

1 Penultimate version. To appear in *MIND*.

4 **Virtue Epistemology for the Zetetic Turn**

5 J. ADAM CARTER & CHRISTOPHER WILLARD-KYLE

6 Cogito, University of Glasgow, UK & University of Kentucky, USA

7 adam.carter@glasgow.ac.uk; chris.willard.kyle@gmail.com

8
9
10 Abstract: This paper develops and argues for a virtue epistemology that includes performance-
11 normative evaluations of interrogative attitudes (IAs) and not just beliefs. In this way, it brings
12 virtue epistemology to interrogative epistemology. The motivating thought behind our proposal is
13 an analogy: as the semantic property of truth is to answering attitudes, so the semantic property of
14 soundness (the property questions have when they admit of true, direct answers) is to questioning
15 attitudes. With this notion of basic success for IAs in hand, we go on to develop notions of
16 competence and aptness for IAs. Roughly, our view says that (1) IAs are *competent* when they
17 manifest (rightly situated) dispositions to form IAs in sound questions reliably enough, and (2) IAs
18 are *apt* when their being directed towards a sound question manifests the relevant competence. A
19 complete virtue epistemology that treats both beliefs and IAs is, we argue, better positioned to
20 account for the way that success in questioning sets up yet further success in answering. Moreover,
21 our complete virtue epistemology generates evaluative norms on IAs not yet represented in the
22 emerging zetetic literature.

23
24 **Keywords:** inquiry; zetetic epistemology; virtue epistemology; questions

26 **1. Introduction**

27
28 Virtue Epistemology has earned its place among the dominant approaches to epistemology.¹
29 Traditionally, virtue epistemologies have focused on what it means for *beliefs* to be virtuously
30 formed. Within the broadly Sosa, reliabilist tradition, virtue epistemologies begin by identifying
31 a correctness condition (truth)² and then building performance-normative evaluations around
32 that correctness condition. On Sosa's (2007) view, for instance, beliefs are correct when they are
33 *true*, and this basic success can be a constituent of more complex kinds of success, *e.g.*
34 knowledgeable belief, which is not only accurate, but also *adroit* and *apt*.

35 A strength of Virtue Epistemology is the generality of its performance-normative
36 approach.³ Knowledge is just one successful performance among many: the knower's apt belief is
37 like the archer's apt shot.⁴ Accordingly, Virtue Epistemology has been illuminatingly applied to
38 phenomena besides knowledge. There are virtue theoretic treatments of trust (Carter 2020;
39 2022b; 2022a; 2023a), suspension of judgement (Sosa 2021), blame (Simion 2021), and epistemic
40 basing (Miracchi Titus and Carter 2023).

¹ See, *inter alia*, Sosa (1980; 1991; 2007; 2015), Greco (2009; 2010), Zagzebski (1996), Carter (2021; 2023b), Broncano-Berrocal (2017), Kelp (2013; 2018), and Miracchi (2015).

² For simplicity, we focus on standard versions of virtue epistemology that operate only with 'truth-aimed' belief.

³ For recent evaluations of performance-theoretic epistemologies, see Fernandez (2016).

⁴ Though see Chrisman (2012).

41 Missing, however, from the contemporary landscape is a virtue theoretic treatment of
42 *interrogative attitudes* (IAs). IAs are question-directed attitudes like being curious about that are
43 possessed by inquirers in paradigmatic cases of inquiry.⁵ IAs are variously theorized—some have
44 thought that suspension of judgment (Friedman 2017) or mere consideration of a question
45 count. We tend to be more restrictive: IAs towards *Q* are only those attitudes whose possession
46 or manifestation makes it appropriate to say that an agent *wonders* about *Q*.⁶ By focusing on IAs,
47 we focus on the subject’s mental state rather than acts of inquiry that such a state may motivate.
48 Whereas judgments (or beliefs) mentally put forward a proposition, IAs mentally put forward a
49 question. IAs are inner questionings.

50 Questioning and answering are activities with an intimate, dyadic relationship: questions
51 call for answers and answers resolve questions. There’s theoretical reason, then, to treat both as
52 elements of a common ecosystem. Moreover, given the recent ‘zetetic turn’ in epistemology
53 (Kelp 2021a; 2021b; Friedman 2020; forthcoming), it’s of special interest whether virtue
54 epistemology can be extended to interrogative attitudes.

55 There are reasons, however, it hasn’t been done before. Indeed, we see three obstacles to
56 even attempting a performance-normative treatment of IAs.

57 First, performance-normative evaluations within virtue epistemology begin with
58 identifying a basic kind of correctness or success. But plausibly, the kind of correctness most
59 intimately associated with *asking* a question is correctly *answering* it!⁷ If that’s right, then Virtue
60 Epistemology *already* addresses the kind of success that’s relevant to IAs by offering a
61 performance-normative evaluation of agents’ *answers*.

62 Second, the most popular norm on IAs is the *ignorance norm*. It says, roughly, that one
63 shouldn’t inquire into a question unless ignorant of its answer.⁸ This exacerbates the first worry:
64 not only are the successes of questioning attitudes properly found in *answerings*, but, by itself,
65 what having an IA encodes is *lack* of epistemic success. Both problems evoke the view that
66 questionings are mere set-up moves for things that can enjoy epistemic success (namely,
67 *answerings*), but not the sorts of things that can be epistemically successful performances
68 themselves.

69 Third, Sosa (2021, Ch. 2) distinguishes *gnoseology* from *intellectual ethics*.⁹ Gnoseology
70 studies the performance normativity that determines whether a particular affirmation counts as a
71 (knowledgeable) success. But factors other than whether an affirmation is knowledgeable may
72 bear on whether a proposition is sensible to affirm: some propositions are trivial, or
73 pragmatically bad to believe. Plausibly, such factors don’t immediately affect whether we know

⁵ See Friedman (2013; 2017).

⁶ Our notion of IAs is otherwise flexible. There is, however, a puzzle about how to think about what can (in English) be substituted for *Q* in expressions like ‘S wonders about *Q*’. Substitutions for *Q* tend to be natural when the expressed question is embedded in a wh-complement as in ‘Sam wonders about *whether* there is life on Mars’. But substitutions are sometimes unnatural when questions are expressed in the normal way for independent interrogative clauses, e.g., ‘Sam wonders about *is there life on Mars*’. So when reading ‘IA towards *Q*’, the substitution for *Q* should be introduced with a wh-complement. Alternatively (and as the Editors helpfully suggested), one could substitute for *Q* something like: ‘the question: is there life on Mars?’, and likewise for other questions. (Thanks to the Editors for identifying both the puzzle and some solutions.)

⁷ Cf. Kelp (2014; 2021b) on knowledge as the goal of inquiry.

⁸ See Whitcomb (2010; 2017), Friedman (2017), Millson, (2021, p. 685), Sapir & van Elswyk (2021), Willard-Kyle (2021, 2023a), and Haziza (2023). Against, see Archer (2018, 2021), Palmira (2020, p. 4959), Falbo (2021; 2023), and Woodard (2022).

⁹ Sosa (2021, ch. 2).

74 some proposition—I can successfully know how many blades of grass are in a field by counting,
75 however trivial the resulting belief. Such factors belong not to gnoseology but to intellectual
76 ethics.

77 The problem is that, initially, normative considerations governing IAs seem to fall into
78 the ‘intellectual ethics’ bucket.¹⁰ We choose which questions to inquire into based on our
79 interests and practical goals. These seem to bear on intellectual ethics, not the stuff of accuracy-
80 driven performance normativity. The epistemic evaluation of IAs can be subsumed under the
81 study of intellectual ethics.

82 These are important challenges; nevertheless, we argue that there *are* successful,
83 competent, and (epistemically) apt interrogative attitudes. Questionings are set-up moves to
84 answers, but they are not *mere* set-up moves: rather, the kind of epistemic accomplishment that’s
85 distinctive to IAs explains how successful IAs in the domain of questionings set up their
86 inquirers for yet further successes in the domain of answerings.

87 In §2, we’ll articulate a notion of basic success or correctness for IAs. Whereas truth is
88 the basic ingredient in constructing a reliabilist performance normativity for doxastic attitudes,
89 *soundness* (roughly, having a true answer) is the basic ingredient in constructing a reliabilist
90 performance normativity for interrogative attitudes. In §3, we’ll construct a notion of competent
91 (*i.e.*, adroit) and apt IAs using this notion of correctness. Then in §4, we’ll consider the rightful
92 place of interrogative attitudes in virtue epistemology.

93

94 **2 Soundness as success**

95

96 *2.1 Truth and soundness*

97 Beliefs are correct insofar as their contents are true. Virtue Epistemology for affirmations starts
98 with this notion of correctness for beliefs and builds a notion of *apt* belief on its basis. Truth is
99 thus the first ingredient in the virtue-theoretic recipe for apt belief. What is the first ingredient in
100 the virtue-theoretic recipe for apt IAs?

101 Our answer is *soundness*. IAs are correct insofar as their contents are sound. ‘Soundness’ is
102 the semanticist’s term for, roughly, the property questions have when they have true answers. To
103 be correct, qua IA, IAs must pick out a question that presents a genuine target. As truth is to
104 answers, so soundness is to questions.

105 The idea of treating soundness as analogous to truth goes back to early work on the
106 semantics of questions. For example, Belnap and Steel argue that each question presupposes that
107 it has a direct answer and write this:

108

109 To ask something with a false presupposition ...is very much like making a false
110 statement. One can do it knowingly and maliciously, and be exactly on a par with a liar...
111 For this reason we hereby make it a proposal to call a question ‘true’ or ‘false’ according
112 ...as it does or does not have some true answer. (Belnap & Steel 1976, p. 115–16)

113

¹⁰ Sosa (2021, p. 38).

114 Other early proponents of erotetic logic made similar claims.¹¹ While we don't have the nerve to
115 apply the labels 'true' and 'false' to questions, we do propose to take seriously Belnap and Steel's
116 analogy between truth and true answerability (or soundness).

117 The nature of questions is contested territory in semantics. But since at least Hamblin
118 (1976), there is consensus that questions are intimately associated with the propositions that are
119 their candidate answers.¹² The question whether it's raining has the answers *yes* (it is raining) and
120 *no* (it's not raining). The question whether Annie, Beth, or Claire won the race has as its
121 candidate answers that *Annie* won the race, that *Beth* won, and that *Claire* did. And so on.

122 Some questions have true answers; others don't.

123

Sound Questions	Unsound Questions
What colours are on the Scottish flag?	What colour is the hair of the king of France?
Is there a God?	How come the UK hasn't elected a woman as PM?
Was it Washington or (rather) Jefferson who was the first US president?	Was it Lincoln or (rather) Jefferson who was the first US president?

124

125 Only the questions in the first box have true answers. The colours on the Scottish flag are white
126 and blue, but there is no true answer to the question about the haircolour of the king of France.
127 Our account of correctness for IAs says that IAs are correct insofar as they target questions that
128 belong in the first box.

129 The distinction between sound and unsound questions is *not* about whether anyone will
130 ever *know* the answer to the relevant questions. Some people have thought we'll never know
131 whether God exists. The question whether there is a God still *has* an answer even if the answer
132 turns out to be unknowable. Soundness does not require resolvability.

133 Strictly speaking, sound questions must have true *direct* answers. 'Direct answerhood' is a
134 dialectically flexible notion. Here's a gloss: direct answers provide all and only the information
135 that a question calls for. It completely answers the question, 'but just completely' (Belnap & Steel
136 1976, p. 13).¹³ For example, consider:

137

- 138 1. (a) Who won the race: Annie, Beth, or Claire?
139 (b) Annie won the race.
140 (c) Either Annie or Beth won the race.
141 (d) It wasn't Beth who won the race.
142 (e) Annie won the race, and she has red hair.
143 (f) Actually, the race was cancelled, so there wasn't a winner.

144

¹¹ Cf. Leonard (1959), Harrah (1961), and Wiśniewski (1991; 1995). For more recent work on the role of soundness in interrogative epistemology, see Willard-Kyle (2023b, 2024), Whitcomb & Millson (2024), Rosa (2025), and Willard-Kyle, Millson, & Whitcomb (forthcoming). For critique, see Falbo (2024).

¹² See also Karttunen (1977) and Groenendijk and Stokhof (1984).

¹³ For more on 'direct answerhood', see Hintikka (1976), Ginzburg (1995a; 1995b) and Whitcomb & Millson (2024).

145 Among these, only (1b) is a direct answer to (1a). Being told that Annie won the race provides all
146 and only the information that the question <Who won the race: Annie, Beth, or Claire?> calls
147 for. (1c) only *partially* answers the question, and (1d) *eliminates* a wrong answer. Neither provides
148 *all* the information (1a) calls for. In contrast, (1e) says too much: it does not provide *only* the
149 information than the question calls for.

150 The final response, (1f), is different. This response, if true, would make it appropriate to
151 cease inquiring into the opening question. But it doesn't *resolve* the question in the ordinary way.
152 It doesn't provide the information that the question calls for; rather, it denies that the
153 information the question called for was semantically appropriate.

154 Many semanticists¹⁴ don't even consider (1f) an *answer* strictly speaking: it is a mere
155 response. Mere responses can be *issue-dispelling* even if they are not (like direct answers) *issue-*
156 *resolving* (Isaacs & Rawlins 2008, p. 283). But even those who call them answers, e.g. Van
157 Fraassen (1980, p. 140), call them *corrective* answers. Corrective answers deny one of the
158 question's presuppositions and indicate that the initial question is defective in virtue of having no
159 true, *direct* answer. In any case, the very category of *corrective* responses (or answers) vindicates the
160 project of identifying a notion of correctness for IAs. For there to be corrective responses to
161 questions presupposes that questions can be *corrected*. And if questions can be *corrected*, then
162 questions can be correct. So while questions do not have truth conditions in the way that
163 propositions do, they still have a standard of correctness.

164 165 2.2 *Varieties of Success*

166 As beliefs inherit the correctness conditions of the propositions they are directed toward, so IAs
167 inherit the correctness conditions of the questions they are directed toward. As beliefs are
168 correct insofar as their contents are *true*, IAs are correct insofar as their contents are *sound*. That's our
169 view.¹⁵

170 In defending the idea that soundness is the standard of correctness for IAs, we've taken
171 as the salient contrast the idea that IAs don't have a standard of correctness *at all*. But now that
172 we have a candidate notion of IA correctness on the table, we've opened ourselves up a separate
173 criticism: that the standard of correctness is something *other* than soundness! After all, having an
174 IA towards *Q* involves being in a motivational state that is something like wanting (or aiming) to

¹⁴ E.g., Belnap and Steel (1976), Groenendijk & Stokhof (1984),

¹⁵ One might think that even if the question 'Is the present King of France bald?' is defective *qua* question (since unsound), it isn't a mistake to simply *consider* whether the present king of France is bald. But isn't considering also an IA? If so, is soundness really a standard of correctness for IAs? We think that not all ways of considering are ways of having an IA. Think, for instance, of 'considering' in thought experiments or fiction. I can consider for such purposes, e.g., the moon's being made of cheese without having any IA about this. There is, for all we've said, nothing defective about considering unsound content in this sense (when no IA is present). This is compatible with the idea that IAs are correct iff directed towards sound questions.

But perhaps *some* considerings are IAs, and perhaps *those* considerings are not defective even when directed toward unsound questions. Suppose I (wrongly thinking France has a king) consider whether he is bald, in a way that makes it true to say I am so wondering. In this case, we are happy to say that the relevant considering *is* defective. For one thing, on the assumption this is a genuine IA, we can expect one who considers whether there is presently a King of France to acknowledge a mistake upon acquiring evidence that there is no King of France.

Second, we want to distinguish between different dimensions of appropriateness. If someone lied to you about the French monarchy, it might then be epistemically justified to consider whether the King of France is bald; but this kind of appropriateness would then come apart from the separate kind of appropriateness applicable to the IA as such. (Compare: you could be practically justified in firing an arrow at a target by throwing it backwards—perhaps an expert misled you about proper form—but it's still a mistake (incorrect) if it misses the target, and still incompetent when thrown in such a way. Thanks to a referee at MIND for discussion.)

175 know Q 's answer (cf. Friedman 2017, p. 307–8, Deigan forthcoming). So perhaps the central aim
176 of an IA—and the condition for its success—is success in answering.¹⁶

177 Though tempting, this proposal is mistaken. First, it conflates the correctness of
178 questioning and answering attitudes. Suppose Sam has long had an IA towards the question of
179 how black holes form. Even if Sam never *answers* the question how black holes form, intuitively,
180 Sam's *question* will have been a good one. And so, we think, their IA (*qua* IA) would have been
181 correct even if unaccompanied by a correct answering attitude.

182 Second, we can grant that having an IA towards Q involves being in a motivational state
183 aimed at knowing Q 's answer. But just because having an IA involves having a certain aim
184 doesn't settle whether that aim is a correct or appropriate one to have. Our claim is: to the extent
185 that having an IA towards Q involves having an aim to answer Q , what makes having *that* aim
186 (epistemically) appropriate is Q 's being sound.

187 Nevertheless, the objection contains an insight. Obviously, the correctness conditions for
188 questioning and answering attitudes are intimately related. We don't think the relationship is
189 *identity*, but our account should explain their deep connection. To that end, (inspired by Roberts
190 1996), we think it's productive to think of IAs as *set-up moves* for answering attitudes. Set-up
191 moves are doubly evaluable for success. Set-up moves can be evaluated for whether they lead to
192 a success of the sort of move it is a set-up move *for*, but also for whether they are successful
193 performances of the kind to which the set-up move belongs.

194 Here's an example. In tennis, one can make a successful serve. But here's another thing
195 one can successfully do: make a good pre-serve *toss*. The toss is a set-up move for the serve. In
196 that sense, the *ultimate* aim of making a toss is to facilitate making a good serve.

197 But a toss can also be evaluated *as a toss*. I can try to make a good serve, but I can also try
198 *to make a good toss* in pursuit of that further goal. Other things being equal, a good toss is slightly
199 ahead of the baseline and neither too low nor too high.

200 A toss can be good or bad, *regardless of whether the resulting serve is*. Sometimes a player
201 makes a good toss but serves poorly anyway; sometimes a player overcomes a bad toss to make a
202 good serve. Good tosses tend to increase the odds that the player will serve well, although not
203 necessarily so (as when environmental conditions are bad). Success in tossing and success in
204 serving are related but distinct.

205 Or switching sports, it's one thing for a pass to constitute an assist (a pass which leads to
206 a successful shot) and another thing for a pass to be successful as a pass (regardless of whether
207 the sequence of play results in a goal). More generally, when an action φ of type M is a set-up
208 move for an action of type M', φ can be evaluated both (1) by whether φ led (in the right sort of
209 way) to a successful performance, ψ , of type M', and (2) by whether φ was successful as a
210 performance of type M. When we say that successful IAs are those directed towards sound
211 questions, we are engaged in the second kind of evaluation.

212 By thinking about IAs as set-up moves for answering attitudes, we can explain the
213 thought that, in an important sense, the ultimate aim of an IA is realised by a correct answering
214 attitude towards the question posed. And so, in addition to being evaluated for soundness, IAs
215 can be evaluated with respect to whether they lead (in the right sort of way) to successful
216 answering attitudes. Such IAs constitute interrogative assists. But that shouldn't stop us from

¹⁶ Thanks especially to Matt McGrath for conversation on related points.

217 asking whether an IA is successful *qua* IA. Indeed, this dual-aspect evaluation is typical for set-up
218 moves.

219 If IAs are correct insofar as they are directed towards sound questions, in what sense are
220 they set-up moves for answering attitudes? Correct IAs guarantee that there is a correct
221 answering attitude to the question posed by the IA. The correctness conditions for questioning
222 attitudes thus *enable* the correctness conditions for answering attitudes. Conversely, IAs are
223 incorrect insofar as they pose questions that are not correctly answerable. Although soundness is
224 distinct from truth as a standard of correctness, it is essentially truth-involving, since questions
225 are sound insofar as they have true, direct answers. In that way, the standards of correctness for
226 questioning and answering attitudes are both, broadly speaking, accuracy relevant.

227 Moreover, the relevance of accuracy to both IAs and beliefs is further evidenced by the
228 plausibility of a kind of ‘shared transparency’ we find across the belief/truth connection and the
229 IA/soundness connection. Consider that, at least according to one familiar proposal (e.g., Shah
230 2003, p. 447), believing that p normatively requires, and might even metaphysically necessitate (at
231 least for those who possess the concept of truth), believing that p is true. Likewise, it is plausible
232 that having an IA towards Q seems to normatively require, and may even metaphysically
233 necessitate (at least for those who possess the concept of soundness), believing that Q is sound.¹⁷
234 Linguistic data seem to support the analogy. There’s a *prima facie* Moorean oddity to asking a
235 question and wondering if the question is sound.¹⁸ Likewise, there’s a *prima facie* Moorean
236 oddity about believing something and wondering if the proposition believed is true.¹⁹ Such
237 observations deepen the analogous relationship we will explore between success in asking and
238 success in answering.

239 Another proposal for interrogative-relevant success in the literature, one which is
240 epistemic rather than semantic, comes from Rosa (2023), who argues that suspension of
241 judgment may have a kind of success analogous to knowledge (2023, p. 2006–7). According to
242 Rosa, suspending judgment on p is correct iff p and $\neg p$ each might be true given the subject’s
243 epistemic circumstances (Rosa 2023, p. 2004). If combined with Friedman’s (2017) view that
244 suspension of judgment is an IA, this proposal modulates into a standard of correctness for (at
245 least a class of) IAs.

246 Here, we simply note that we deny Friedman’s thesis that suspension of judgment is an
247 IA. We think it’s possible to suspend judgment on a question without wondering what the
248 answer is, for the sort of reasons articulated in McGrath (2021, p. 476–9). So our thesis is
249 compatible with Rosa’s. Notably, the notions of correctness involved are different. Ours requires
250 of an IA that it be directed towards a question that has a (true, direct) answer, whereas Rosa’s
251 requires that one’s epistemic circumstances leave open multiple answers as possible. A question
252 could have multiple answers that might be true given the agent’s circumstances even though
253 none of them is, and a question might be sound even though only one answer (the correct one)
254 might be true given the agent’s epistemic circumstances.

¹⁷ Thanks to a referee at MIND for suggesting this comparison.

¹⁸ Cf. Friedman (2017, p. 315–16).

¹⁹ How well does this transparency analogy hold up? One might point out that good grounds for believing that p is true are, *ipso facto*, good grounds for believing that p . But good grounds for believing that Q is sound are not *ipso facto* good grounds for having an IA towards Q . We grant this, but point out that the analogy is easily regained via negation: good grounds for *not* believing that Q is sound are *ipso facto* good grounds for *not* having an IA towards Q . Thanks to a referee at MIND for prompting discussion here.

255 With a standard of correctness for IAs in hand, we can now consider whether IAs are
256 performance-normatively evaluable. Can IAs, like beliefs, be not just correct but *competent* and
257 *apt*? And, if so, do such IAs mark accomplishments that we find valuable?

258

259 3. Competent and apt IAs

260

261 Recall the question, ‘Who won the race: Annie, Beth, or Claire?’, for which ‘Annie won the race’
262 was the true, direct answer. IAs into this question are *successful* in the sense that they meet the
263 standard of correctness (*viz.*, being directed towards a sound question) for performances of that
264 type.

265 Note that a given IA would not lose its status as a successful even if an inquirer were to
266 ask a sound question by luck. Suppose that Ian is simply posing as a racing fan in order to gain
267 social access. Ian can’t remember exactly who was competing in the race at hand (a race with just
268 three participants), while nonetheless hoping to appear informed about the race. Suppose Ian
269 prepared for his ploy by memorising the names of 20 people who hoped initially to qualify for
270 this race, while remaining—unfortunately for him—unsure exactly which of the 20 are actually
271 running. Picking three names randomly, he asks ‘Who won the race: Annie, Beth, or Claire?’

272 One thing we can say about Ian’s posing of the question, is that—qua performance—it is
273 successful or correct in the way a luckily true belief is still correct (despite any other performative
274 failing) qua belief, and in the way a basketball shot that goes through the hoop is still successful
275 as a shot even if heaved at the basket with one’s eyes closed.²⁰

276 But isn’t Ian’s IA, in a sense, also obviously criticisable? Yes, and this despite the
277 question’s soundness. That you can form an IA badly yet successfully is analogous to what
278 happens in other domains of performance (*e.g.*, lucky guesses and ‘Hail Mary’ shots that happen
279 to go in). Here, an analogy to the value of knowledge literature is helpful: epistemologists care
280 about more than merely true (*i.e.*, successful) belief; unsurprisingly, the value of knowledge
281 (roughly: a belief whose correctness is in some way through ability) is widely taken to exceed the
282 value of mere true belief.²¹ Not implausibly, there is a species of IAs the value of which is greater
283 than *merely successful* IAs. If the value of knowledge debate is any guide, we’ll locate such value by
284 first looking at what questioning well, *viz.*, forming IAs *competently*, involves, and then by
285 considering how successful IAs can be connected to competent IA-formations.

286

287 3.1. *Competent IAs*

288 A competence is, roughly, an agent’s disposition to perform well in a given domain of
289 endeavour.²² What counts as performing well, or reliably enough, depends on the domain of
290 performance. Possessing a bomb-defusing competence might require a very high level of

²⁰ We think that Rosa (2023) is best interpreted as not giving a view about correctness for IAs because we think suspension isn’t an IA. But suppose it were. Notice that an IA can apparently be incorrect even if it meets Rosa’s (2023) condition. For suppose that it’s consistent with my epistemic situation that any of the three racers (Annie, Beth, or Claire) won, and, moreover, Claire in fact did. And suppose I have an IA towards (merely) whether Annie or Beth won (omitting Claire). There is a sense in which an IA towards that question is incorrect, even though multiple answers to my question (namely, that Annie won and that Beth won) might be true given my epistemic circumstances.

²¹ See Kvanvig (2003); Sosa (2000); Pritchard (2009a; 2009b); Carter, Pritchard, and Turri (2018); Riggs (2002); Olsson (2007).

²² See Sosa (2010).

291 reliability at succeeding when one tries; though, one might be a competent pinch hitter in
292 baseball even with a batting average of .250. These points are familiar from standard virtue
293 epistemology, which takes doxastic competences to require reliability but not infallibility.

294 The term ‘competence’ is often used to track two distinct kinds of dispositions. Consider
295 flammability. In one sense, a match remains flammable (as a cube of ice is not) even when the
296 match is covered in sand, deprived of oxygen. It remains true in those circumstances of the
297 match that it *would catch fire* if struck in normal conditions; the same is not true of the ice cube.
298 That the match is flammable in this ‘general’ sense is compatible with it being true that (here and
299 now, under sand without oxygen) the match would not light if struck. A more demanding sense
300 of ‘flammable’ requires that the match *is presently in appropriate conditions*, such that the match
301 would now (likely enough) catch fire if struck.

302 What goes for dispositions goes for competences (an agent’s dispositions to perform
303 well). We can distinguish (i) an agent’s innermost competence (or *skill*) from (ii) her *complete*
304 *competence*. A performer in any domain of endeavour, D, might retain the skill to φ_D reliably, just
305 so long as it’s true of her that, *were she in proper shape and situation* (for φ_D -ing) she’d succeed
306 reliably enough. An expert ornithologist retains a bird identifying skill even when not in a
307 position to exercise it, and likewise, an informed and skilled *questioner* might possess a skill at
308 forming IAs even if in the wrong conditions to exercise it. The questioner might be asleep or
309 drugged, and so not in the right shape to reliably ask sound questions. A questioner might also
310 not be in the appropriate circumstances to exercise her questioning skill if she is misled *ex ante*.
311 For instance, we shouldn’t evaluate questioning skill by assessing how reliably an inquirer would
312 form IAs in sound questions *in circumstances leading up to a surprise party*—that is, circumstances in
313 which the risk of asking an *unsound* question is unusually high. Evaluating for questioning skill in
314 *that* kind of social-epistemic predicament would be like assessing one’s driving skill by asking
315 how reliably they’d perform on unusually slick roads.

316 Having a *complete competence* asymmetrically entails having a skill. To have a complete
317 competence (in D), the agent must not only possess the relevant D-skill but also be situated in
318 conditions appropriate for its exercise. In the case of forming IAs, then, we can distinguish
319 between *erotetic* (or questioning) *skill*, which is a disposition to form IAs in sound questions
320 reliably enough, when in proper shape and situation, and (complete) *erotetic competence*, which
321 requires one have an erotetic skill and being in appropriate conditions.

322 We can now distinguish *successful* from *competent* IAs:

- 323
- 324 ● An IA is successful (qua IA) iff the question towards which the IA is directed is sound,
325 *viz.*, if it admits of a direct, true answer, and regardless of how it was formed;
 - 326 ● An IA is *competent* (qua IA) iff its formation manifests a complete erotetic competence,
327 and regardless of whether the IA’s question is sound.²³
- 328

329 But what do competently (or incompetently) formed IAs look like? We start with some *bad*
330 examples:

²³ These are instances of *evaluative norms* (McHugh 2012, p. 9; Simion et al 2016, p. 385-7). Evaluative norms regulate what it takes for a token of a particular type of thing to be good or bad with regard to its type; the ‘goodness’ here is *attributive* in Geach’s (1956) sense; this is the sense in which a sharp knife is a good knife, *qua* knife, regardless of whether good or bad *simpliciter*.

331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363

TABLOIDS: A recent British tabloid headline reads: ‘Why does Meghan Markle hate the royal family?’ Despite background knowledge that the track-record of such tabloids is abysmal, a reader begins to wonder: Why does Meghan Markle hate the royal family?

BIASED BOSS: A company is trying to execute a difficult merger, and Brad tasks three of his employees, Angie, Bill, and Callum, with developing a successful strategy. Later in the week, Brad asks his employees for their best plan, and the team explains plan P. Brad is impressed with the strategy and wonders: Did Bill or Callum originally come up with plan P? Because of his implicit bias, Brad’s question fails to make room for the possibility that *Angie* originally came up with plan P.

Tabloids aren’t any more reliable at *querying* than they are at *testifying*, and implicit bias can infect our interrogative attitudes no less than our doxastic attitudes. The IAs in the examples *might* be directed towards sound questions—maybe Callum *did* come up with the relevant plan, for instance. But the inquirers do not manifest a competence to wonder about questions only if they are sound. These cases recommend thinking about competences to form IAs towards sound questions by analogy with competences to believe true propositions.

In contrast, here’s a competently formed IA:

CLUEDO: Rowan is playing a game of Cluedo. At the outset of the game, they know that the murderer must be either: Colonel Mustard, Mrs Peacocke, Miss Scarlet, Revd Green, Mrs White, or Professor Plum. Rowan looks at their hand and sees that they have the following cards: Mrs Peacocke, Revd Green, and Mrs White. (No character in a player’s hand can be the murderer.) So Rowan wonders: Is the murderer Colonel Mustard, Miss Scarlet, or (rather) Professor Plum?

As CLUEDO (positively) and BIASED BOSS (negatively) illustrate, agents often manifest competence through their IAs by being disposed to have IAs towards questions that exhaust all (plausible enough) alternatives.²⁴ Rowan’s IA—unlike the biased boss’s—manifests a disposition to reliably enough cover all the relevant alternatives. Dispositions to cover all relevant alternatives help make it likely that the resulting questions will be sound.²⁵ Of course, these examples do not exhaust the ways that IAs can (fail to) manifest an erotetic competence. But

²⁴ cf. Millson (2021a, p. 225–27).

²⁵ It’s worth noting that IAs can be incompetently formed, and thus unjustified on VE, even when the inquirer possesses sufficient evidence to know the question is sound. Imagine a modified version of BIASED BOSS. In the original case, the boss incompetently forms an IA towards whether Bill or Callum came up with the plan *and also* it is propositionally unjustified for the boss to believe that the relevant question is sound. In the new version, suppose Bill is highly biased and relies on the availability heuristic (e.g., Tversky & Kahneman 1973), choosing to focus on those who most recently come to mind when querying responsibility for outcomes. Since he’s just thought about Bill, Callum, and Angie, the bias leads him to ask ‘Did Bill, Callum, *or Angie* originally come up with plan P?’ Manifesting this bias is an unreliable method of forming such questions—after all, one could easily pass over an alternative merely because it wasn’t recently in mind. Yet, we can stipulate that in this specific instance, the IA triggered by this bias aligns with the available evidence. According to the VE approach we favour, this IA would be deemed unjustified due to its incompetent formation, despite being propositionally justified by, in this case, luck. Thanks to a referee for encouraging us to explore such cases. For further discussion on how virtue-theoretic approaches in epistemology address the significance of evidence, see, e.g., Sylvan and Sosa (2018), Sylvan (2021), and Neta (2022).

364 they give us a handle on the conceptual point that since it is possible for IAs to be successful, it
365 is also possible for IA-formation to manifest a competence to be successful.

366 Two points of clarification. First, the reason *why* a given IA can be competent *even if*
367 *unsuccessful* lines up symmetrically with what Virtue Epistemology maintains about belief. Visual-
368 perceptual competences, for example, are dispositions to affirm correctly *reliably enough* but not
369 infallibly. What *threshold* of reliability we should demand of questioning competence is a separate
370 matter; we think it's plausible that—given that both IAs and beliefs are moves within inquiry
371 more broadly—the level of reliability required for a questioning competence will roughly match
372 the level of reliability required for doxastic competences.

373 Second, while the success and competence norms of IAs are such that their satisfaction is
374 logically independent of each other, we might nonetheless be tempted to think that whenever
375 both norms *are* satisfied by a given inquirer, in the case of a given question asked, the IA is
376 thereby, 'in the clear'. But as we'll see, this isn't quite right.

377

378 3.2. *Apt IAs*

379 Consider the following case:

380

381 GETTIERED IA: Edmund is a responsible race enthusiast who has spent the week
382 fretting over which of the three competitors in the upcoming race, Annie, Beth, or
383 Claire, is likely to win. Unfortunately, Edmund's television loses power right before the
384 race. Unable to fix it, he calls his friend at the racetrack and asks 'So who won the race:
385 Annie, Beth, or Claire?' His friend offers the true, direct answer 'Annie won the race.'
386 Unbeknownst to Edmund, the race was called off moments before its scheduled start,
387 due to a security breach; however, shortly after the race was called off (and before
388 everyone left) the security breach was luckily resolved (a disgruntled former employee
389 had stolen the starter gun, and just so happened to be caught). And so, the race
390 proceeded.

391

392 The above case exhibits lucky questioning, in which the relevant luck taints our assessment of
393 Edmund's questioning successfully. But the luck does not resemble (in epistemology) a
394 completely wild guess or (in an athletic performance) a shot performed blindly, and that's
395 because the IA is a *competent* one.

396 Edmund's case is something like a 'justified, successful IA' where the source of the
397 justification (or competence) has little to do with why the trusting is successful. This suggests
398 that even when one forms an IA in a way that manifests competence and success, something still
399 might fall short. One's forming an IA successfully might still not *manifest* one's competence to
400 question well.

401 On a performance normativity model, when a performance's success issues (non-
402 deviantly) from a complete competence, the performance is not only successful and competent
403 but *apt*. There is something better than forming an IA (merely) successfully and competently: one
404 can form an IA aptly, where one's doing so successfully manifests one's complete questioning
405 competence. Edmund's IA falls short by being inapt. We can thus add an *aptness* norm, which
406 asymmetrically entails the other two:

407

- 408 • An IA is *apt* (qua IA) iff its success (namely, being directed towards a sound question)
409 manifests a complete erotetic competence.

410

411 By referencing the aptness norm, we can assess the quality of Edmund’s IA in a way that is
412 sensitive to how it comes up short in just the way a Gettiered belief does. A Gettiered belief is
413 not *apt* in the way knowledge (arguably) requires, and by parity of reasoning, neither is a
414 ‘Gettiered IA’.²⁶

415 This observation maps well to our earlier discussion on the *value* of knowledge. Note the
416 Virtue Epistemologist’s reasoning for why knowledge is more valuable than merely true belief.
417 Achievements are successes that are *because* of ability. Knowledge, given VE, is a cognitive
418 success that is because of cognitive ability.²⁷ Therefore, knowledge is a (species of) cognitive
419 achievement. Achievements are finally valuable in a way that isn’t swamped by the value of the
420 success alone; therefore, knowledge (qua achievement) is finally valuable.

421 The Virtue Epistemologist is equipped to assess the quality of IAs in like manner. Apt
422 IAs, on Virtue Epistemology, are erotetic successes that are successful *because* of an exercise of
423 erotetic ability. Thus, apt IAs are an erotetic achievement. Achievements are finally valuable;
424 therefore, apt IAs (*qua* erotetic achievement) are finally valuable.

425 We began this section by positing a parallel between beliefs and IAs: each has an internal
426 ‘aim’ or standard of correctness. In slogan form, truth is to answering as soundness is to asking.
427 This section has drawn out further consequences of this parallel within a virtue-theoretic
428 framework. Just as beliefs can be not only successful but competent and apt, so too can IAs.
429 What results is a three-fold symmetry of evaluative²⁸ norms governing believing and questioning:
430

virtue-theoretic erotetic norms	analogue belief norms
<p>Success norm: An IA is successful (qua IA) iff the question towards which the IA is directed is sound, <i>viz.</i>, it admits of a direct, true answer, and regardless of how it was formed.</p>	<p>Success norm: A belief is successful (qua belief) iff it is true, and regardless of how it was formed.</p>
<p>Competence norm: An IA is <i>competent</i> (qua IA) iff its formation manifests a complete erotetic competence, and regardless of whether the IA’s question is sound.</p>	<p>Competence norm: a belief is <i>competent</i> (qua belief) iff its formation manifests a complete belief-forming competence, and regardless of whether it is true.</p>
<p>Aptness norm: An IA is <i>apt</i> (qua IA) iff its success (namely, being directed towards a</p>	<p>Aptness norm: a belief is <i>apt</i> (qua belief) iff its success manifests a complete belief-forming competence.</p>

²⁶ Bi-level virtue epistemologists will be attracted to a further parallel between beliefs and IAs—beyond mere aptness—which countenances *full aptness* of both beliefs and IAs. Fully apt IAs are guided to aptness by a reflectively apt risk assessment. See Sosa (2015, ch. 3).

²⁷ See Greco (2010, ch. 6), Pritchard (2009a; 2009b) and Pritchard, Turri, and Carter (2022, ch. 3).

²⁸ This contrasts with the prescriptivist approach to soundness-side erotetics in Willard-Kyle (2023b) and dodges critiques in Palmira (ms).

sound question) manifests a complete erotetic competence.	
---	--

Table 2. Virtue theoretic symmetries: norms on questioning and believing

4. Virtue Epistemology Expanded

Our paper has developed a performance-normative treatment of IAs that mirrors the way virtue epistemology handles belief. The various levels of interrogative success are built on a notion of question soundness. But the reader might wonder: isn't this picture incomplete? After all, much that intuitively matters for assessing IAs apparently has nothing to do with soundness. Whitcomb (2010, 2017) and Friedman (2017), *inter alia*, argue that IAs are only proper insofar as their inquirers are *ignorant* of the answer. And many have argued that some questions are just more *significant* than others in a distinctively epistemic way: *e.g.*, it's just more interesting (epistemically) what the fundamental scientific laws are than where my dentist's office is.²⁹ Why hasn't our account said anything about *these* considerations?

In response, recall Sosa's distinction between gnoseology and intellectual ethics. Gnoseology is the theory of knowledge. Its domain only concerns things that are connected to the performance-normative evaluation of beliefs in pursuit of truth:

Here we are in the realm of telic assessment, and, in the case of judgment..., the aim determines how it would be proper to assess the pertinent attempts. For such assessment, the aim that matters is that of getting it right aptly with one's alethic affirmation. All such gnoseological assessment is then by reference to that aim, so that practical, extraneous aims are irrelevant. (Sosa 2021, p. 39)

That does not mean, Sosa stresses, that other aims are irrelevant to epistemology generally. When considerations bear on what we should believe but not on gnoseology, they are a part of 'intellectual ethics'.

Sosa draws this distinction to preserve the special character of performance-normative evaluations of beliefs with respect to the goal of truth while liberally allowing that a wide range of aims may be relevant to normative epistemology generally. So, suppose I carefully count the blades of grass in my garden and come to know that it contains 121,544,382 blades of grass. That belief gets full marks for aptness in the domain of gnoseology—it was, by hypothesis, knowledgeable—even if the most pressing thing to say about my belief is that it's criticisable for being epistemically insignificant and a waste of cognitive resources. Notice that since gnoseology only concerns evaluating whether beliefs are successful *in ways that relate to the aim of truth*, gnoseology is always silent on broader issues of belief *selection* (that is, since a belief can be, *e.g.*, apt, even if insignificant or otherwise problematic).³⁰ We submit that Sosa's distinction between gnoseology and intellectual ethics regarding the epistemology of beliefs mirrors an analogous

²⁹ This example is from Grimm (2011, p. 521), but there is a large literature on epistemic significance. See Locke (1706/1996, §43–44), Hurka (1993, p. 100), Goldman (1999, p. 88, 94–96), Bishop & Trout (2005, p. 93–103), Whitcomb (2007), Roberts & Wood (2007, p. 153–82), Brady (2009), and Baril (2010).

³⁰ Nor do we impose the requirement that a belief has value in the domain of the ethics of belief only if it is true.

469 distinction between erotetics and interrogative ethics regarding the epistemology of IAs.
470 Erotetics, as we use the term, involves the telic assessment of IAs with reference to the aim of
471 soundness. Erotetics does *not* consider how well an IA does epistemically with respect to other
472 aims. Since erotetics only concerns whether IAs have been successful in ways that relate to
473 soundness, erotetics is silent on broader issues of question selection.³¹

474 So, suppose that I wonder how many blades of grass are in my garden. This is an
475 insignificant question. Nevertheless, supposing I have a garden, this question is sound. An IA
476 towards this question could get full marks in the domain of erotetics. And that's so even if the
477 most urgent thing to say about such IAs is that they are (epistemically) criticisable for being
478 trivial. Having IAs about the number of blades of grass in your garden manifests bad *question*
479 *selection*, but it doesn't make one's questions (even disposed to be) unsound.

480 Something similar is true of the Ignorance Norm. Inquiring into a question you already
481 know the answer to evinces poor question selection. Here, the issue is not that the questions
482 about what one already knows are insignificant but that they have a kind of agent-relative
483 redundancy. IAs into questions to which one already knows the answer never yield new
484 information of the sort called for by the relevant question. In that sense, such questions are
485 always agent-relatively trivial.

486 Inquiries into questions with (already) known answers feel pointless in a way that inquiries into
487 merely insignificant questions do not. We can explain this by recalling that IAs are set-up moves
488 for an agent to acquire a correct answering attitude. Anyone who knows the answer to *Q* already
489 has achieved the thing that having an IA towards *Q* would be a set-up move for. And so, we can
490 explain why the Ignorance Norm has (deservedly) been of special interest to the recent zetetic
491 turn in epistemology. But that doesn't mean ignorance bears on erotetics in our narrow sense.

492 Indeed, now that we've articulated our account of erotetics, we can build on Sosa's
493 binary distinction between gnoseology and intellectual ethics to create a fourfold distinction that
494 maps different domains of epistemic normativity. On the performance-normative side,
495 gnoseology mirrors erotetics. We split intellectual ethics in twain. The ethics of belief pairs with
496 gnoseology, and interrogative ethics pairs with erotetics. Both fall under the broader category of
497 intellectual ethics—those aspects of epistemic normativity that are not performance-normatively
498 evaluable in terms of the basic goals of truth (gnoseology) or soundness (erotetics). Taken
499 together, this fourfold distinction maps the regions of epistemic normativity within virtue
500 epistemology as regards both questioning and answering attitudes. Importantly, nothing we've
501 said commits us to thinking that interrogative ethics is less a part of 'real epistemology' than the
502 performance-normative evaluations central to erotetics. Our goals are taxonomical not
503 hierarchical.³²

504 The motivating thought behind our proposal is an analogy: as truth is to answering
505 attitudes, so soundness is to questioning attitudes. These semantic concepts of truth and
506 soundness are interrelated. *Sound questions* are those that have *true answers* of the right sort. This

³¹ Nor do we impose the requirement that an IA has value in the domain of interrogative ethics only if it is directed towards a sound question.

³² We confess that we've not fleshed out a full picture here and that our account incurs certain explanatory burdens that, for reasons of length, we shall not attempt to fully shoulder. Of particular interest is whether there isn't *some* sense in which truth is the primary dimension of assessment in the domain of belief. And if that is so, can our account explain how soundness is the primary dimension of assessment for IAs in roughly the same way? Or is there, rather, a disanalogy here that stands in need of explanation? (We're grateful to conversation with the Editors for putting these pressing questions to us.)

507 tight connection between semantic concepts is reflected by the connection between beliefs and
508 IAs. IAs are successful whose contents are sound: beliefs are successful whose contents are true.
509 Accordingly, successful IAs set up inquirers for the possibility of successful affirmations. These
510 notions of basic success can both be treated according to the performance normativity
511 distinctive to virtue epistemology. Like beliefs, IAs can be not only successful but competent and
512 apt. A complete virtue epistemology treats affirmative and interrogative attitudes as part of a
513 common ecosystem. IAs play a preparatory role, and perhaps for this reason, their contributions
514 to epistemology have sometimes been overlooked. But IAs are not *mere* set-up moves: they have
515 their own distinctive kind of epistemic, virtue-theoretically evaluable success.³³

³³ We're grateful to Carolina Flores, Arianna Falbo, Giada Fratantonio, Tim Kearl, Matt McGrath, Lilith Newton, Michele Palmira, Ernie Sosa, Josh Thorpe, audiences at the Scottish Epistemology Early Career Research Group, Autonomous University of Madrid, and the 2023 Pacific APA, our referees, and the Editors at MIND for thoughtful feedback on this project. Willard-Kyle's research on this project was supported by Therme Group; Carter's research is supported by AHRC Grants No. AH/W005077/1 and AH/W008424/1.

516 **References**

- 517 Archer, Avery 2018, ‘Wondering About What You Know’, in *Analysis* 78
- 518 ———2021, ‘The Aim of Inquiry’, in *Disputatio* 13
- 519 Baril, Anne 2010, ‘A Eudaimonist Approach to the Problem of Significance’, in *Acta Analytica* 25
- 520 Belnap, Nuel & Thomas Steel 1976, *The Logic of Questions and Answers* (New Haven: Yale University Press)
- 521 Bishop, Michael A. and Trout, J.D. 2005, *Epistemology and the Psychology of Human Judgment* (Oxford: Oxford
- 522 University Press).
- 523 Brady, Michael 2009, ‘Curiosity and the Value of Truth’, in A. Haddock, A. Millar and D. Pritchard (eds.),
- 524 *Epistemic Value* (Oxford: Oxford University Press)
- 525 Broncano-Berrocal, Fernando 2017, ‘A Robust Enough Virtue Epistemology’, in *Synthese* 194
- 526 Carter, J. Adam 2020, ‘On Behalf of a Bi-Level Account of Trust’, in *Philosophical Studies* 177
- 527 ———2021, ‘De Minimis Normativism: A New Theory of Full Aptness’ in *Philosophical Quarterly* 71
- 528 ———2022a, ‘Trust and Trustworthiness’, in *Philosophy and Phenomenological Research* 107
- 529 ———2022b, ‘Trust as Performance’, in *Philosophical Issues* 32
- 530 ———2023a, *A Telic Theory of Trust* (Oxford: Oxford University Press).
- 531 ———2023b, *Stratified Virtue Epistemology: A Defence* (Cambridge: Cambridge University Press)
- 532 Carter, J. Adam and Duncan Pritchard 2015, ‘Knowledge-How and Epistemic Luck’, in *Noûs* 49
- 533 Carter, J. Adam, Duncan Pritchard, and John Turri 2018, ‘The Value of Knowledge’, in *Stanford*
- 534 *Encyclopedia of Philosophy* <[https://plato.stanford.edu/archives/fall2022/entries/knowledge-](https://plato.stanford.edu/archives/fall2022/entries/knowledge-value/)
- 535 [value/](https://plato.stanford.edu/archives/fall2022/entries/knowledge-value/)>
- 536 Chrisman, Matthew 2012, ‘The normative evaluation of belief and the aspectual classification of belief
- 537 and knowledge attributions’, in *Journal of Philosophy* 109
- 538 Deigan, Michael forthcoming, ‘Wondering as Appetitive Desire’, in *Analysis*
- 539 Falbo, Arianna 2021, ‘Inquiry and Confirmation’, in *Analysis* 81
- 540 ———2023, ‘Inquiring minds want to improve’, in *Australasian Journal of Philosophy* 101
- 541 Falbo, Arianna 2024, ‘Inquiry for the Mistaken and Confused’, in *Philosophy and Phenomenological*
- 542 *Research* 109
- 543 Fernandez, Miguel Angel 2016, *Performance Epistemology: Foundations and Applications* (Oxford:
- 544 Oxford University Press)
- 545 Friedman, Jane 2013, ‘Suspended Judgment’ in *Philosophical Studies* 162
- 546 ———2017, ‘Why Suspend Judging?’ in *Noûs* 51
- 547 ———2019, ‘Inquiry and Belief’ in *Noûs* 53
- 548 ———2020, ‘The Epistemic and the Zetetic’ in *Philosophical Review* 129
- 549 ———forthcoming, ‘Zetetic Epistemology’, in A.K. Flowerree and B. Reed (eds.), *Towards an Expansive*
- 550 *Epistemology: Norms, Action, and the Social Sphere* (New York: Routledge)
- 551 Goldman, Alvin 1999, *Knowledge in a Social World* (Oxford: Oxford University Press)
- 552 Ginzburg, Jonathan 1995a, ‘Resolving Questions, I’, in *Linguistics and Philosophy* 18
- 553 ———1995b, ‘Resolving Questions, II’, in *Linguistics and Philosophy* 18
- 554 Greco, John 2009, ‘Knowledge and Success from Ability’, in *Philosophical Studies* 142
- 555 ———2010, *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge
- 556 University Press)
- 557 Grimm, Stephen 2011, ‘What is interesting?’ in *Logos and Episteme* 2
- 558 Jeroen Groenendijk and Martin Stokhof 1984, *Studies on the Semantics of Questions and the Pragmatics of*
- 559 *Answers*, (Amsterdam: University of Amsterdam)
- 560 Hamblin, Charles L. 1976, ‘Questions in Montague English’, in B.H. Partee (ed.), *Montague Grammar*
- 561 (Exeter: Imprint)
- 562 Harrah, David 1961, ‘A logic of questions and answers’, in *Philosophy of Science* 28
- 563 Haziza, Eliran 2023, ‘Questioning and addressee knowledge’, in *Synthese* 201
- 564 Hintikka, Jaakko 1976, *The Semantics of Questions and the Questions of Semantics* (Ithaca: Cornell University Press)
- 565 Hurka, Thomas 1993, *Perfectionism* (New York: Oxford University Press)
- 566 Isaacs, James and Kyle Rawlins 2008, ‘Conditional Questions’, in *Journal of Semantics* 25
- 567 Karttunen, Lauri 1977, ‘Syntax and Semantics of Questions’, in *Linguistics and Philosophy* 1
- 568 Kelp, Christoph 2013, ‘Knowledge: The Safe-Apt View’, in *Australasian Journal of Philosophy* 91
- 569 ———2018, *Good Thinking: A Knowledge First Virtue Epistemology* (London: Routledge)
- 570 ———2021a, ‘Theory of Inquiry’, in *Philosophy and Phenomenological Research* 103

- 571 ———2021b, *Inquiry, Knowledge, and Understanding* (Oxford: Oxford University Press)
- 572 Kvanvig, Jonathan L. 2003, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge
- 573 University Press).
- 574 Leonard, Henry 1959, 'Interrogatives, Imperatives, Truth, Falsity, and Lies', in *Philosophy of Science* 26
- 575 Locke, John 1706/1996, 'Of the Conduct of the Understanding', in R.W. Grant & N. Tarcov (eds.) *Some*
- 576 *Thoughts Concerning Understanding and Of the Conduct of the Understanding* (Indianapolis: Hackett
- 577 Publishing Company)
- 578 McGrath, Matthew 2021, 'Being Neutral: Agnosticism, Inquiry, and the Suspension of Judgment', in *Noûs*
- 579 55
- 580 McHugh, Conor 2012, 'The Truth Norm of Belief', in *Pacific Philosophical Quarterly* 93
- 581 Millson, Jared A. 2021a, 'A Defeasible Calculus for Zetetic Agents', in *Logic and Logical Philosophy* 30
- 582 ———2021b, 'Seeking Confirmation: A Puzzle for Norms of Inquiry', in *Analysis* 80
- 583 Miracchi, Lisa 2015, 'Competence to Know', in *Philosophical Studies* 172
- 584 Miracchi Titus, Lisa and J. Adam Carter 2023, 'What the Tortoise Should Do: A Knowledge-First Virtue
- 585 Approach to the Basing Relation', in *Noûs* 58
- 586 Moyer, Morgan and Kristen Syrett 2019, 'The Semantics of Questions', in *WIREs Cognitive Science* e1513
- 587 Olsson, Erik J. 2007, 'Reliabilism, Stability, and the Value of Knowledge', in *American Philosophical Quarterly*
- 588 44
- 589 Palmira, Michele 2020, 'Inquiry and the Doxastic Attitudes', in *Synthese* 197
- 590 ———ms, 'The Puzzle of Defective and Permissible Inquiry'
- 591 Pritchard, Duncan 2009a, 'Knowledge, Understanding and Epistemic Value', in *Royal Institute of Philosophy*
- 592 *Supplement* 64
- 593 ———2009b, 'The Value of Knowledge' in *The Harvard Review of Philosophy* 16
- 594 Pritchard, Duncan, John Turri, and J. Adam Carter 2022, 'The Value of Knowledge', in *Stanford*
- 595 *Encyclopedia of Philosophy* <[https://plato.stanford.edu/archives/fall2022/entries/knowledge-](https://plato.stanford.edu/archives/fall2022/entries/knowledge-value/)
- 596 [value/](https://plato.stanford.edu/archives/fall2022/entries/knowledge-value/)>
- 597 Riggs, Wayne D 2002, 'Reliability and the Value of Knowledge', in *Philosophy and Phenomenological Research*
- 598 64
- 599 Roberts, Craige 1996, 'Information Structure in Discourse', in *Semantics and Pragmatics* 5
- 600 Roberts, Robert C. and Wood, W. Jay 2007, *Intellectual Virtues: An Essay in Regulative Epistemology* (Oxford:
- 601 Clarendon Press)
- 602 Rosa, Luis 2023, 'Suspending Judgment the Correct Way', in *Inquiry* 66
- 603 ———2025, *Rules for the Inquiring Mind: A Unified Framework of Norms of Inquiry* (New York:
- 604 Routledge)
- 605 Sapir, Yasha, and Peter van Elswyk 2021, 'Hedging and the Ignorance Norm on Inquiry', *Synthese* 199.
- 606 Shah, Nishi 2003, 'How Truth Governs Belief', in *The Philosophical Review* 112
- 607 Simion, Mona 2021, 'Blame as Performance', in *Synthese* 199
- 608 Simion, Mona, Christoph Kelp, and Harmen Ghijsen 2016, 'Norms of Belief', in *Philosophical Issues* 26
- 609 Sosa, Ernest 1980, 'The Raft and the Pyramid: Coherence versus Foundations in the Theory of
- 610 Knowledge', in *Midwest Studies in Philosophy* 5
- 611 ———1991, *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge: Cambridge University Press)
- 612 ———2000, 'For the Love of Truth?', in L. Zagzebski and A. Fairweather (eds.), *Virtue Epistemology: Essays*
- 613 *on Epistemic Virtue and Responsibility* (Oxford: Oxford University Press)
- 614 ———2007, *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I* (Oxford: Oxford University
- 615 Press)
- 616 ———2015, *Judgment and Agency* (Oxford: Oxford University Press)
- 617 ———2017, *Epistemology* (Princeton: Princeton University Press)
- 618 ———2021, *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains* (Oxford: Oxford
- 619 University Press)
- 620 van Fraassen, Bas C. 1980, *The Scientific Image* (Oxford: Oxford University Press)
- 621 Whitcomb, Dennis 2007, *Intellectual Goods: An Epistemic Value Theory*, dissertation (New Brunswick:
- 622 Rutgers University).
- 623 ———2010, 'Curiosity Was Framed', in *Philosophy and Phenomenological Research* 81
- 624 ———2017, 'One Kind of Asking', in *Philosophical Quarterly* 67
- 625 Whitcomb, Dennis & Jared A. Millson 2024, 'Inquiring Attitudes and Erotetic Logic: Norms of
- 626 Restriction and Expansion', in *Journal of the American Philosophical Association* 10

- 627 Willard-Kyle, Christopher 2021, 'Surprising suspensions: the epistemic value of being ignorant',
628 dissertation (New Brunswick: Rutgers University)
629 ———2023a, 'Valuable Ignorance: Delayed Epistemic Gratification', in *Philosophical Studies* 180
630 ———2023b, 'The Knowledge Norm for Inquiry', in *Journal of Philosophy* 120
631 ———2024, 'Ignorance, Soundness, and Norms of Inquiry', in *Philosophical Studies* 181
632 Willard-Kyle, Christopher, Jared Millson & Dennis Whitcomb forthcoming, 'Evoked Questions and
633 Inquiring Attitudes', in *Philosophical Quarterly*
634 Wiśniewski, Andrzej 1991, 'Erotetic Arguments: A Preliminary Analysis.' in *Studia Logica* 50
635 ———1995, *The Posing of Questions: Logical Foundations of Erotetic Inferences* (Dordrecht: Springer)
636 Woodard, Elise 2022, 'The Ignorance Norm and Paradoxical Assertions', in *Philosophical Topics* 49
637 Zagzebski, Linda T. 1996, *Virtues of the Mind* (Cambridge: Cambridge University Press)