

Why Ideal Epistemology?

1
2
3
4
5

6 1 Ideal and nonideal epistemology

7 What are ideal and nonideal epistemology? We can begin by gesturing toward some loose sociological
8 trends:

nonideal epistemology	ideal epistemology
informal, non-mathy	formal, mathy ¹
talks about beliefs	talks about credences ²
emphasizes relation to knowledge	emphasizes relation to decision theory
uses “justified”	uses “rational”
uses “reasons” language	doesn’t
unimpressive babies	superbabies
written in Word	written in LaTeX
⋮	⋮

9
10 Ideal epistemologists are concerned with questions about what perfectly rational, cognitively ide-
11 alized, computationally unlimited believers would believe. (Note: I use the term “belief” broadly for any
12 doxastic state, including binary belief, credences, comparative confidence, etc.; mutatis mutandis for
13 “believe” and “believer”.) Often this involves presupposing or defending epistemic norms that, arguably,
14 no actual humans can satisfy: norms mandating

- 15 ○ logical omniscience;

¹ Note that despite this trend, the ideal/nonideal distinction in epistemology is strictly orthogonal to the formal/informal distinction. Theorists of bounded rationality often pursue formal nonideal epistemology; I personally often do informal ideal epistemology.

² Again, despite the trend, this distinction is orthogonal to the ideal/nonideal distinction: AGM models of belief revision are a form of ideal epistemology, for example (Alchourrón et al., 1985). And much of the bounded rationality literature focuses on credence.

- 1 o consistency and closure of binary beliefs;
- 2 o infinitely precise credences;
- 3 o credences satisfying the Kolmogorov probability axioms;
- 4 o updating by conditionalization;
- 5 o immediate update (rather than temporally extended reasoning);
- 6 o closure of doxastic attitudes under boolean operations;

7 ...and so on.³ None of these norms is universally accepted by ideal epistemologists. But where ideal
8 epistemologists reject a norm on this list, their reasons for doing so don't stem from the cognitive
9 demandingness of the norm.

10 By contrast, nonideal epistemologists are concerned with questions about epistemic norms that are
11 satisfiable by most humans much of the time. The distinction between ideal and nonideal epistemology
12 is vague. There are philosophers who work in both research programs, and there are philosophers who
13 work at the borderline.

14 Ideal epistemologists are often asked to defend their research program. We're often asked: why
15 spell out norms that only apply to fictional creatures in distant possible worlds? The question rests
16 on a false presupposition: the norms described by ideal epistemology are meant to apply to actual
17 humans. We just fail to satisfy them. Still, the underlying concern might be cashed out in one or more
18 of the following challenges:

19 *Challenge #1: Changing the subject.* By the standards of ordinary epistemic evaluations,
20 many of us are rational. Ideal epistemology claims that none of us is rational. So ideal
21 epistemology is either incompatible with our ordinary epistemic intuitions (and so re-
22 quires an error theory), or else uses "rational" in some other sense and so changes the
23 subject from ordinary epistemic evaluation.

24 *Challenge #2: "Ought" implies "can".* We humans *cannot* satisfy norms of ideal epistemology,
25 given our cognitive limitations: our limited computational power, informational storage
26 space, processing speeds, integration of different cognitive systems, information reten-
27 tion, and so on. But intuitively, "ought" implies "can". So the ideal epistemologist must
28 explain how it's possible for the norms of ideal epistemology to apply to us.

³ In a general theory of ideal rationality, we would also include ideal *practical* constraints, e.g.:

- o weak preferences that are transitive and complete, satisfy the Sure Thing Principle, etc. (i.e., the kinds of constraints assumed in representation theorems);
- o preferences that maximize expected utility relative to the agent's credence function and utility function.

This paper focuses on the case of epistemic rationality, but some of its observations extend to ideal and nonideal practical rationality.

1 *Challenge #3: Pointlessness.* Ideal epistemologists research norms that humans universally
2 fail to satisfy. These norms don't distinguish between reasonable, conscientious evalua-
3 tors of evidence and conspiracy theorists in tinfoil hats—or more mundanely, between
4 our reasoning at our soberest and our reasoning on reason-distorting drugs. We're all
5 always irrational. So the ideal epistemologist must explain why these norms are worth
6 discussing.

7 This paper addresses each of these challenges. I aim to offer a unified interpretation of the ideal and
8 nonideal epistemic “ought”,⁴ “rational”, and so on. I argue that ideal and nonideal epistemic evaluations
9 are governed by a shared ordering over epistemic states, and differ with respect to which states they
10 treat as open possibilities. In nonideal epistemology, “ought” implies “can” in some substantive sense
11 (that is, one that is sensitive to at least some human cognitive limitations, e.g., our limited computational
12 power, informational storage, processing speeds, integration of different cognitive systems, lifespan,
13 and so on). In ideal epistemology, it doesn't.

14 Why, then, is ideal epistemology worth so much spilled ink? I argue that only ideal epistemic evalua-
15 tions are **normatively robust**: that is, they are both *non-conventional* and not *seriously context-sensitive*,
16 in senses to be explained. Nonideal epistemic evaluations are both conventional and seriously context-
17 sensitive. Nonideal epistemology depends not merely on what's epistemically valuable and how to
18 pursue it effectively, but also on *modally contingent* epistemic conventions and *contextually contingent*
19 constraints on our epistemic options. It doesn't carve nature at its normative joints. If we want a nor-
20 matively robust theory of epistemic rationality, ideal epistemology is the only game in town. But this
21 doesn't detract from the importance of nonideal epistemology. Instead, it motivates ideal epistemology
22 as an ineliminable component within a shared research program.

23 The paper proceeds as follows: §2 introduces core concepts. §3 argues that evaluations are sensitive
24 to both an *ordering* over possibilities and a set of *fixed circumstances*. In §4, I argue that the ordering
25 used in nonideal epistemic evaluations is normatively non-robust, governed by convention. In §5, I
26 argue that the fixed circumstances in nonideal epistemic evaluations are also normatively non-robust:
27 they're conventional, or seriously context-sensitive, or both. §6 explains how normatively robust ideal
28 epistemology and non-robust nonideal epistemology can happily coexist.

29 2 **Stage-setting**

30 2.1 *“Ideal”*

31 In what sense is ideal epistemology “ideal”? It's helpful to distinguish two senses:⁵

32 *Descriptive idealization:* ideal epistemology involves “idealizations” in the sense of *simplifying as-*
33 *sumptions*. In physics, it's sometimes useful to talk of the behavior of point particles on a frictionless

⁴ N.B. I use “epistemic ‘ought’” for the *deontic* necessity modals used in epistemology. The expression is unfortunately ambiguous, and sometimes used for (non-deontic) necessity given one's knowledge, e.g. “With this warm front moving in, it ought to rain tonight.”

⁵ I owe this observation to [omitted.]

1 plane. Similarly, in epistemology it's useful to talk about the doxastic states of believers who are com-
2 putationally unlimited and capable of instantaneous massive belief revisions.

3 *Normative idealization*: ideal epistemology describes norms of "ideal", in the sense of *perfect* or *max-*
4 *imal*, rationality. It doesn't concern rationality that sacrifices, or imperfect or bounded rationality.

5 On my usage, ideal epistemology is "ideal" in both senses. First, like descriptive idealizations in sci-
6 entific modeling, it allows researchers to abstract away from complicating extraneous factors, to zero
7 in on the fundamental nature of a specific concept or phenomenon (Ismael, 2016). But this requires a
8 caveat: the norms of ideal epistemology aren't mere approximations of the genuine epistemic norms.
9 They are themselves genuine epistemic norms. Second, ideal rationality is also normatively ideal, epis-
10 temically perfect. But this too requires a caveat: as we'll see, minimizing distance from the ideal may
11 not always be better.

12 2.2 "Normatively robust"

13 I introduce "normatively robust" as a term of art. A kind of epistemic evaluation is *normatively robust*
14 just in case evaluations of that kind are both *non-conventional* and not *seriously context-sensitive*.

15 By "epistemic evaluations", I mean claims about what a believer "should" or "ought" to believe, as
16 well as attributions like "(ir)rational", "(un)justified", "(un)reasonable", "(il)logical", "wise", "stupid", and so
17 on. Epistemic evaluations target believers, doxastic states, epistemic rules (functions from evidence to
18 doxastic states), and epistemic standards (individuals' total set of epistemic rules). A believer's epistemic
19 standards characterize her dispositions to respond to her epistemic situation. We might think of them
20 as Gibbardian hyperplans for her doxastic states.⁶ They need not be luminous; when adopted, they need
21 not be consciously so; they even may conflict with the standards that the agent explicitly endorses.

22 By "non-conventional", I mean that the standards on which the evaluation is based are not standards
23 arrived at by convention. *Conventions*, I stipulate, are regularities in behavior that serve some coordina-
24 tive function, where some alternative regularity in behavior would have served the same coordinative
25 function equally well, if widely adopted.⁷

26 A kind of evaluation is *seriously context-sensitive* if and only if there's no normatively privileged res-
27 olution of one or more of the context-sensitive parameters for evaluations of that kind. Some norma-
28 tive expressions are context-sensitive but not *seriously* context-sensitive: they admit of a normatively
29 privileged interpretation. For example, ethicists might accept different moral theories. But if there's
30 a unique true moral theory, then it can generate a normatively privileged interpretation of the moral
31 "ought". When people disagree about what one "morally ought" to do, they use this privileged interpre-
32 tation of "ought", on pain of talking past each other—even if they don't know what the correct moral
33 theory says (Carr, 2015). If no such privileged resolution exists for one or more of the contextual pa-
34 rameters of some kind of evaluation, then that kind of evaluation counts as *seriously* context-sensitive.
35 I explain the contextual parameters associated with epistemic "ought" in §3.1 below.

36 Why does normative robustness matter? If you go in for a certain kind of metaphysics, then you

⁶ Gibbard (2003).

⁷ This is a simplification of Lewis (1969).

1 can think of normatively robust evaluations as latching onto *normatively fundamental* properties. They
2 carve nature at its normative joints. Conventions, by contrast, depend on historical contingencies that
3 are independent of these normative joints; and if there are fundamental properties, they ought to be
4 describable in language that isn't seriously context-sensitive: they ought to provide evaluations with a
5 privileged interpretation, a privileged resolution of contextual parameters.

6 2.3 "Ought"-implies-"can" principles in epistemology

7 **Claim:** the fundamental distinction between ideal and nonideal epistemology is that nonideal episte-
8 mology holds that epistemic "ought" implies some substantive "can"; ideal epistemology does not.

9 *Prima facie*, this might seem like a problem for ideal epistemology, since "ought"-implies-"can" (OIC)
10 principles are so widely accepted in other areas of normative philosophy. But OIC principles have a
11 vexed history in epistemology. (E.g., Alston, 1988; Feldman, 2001.) One worry that has been raised for
12 epistemic OIC principles is that they are incompatible with doxastic involuntarism, the popular view
13 that beliefs are not under our voluntary control. Even if one ought to hold some belief, it doesn't seem
14 to follow that one *can* voluntarily induce the belief.

15 A somewhat extreme example: somatoparaphrenia is the delusion that one's limb belongs to some-
16 one else. Presented with proof that a limb is their own, people with somatoparaphrenia produce elabo-
17 rate confabulations about who the limb really belongs to. Intuitively, someone with somatoparaphrenia
18 *cannot* have rational beliefs about which limbs are theirs. But they're still irrational for failing to do so.
19 Indeed, they seem to provide a canonical real-world case of epistemic irrationality. In other words,
20 they *ought* to have different beliefs; but, in some sense, they *cannot*. So, one might think, even nonideal
21 epistemologists reject epistemic OIC.⁸

22 But not all OIC principles entail that "ought φ " implies "can voluntarily φ " or even implies "can
23 φ , given one's brain lesion". There's certainly *some* sense of "can" in which someone with somatopara-
24 phrenia *can* believe what their evidence supports. For example, insofar as they are capable of attributing
25 limbhood and first-personal possession, and of using demonstratives, there's a sense in which they have
26 all the necessary ingredients for forming the rational belief: *This is my limb*. So the example doesn't rule
27 out *all* epistemic OIC principles. After all, "can" is wildly flexible in its interpretation, as §3.1 explains.

28 Nonideal epistemologists hold that what one epistemically ought to do is *somehow or other* sensitive
29 to our cognitive limitations. And these sorts of sensitivity can be fleshed out as entailing corresponding
30 OIC principles, even if those principles are much weaker than "ought"-implies-"can voluntarily". Many
31 have epistemic intuitions that reflect a commitment to some form of epistemic OIC. (For example,
32 that forgetting isn't irrational because it's psychologically inevitable.) These intuitions shouldn't be
33 dismissed. Attempting to systematize a theory that accommodates such intuitions is central to the
34 project of nonideal epistemology.

35 *Objection:* ideal epistemology is also committed to some substantive OIC principle. After all, the
36 true epistemic ideal would be omniscience; but even ideal epistemologists don't treat non-omniscience

⁸ This is a worry for a *substantive* epistemic OIC principle. But even *non-substantive* epistemic-"ought"-implies-metaphysical-"can" principles may have controversial consequences: for example, that skepticism is impossible (Greco, 2012), or that it's not the case that one ought to believe, or even aim to believe, the truth (Bykvist & Hattiangadi, 2007; Castro & Vassend, 2017).

1 as irrational.⁹ Without OIC, why isn't ideal epistemology committed to saying that rationality requires
 2 us to be certain of all truths? *Two responses*: first, epistemic internalists have the option of holding
 3 that omniscience can't be a norm of rationality simply because rational statuses supervene on what's
 4 internal to the agent and omniscience doesn't. Second and more generally: surprisingly enough, there's
 5 no acceptable interpretation of such a norm, as Bykvist & Hattiangadi (2007) and Carr (manuscript)
 6 argue. For example: there's some true proposition p that you're uncertain of. If you ought to believe all
 7 actual truths, then you ought to believe p and to believe that you're uncertain of p . But in doing so, you'd
 8 have a false belief, and so violate the norm. Alternatively, if the norm is that you ought to be omniscient
 9 but not necessarily believe all *actual* truths, then you ought to believe you're actually omniscient—even
 10 though that's false. Similar problems arise for other interpretations of the truth norm. For this reason,
 11 it ought to be rejected.

12 3 Unifying ideal and nonideal epistemic “ought”

13 3.1 What “ought” and “can” mean

14 One form or another of OIC is widely accepted within moral theories and theories of practical ratio-
 15 nality. OIC also plays a role in the standard semantics for natural language modals like “ought”, often
 16 called “Kratzer semantics” (Kratzer, 1977, 1981, 1991, 2012). To understand the relation between the
 17 ideal and nonideal epistemic “ought”, it's worth rehearsing a simplified summary of Kratzer semantics.

18 “Ought” is context-sensitive.¹⁰ When context determines that it's being used *deontically*, it is sensi-
 19 tive to two features of context:

- 20 ○ *normative parameter*: an ordering of worlds by ideality;
- 21 ○ *descriptive parameter*: a set of circumstances that are treated as fixed.¹¹

22 Kratzer semantics formalizes the intuitive thought: *you ought to do the best thing possible, given the cir-*
 1 *cumstances*. The fixed circumstances tell you what's possible; the ordering tells you what's best.

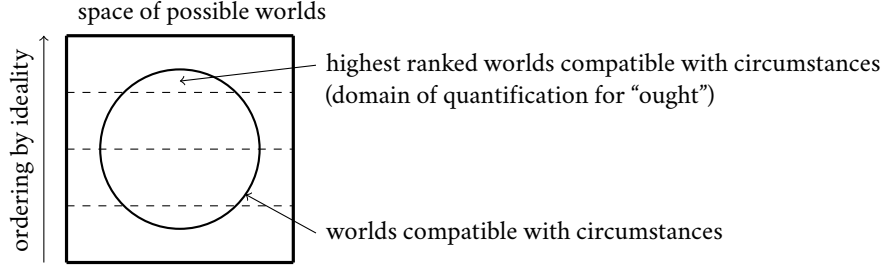
⁹ Thanks to [omitted].

¹⁰ N.B. I don't presuppose Kratzer's contextualism; what follows is compatible with relativism.

¹¹ In the lingo: the first contextual parameter is determined by an “ordering source”: a set of possible worlds propositions that characterize ideals. A weak partial ordering over possibilities can be projected from this set: a world w is at least as ideal as a world v , relative to an ordering source os , if and only if the set of os propositions that w satisfies includes the set of os propositions that v satisfies. Formally: let $\ulcorner w \succeq_{os} v \urcorner$ be interpreted as: *w is at least as ideal as v according to os*. Project the ordering \succeq_{os} from os as follows: for all worlds w, v ,

$$w \succeq_{os} v \leftrightarrow \{p \in os : w \in p\} \supseteq \{p \in os : v \in p\}$$

The second contextual parameter is the “modal base”: also a set of propositions, the intersection of which is the set of circumstantial possibilities: what's possible, given the circumstances.



2

3 $\lceil \text{Ought } \varphi \rceil$ is true (relative to a world, ordering, and set of fixed circumstances) just in case φ is true
 4 in all of the best (highest ordered) possibilities compatible with the circumstances.

5 The set of fixed circumstances determines what the agent “can” do: which possibilities for their
 6 behavior are left open and which are foreclosed. Since the set of worlds that “ought” universally quan-
 7 tifies over must be compatible with fixed circumstances, constraining what “can” happen in context,
 8 “ought” implies “can”.

9 In different contexts, the ordering may be provided by, e.g., state and federal laws, or a board game
 10 manual, or my friend Saba’s preferences, or the Objective Moral Truths. In the first example: ideal
 11 worlds might be ones where the laws are never violated; the best worlds possible, given the circum-
 12 stances, will include worlds where the laws have been violated but the violations have been properly
 13 sanctioned or rectified.

14 Which circumstances are held fixed varies from context to context, even in evaluating one and the
 15 same event. Consider the following discourse:

- 16 (1) Kerry ought to return her library books by 6:00 p.m.
 17 (2) But given that it’s 5:40 p.m. and she’s an hour away, Kerry ought to return her library books
 18 tomorrow and pay her library fines.

19 One of these sentences describes a requirement to turn in books before a date, while the other de-
 20 scribes a requirement to turn in books after the same date—incompatible requirements. But (1) and (2)
 21 are clearly compatible; they can be uttered one after the other without contradiction. Kratzer explains
 22 this compatibility via the context-sensitivity of the fixed circumstances. In (1), we don’t rule out Kerry’s
 23 turning in her library books on time. But the “given that” clause in (2) adds to the relevant fixed cir-
 24 cumstances that Kerry will not return her library books on time. Because the context updates between
 25 the two utterances, (1) and (2) don’t contradict each other.

26 Notice that by holding more circumstances fixed, (2) shrinks the set of open possibilities to those
 27 where Kerry does not return her library books on time. It excludes the very best worlds, by library
 28 norms. And so it *lowers the bar* for the library “ought”. In general, holding fixed more circumstances
 29 will tend to lower the evaluative bar.

30 Circumstantial modality is one of the least understood and least researched flavors of modality in
 31 natural language. Which contextually salient circumstances are held fixed for circumstantial “can” is
 1 extraordinarily flexible: there are clearcut senses in which both (3) and (4) are true:

- 2 (3) I physically **can** grasp my ankles from a standing backbend.
 3 (Unlike some people, my skeletal structure permits it; I just need to do regular stretches for a few years!)
- 4 (4) I physically **cannot** grasp my ankles from a standing backbend.
 5 (If I tried right now, I would fall on my face.)

6 The latter holds fixed my present muscle tightness; the former does not.¹²
 7 It's no surprise, then, that OIC principles in epistemology are hard to formulate. As I argued in
 8 §2.3, the difficulty involves determining which kinds of human limitations are fixed for the relevant
 9 circumstantial modal "can". These might be as variable across contexts as circumstantial "can" in gen-
 10 eral.

11 3.2 *Unified semantics and post-semantics*

12 I've suggested that we should think of nonideal epistemology as investigating the epistemic "ought"
 13 that implies a substantive "can", while ideal epistemology investigates the epistemic "ought" that does
 14 not imply a substantive "can". In that case, should we think of epistemic "ought", and other language
 15 used for epistemic evaluation—"(ir)rational", "(un)justified", etc.—as *ambiguous*? Are ideal and nonideal
 16 epistemologists merely talking past each other? Are they concerned with two distinct subject matters:
 17 rationality for idealized believers and, say, *reasonability* for ordinary believers?¹³

18 I doubt it. Semantically, "rational", and the "ought" of rationality, behave like other normative ex-
 19 pressions: they are context-sensitive, and in particular sensitive to contextually salient orderings of
 20 possibilities and fixed circumstances. In Kratzer semantics, ideal and nonideal epistemic "ought" have
 21 the same semantics—and the same semantics as the "ought" of morality and the "ought" of Pictionary.

22 In that case, one might protest, a unified semantics doesn't rule out the possibility that ideal and
 23 nonideal epistemologies have different subject matters, just as they both have different subject matters
 24 from morality and Pictionary.

25 My suggestion: ideal and nonideal epistemology are unified not only in their semantics but in some
 26 elements of their *post-semantics* (the resolution of their contextual parameters): I suggest they involve
 27 *the very same ordering*, and differ only in the circumstances they hold fixed. The ideal epistemic "ought"
 28 is expressed in contexts where the ordering is epistemic and the fixed circumstances don't include
 29 cognitive limitations. The nonideal epistemic "ought" is expressed in contexts where the ordering is
 30 epistemic but the fixed circumstances include some of the believer's cognitive limitations, ruling out,
 31 e.g., logical omniscience. As these fixed circumstances are pared down, as new possibilities open up,
 32 the nonideal epistemic "ought" converges to the ideal epistemic "ought".

33 What does this unity amount to? Consider the following two sentences, by analogy with (1) and (2)
 34 above:

- 1 (5) You ought to satisfy conditionalization, and so never forget any of your evidence.

¹² See also Lewis's (1976) famous "I can speak Finnish" example.
¹³ In discussing forgiving epistemic standards for ordinary believers, both Lasonen-Aarnio (2010) and Schoenfield (2012a) use the terminology of (nonideal) "reasonability".

2 (6) But if you do forget something, then you ought to have credences that cohere with your total
3 remembered evidence.

4 The first is an ideal epistemic evaluation, the second nonideal. But as with the library books, the two
5 sentences can be uttered truly, with no equivocation, no change in the relevant ordering. The epistemic
6 “ought” they use is, in an important sense, univocal.

7 3.3 Responses to challenges

8 In §1, I discussed three challenges to ideal epistemology: (i) that ideal epistemology “changes the sub-
9 ject”; (ii) that it’s false because “ought” implies “can”; and (iii) that, even if true, it’s pointless.

10 I’ve responded to the first: ideal and nonideal epistemology share part of their subject matter (an
11 ordering of worlds by epistemic ideality) and diverge in part of it (the circumstances they treat as fixed).

12 For the second, I’ve argued that only very weak versions of OIC apply in both nonideal and ideal
13 epistemology. In nonideal epistemology, “ought” doesn’t apply voluntary control. And in ideal episte-
14 mology, it might imply little more than metaphysical possibility. But the norms of nonideal and ideal
15 epistemology still apply to us nonideal agents—even if in many salient senses, we *can’t* satisfy them.

16 The third worry is that discovering the norms of ideal rationality is pointless: humans are all
17 doomed to irrationality, so ideal epistemology has no informative instructions to give them. This is
18 like saying that if you’re not an ideal library patron—you already have an overdue book—then the li-
19 brary rules have nothing to say to you. This is obviously false. Suppose the library rules say:

- 20 1. Return books by their due dates;
- 21 2. If you have overdue books, pay your library fines;
- 22 3. If you have unpaid library fines, answer calls from the collections agency;
- 23 ⋮

24 This set of norms generates an ordering of comparative library rule-abidingness. “Ideal rule-abidingness”
25 amounts to obeying *all* of these (where you trivially satisfy everything after 1 by virtue of satisfying 1).
26 Still, if you *can’t* obey all of the rules, you’re required to obey 2, 3, ...; and if you *can’t* obey 2, then you’re
27 required to obey 3, etc. What the library rules require of you in particular cases is relativized to which
28 circumstances we treat as fixed. If these circumstances entail it’s impossible to satisfy all of the rules,
29 then you can’t be *ideally* rule-abiding. But the library itself has further dictates for those who violate its
30 rules.

31 We can represent the norms of rationality as having a similar structure. An exaggeratedly simple
32 example:

- 33 1. Satisfy all of the norms of subjective bayesianism with rationally permissible priors;
- 34 2. If you have impermissible priors, violate conditionalization so that your new credences are the
1 result of conditionalizing permissible priors on your total evidence;

- 2 3. If you can't remember your total evidence, conditionalize on your total remembered evidence;
3 ⋮
4 *i.* If you don't satisfy all of *such-and-such set of ideal epistemic norms*, then satisfy all of *such-and-such*
5 *set of nonideal epistemic norms*;¹⁴
6 ⋮

7 So, like the library "ought", we can think of the epistemic "ought" as *relative*—in particular, relative to
8 some contextually salient circumstances, including cognitive limitations. Ideal rationality is a special
9 case of relative rationality, where no cognitive limitations are treated as fixed.

10 Still, one might respond: even if ideal and nonideal epistemic norms are merely different regions
11 on the same scale, looking to the unreachable top of the scale is pointless. What matters is comparative
12 values on the nonideal stretch of the scale. As Sen (2006) argues against ideal theory in political phi-
13 losophy, a hiker needn't know which mountain is the tallest in the world before she can compare the
14 heights of two mountains before her.

15 To address this challenge, we need to understand the shared scale—the epistemic ordering—better.
16 I'll argue that only its pinnacle, the epistemic ideal, is normatively robust. Only within ideal epistemol-
17 ogy can we learn about the fundamental nature of epistemic goodness and rightness.

18 In the remainder of the paper, I'll argue that nonideal epistemic evaluations are not normatively
19 robust. In fact, they're really, *really* normatively non-robust. In §4, I'll argue that the norms of nonideal
20 epistemology, governed by the non-maximal part of the epistemic ordering, are conventional. In §5, I'll
21 argue that there's no privileged set of cognitive limitations to treat as fixed circumstances; as a result,
22 nonideal epistemic evaluations are conventional, seriously context-sensitive, or both.

23 I identify five distinct sources of non-robustness. You might not accept all of them. (Those in §4
24 may be especially unpalatable.) But if I'm right about even one of the five, that's enough to establish
25 that nonideal epistemic evaluations are not normatively robust.

26 4 Non-robust ordering

27 What determines the ordering for epistemic evaluations? In this section, I give reasons to believe that
28 the epistemic ordering is, to a significant degree, *conventional*. I present a general metaepistemological
29 picture and argue that, on this picture, three sources of conventionality arise in the determination of
30 which epistemic rules are rational. The latter two sources of conventionality have been argued rigor-
31 ously and in great depth in the work of Dogramaci (2012, 2015), though he puts these observations to
1 very different work; indeed, his conclusions from these observations are incompatible with mine.

¹⁴ Notice that this doesn't assume that satisfying more of the ideal epistemic rules is always better than satisfying fewer, or more generally, that the ideal rules determine the ordering anywhere below the apex. Indeed, I argue below that the "back-up norms" determining the ordering for the nonideal epistemic "ought" are sometimes incompatible with ideal epistemic rules.

2 4.1 *A meta-epistemological picture*

3 Here is a simple picture of the purpose of epistemic evaluation. The ultimate epistemic end is alethic:
4 true belief or accurate credence. Rationality is valuable only instrumentally, because it conduces to this
5 end.¹⁵ There are two broad means of promoting the end of truth. Epistemically: we should adopt truth-
6 conducive epistemic rules. And practically: we should improve our epistemic situations by developing
7 tools to collect information as input for epistemic rules.

8 One tool for improving our epistemic situations is other people. If other people's epistemic stan-
9 dards are truth-conducive, then trusting others' sincere testimony is truth-conducive. If trusting others
10 is truth-conducive, then (defeasibly) trusting improves one's epistemic situation. The purpose of epis-
11 temic evaluations, on this picture, is to encourage others to adopt truth-conducive epistemic standards,
12 to make trust in others a reliable source of information. This picture has been defended forcefully by
13 [Dogramaci \(2012, 2015\)](#) and [Greco & Hedden \(2016\)](#).

14 Supposing that's correct, the best possible means to our epistemic ends will involve the following
15 two strategies: first, we should adopt all truth-conducive epistemic rules. Second, we should encourage
16 others to adopt all truth-conducive epistemic rules.

17 4.2 *Source #1: incompatible truth-conducive rules*

18 There's a problem for this strategy: it's not possible to adopt *all* truth-conducive epistemic rules. While
19 some epistemic rules are deductive, many are not: rules governing inductive reasoning, inference to the
20 best explanation, and so on. Good non-deductive epistemic rules don't guarantee truth-preservation:
21 they merely probabilify it. They are truth-conducive, but only imperfectly. And so these rules offer no
22 guarantee of maximizing accuracy. At most they maximize *expected* accuracy. But expected accuracy is
23 necessarily relativized to some perspective, and there may be no objectively privileged perspective. So
1 we may expect ties between incompatible epistemic rules, each with an equal claim to truth-conducivity

¹⁵ This view has many proponents in various areas of epistemology. Some representative quotations:

"[A] satisfactory defense of a particular standard of epistemic justification must consist in showing it to be truth-conducive, i.e. in showing that accepting beliefs in accordance with its dictates is likely to lead to truth (and more likely than any proposed alternative)" ([BonJour 1978](#), 5).

"Justification is a kind of guide to the truth. We seek to form justified beliefs as a means to forming true beliefs" ([White 2006](#), 539).

The view is central to the popular research program often called "accuracy-first epistemology", spearheaded by [Joyce \(1998, 2009\)](#). For criticism of this form of epistemic consequentialism, see e.g. [Berker \(2013a,b\)](#) and [Carr \(2017\)](#). I assume accuracy-based epistemic consequentialism for simplicity; the metaepistemic picture presented should be generalizable to other theories of epistemic value and epistemic rightness.

2 (relative to different perspectives).^{16, 17}

3 A *prima facie* solution: perhaps rationality should instead require adopting *some* maximal consist-
4 tent set of epistemic rules from those tied for equal claim to maximizing truth-conducivity—where
5 competing standards might, for example, countenance different degrees of inductive risk. (I need to
6 see 23 quokkas with black noses, and none without, before I reach 80% confidence that all quokkas
7 have black noses; you only need to see 19.)

8 The problem: in order for trust in others to be truth-conducive, we must *agree* in our epistemic
9 standards.¹⁸ Your epistemic standards maximize expected accuracy from your perspective. (Otherwise,
10 why would you have those standards?) But trust in others only maximizes expected accuracy from your
11 perspective if others share your epistemic standards. Otherwise, they'll be disposed to sometimes draw
12 conclusions from their evidence that you regard as unreliable. So to learn from each other's testimony,
13 we should all prefer that we share epistemic standards—but we might disagree about which standards
14 to coordinate on. If we're caught in a deadlock, we can't learn from each other. Communities with more
15 broadly shared epistemic standards trust each other's testimony more. Their members individually and
16 collectively benefit from pooling their information, dividing the evidence-gathering labor, promoting
17 the community's success. So we should expect social (and perhaps even evolutionary) pressure toward
18 consensus—even if the consensus epistemic standards have no unique claim to truth-conducivity.¹⁹ If
19 the uniquely rational standards did not exist, it would be necessary to invent them.

20 In order to coordinate on epistemic standards, there must be a selection among maximal consistent
21 sets of epistemic rules that are tied in their claims to truth-conducivity. With respect to our epistemic
1 ends, then, this choice is arbitrary; coordinating on standards is a matter of convention.^{20, 21}

¹⁶ I assume that rationality requires some immodesty, in the sense of [Gibbard \(2008\)](#); [Horowitz \(2013\)](#). So the sense in which different epistemic standards have an equal claim to truth-conducivity is delicate. The truth-based considerations that I can offer in favor of my standards—that they maximize expected accuracy (from the perspective I inhabit)—are equally available to you—that yours maximize expected accuracy (from the perspective you inhabit). But that doesn't mean I should give up my epistemic standards: to do so would be to move to some alternative standards that I genuinely expect to lead me away from the truth! For discussion, see [Schoenfield \(2012b\)](#).

¹⁷ [Pettigrew \(2016\)](#) argues that accuracy-based considerations may privilege a unique body of epistemic standards for credences, determined by the Principle of Indifference together with a partition over logical space. (Though this partition may also be conventional.) For this and other reasons, I'm not convinced of the claim I assert in this subsection. If it is incorrect, then this section will only correctly identify two ways in which nonideal epistemic evaluations are conventional.

¹⁸ For compelling arguments, see [Greco & Hedden \(2016\)](#) and [Dogramaci & Horowitz \(2016\)](#).

¹⁹ Notice: we can regard our epistemic standards as most truth-conducive, even if we recognize that had we landed on other standards, we could have given perfectly parallel reasons for regarding those other standards as most truth-conducive.

²⁰ Both [Greco & Hedden \(2016\)](#) and [Dogramaci & Horowitz \(2016\)](#) argue that for the purposes of coordination in epistemic standards, epistemic evaluations must presuppose the Uniqueness principle ([Feldman, 2006](#)). But these arguments give no reason to expect that one body of epistemic rules has a unique claim to greater truth-conducivity. It may then be that if Uniqueness is true, the uniquely permissible epistemic standards are conventionally determined. In this case, then, theories like objective bayesianism might fall partly within the purview of nonideal epistemology. (!)

²¹ *Objection:* as [Kuhn \(1977\)](#) and [Kitcher \(1990\)](#) have argued in the context of science, diversity in epistemic standards, rather than uniqueness, may better promote our alethic ends. So it may be serve us better *not* to coordinate our epistemic standards. (Thanks to [omitted].)

Replies: First, our actual conventions might not be the optimal conventions. Luckily, they seem to work (more or less).

2 4.3 Source #2: coordinating around cognitive limitations

3 A sociological observation: there are some epistemic rules that are truth-conducive but that we regard
4 as irrational. (See, e.g., Bonjour's (1980) clairvoyant or Lehrer's (1990) TrueTemp.)

5 Focus on the case of *valid* epistemic rules. Valid epistemic rules are maximally truth-conducive and
6 mutually compatible. So all maximal consistent truth-conducive epistemic standards must include all
7 valid epistemic rules.

8 But are all valid epistemic rules intuitively *rational*? Dogramaci (2012, 2015) argues that many are
9 not. For example:²²

10 **Fictional Lagrange.** Fictional Lagrange offers the following argument for Lagrange's
11 four-square theorem (LFST): from the Peano axioms, he immediately infers LFST. The
12 argument includes no intermediary steps—just a direct inferential leap from basic arith-
13 metic to LFST.

14 Is Fictional Lagrange's inference rational? In traditional nonideal epistemology, the intuition is gener-
15 ally that it's not: that this inference would have to constitute some sort of lucky guess. What explains
16 this?

17 This example illustrates a general sociological fact: of the infinite number of valid inference rules,
18 we ordinarily only accept a few as capable of rationalizing beliefs. In many cases, a believer may
19 rationally believe a set of premises and may, on that basis, immediately infer a proposition entailed
20 by those premises—without the inferred belief being rational. A student first learning propositional
21 logic who immediately infers the axiom of conjunction elimination $((p \wedge q) \rightarrow p)$ from no premises
22 is intuitively rational. But, Dogramaci argues, if the same student immediately infers Pierce's Law
23 $((p \rightarrow q) \rightarrow p) \rightarrow p$ from no premises or intermediate steps, her belief is intuitively unjustified. If
24 it's correct that epistemic rationality derives its value somehow or other from its relation to true belief,
25 this sociological fact should be surprising. It stands in need of explanation.

26 The research program of nonideal epistemology aims to respect some substantive OIC principle.
27 In some relevant sense of "can", ordinary human believers *can't* adopt all of the infinitely many valid
28 epistemic rules. At best, a believer can adopt a small subset of rules that are as powerful as she can
29 manage. There are many sets of epistemic rules that are equally small and equally powerful, or on a par
30 with respect to these two properties weighed against each other. No alethic considerations tell these
31 apart.

32 Can we each choose whichever small, powerful subsets we like? The purpose of epistemic evalua-
33 tion, I argued, is to make trust in others more truth-conducive. But as I argued in §4.2, trust is expect-
34 edly truth-conducive only if epistemic standards are widely shared. To serve our epistemic ends, we

Second: there may be advantages in sustaining a convention whereby scientists are convinced that others *ought* to agree with them (if, e.g., it spurs them to present better arguments to try to convince others from shared evidence). Kuhn and Kitcher each argued that there may be conflicts in individual rationality and group rationality. This in fact *supports* the idea that our individual epistemic evaluations may mandate conformity to shared epistemic standards, even when this conflicts with group-level epistemic aims.

²² Boghossian (2003) provides a version of this example.

1 need to *coordinate* in our epistemic standards. And this requires coordinating on a small, computation-
2 ally tractable set of shared epistemic rules, rejecting other rules as irrational.

3 Many of these rejected rules are truth-conducive, even perfectly so: for example, Fictional La-
4 grange’s rule. Because members of the evaluative community are cognitively limited, they aren’t in
5 a position to evaluate the accuracy of all individual inferences. (Many of us can’t assess Fictional La-
6 grange’s inference.) For a community to coordinate, we should expect it to err on the side of impermis-
7 siveness, allowing only a tractable subset of the epistemic rules that are recognized as truth-conducive.
8 Here, then, cognitive limitations affect how many rules evaluators are capable of sanctioning, or how
9 complex these rules may be. They will sometimes require excluding rules of ideal epistemology, sub-
10 stituting in more tractable alternatives.

11 The peculiar result is that Fictional Lagrange’s inference isn’t merely unrequired; it’s prohibited.
12 Common cognitive limitations don’t strictly loosen the requirements of nonideal rationality, permit-
13 ting us to ignore some hidden avenues to the truth. In some cases, nonideal rationality is *more* stringent
14 than ideal rationality, blocking these avenues to the truth.

15 Whichever small, tractable, powerful set of epistemic rules we coordinate on is, from an alethic
16 perspective, *arbitrary*. Our epistemic evaluations hew to some such set of rules only as a matter of con-
17 vention. And these conventions will treat some arbitrary set of truth-conducive epistemic rules—e.g.
18 Fictional Lagrange’s rule—as *irrational*, even though from an alethic perspective, they are unimpeach-
19 able.

20 4.4 Source #3: *facilitating conventions*

21 The second source of non-robustness stemmed from our *individual* cognitive limitations. Another
22 stems from challenges in coordinating as a group. The larger, more complex, and more arbitrary the
23 set of rules, the harder it is to establish a convention. This is presumably why, for example, Pig Latin
24 caught on as a popular code rather than letter-by-letter ciphers. Individuals may all be able to conform
25 to more complex epistemic standards than would be feasible for populations to coordinate on.

26 *Objection:* this depends on the pretense that epistemic conventions are socially constructed and
27 must be learned from others—which is obviously not literally true.

28 *Reply:* epistemic conventionalism is not committed to the claim that alternative sets of epistem-
29 ically adequate norms—norms that would serve the functional roles of epistemic evaluation equally
30 well—would all be equally easy to adopt, given our psychological makeup. We may be hardwired to
31 reason by modus ponens and not by Fictional Lagrange’s rule. [Dogramaci \(2015\)](#) suggests that, in our
32 evolutionary history, the need for epistemic coordination may have involved an arbitrary choice among
33 sets of epistemic rules, not by humans (individually or collectively), but by natural selection.

34 But even if epistemic conventions are established by natural selection, it’s still the case that larger
35 sets of independent rules are harder to establish as conventions. Compare universal grammar: there’s a
36 reason why we didn’t evolve a *total, shared innate language*, even if doing so would better serve language’s
37 communicative purposes. And there’s a reason why languages contain a finite vocabulary and a small
1 finite number of compositional rules.

2 Indeed, we may have evolved an “universal epistemic grammar”, generating some innate constraints

3 on epistemic standards while allowing social conventions to determine others.²³ A universal epistemic
4 grammar might require sound and complete collections of epistemic rules as well as reliable induc-
5 tive epistemic rules, but with flexibility in which specific epistemic rules we can adopt, where many
6 conventions are socially inherited rather than hardwired. On this hypothesis, our hardwiring doesn't
7 accommodate the acquisition of epistemic rules for inferring abstruse mathematical or logical theo-
8 rems immediately from basic axioms of arithmetic or logic. But it is flexible enough to accommodate
9 acceptance of subtly different, overlapping collections of epistemic rules. This may explain the diver-
10 gences in "intuition" we sometimes see among students, and among established philosophers.

11 To facilitate coordination in our epistemic rules, then, the set of epistemic rules we coordinate
12 on must exclude some of the epistemic rules that survived the winnowing processes in the previous
13 two subsections, and introduce rules that are necessarily less truth-conducive but more tractable. But
14 again, these exclusions and additions are arbitrary: alternatives would have served just as well for the
15 purposes of attaining our epistemic ends.

16 From the conventionality of epistemic evaluations, [Dogramaci](#) argues that *there can be no theory*
17 *of epistemic rationality*—a potentially alarming conclusion for us career epistemologists! But it's worth
18 noting that in *ideal* epistemology, there's no barrier to including *all* valid rules among one's epistemic
19 standards. Logical omniscience is broadly accepted as a norm of ideal rationality. So examples like
20 Fictional Lagrange provide no reason for conventionality, and no barrier to constructing an *a priori*
21 theory of ideal epistemic rationality. I'll discuss this contrast further in §6.

22 4.5 Comparison with other accounts of nonideal epistemology

23 The view described above doesn't entail that there's a set of ideal epistemic rules that determine an
24 epistemic ordering. The ideal epistemic rules only characterize what's at the top of the ordering. One
25 might think that the rest of the ordering could be projected from the ideal rules, either in terms of
26 nonideal states' divergence from the ideal or in terms of the number of ideal rules satisfied. Or more
27 generally, one might hold that we should adopt the most truth-conducive standards possible, given our
28 cognitive limitations. Notice, though, that none of these orderings would be conventional.

29 I've argued that the nonideal tail of the epistemic ordering is conventional. The conventions aim
30 for truth-conducivity but are constrained by implementability. They may sacrifice individual-level
31 truth-conducivity in favor of population-level truth-conducivity. Standards that are not ideally truth-
32 conducive, but are more tractable than the ideal, can develop into conventions. These conventions
33 allow for practices of epistemic evaluation that encourage more overall truth-conducive standards in
34 a population. (This enables information pooling, which more than compensates individuals for the
35 alethic costs of ruling out some truth-conducive rules.) Sometimes, then, conforming to ideal rules
36 will *violate* the conventions, and will be ruled irrational. Fictional Lagrange is an example.

1 The conventionalist picture differs substantively from, e.g., [Staffel \(forthcoming b\)](#) or [Williams](#)
2 ([manuscript](#)) in how it ranks nonideal epistemic states for comparative rationality. On these views, the

²³ Compare with Chomskyan universal grammar, which determines innate structural constraints on human languages but allows cross-linguistic flexibility in vocabulary, parts of syntax, and so on.

3 ordering is determined by the epistemic state's distance from some ideally rational epistemic state: for
4 example, the distance between the different degrees of credence that two credence functions assign to
5 various propositions, by some divergence like the Brier score. These divergences will respect principles
6 like Rationality-Directedness. Let I be the ideal credence function in some context, c and c' be two
7 credence functions, and F be the algebra of propositions over which I , c , and c' are defined:

8 **Rationality-Directedness.** If for all $p \in F$, either $c'(p) \geq c(p) \geq I(p)$ or $c'(p) \leq c(p) \leq I(p)$, and
9 for some $p \in F$, $c'(p) > c(p) \geq I(p)$ or $c'(p) < c(p) \leq I(p)$, then c is epistemically better than c' .

10 In English: if c 's credences are uniformly at least as close to the ideal as c' 's and sometimes closer, than
11 c is epistemically better than c' .

12 While the conventionalist picture is capable of accommodating orderings like these, it also accom-
13 modates very different kinds of orderings: orderings in which, for example, greater incoherence may
14 be preferable, even required. This is important: it allows the conventionalist picture to capture the
15 claims of traditional nonideal epistemology. Examples like Fictional Lagrange show that traditional
16 nonideal epistemology sometimes *requires* agents to forego greater accuracy, expected accuracy, and
17 coherence—to choose options that are quantitatively farther from the ideal than some available alter-
18 natives. On the conventionalist picture, second-best need not be what's quantitatively closest to best.

19
20 To summarize: nonideal epistemology embraces (perhaps tacitly) some form of substantive OIC
21 principle. Epistemic OIC principles require coordinating on some arbitrary selected set of epistemic
22 rules. At each stage, truth-conducive epistemic rules—even maximally truth-conducive rules—are
23 ruled out as irrational, purely as a matter of convention.²⁴

24 5 Non-robust fixed circumstances

25 In §3, I argued that evaluative language is sensitive to an ordering over possibilities and a set of fixed
26 circumstances. In §4, I argued that the ordering used in nonideal epistemic evaluations is norma-
27 tively non-robust, because convention-bound. In this section, I'll argue that the fixed circumstances
28 in nonideal epistemic evaluations are also normatively non-robust: they're conventional, or seriously
1 context-sensitive, or both. There is no normatively privileged set of cognitive limitations to treat as
2 fixed, determining what we "can" do and therefore constraining what we "ought" to do.

²⁴ Epistemic conventionalism might be motivated independently of the metaepistemology presented in section 4.1. We may, for example, have impermissivist intuitions, on which one can never rationally regard some other response to one's evidence as just as rational as one's own. But we might equally worry that epistemically significant considerations—whether these are reducible to accuracy or not—can't possibly whittle down the set of permissible standards to a singleton. (In the Bayesian context, e.g., one might worry that epistemically significant considerations can't solve the problem of the priors.) One might resolve this tension with the form of conventionalism described in 4.2. Alternatively, we might worry that epistemically significant considerations won't be sufficient to explain or systematize various nonideal epistemic intuitions, e.g. intuitions prohibiting Fictional Lagrange's (accuracy-promoting) inference; this worry fits well with source #2 of conventionality.

3 5.1 Source #4: Why these cognitive limitations?

4 I argued that nonideal epistemology embraces some form of OIC, sensitive to cognitive limitations that
5 are treated as fixed circumstances. Nonideal epistemology tacitly divides cognitive limitations into two
6 kinds:

7 1. *Cognitive limitations that lower the bar for nonideal rationality*: our limited computational power,
8 informational storage, processing speeds, integration of different cognitive systems, information
9 retention, ...

10 2. *Cognitive limitations that don't*: our dispositions toward implicit biases, unreliable heuristics,²⁵
11 delusional reasoning, misinterpreting statistical phenomena as having causal explanations, in-
12 flated sense of one's own driving ability, over-optimism/pessimism, overestimating the moral
13 superiority of one's own side in a fight with a spouse, family member, or departmental faction,
14 ...

15 What justifies dividing things in this particular way? Why do some cognitive limitations *lower the bar*
16 *for rationality* while others *constitute irrationality*?

17 *Hypothesis*: the former are universal, not culturally contingent in the way that, say, racial bias may
18 be culturally contingent.

19 *Problem*: many of the latter limitations are just as universal among humans as the former; and
20 if they weren't, it's not clear why that would change their categorization. Learning that one of the
21 former limitations was not universal intuitively wouldn't change its categorization. If it turned out
22 that Fictional Lagrange were real, it's hard to imagine that his existence would render the rest of us
23 irrational (within nonideal epistemology) for our mathematical ignorance.

24 *Hypothesis*: only the latter dispositions come in degrees, leaving room for improvement and there-
25 fore room for criticism. We can, e.g., ameliorate our propensity toward confirmation bias, but so long
26 as we're finite beings, we'll have finite information storage and processing speeds.

27 *Problem*: we can also improve on many of the former limitations. We're certainly capable of im-
28 proving our deductive abilities (by concentrating, practicing math and logic exercises, consuming stim-
29 ulants...), information retention (by repeating facts to ourselves, developing mnemonics...), etc. We
30 don't expect that ordinary believers can infer Lagrange's four-square theorem from the Peano axioms.
31 We do expect that ordinary believers can infer conjuncts from conjunctions, and are sometimes irra-
32 tional for failing to do so. Between these two extremes, there are a range of cases and no obvious cutoff
33 points.

34 I don't claim that there are no differences between the two kinds of limitation. Instead, I claim
35 that the line between them is blurry. Whatever sharp line nonideal epistemologists draw, it could have
36 been drawn elsewhere. This suggests that the distinction is conventional. Given that we have a large
1 array of cognitive imperfections and can't continually consciously mitigate all of them, it makes sense
2 to implement some conventions determining which to try to mitigate and which to accept as our fate.

²⁵ For example, the representativeness and availability heuristics (Kahneman & Tversky, 1972; Tversky & Kahneman, 1973).

3 We then treat the former as constituting irrationality and the latter as lowering the bar for rationality.
4 Epistemic evaluations are more efficacious if they're predictable.

5 But it seems to me that outside the epistemologists' den, the borderline between these two cate-
6 gories in ordinary epistemic evaluations is far more flexible. In ordinary epistemic evaluations, one can,
7 for example, describe someone's beliefs as *rational given their somatoparaphrenia*, *rational given their in-*
8 *flated estimate of their own driving skills*, etc. So the distinction may instead, or also, be seriously context-
9 sensitive.

10 There's no clearly normatively privileged boundary between which limitations we treat as fixed
11 circumstances and which we suppose are surmountable: which constitute epistemic failures and which
12 constitute epistemic excuses. Whether the boundary relevant for an epistemic evaluation is determined
13 by convention or by context, the evaluation is not normatively robust.

14 5.2 Source #5: irrelevance of individual variation in ability

15 A sociological observation: the "can" implied by the traditional nonideal epistemologists' "ought" is
16 *not sensitive to individual variations in ability*. It makes no special allowances for particular dispositions
17 toward biases; no dispensations for somatoparaphrenics or Alex Jones-style conspiracy theorists.

18 And yet the following seems equally obvious: if there were a non-arbitrary, *normatively privileged*
19 "can" for epistemic OIC, it would be sensitive to individual variations in ability. This is a bare intuition.
20 But we do accept this bare intuition for the "can" of OIC in ethics and practical rationality: these are
21 uncontroversially sensitive to individual variations in ability. (Jeff Bezos is morally obligated to do-
22 nate \$123 billion to charity; you and I are not, because we cannot.) It's hard to see what could justify
23 insensitivity to individual variations in the practical case and not the epistemic.

24 On the picture described above, this fact is easy to explain. Coordination in epistemic standards
25 would be near impossible to achieve if the applicability of epistemic evaluations depended on knowl-
26 edge of idiosyncratic psychological features of individual believers.²⁶

27 *Whose* limitations we treat as fixed is arbitrary, and so plausibly conventional. The "can" implied
28 by nonideal epistemic "ought" draws an arbitrary, normatively non-robust distinction: some people's
29 limitations constitute irrationality; others' are treated as faultless.

30 6 Conclusions

31 6.1 Why ideal epistemology?

32 Nonideal epistemic evaluations are infected with normative non-robustness from all angles. *Ideal* epis-
33 temic evaluations, by contrast, are free from these sources of normative non-robustness.

1 Consider Dogramaci's observation that our (ordinary, nonideal) epistemic intuitions exclude many
2 valid epistemic rules as irrational. This forms the basis for his argument that there can be no theory of

²⁶Note that this is compatible with there being some serious context-sensitivity in cases where we do accommodate, e.g., so-
matoparaphrenia, as hypothesized in §5.1.

3 epistemic rationality. But Dogramaci doesn't discuss ideal epistemology—which often involves accep-
4 tance of logical omniscience as a norm for ideally rational believers. It follows from deductive closure
5 constraints on binary beliefs and from probabilistic constraints on rational credences. So ideal episte-
6 mologists generally accept that rationality permits reasoning by Lagrange's rule. Indeed, they generally
7 accept that if an idealized believer has any doxastic attitude toward LFST whatsoever, it's rationally
8 required to be full belief or maximal confidence. More generally, the ideal epistemologist will reject
9 Dogramaci's foundational premise: that the rational deductive epistemic rules don't include all of the
10 valid rules. For idealized believers, there's no barrier to adopting all valid deductive rules. Without the
11 need for an arbitrary selection, there's no barrier to the possibility of a unifying theory of the deductive
12 epistemic rules governing ideally rational believers' inferences. So Dogramaci's argument presents no
13 challenge to the possibility of a unified theory of *ideal* rationality. It only poses a challenge for theo-
14 ries of nonideal rationality. Indeed, candidate unifying theories of ideal epistemic rationality have long
15 been in circulation.²⁷

16 If there is any normatively robust epistemic "ought", it's the "ought" researched by *ideal* episte-
17 mologists. And this matters: only normatively robust evaluations carve nature at its normative joints.
18 Conventions are not, in the appropriate sense, normatively fundamental: they depend on historical
19 contingencies that have nothing to do with the Great Normative Truths in the Sky. And if these nor-
20 mative truths exist, it ought to be possible to express them with language that isn't seriously context-
21 sensitive. For theorizing about fundamental epistemic normative properties, our only option is ideal
22 epistemology.

23 A staunch ideal epistemology skeptic might object: if ideal epistemology is the only option for a
24 normatively robust theory, so much the worse for normative robustness! Why not reject ideal episte-
25 mology in favor of thoroughgoing epistemic conventionalism?

26 *Reply:* Let's consider two kinds of ideal epistemology skeptic: one who accepts that the purpose of
27 these conventions is to help us get closer to the truth and one who doesn't.

28 The first kind of skeptic should recognize that if we had fewer cognitive limitations, other conven-
29 tions would better serve this purpose. (If forgetting were under our control, for example, it'd be better
30 to have conventions that sometimes made information loss impermissible.) In the limit, these counter-
31 factuals lead the case where believers are cognitively unlimited and epistemic norms no longer count as
32 conventional because there are no alternatives that serve our epistemic purposes equally well—in other
33 words, to norms of ideal epistemology. There are epistemic standards that are logically guaranteed to
34 serve this purpose better than any alternatives (as accuracy-first epistemology has shown). If idealized
1 believers are epistemically evaluable, then these evaluations fall in the purview of ideal epistemology.²⁸

²⁷ These include theories formulated in terms of pragmatic benefit (Dutch book arguments for bayesian epistemic norms (Ramsey et al., 1931; de Finetti, 1937), Regularity (Kemeny, 1955; Shimony, 1955), the Principal Principle (Howson, 1992), the Reflection Principle (van Fraassen, 1984), and Jeffrey Conditionalization (Armendt, 1980)), and in terms of truth- or accuracy-conducivity (accuracy-based arguments for bayesian epistemic norms (Joyce, 1998, 2009; Greaves & Wallace, 2006; Leitgeb & Pettigrew, 2010a,b), the principal principle (Pettigrew, 2012, 2013), the principle of indifference (Pettigrew, 2016), binary belief single-premise closure and pairwise consistency (Easwaran, 2016), etc.).

²⁸ One may deny that they are epistemically evaluable; but that would need motivating.

2 So there are questions for ideal epistemology to answer.²⁹

3 The other form of ideal-epistemology skeptic holds that once we go conventionalist, we give up
4 on the idea that truth (knowledge, understanding...) has any value that's essential to the *real* deontic
5 epistemic theory. But conventions that float free of truth (etc.), even conventions governing belief, do
6 not generate *epistemic* norms. So this form of skepticism seems to give up on epistemology altogether.

7 6.2 *Whence nonideal epistemology?*

8 If nonideal epistemology is normatively nonrobust, that isn't an objection to it! Instead, its non-
9 robustness reframes the research program. If nonideal epistemology is normatively non-robust, this
10 suggests a fruitful role for ideal epistemology. It describes the standard that ordinary rational believers
11 *approximate*, to greater or lesser degrees, in convention-governed ways, held back or excused by our
12 cognitive limitations.

13 The role for ideal rationality, as an ideal that we improve our rationality by approaching, is often
14 presupposed in the literature on bounded rationality. But while theories of bounded rationality take
15 into account various of our cognitive limitations, they ignore the normative non-robustness of non-
16 ideal epistemic evaluations—the conventionality and serious context-sensitivity in how we rank non-
17 ideal states (the ordering) and which cognitive limitations we accommodate (the fixed circumstances).

18 To build a complete, predictive nonideal epistemology, we need three forms of input:

- 19 1. *Normatively fundamental factors*: a theory of the epistemically ideal;
- 20 2. *Contingent empirical factors*: a psychologically sophisticated understanding of contingent human
21 cognitive limitations (including variations across individuals);
- 22 3. *Non-robust factors*:
 - 23 (a) sociological/evolutionary psychological facts about which sets of epistemic rules we've co-
24 ordinated on,
 - 25 (b) sociological/linguistic facts about which cognitive limitations loosen the requirements of
26 rationality, where these may turn out to be conventional or seriously context-sensitive.

27 This first input just is ideal epistemology. And so a complete nonideal epistemology will depend on,
28 and justify, the program of ideal epistemology.

29 Nonideal epistemology is, from this perspective, partly empirical. This doesn't exclude it from the
30 domain of philosophy: philosophy of language and philosophy of mind have always been partly em-
31 pirical; applied ethics and philosophies of the special sciences are, by definition, empirically informed.

²⁹ It's a separate question whether the norms for idealized believers apply to us. But, again, nothing in the *language* of ideal epistemic evaluation blocks its applicability to us (or, for that matter, to our pets). We can measure ourselves against the standards of ideal rationality, even if we inevitably find ourselves lacking—just as we can measure ourselves against celebrity beauty standards, even if we inevitably find ourselves unlovely. So long as ideal evaluation is doesn't preclude nonideal evaluation, why not let a thousand evaluative flowers bloom?

32 These areas of philosophy bring together empirical information and insights from areas of philoso-
1 phy at a greater theoretical remove—logic, metaphysics, general normative ethics, ideal epistemology,
2 etc.—using each to shed light on the other. One role of the nonideal epistemologist is to systematize
3 our nonideal intuitions, to expose the structure of the conventions we in fact have. There are some
4 topics about which we have vast implicit knowledge and extremely limited explicit knowledge. (The
5 fact that linguistics exists as a field of research demonstrates this.) Nonideal epistemology helps us to
6 better understand the norms that guide us. Another role for the nonideal epistemologist may be *ame-*
7 *liorative*:³⁰ nonideal epistemology can reveal the shortcomings of our conventions and propose better
8 alternatives.³¹

9 Ideal epistemology offers a theory of the normatively robust. Nonideal epistemology should couple
10 it with empirical research into our cognitive limitations, epistemic conventions, and context-sensitive
11 expressions. Only then will we explain why, how, and in what contexts we settle for approximating the
12 ideal.

13 References

- 14 Alchourrón, Carlos E., Peter Gärdenfors & David Makinson (1985). “On the Logic of Theory Change:
15 Partial Meet Contraction and Revision Functions.” *Journal of Symbolic Logic* 50(2), 510–530. doi:
16 10.2307/2274239.
- 17 Alston, William (1988). “The Deontological Conception of Epistemic Justification.” *Philosophical Per-*
18 *spectives* 2. 257–299. doi:10.2307/2214077.
- 19 Armendt, Brad (1980). “Is There a Dutch Book Argument for Probability Kinematics?” *Philosophy of*
20 *Science* 47(4), 583–588. doi:10.1086/288958.
- 21 Berker, Selim (2013a). “Epistemic Teleology and the Separateness of Propositions.” *Philosophical Review*
22 122. 337–393.
- 23 Berker, Selim (2013b). “The Rejection of Epistemic Consequentialism.” *Philosophical Issues* 23(1), 363–
24 387. doi:10.1111/phis.12019.
- 25 Boghossian, Paul (2003). “Blind Reasoning.” *Aristotelian Society Supplementary Volume* 77(1), 225–248.
- 26 Bonjour, Laurence (1978). “Can Empirical Knowledge Have a Foundation?” *American Philosophical*
27 *Quarterly* 15(1), 1–14.

³⁰In the sense of Haslanger (2006).

³¹For this reason, I view the conventionalism in this paper as compatible with Staffel’s (2019) project. The conventionalist picture aims to characterize norms that we in fact use in epistemic evaluation. Staffel’s proposal might be thought of as providing a system of norms that we ought instead to adopt: one that is normatively robust (albeit one that might fall afoul of traditional nonideal epistemic OIC principles).

- 28 BonJour, Laurence (1980). "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy*
1 5(1), 53–73. doi:10.1111/j.1475-4975.1980.tb00396.x.
- 2 Bykvist, Krister & Anandi Hattiangadi (2007). "Does Thought Imply Ought?" *Analysis* 67(4), 277–285.
- 3 Carr, Jennifer Rose (2015). "Subjective Ought." *Ergo* 2(27), 678–710.
- 4 Carr, Jennifer Rose (2017). "Epistemic Utility Theory and the Aim of Belief." *Philosophy and Phenomeno-*
5 *logical Research* 95(3), 511–534.
- 6 Carr, Jennifer Rose (manuscript). "Should You Believe the Truth?"
- 7 Castro, Clinton & Olav Vassend (2017). "Ideal Counterpart Theorizing and the Accuracy Argument
8 for Probabilism." *Analysis* 78(2), 207–216.
- 9 Dogramaci, Sinan (2012). "Reverse Engineering Epistemic Evaluations." *Philosophy and Phenomenolog-*
10 *ical Research* 84(3), 513–530.
- 11 Dogramaci, Sinan (2015). "Communist Conventions for Deductive Reasoning." *Noûs* 49(4), 776–799.
- 12 Dogramaci, Sinan & Sophie Horowitz (2016). "An Argument for Uniqueness About Evidential Support."
13 *Philosophical Issues* 26(1), 130–147.
- 14 Easwaran, Kenny (2016). "Dr. Truthlove: or, How I Learned to Stop Worrying and Love Bayesian
15 Probabilities." *Noûs* 50(4), 816–853.
- 16 Feldman, Richard (2001). "Voluntary Belief and Epistemic Evaluation." In Matthias Steup (ed.), *Knowl-*
17 *edge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, 77–92. Oxford Uni-
18 versity Press.
- 19 Feldman, Richard (2006). "Reasonable Religious Disagreements." In Louise Antony (ed.), *Philosophers*
20 *Without Gods: Meditations on Atheism and the Secular Life*, 194–214. Oxford University Press.
- 21 de Finetti, Bruno (1937). "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut*
22 *Henri Poincaré* 17. 1–68.
- 23 van Fraassen, Bas C. (1984). "Belief and the Will." *Journal of Philosophy* 81(5), 235–256. doi:10.2307/
24 2026388.
- 25 Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.
- 26 Gibbard, Allan (2008). "Rational Credence and the Value of Truth." In *Oxford Studies in Epistemology*,
27 vol. 2, 143–164. Oxford University Press.
- 28 Greaves, Hilary & David Wallace (2006). "Justifying Conditionalization: Conditionalization Maximizes
29 Expected Epistemic Utility." *Mind* 115(459), 607–632.

- 30 Greco, Daniel (2012). "The Impossibility of Skepticism." *Philosophical Review* 121(3), 317–358.
- 1 Greco, Daniel & Brian Hedden (2016). "Uniqueness and Metaepistemology." *Journal of Philosophy*
2 113(8), 365–395. doi:10.5840/jphil2016113825.
- 3 Haslanger, Sally (2006). "What Good Are Our Intuitions: Philosophical Analysis and Social Kinds."
4 *Aristotelian Society Supplementary Volume* 80(1), 89–118. doi:10.1111/j.1467-8349.2006.00139.x.
- 5 Horowitz, Sophie (2013). "Immoderately Rational." *Philosophical Studies* clxvii(1), 41–56.
- 6 Howson, Colin (1992). "Dutch Book Arguments and Consistency." *PSA: Proceedings of the Biennial*
7 *Meeting of the Philosophy of Science Association* 1992. 161–168.
- 8 Ismael, Jenann (2016). "A Philosopher of Science Looks at Idealization in Political Theory." *Social*
9 *Philosophy and Policy* 33(1-2), 11–31. doi:10.1017/s026505251600039x.
- 10 Joyce, James M. (1998). "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65(4),
11 575–603.
- 12 Joyce, James M. (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Be-
13 lief." In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*, vol. 342, 263–297. Synthèse.
- 14 Kahneman, Daniel & Amos Tversky (1972). "Subjective probability: A judgment of representativeness."
15 *Cognitive Psychology* 3(3), 430–454.
- 16 Kemeny, John G. (1955). "Fair Bets and Inductive Probabilities." *Journal of Symbolic Logic* 20(3), 263–
17 273. doi:10.2307/2268222.
- 18 Kitcher, Philip (1990). "The Division of Cognitive Labor." *Journal of Philosophy* 87(1), 5–22. doi:
19 jphil199087178.
- 20 Kratzer, Angelika (1977). "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1(3),
21 337–355. doi:10.1007/BF00353453.
- 22 Kratzer, Angelika (1981). "The Notional Category of Modality." In H.-J. Eikmeyer & H. Rieser (eds.),
23 *Words, Worlds, and Contexts: New Approaches in World Semantics*, Berlin: de Gruyter.
- 24 Kratzer, Angelika (1991). "Modality." In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: An*
25 *International Handbook of Contemporary Research*, 639–650. Berlin: de Gruyter.
- 26 Kratzer, Angelika (2012). *Modals and Conditionals*. Oxford: Oxford University Press.
- 27 Kuhn, Thomas S. (1977). "Objectivity, Value Judgment, and Theory Choice." In *The Essential Tension: Selected Studies in Scientific Tradition and Change*, 320–39. University of Chicago Press.
- 28
29 Lasonen-Aarnio, Maria (2010). "Unreasonable Knowledge." *Philosophical Perspectives* 24(1), 1–21.

- 30 Lehrer, Keith (1990). *Theory of Knowledge*. Routledge.
- 1 Leitgeb, Hannes & Richard Pettigrew (2010a). "An Objective Justification of Bayesianism I: Measuring
2 Inaccuracy." *Philosophy of Science* 77(2), 201–235.
- 3 Leitgeb, Hannes & Richard Pettigrew (2010b). "An Objective Justification of Bayesianism II: The Con-
4 sequences of Minimizing Inaccuracy." *Philosophy of Science* 77(2), 236–272.
- 5 Lewis, David (1969). *Convention: A Philosophical Study*. Wiley-Blackwell.
- 6 Lewis, David (1976). "The Paradoxes of Time Travel." *American Philosophical Quarterly* 13(2), 145–152.
- 7 Pettigrew, Richard (2012). "Accuracy, Chance, and the Principal Principle." *Philosophical Review* 121(2),
8 241–275.
- 9 Pettigrew, Richard (2013). "A New Epistemic Utility Argument for the Principal Principle." *Episteme*
10 10(1), 19–35.
- 11 Pettigrew, Richard (2016). "Accuracy, Risk, and the Principle of Indifference." *Philosophy and Phe-
12 nomenological Research* 92(1), 35–59.
- 13 Ramsey, Frank Plumpton, R. B. Braithwaite & G. E. Moore (1931). "The Foundations of Mathematics
14 and Other Logical Essays." *Mind* 40(160), 476–482.
- 15 Schoenfield, Miriam (2012a). "Chilling Out on Epistemic Rationality: a Defense of Imprecise Cre-
16 dences (and Other Imprecise Doxastic Attitudes)." *Philosophical Studies* 158. 197–219.
- 17 Schoenfield, Miriam (2012b). "Permission to Believe." *Noûs* 48(2), 193–218.
- 18 Sen, Amartya (2006). "What Do We Want From a Theory of Justice?" *Journal of Philosophy* 103(5),
19 215–238. doi:10.5840/jphil2006103517.
- 20 Shimony, Abner (1955). "Coherence and the Axioms of Confirmation." *Journal of Symbolic Logic* 20(1),
21 1–28. doi:10.2307/2268039.
- 22 Staffel, Julia (2019). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press.
- 23 Staffel, Julia (forthcoming b). "Bayesian Norms and Non-Ideal Agents." In Maria Lasonen-Aarnio &
24 Clayton M. Littlejohn (eds.), *Routledge Handbook of the Philosophy of Evidence*, Routledge.
- 25 Tversky, Amos & Daniel Kahneman (1973). "Availability: A Heuristic for Judging Frequency and Prob-
26 ability." *Cognitive Psychology* 5(2), 207–232.
- 729 White, Roger (2006). "Problems for Dogmatism." *Philosophical Studies* 131(3), 525–557.
- 730 Williams, J. Robert G. (manuscript). "A Nonpragmatic Dominance Argument for Conditionalization."