge under a practical mode of presentation, or…) to flip successfully to count as having ability-constituting knowledge of how to flip eggs.

This little vignette is just meant to bring out another consideration one might cite on behalf of appeals to inferential knowledge of what one is doing as they do it. To wit, it enables rational skill-acquisition via course-correction. But of course, for Mike to infer by way of observation and testimony that what he is doing--grabbing and flicking in whatever way--does not count as a way to successfully flip eggs, he must already know what he is doing as he does it in another, distinctively non-inferential way. It cannot simply be inferential knowledge all the way down.

5. Lingering objections

For a moment, consider what a die-hard inferentialist might say in reply to section 3: "I agree that inferential knowledge of action is valuable in that it enables skill acquisition and course-correction. What I find contentious is the claim that course-correction by inference presupposes non-inferential knowledge of what one is doing as one does it. This seems to (i) beg the question against proponents of IKA in favor of AKA, and (ii) invite a quasi-Cartesian reading of AKA."

In response to (i), I have comparatively little to say. My aim is not, in the first place, to put forward a view of intentional action that follows from premises anyone could accept. Instead, it is to engage with a set of fairly specific positions and the cases they take to be inconsistent with an epistemic constraint on intentional action. Those fairly specific positions and cases, I argued, are only problematic for a narrow band of views, ones which are rightly criticized as overly intellectual. If we opt for a more deflationary reading of the relevant knowledge manifested in action (and inference), then what were once hard cases become much easier to explain. Moreover, appeals to

the more deflationary kind of knowledge are not *ex nihilo*; they are motivated by meta-normative considerations—drawing on theories of epistemic akrasia and inadvertent moral worth.

I am much more concerned with objection (ii). Here, the thought seems to be that ability-constituting knowledge is still a kind of *knowledge*, and we, good contemporary epistemologists, are staunchly *anti-Cartesian*. That is, we think that all knowledge is defeasible. Billy's case (from **Bike**), for instance, is only compelling if we think that ability-constituting knowledge is *indefeasible*, so that case is not compelling.

This objection is helpful because it forces us to think hard about the relationship between having ability-constituting knowledge of action, on the one hand, and the evidence one acquires in skill-acquisition, on the other. Moreover, the objection is helpful because it is partially correct. Billy's case is only compelling if we think that ability-constituting knowledge is insensitive to or insulated from certain kinds of evidence, and so, in some weak sense "indefeasible". However, this "indefeasibility" is not objectionable; it follows from a fairly natural picture of our agency.

The first thing to say in response to (ii) is that it is too strong. For all I've said, it is entirely possible for an agent to lose *some* kinds of ability-constituting knowledge in the face of (perhaps misleading) evidence. Most clearly, some of our ability-constituting knowledge is itself constituted by background evidence—the background evidence accumulated in acquiring the ability itself. Knowing how to hail a taxi involves acquiring and retaining a good deal of evidence about what to do and when to do it. To the extent that one's knowledge of how to hail a taxi is constituted by certain background evidence that can be defeated, so too can one's ability-constituting knowledge of how to hail a taxi.

The ability-constituting knowledge one manifests in, for instance, raising one's arm above one's head is importantly different from the knowledge one manifests in hailing a taxi. (Statistically) normal, bodied agents do not know how to raise their arms

93

by doing other things. It follows that such agents do not know how to raise their arms by learning which other things count as arm-raising under which conditions. Moreover, unlike the knowledge of how to hail a taxi, the knowledge of how to raise one's arm cannot be defeated by defeating the background evidence that partly constitutes the ability itself, since the ability itself is not even partly constituted by background evidence in the first place.

With the distinction between basic and non-basic ability constituting knowledge in place, I don't need to go as far as claiming that ability-constituting knowledge of action is *indefeasible*. What I claim is much more modest: basic ability-constituting knowledge is *insulated* from defeating evidence, while non-basic ability-constituting knowledge is not.

One peculiar case from the action theory literature brings out this phenomenon of insulation fairly well:

> **Temporary paralysis**: Susie's arm has been temporarily paralyzed for a medical procedure. As she is leaving the doctor's office, she is informed that the paralysis should subside after two to four hours. At home, she sits anxiously, often checking her watch, which indicates that only 90 minutes have passed. Susie firmly believes that she is presently unable to raise her arm, and this seems rational in light of the information she received at the doctor's office. On a whim, Susie thinks to herself, *what harm is there in trying?* She proceeds to raise her arm. (Paul 2009; Setiya 2009)

Susie's knowledge of how to raise her arm is insulated from the evidence that might be thought to defeat it. In a sense, Susie is rather lucky; her paralysis wears off much earlier than she should have expected, given the information from the doctor's office. Her judgments are, rightly to mind my, guided by this expert information, and this explains why she believes that she cannot raise her arm, even as she begins to raise it. Nevertheless, this is not the sort of luck that makes us think Susie merely

94

unintentionally raised her arm, nor is it the sort of luck that makes us think Susie didn't *act* at all. Rather, Susie receives strong empirical evidence that defeats whatever knowledge she would express by *explaining* what she is doing as she does it, but she surely retains and manifests knowledge of another sort. This is precisely what makes her arm-raising an *intentional* action (albeit an inadvertent one), as opposed to merely unintentional action or mere behavior.

Still, one might hope, instead of simply being given cases in which it is natural to think that an agent's basic-ability constituting knowledge is insulated from defeating evidence, to be given a general argument that basic ability-constituting knowledge—or *any* knowledge, for that matter—is ever *insulated* from defeating evidence.

Here, I can only offer the outline of a general argument, which proceeds by drawing on theories of inadvertent epistemic virtue. For instance, Brian Weatherson has suggested that, in cases of inadvertent epistemic virtue, an agent's first-order evidence "evidentially screens-off" her judgments *about* her first-order evidence.[83] That is not to deny that, in general, such judgments are evidentially relevant, nor is it to deny that "evidence of evidence is evidence"[84]; rather, it is to give voice to the thought that one's first-order evidence has priority over one's judgments about one's evidence in determining what one ought to believe. Aki the akratic, convinced of testimonial skepticism, nevertheless finds herself with what is in fact strong first-order testimonial evidence to believe that P. Her strong first-order testimonial evidence that P evidentially screens-off her judgments about what her evidence supports. This is how first-order evidence can be "insulated" from judgments about what one's first-order evidence supports. To borrow a phrase from Nomy Arpaly, Aki's judgments *about* what

---

[83] He says: "facts about one's own judgment are not evidentially relevant to what judgment one makes, provided one has access to the evidence behind that judgment. And that suggests that the judgment should really just be judged on how it tracks the evidence." (2019, p. 198)
[84] Comesaña and Tal (2015; 2017)

her evidence supports play a rather superficial role in the drama of her rational inference.[85]

This metaphor can be extended to the action theoretic cases we've considered. In cases like **Temporary paralysis,** the explanation of Susie's intentionally raising her arm does not appeal at all to her rational, false judgments about her ability to raise her arm, or to her disposition to make false judgments about what she is doing as she begins to raise her arm. If, instead, Susie were to have never considered the question of whether her arm was still paralyzed, were she to have mistakenly thought that more than 90 minutes had passed, or were she to have disregarded the doctor's information on entirely irrational grounds, her raising her arm would still have been an intentional action. This suggests that these judgments, and the dispositions underwriting them, seem to play a rather superficial role in the drama of her intentional action. Even if such judgments are not, strictly speaking, "evidentially screened-off" by one's basic ability-constituting knowledge, there is clearly a sense in which they are screened-off explanatorily.[86]

---

[85] Arpaly suggests that one's judgments about what one ought to do sometimes play a "rather superficial role in a drama in which a person is motivated by sinister [or virtuous, as the case may be] desires…" (2002, p. 106).

[86] I suspect that offering a fuller account of how basic ability-constituting knowledge is insulated from defeating evidence would require adopting a particular metaphysics of ability-constituting knowledge, which I have tried to avoid doing. In part, that's because it seems plausible to think that insulation reflects differences between ability-constituting knowledge and other kinds of knowledge. Whether that difference is ultimately cashed out in terms of differing modes of presentation, in terms of dispositional versus propositional knowledge, or whatever, is an issue for another time. This is, of course, as much a promissory note for an account of insulation as it is a defense of that notion. Nevertheless, I hope to have shown that appeals to insulated knowledge are not objectionable in and of themselves.

6. Summary

I have argued that one distinguishing mark of intentional action is that the agent knows what she is doing as she does it. The sort of knowledge in question is captured by AKA:

> **Ability-constituting knowledge of action ("AKA"):** When an agent X-s intentionally, she manifests certain non-inferential ability-constituting knowledge as she X-s.

Ability-constituting knowledge is not the kind of knowledge that serves as a reason to act or believe; rather, it guides these sorts of reason-based performances. For instance, competent cyclists often lack whatever knowledge enables them to correctly explain how to steer, while nevertheless knowing that one initiates a right turn by steering left; having this knowledge is part of what their competence consists in. That very knowledge may be merely *de re* or sub-personal, inconsistent with one's *de dicto*, person-level judgments.

If my arguments are correct, I have identified a form of knowledge the manifestation of which is necessary for different kinds of reasons-based performances (in action and judgment) to count as "successes". Moreover, I have shown that the manifestation of that knowledge is what explains why cases of inadvertent reasons-based performances are successes (in action and judgment), rather than things that merely *happen* to the agent. I have offered a way to think about ability-constituting knowledge in terms of knowledge-how, but this invites a number of questions about proprietary distinctions within extant accounts of knowledge-how, and I have, admittedly, only gestured at a complete answer.

Without the space to fully explore these ideas here, in closing I want to note two advantages of accepting AKA. The first should appeal to philosophers regardless of traditional commitments: AKA supplies principled reason to admit of an epistemic constraint on intentional action that side-steps familiar problems surrounding the nature of intention itself. The position developed here does not depend on any positive

view about the nature of intentions, whether they are belief-like, desire-like, a hybrid belief-like and desire-like attitude, for instance.[87] Nor does it depend on whether and to what extent the contents of one's intentions must relate to (a particular description of) one's actions in order to render those actions intentional (under that description). These are extremely important and interesting issues that any complete theory of action must address. But I hope to have shown that, whatever the final verdict on the nature and content of intentions *per se*, one must appeal to something like AKA to explain how intentions—however understood—could perform their characteristic *guiding* role.

The second advantage of accepting AKA should appeal, I think, primarily to causalists. Namely, AKA is consistent with views in the causalist tradition. This may come as a surprising result, given the way that causalist and Anscombian views are often pitted against one another. The apparent incompatibility of these traditions seems to be driven in no small part by an assumption that debates around an epistemic condition on intentional action are settled by considerations of the nature of intention *per se*. If my remarks above are correct, I have cast doubt on that assumption and opened up a section of previously unoccupied theoretical space, one that ostensibly inherits the benefits of both causal and Anscombian theories of action. Thus, to the extent that reconciliation possible, I have tried to reconcile causalist and Anscombian theories of action.

Why should this be appealing? Causal theories of action are *incomplete,* and just about everyone admits it; even if intentional actions necessarily have certain

---

[87] An anonymous referee wonders how we might distinguish *bona fide* intentional actions from merely highly routinized behavior. Strictly speaking, I am offering only a necessary condition on intentional action, so further conditions likely distinguish *bona fide* intentional actions from merely routine or automatic ones, many of which might manifest ability-constituting knowledge (a basketball player's off-ball movement in a zone defense, for instance). Without the space to go into detail, I believe that an account of intentions, which I have tried to avoid discussing, will supply the missing conditions.

distinguished mental causal antecedents, causalists notoriously have problems with cases of *causal deviance*. What does its deviance consist in, or to put it differently, what would have to be added back into cases of deviance to turn them into cases of *bona fide* intentional action? On the simplest version of this causalist-*cum*-Anscombian view, an agent's behavior counts as an intentional action only if her ability-constituting knowledge is itself a cause of that behavior. We might then say, for instance, that in **Bike**, Billy turns right intentionally because not only are his antecedent beliefs, desires, and intentions are a cause of his turning right, so too is Billy's ability-constituting knowledge of how to turn right. On more complex versions of this kind of view, the ability-constituting knowledge is not itself a cause of Billy's turning right, but it nevertheless grounds the relevant causal relations. We might then say that Billy's ability-constituting knowledge grounds the fact that his intention to turn right causes his turning right. Whatever the details of the particular causal theory, AKA can help us say more about what it is to have the right kind of cause, rather than compete with such theories for an explanation of what it is to act intentionally at all.[88]

References for *Good Thinking*

Alston, William, 1980, "Level-Confusions in Epistemology," *Midwest Studies in Philosophy* 5 (1):135-150.

---------, 1983, "What's Wrong with Immediate Knowledge?" *Synthese* 55: 73-96.

Anscombe, Elizabeth, 1957, *Intention*, Cambridge, MA: Harvard University Press.

Arpaly, Nomy, 2002, *Unprincipled Virtue: An Inquiry into Moral Agency*, Oxford University Press.

Arpaly, Nomy and Timothy Schroeder, 2013, *In Praise of Desire*, Oxford University Press.

Boult, Cameron, 2016, "Epistemic Normativity and the Justification-Excuse Distinction," *Synthese*, DOI 10.1007/s11229-016-1127-8.

---------, 2019, "Excuses, Exemptions, and Derivative Norms," *Ratio* 32 (2): 150-158

Bratman, Michael, 1984, "Two Faces of Intention", *Philosophical Review*, 101: 375-405.

---------, 1987, *Intentions, Plans, and Practical Reason*, Cambridge, MA: Harvard University Press.

Buckwalter, Wesley and John Turri, 2018, "Inability and Obligation in Intellectual Evaluation", *Episteme*, DOI 10.1017/epi.2018.49.

Carroll, Lewis, 1895, "What the Tortoise Said to Achilles", *Mind*, Volume 4, 14: 278-280.

Chomsky, Noam, 1965, *Aspects of a Theory of Syntax*, MIT Press.

Chudnoff, Elijah, 2020, *Forming Impressions: Expertise in Perception and Intuition*, Oxford University Press.

Cohen, Stewart, 1984, "Justification and Truth," *Philosophical Studies* 46 (3): 279-95.

---------, 2002, "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65 (2):309-329.

---------, 2005, "Why Basic Knowledge is Easy Knowledge," *Philosophy and Phenomenological Research* 70 (2):417-430.

--------, 2010, "Bootstrapping, Defeasible Reasoning, and a priori Justification," *Philosophical Perspectives* 24 (1):141-159.

--------, 2016, "Theorizing about the Epistemic," *Inquiry: An Interdisciplinary Journal of Philosophy* 59 (7-8): 839-857

Comesaña, Juan, 2020, *Being Rational and Being Right*, Oxford University Press.

Comesaña, Juan and Carolina Sartorio, 2014, "Difference-making in Epistemology," *Nous* 48:368-387.

Comesaña, Juan and Eyal Tal, 2015, "Evidence of evidence is evidence (trivially)," *Analysis*, Volume 75, 4: 557-559.

--------, 2017, "Is Evidence of Evidence Evidence?" *Nous*, Volume 51, 1: 95-112.

Conee, Earl and Richard Feldman, 1985, "Evidentialism," *Philosophical Studies* 48 (1):15-34.

--------, 2001, "Internalism Defended," *American Philosophical Quarterly* 38 (1):1-18.

Davidson, Donald, 1971, "Agency," reprinted in *Essays on Actions and Events*, Oxford: Oxford University Press, 1980.

--------, 1978, "Intending," reprinted in *Essays on Actions and Events*, Oxford: Oxford University Press, 1980.

Edgington, Dorothy, 1991, "The Mystery of the Missing Matter of Fact," *Proceedings of the Aristotelian Society*, Volume 65: 185-209.

--------, 2009, "Conditionals, Truth and Assertion," in I. Ravenscroft (ed.), *Minds, Ethics, and Conditionals*, Oxford: Oxford University Press.

Elga, Adam and Agustin Rayo, "Fragmentation and information access," forthcoming in C. Borgoni, D. Kindermann, and A. Onofri, eds. *The Fragmented Mind*, Oxford: Oxford University Press.

Fantl, Jeremy and Matthew McGrath, 2009, *Knowledge in an Uncertain World*, Oxford University Press. doi:10.1093/acprof:oso/9780199550623.001.0001

Field, Hartry, 2009, "What is The Normative Role of Logic?" *Proceedings of the Aristotelian Society* Supplementary Volume 83: 251-268.

Fischer, John and Martin Ravizza, 1998, *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge University Press.

Foley, Richard, 1993, *Working Without a Net*, Oxford University Press.

Frankfurt, Harry, 1971, "Freedom of the Will and the Concept of a Person," *Journal of Philosophy* 66: 5-20.

Fumerton, Richard, 1995, *Metaepistemology and Skepticism*, Rowman & Littlefield.

Grice, H. P., 1971, "Intention and Uncertainty," *Proceedings of the British Academy*, 5: 263-279.

Hall, Ned, 1994, "Correcting the Guide to Objective Chance," *Mind* 103 (412):505-518.

Harman, Gilbert, 1976, "Practical Reasoning," *Review of Metaphysics* 79: 431-63.

---------, 1986, *Change in View: Principles of Reasoning*, MIT Press.

---------, 1986, "Willing and Intending." In R. Grandy and R. Warner, eds., *Philosophical Grounds of Rationality*, Oxford: Clarendon Press.

Hasan, Ali, 2013, "Phenomenal Conservatism, Classical Foundationalism, and Internalist Justification," *Philosophical Studies* 162 (2):119-141.

Huemer, Michael, 2006, "Phenomenal Conservatism and the Internalist Intuition," *American Philosophical Quarterly* 43, 147-158.

---------, 2013, "Phenomenal Conservatism Uber Alles," In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Chris Tucker (ed.), Oxford University Press.

Ichikawa, Jonathan and Benjamin Jarvis, 2013, *The Rules of Thought*, Oxford University Press.

Kenny, Anthony, 1975, *Will, Freedom, and Power*, Oxford: Blackwell Publishing.

Kornblith, Hilary, 1983, "Justified Belief and Epistemically Responsible Action," *The Philosophical Review* 92 (1):33-48.

Lasonen-Aarnio, Maria, 2014, "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research* 88 (2):314-345.

---------, "Virtuous Failure and the Victims of Deceit," Forthcoming in Dutant and Dorsch (eds), *The New Evil Demon*, Oxford University Press.

Levy, Neil, 2007, "Doxastic Responsibility", *Synthese* 155: 127-155

Lewis, David, 1976, "The Paradoxes of Time Travel", *The American Philosophical Quarterly* 13 (2): 145-152.

Littlejohn, Clayton, 2012, *Justification and the Truth Connection*, Cambridge University Press.

---------, "A Plea for Epistemic Excuses", Forthcoming in Dutant and Dorsch (eds), *The New Evil Demon*, Oxford University Press.

Maier, John, 2013, "The Agentive Modalities", *Philosophy and Phenomenological Research*. 90 (1): 113-134.

Markie, Peter, 2013, "Searching for True Dogmatism," In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, Chris Tucker (ed.), Oxford University Press.

---------, 2015, "The Special Ability View of knowledge-how," *Philosophical Studies* 172: 3191-3209.

Martin, C.B., 1994, "Dispositions and Conditionals," *The Philosophical Quarterly* 44 (147): 1-8.

Mason, Elinor, 2016, "Moral Incapacity and Moral Ignorance," *Perspectives on Ignorance From Moral and Social Philosophy*, Rik Peels, ed. Routledge Press.

McHugh, Conor, 2013, "Epistemic Responsibility and Doxastic Agency," *Philosophical Issues*, 23: 132-157

McKenna, Michael, 2012, *Conversation and Responsibility*, Oxford University Press.

---------, 2013, "Reasons-Responsiveness, Agents, and Mechanisms," In *Oxford Studies in Agency and Responsibility*, David Shoemaker, ed. Oxford University Press: 151-84.

Mele, Alfred, 1992, "Recent Work on Intentional Action," *American Philosophical Quarterly*, Volume 29, 3: 199-217.

---------, 1992, *The Springs of Action*, New York: Oxford University Press.

---------, 2001, *Autonomous Agents*, Oxford: Oxford University Press.

---------, 2002, "Agents' Abilities," *Noûs*, 37: 447–470.

Mele, Alfred and Paul Moser, 1994, "Intentional Action", *Nous*, Volume 28, 1: 39-68.

Nelkin, Dana, 2011, *Making Sense of Freedom and Responsibility*, Oxford University Press.

O'Shaughnessy, Brian, 1980, *The Will*, Cambridge: Cambridge University Press.

Owens, David, 2000, *Reason Without Freedom: The Problem of Epistemic Normativity*, Routledge.

Paul, Sarah, 2009, "How We Know What We're Doing," *Philosopher's Imprint*, Volume 9, 11: 1-24.

Pereboom, Derk, 2001, *Living Without Free Will*, Cambridge University Press.

---------, 2014, *Free Will, Agency, and the Meaning of Life*, Oxford University Press.

Plantinga, Alvin, 1983, "Reason and Belief in God," In *Faith and Rationality: Reason and Belief in God*, Alvin Plantinga & Nicholas Wolterstorff (eds.). University of Notre Dame Press.

Pollock, John, 1971, "Perceptual Knowledge," *The Philosophical Review* 80 (3):287-319.

---------, 1987, "Defeasible Reasoning," *Cognitive Science* 11: 481-517.

Pryor, James, 2000, "The Skeptic and the Dogmatist," *Noûs* 34(4): 517–549.

Ramsey, Frank, 1929, "General Propositions and Causality," in D. H. Mellor (ed.), *F.P. Ramsey, Philosophical Papers*, Cambridge: Cambridge University Press.

Ryle, Gilbert, 1949, *The Concept of Mind*, London: Hutchinson.

Scanlon, Thomas, 2008, *Moral Dimensions: Permissibility, Meaning, Blame*, Harvard University Press.

Schaffer, Johnathan, 2010, "The Debasing Demon", *Analysis* 70 (2): 228-237.

Setiya, Kieran, 2008, "Practical Knowledge", *Ethics*, Volume 118, 3: 388-409.

--------, 2009, "Practical Knowledge Revisited", *Ethics*, Volume 120, 1: 128-137.

--------, 2017, *Practical Knowledge: Selected Essays*, Oxford: Oxford University Press.

Shope, Robert, 1978, "The Conditional Fallacy in Contemporary Philosophy", *The Journal of Philosophy* 75 (8): 397-413.

Siegel, Susanna, 2010, *The Contents of Visual Experience*, Oxford: Oxford University Press.

Siegel, Susanna and Nicholas Silins. "The Epistemology of Perception" (2015). In *The Oxford Handbook of Philosophy of Perception*, Mohan Matten (ed.). Oxford University Press.

Silva, Paul, 2017, "Knowing How to Put Knowledge First in the Theory of Justification", *Episteme,* 14 (4): 393-412.

Smith, Angela, 2012, "Attributability, Answerability, and Accountability: In Defense of a Unified Account", *Ethics* 122 (3): 575-589.

Sosa, Ernest, 2007, *Apt Belief and Reflective Knowledge, Volume I: A Virtue Epistemology*, Oxford: Oxford University Press.

--------, 2010, *Knowing Full Well,* Princeton: Princeton University Press.

Stanley, Jason, 2005, *Knowledge and Practical Interests,* New York: Oxford University Press. doi:10.1093/0199288038.001.0001

--------, 2011, *Know How*, Oxford: Oxford University Press.

Stanley, Jason and Timothy Williamson, 2001, "Knowing How," *The Journal of Philosophy*, 98(8): 411-444.

Strawson, Peter, 1962, "Freedom and Resentment," In *Proceedings of the British Academy* 48: 1-25.

Sturgeon, Scott, 2020, *The Rational Mind*, Oxford University Press.

Sylvan, Kurt, 2015, "What Apparent Reasons Appear to be," *Philosophical Studies* 172 (3): 587-606.

Thagard, Paul, 1982, "From the Descriptive to the Normative in Psychology and Logic," *Philosophy of Science* 49 (1):24-42.

Titelbaum, Michael, *Fundamentals of Bayesian Epistemology*, forthcoming in Oxford University Press.

Turri, John, 2010, "On the Relationship Between Propositional and Doxastic Justification," *Philosophy and Phenomenological Research* 80 (2):312-326.

Tversky, Amos and Daniel Kahneman, 1983, "Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment," *Psychological Review* 90(4): 293-315.

Van Inwagen, Peter, 1983, *An Essay on Free Will*, Oxford University Press.

---------, 2015, "Some Thoughts on *An Essay on Free Will*," *The Harvard Review of Philosophy* 22: 16-30.

Velleman, David, 1989, *Practical Reflection*, Princeton University Press.

Vihvelin, Kadri, 2004, "Free Will Demystified: A Dispositional Account," *Philosophical Topics 32*: 427-450.

---------, 2013, *Causes, Laws, and Free Will: Why Determinism Doesn't Matter*, Oxford University Press.

Vogel, Jonathan, 2000, "Reliabilism Leveled," *The Journal of Philosophy* 97 (1):602-623.

Wallace, R Jay, 1994, *Responsibility and the Moral Sentiments*, Harvard University Press.

Weatherson, Brian, 2019, *Normative Externalism*, Oxford University Press.

Wedgewood, Ralph, 2011, "Justified Inference," *Synthese* 189: 273-295.

---------, 2013, "A Priori Bootstrapping," In *The A Priori in Philosophy*, Albert Casullo and Joshua C. Thurow (eds.), Oxford University Press.

Weisberg, Jonathan, 2010, "Bootstrapping in General," *Philosophy and Phenomenological Research* 81 (3):525-548.

White, Roger, 2006, "Problems for Dogmatism," *Philosophical Studies* 131 (3):252-557.

Whittle, Anne, 2010, "Dispositional Abilities," *Philosopher's Imprint* 10(12): 1-23.

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

---------, 2011, "Improbable Knowing," In *Evidentialism and its Discontents*, Trent Dougherty (ed). Oxford University Press.

---------, 2014, "Very Improbable Knowing," *Erkennetis* 79: 971-999.

---------, "Justification, Excuses, and Skeptical Scenarios," Forthcoming in Dutant and Dorsch (eds), *The New Evil Demon*, Oxford University Press.

Wolf, Susan, 1990, *Freedom Within Reason*, Oxford University Press.