ORIGINAL RESEARCH

Greco's explanatory salience contextualism revisited

Sven Bernecker[1,2]

Abstract
According to Greco's early explanatory salience contextualism, *S* knows that *p* if and only if *S*'s cognitive abilities are the salient factor in a causal explanation of why *S* holds a true belief rather than a false belief or no belief at all. Greco abandoned this view because it cannot handle fake barn cases and because it proves impossible to analyze knowledge in terms of a quantitative characterization of explanatory salience. The paper argues that the core idea of explanatory salience contextualism could have been saved if only the causal notion of explanation had been replaced by an abductive notion.

Keywords John Greco · Virtue reliabilism · Explanationism · Gettier · Knowledge · Abduction

Virtue reliabilism holds that a person knows that *p* when they hold a true belief because they have exercised an intellectual ability, skill, competence, disposition, or virtue. Knowledge is epistemic success that is due to intellectual ability or simply: knowledge is success from ability. A pioneer of this approach, Ernest Sosa, formulates the core insight as follows: "knowledge is true belief... that turns out right by reason of the virtues and not just by coincidence" (Sosa, 1991, p. 277). In Greco's words "*S* knows that *p* if and only if *S* believes the truth with respect to *p* because S's belief that *p* is produced by intellectual ability" (Greco, 2010, pp. 12, 71; see also Greco, 2019, p. 125).

The virtue reliabilist account of knowledge has three parts—the success, the ability, and the connection between them. The epistemic success consists in the formation of a true belief as opposed to false belief or no belief at all. An intellectual ability is a stable disposition to form a true belief under certain external and internal conditions.

✉ Sven Bernecker
s.bernecker@uci.edu

1 Department of Philosophy, University of Cologne, Albertus-Magnus-Platz, 50923 Cologne, Germany

2 Department of Philosophy, University of California, Irvine, Irvine, CA 92697-4555, USA

If knowledge is tied to intellectual abilities, true beliefs resulting from 'strange and fleeting cognitive processes' do not count as knowledge (Greco, 2010, ch. 9). The third element of the virtue reliabilist account of knowledge is the connection between the ability and the success. The epistemic success must be due to, or produced by, the intellectual ability. Knowledge is said to be cognitive success that exists *because* of the exercise of a person's cognitive ability.[1] Greco understands the 'because'-clause as referring to a *causal-explanatory* relation. "[I]n cases of knowledge, the fact that *S* has a true belief is explained by the fact that *S* believes from ability" (Greco, 2010, p. 71). Greco (2008, p. 421) dubs this position *explanatory salience contextualism*, or *ESC* for short.

When someone believes the truth because of their intellectual abilities, they have accomplished something for which they deserve credit. It is the hallmark of knowledge that we can give the knower credit for their epistemic success. This is why Greco (2019, p. 125) treats the statement "*S* has a true belief *because* she believes from intellectual virtue" as interchangeable with the statement "*S*'s believing the truth is *attributable* to her virtuous intellectual agency".

ESC was short-lived. In 2012, Greco convinced himself of its untenability and developed a new form of virtue epistemology that does without explanatory salience considerations. The point of this paper is to show that the problems identified with ESC do not speak against the basic idea (viz., that knowledge is epistemic success explainable in terms of intellectual abilities) but only against the causal conception of explanation. If only the causal notion of explanation is replaced with an abductivist notion, the problems with ESC evaporate. Greco could have saved the core insight of ESC by changing the conception of explanation.

Section 1 explains ESC by pointing out how it eliminates classical Gettier cases. Section 2 compares ESC with contextualist and explanationist approaches to knowledge. Sections 3 and 4 analyze the objections brought against ESC and outlines Greco's newest form of virtue epistemology. Section 5 spells out a different way of combining virtue reliabilism and explanationism. The proposed view is immune to the problems that beset ESC.

1 Causal impact and abnormality

ESC has it that one knows only if one's intellectual abilities are the salient factor in a causal explanation of why one holds a true belief, rather than a false belief or no belief at all. This view is said to provide solutions to the lottery problem, the value problem, and the Gettier problem. Gettierized beliefs do not qualify as knowledge because the epistemic success is not explained by the subject's intellectual abilities.

---

[1] Sometimes, Sosa also uses a 'because'-clause to formulate his position. He writes: "Beliefs fall under the AAA structure, as do performances generally. We can distinguish between a belief's accuracy, i.e., its truth: its adroitness, i.e., its manifesting epistemic virtue or competence; and its aptness, i.e., its being true *because* competent" (Sosa 2007, p. 23). The 'because'-clause is also present in Riggs (2002) and Zagzebski (1996, p. 297). Kelp (2009, p. 583) characterizes 'robust virtue epistemology' as follows: "One knows that p [if and only if] the truth of one's belief that p—alternatively, one's cognitive success—is *because* of the exercise of cognitive ability." For a discussion of the 'because'—relation in virtue reliabilism see Fairweather and Montemayor (2017), Littlejohn (2014) and Nimtz (2012).

Consider a classical Gettier case such as Lehrer's (1965) Ford case. Smith does not know that someone in the office owns a Ford because even though his intellectual abilities are responsible for him deriving the true existential generalization from the false premise that Nogot owns a Ford, they do not contribute to the belief in the existential generalization being true. Yet the salient causal factor of the true belief does not have to do with Smith's abilities but with the lucky coincidence that there happens to be another employee in the company who owns a Ford. Smith does not get things right *because* of his intellectual abilities. The conclusion-belief is true and competently formed but it "is not true *because* competent."[2] The intellectual abilities *cause* the conclusion-belief, but they do not *explain* its truth.

When we say that *y* occurs *because x* occurs, we mark out *x*'s occurring as a particularly important or salient part of the causal story behind *x*'s occurring (Greco, 2003, p. 118, 2010, p. 74, 2012, pp. 5–6). Greco acknowledges that we do not have a complete understanding of the rules governing explanatory salience. We know, however, that salience is determined, among other things, by our interest and purposes and by what is abnormal and unusual. For instance, when picking out the cause of smoke coming from a car engine, we pick out the malfunctioning carburetor because we are interested in the part that needs to be replaced (Greco, 2003, p. 119, 2012, p. 9). Similarly, sparks may be said to be the cause of a fire if the presence of sparks in the area is not normal. Sparks would not count as the cause if we were trying to explain the cause of a fire in a welding shop, where sparks are flying all the time (Greco, 2003, p. 118, 2012, p. 9). Picking out the sparks as the cause of a fire require the presence of sparks to be statistically abnormal in the area under the given circumstances.

In Gettier cases, what counts against the intellectual abilities explaining the belief's truth is the unlikelihood of the subject ending up with a true belief given the evidence and given the circumstances (Greco, 2003, p. 131). In the Ford case, for example, the abnormality consists in Smith's evidence being misleading although someone else in the office owns a Ford. Likewise, in Chisholm's (1989, p. 93) sheep in the field case, the abnormality consists in the 'big surprise' that there is a sheep elsewhere in the field although the subject mistakes a dog for a sheep (Greco, 2003, p. 131). The decisive feature of Gettier cases is that "this default salience is trumped by something abnormal in the case" (Greco, 2002, p. 310, 2010, p. 75, 2019, p. 133). Our intuitions about whether someone knows are said to covary with intuitions about whether their abilities are explanatorily salient:

> That is, in cases where it seems that *S* knows, it seems that it is the case that *S*'s cognitive abilities are important in an explanation why *S* believes the truth. And in cases where it seems that *S* does not know, it seems that *S*'s abilities are less important in such an explanation (Greco, 2010, p. 84).

---

[2] Greco (2012, p. 2). Sosa (2007, pp. 96–97) gives the same diagnosis. "The reasoning by way of Nogot does of course help explain why the believer has that belief, but it does not in the slightest help explain its correctness. … The error resides … in supposing that what explains there being my true belief explains why my belief is true".

2 Between contextualism and explanationism

ESC combines elements of contextualism and explanationism. *Epistemic contextualism* is the view that the truth conditions of knowledge ascribing and knowledge denying statements (sentences of the form '*S* knows that *p*' and '*S* does not know that *p*') vary according to the conversational context in which they are uttered. A knowledge attribution might require one epistemic position when made in a 'low standards' context, but a different, stronger epistemic position when made in a 'high standards' context.[3]

Typically, the semantic models underlying epistemic contextualism are ones of indexicals, demonstratives, or gradable adjectives. Greco proposes a different model for the semantic context sensitivity of knowledge attributions. Knowledge attributions are said to involve a species of causal explanation, and causal explanation is said to be context sensitive: "knowledge attributions inherit the context-sensitivity of causal explanations."[4] Changes in the conversational context bring about changes in explanatory salience, which, in turn, affect knowledge attributions.

ESC borrows not only from contextualism but also from explanationism. *Explanationism about knowledge* grew out of Alvin Goldman's (1967) causal theory which maintains that *S* knows that *p* if the fact that *p* is causally connected in the right way to *S*'s believing that *p*. The causal analysis of knowledge is both too weak and too strong. It is too strong because it does not seem to allow for knowledge of abstract objects, given that abstract objects cannot enter into causal relations with beliefs. Causalism about knowledge is too weak because it attributes knowledge in fake barn cases (more on this in Sect. 4). Explanationism promises to overcome the problems of the causal theory. It maintains that the knowledge constituting connection between a belief and the fact that makes it explanatory, not causal. To know that *p*, the fact that *p* must explain the fact that one's belief that *p* is true.[5]

The crucial difference between ESC and explanationism concerns the explanans. Explanationism has it that a subject's truly believing that *p* must be explained by the fact that *p*. According to ESC, however, the explanatory relation links the true belief with the subject's intellectual abilities. A subject is said to know if it is their intellectual abilities that explain why they hold a true belief as opposed to a false belief or no belief. The explanantia are the intellectual abilities rather than the facts that make the belief true.

3 Practical needs and safety

In his 2012 paper "A (Different) Virtue Epistemology," Greco (2012, pp. 9–10) maintains that ESC is 'pretty good' but 'not good enough.' He cites two objections as

---

[3] Proponents of contextualism are Blome-Tillmann (2009), Cohen (1998), DeRose (2009) and Lewis (1996).

[4] Greco (2010, p. 106). For a similar view see Rieber (1998).

[5] Among the proponents of explanationism about knowledge are Harman (1973), Goldman (1984), Rieber (1998), Jenkins (2006) and Bogardus and Perrin (2022). Faraci (2019) and Goldberg (2015) develop explanationist accounts of epistemic coincidence.

having convinced him of the need to develop a different kind of virtue epistemology. The first objection states that explanatory salience virtue epistemology has problems making sense of knowledge acquired from testimony (cf. Lackey, 2007, 2009). Testimonial believers are typically not appropriately creditable for the truth of their beliefs to qualify as knowing. The second objection states that the position whereby knowledge is a kind of success from ability cannot account for knowledge based on extended cognition (cf. Pritchard, 2010a; Vaesen, 2011). When we rely on external instruments (e.g., thermometers, metal detectors, speedometers) in coming to know something, the most salient feature in the causal explanation of our intellectual success are not our intellectual abilities but these instruments.

These objections give rise to a dilemma. According to ESC, for someone to know something either (a) their abilities must be *the most important part* of the explanation for why they have a true belief (Greco, 2003) or (b) their abilities must be *an important part* of the explanation for why they have a true belief (Greco, 2010). If (a), then Greco's account makes it impossible to know from testimony or through extended cognition; and if (b), then the account fails to rule out knowledge in Gettier cases, for in many such cases the subject's abilities *are* an important part of the explanation for their true belief (Greco, 2012, p. 10, 2019, p. 133, 2020b, p. 95–96). The crux with ESC is that it is not possible to specify with any precision the degree of causal contribution of a subject's intellectual abilities vis-à-vis other causal factors involved in the formation of a true belief. This is why Greco abandons the project of characterizing explanatory salience in *quantitative* terms and instead opts for a *qualitative* characterization.

A subject's intellectual abilities are said to explain the belief's truth just in case the abilities contribute to the belief's truth "in a way that would regularly serve relevant purposes" (Greco, 2012, p. 14, 2019, p. 139). What counts as a 'relevant purpose' in the context of knowledge attributions? Greco follows Edward Craig (1990) in claiming that the function of the concept of knowledge is to flag good informants. Craig's idea is combined with a suggestion made by Jason Stanley (2005) and others that knowledge is the norm of action and of practical reasoning. In cases of knowledge, Greco holds, the subject's abilities must explain the belief's truth in a way that would regularly serve informational needs associated with some relevant domain of action and practical reasoning. In other words,

> *S* knows that *p* [if and only if] *S*'s believing that *p* is produced by an exercise of intellectual ability [of the sort that would regularly serve relevant informational needs, both local (actual) and global (typically and/or likely)], and *S*'s belief being so produced contributes [in a way that would regularly serve relevant informational needs, both local and global] to *S*'s having a true belief (Greco, 2012, p. 19; see also Greco, 2019, p. 139).

On Greco's newest form of virtue epistemology, gettierized beliefs fail to qualify as knowledge because the subject's intellectual abilities do not contribute to the belief's truth "in a way that can be regularly exploited, not in a way that is *dependable* or *reliable*" (Greco, 2012, p. 17; my emphasis; cf. Greco, 2019, p. 140, 2020b, p. 97). This marks an important shift from ESC. Whereas before Greco demanded that *S*'s intellectual abilities be a salient part of the causal explanation of her belief's truth, he now demands that *S*'s abilities reliably produce true belief.

The qualitative analysis of explanatory salience resembles the safety account. On the safety view, a belief-forming method that generates a true belief is knowledge-grade only if it could not easily issue false belief.[6] Similarly, the qualitative analysis of explanatory salience claims that intellectual abilities are knowledge-grade if they are 'dependable or reliable.' 'Reliability' can be replaced here with 'safety,' for Greco (2020a, p. 126) writes: "In cases of knowledge, *S* would not easily go wrong when *S* believes from ability (of the right sort)."

4 Practical environment and abilities

Prior to the paper "A (Different) Virtue Epistemology," Greco used ESC to handle classical Gettier cases, but he required a different strategy to deal with failed threat scenarios such as Alvin Goldman's (1976, p. 772–773) fake barn case. The reason is that ESC cannot eliminate fake barn cases. If I truly believe that there is barn in fake barn county, my perceptual abilities *are* causally salient in the formation of my true belief. Once Greco abandoned the quantitative conception of explanatory salience, he was able to unify his strategy to deal with classical Gettier cases and fake barn cases.

The key idea is that the reliability of the intellectual abilities that are responsible for a true belief must be assessed relative to the environment in which the abilities are operative. In fake barn county, knowing that there is a barn requires the true belief being the result of barn-identifying abilities that are appropriate *in fake barn county*. In a normal environment, barn-identifying abilities are appropriate if they allow the subject to distinguish barns from other objects such as cows, trees, and fences. In fake barn county, however, barn-identifying abilities are appropriate if they allow the subject to distinguish real barns from fakes ones. Our intellectual abilities yield barn-knowledge in regular environments, but not in fake barn county. When knowledge is identified not with success through intellectual ability *simpliciter* but with success through intellectual-ability-relative-to-the-environment, we have a way of denying knowledge in fake barn scenarios.[7]

Greco (2010, pp. 79–80) claims that, when determining whether someone has the required intellectual ability to know, the level of generality of an environment should be fixed according to subject's *practical* environment. The practical environment is determined by the subject's interests and purposes. "Knowledge requires belief from ability that would regularly serve *both* the actual *and* the typical or likely informational needs attaching to the relevant practical environment" (Greco, 2012, p. 23). Suppose you are at a working farm where there are several real barns and no barn facades. You are asked to retrieve a shovel from the barn located just ahead. There are barn facades in the vicinity, but you will never encounter them (Greco, 2010, p. 80, 2012, p. 23, 2017, p. 199). Greco maintains that it is appropriate to say that you *know* that there

---

[6] *S*'s belief that *p* is safe if and only if in all nearby worlds where *S* believes that *p* based on the same method as in the actual world, *p* is true. See Pritchard (2005, pp. 156, 168) and Sosa (2002, pp. 267–269, 2007, p. 26).

[7] See Greco (2012, p. 17, 2017, 2019, pp. 140–141, 2020a, p. 135). For similar views see Littlejohn (2014) and Millar (2010).

is a barn, for your "ability to distinguish barns from non-barns is perfectly reliable relative to the practical environment" (Greco, 2017, p. 199).[8]

Greco's newest form of virtue epistemology offers a unified response to classical Gettier cases and fake barn cases. A true belief fails to qualify as knowledge if, given the subject's intellectual abilities and practical environment, the belief does not covary with the truth through a sphere of close possible worlds.

5 Explanatory salience contextualism revisited

After having presented ESC and discussed the problems that led to its abandonment, I will argue that the core insight of ESC could have been saved if only Greco had substituted an abductivist notion of explanation for the causal notion of explanation.

Abduction (or 'inference to the best explanation' as it is also called) claims that the justificatory status of a hypothesis is determined by how well it explains some fact. A hypothesis is justified in virtue of its ability to explain some fact better than some rival hypotheses. The quality of an explanation is measured by the degree to which it manifests virtues of inference. Examples of such virtues are precision, scope, simplicity, and fit with background belief.[9]

Abduction is a ubiquitous type of inference. Suppose you try to understand why the dinosaurs become extinct. Did they become extinct because there were ill-adapted in comparison to their competitors (hypothesis *A*) or because they were destroyed as a result of meteorological effects caused by a meteorite impact (hypothesis *B*)? What speaks in favor of hypothesis *B* is that the cloud of dust caused by a large meteorite impact would have rapidly led to the extinction of the dinosaurs and that meteorite impacts did occur at the appropriate time in history. We know about the meteorite impacts because there are unusually high iridium deposits around the time of the extinction and because iridium is found in meteorites but not in the Earth's crust. The meteorite-hypothesis *B* is superior to the Darwinism-hypothesis *A* because it not only explains the extinction of the dinosaurs but also the iridium anomalies (Bird, 2000, p. 86).

How is abduction related to causal explanation? According to abduction, a hypothesis is justified if it meets the virtues of inference relative to some fact that is in need of explanation. Whether, in addition, the hypothesis meets the constraints of the causal or any other model of explanation is irrelevant with respect to its justificatory work. The justificatory work solely depends on the realization of virtues of inference. As

---

[8] The thesis that the reliability of intellectual abilities is environment-relative has met with criticism. Some claim our regular perceptual abilities *are* reliable in fake barn county because they allow us to form mostly true perceptual beliefs about farmers, horses, pigs, trees etc. The only sense in which perception is unreliable in fake barn county is with respect to distinguishing real barns from barn facades. Yet individuating intellectual faculties this narrowly raises the generality problem (Lackey, 2007, p. 355n17; for discussion see Gardiner, 2018, p. 306). Others object to the idea that one's perceptual abilities become unreliable as soon as one enters fake barn county. They claim that intellectual abilities should be assessed vis-à-vis the subject's typical or normal environment as opposed to their actual environment (Kallestrup & Pritchard, 2011, p. 343, 2012; Pritchard, 2010b, pp. 37–38).

[9] Rival theories of inference are Bayesianism, hypothetico-deductivism, and falsificationism.

long as a hypothesis does well with respect to the virtues of inference relative to a body of evidence, we are justified in believing it (Cabrera, 2020, p. 743).

When the gist of ESC is cast in abductivist terms, we get a theory of knowledge that reads as follows: *S knows that p if and only if, given the circumstances, S's forming a true belief, rather than a false belief or no belief at all is better explained by S's intellectual abilities than by coincidental factors.* Let us refer to this version of explanatory salience contextualism as *abductive salience contextualism*, or *ASC* for short.

The first thing to note is that it is inference to the *better*, not the *best*, explanation that is assumed by ASC. Explanation is contrastive. The intellectual abilities do not have to be the best explanation for the subject's epistemic success; they only have to do a better job of explaining the belief's truth than certain coincidental factors. Examples of such coincidental factors are "someone else in the company happens to own a Ford", "the subject happens to turn my head when the only real barn is in view", or "there happens to be a sheep hidden by the dog that is mistaken for a sheep". It is important to note that the notion of coincidence assumed by ASC is explanatory rather than metaphysical. When *S* does *not* know, the factors that outperform the intellectual abilities as explanantia of *S*'s truly believing are coincidental in the sense that they are unexpected, surprising or anomalous. Appealing to these coincidental factors is necessary because an explanation of the belief's truth in terms of the subject's intellectual abilities alone is gappy or incomplete.

What is the advantage of casting ESC in abductivist terms? We saw that there are, broadly speaking, two kinds of problems with ESC. One problem has to do with it not being possible to specify the degree of causal contribution of a subject's intellectual abilities vis-à-vis other causal factors involved in the formation of a true belief. This problem does not arise in the context of ASC. The reason is that, according to ASC, the intellectual abilities need not be operative in the production of belief. As long as the belief's truth is best explained in virtue of the subject's intellectual abilities as opposed to some coincidental factor, the true belief counts as knowledge.

It might seem that the problem of specifying the degree of contribution to the explanation also applies to ACS.[10] For just as the *causal* contribution that intellectual abilities play in belief formation allow for degrees, so does the *justificatory* contribution that hypotheses play vis-à-vis some fact in need of explanation. Both causal impact and inferential justification can be stronger or weaker. This raises the question of when the justification generated by an explanation of a subject's epistemic success in terms of their intellectual abilities is strong enough for the true belief to qualify as knowledge.

Since ASC assumes a contrastive notion of explanation, there are no absolute standards the explanation-based justification needs meet to be knowledge-grade. Instead, we compare the explanation of the subject's epistemic success in terms of their intellectual abilities with an explanation in terms of coincidental factors and determine which meets more of the virtues of inference. Knowledge attributions are made on the basis of this comparative assessment.

The other main problem with ESC is that it is committed to attributing knowledge in situations where the subject perceptually detects the truth but where they cannot

---

[10] I am grateful to an anonymous referee for pressing me on this point.

distinguish the fact that causes their belief from relevant alternatives. To see this, compare the bad barn scenario with a good barn scenario. In both cases, it is the subject's intellectual abilities that causally explain the truth of the belief that there is a barn. Yet there is an important difference. In *Bad Barn*, the belief is merely coincidentally true and hence does not qualify as knowledge. There does not seem to be a way to capture the knowledge-undermining coincidentality of the bad barn scenario within the framework of causal explanationism.[11]

How does ASC fair with respect to the fake barn scenario? Is the fact that the person driving through fake barn county (let us call him Henry) forms a true barn-belief better explained in terms of his perceptual abilities or in terms of the lucky coincidence that he forms the barn-belief right in the moment when the only real barn is in view? For ASC to yield the desired verdict that Henry does *not* know that there is a barn, the description of the circumstances has to mention that there are fake barns in the vicinity (bad luck), but not that Henry happens to form the belief that there is a barn while looking at the only real barn far and wide (good luck). If we know of the presence of fake barns and know that Henry's perceptual abilities are unable to discriminate between fake and genuine barns, it follows that his epistemic success is not properly explained by his intellectual abilities.

What is a principled reason for including the bad luck but not the good luck in the description of the circumstances in the explanandum? The justificatory import of abduction crucially depends on something being in need of explanation. For the abductive approach to get a foothold, the phenomenon in question must be in need of explanation.[12] Justification is generated only by necessary explanation, not by super-fluous one. In the fake barn case, the explanandum in question is *S*'s truly believing that there is a barn while driving through fake barn county. This fact is in need of explanation only if the description of the circumstances does not already give away why the belief is true—namely, that Henry happens to turn the head when the only real barn is in view. The description of the circumstances must be sufficiently parsimonious for there to be an explanatory gap. In other words, it is the logic of abduction that dictates the exclusion of the good luck from the description of the circumstances.

In sum, ASC has a leg up over ESC not only because it gives the correct verdict in bad barn cases but also because it gives classical Gettier cases and bad barn cases a unified treatment. In both cases it is the failure of the subject's intellectual abilities to properly explain why, given the circumstances, the subject forms a true belief, rather than a false belief or no belief at all. In sum, Greco could have saved the core insight of ESC by replacing the causal notion of explanation with an abductive notion.[13]

---

[11] Bogardus and Perrin (2022, p. 193n24), Goldberg (2015, pp. 288–289) and Rieber (1998, p. 201) are proponents of causal explanationism who attribute knowledge in fake barn cases. Some proponents of causal explanationism employ the no-essential-false-assumption clause (Harman, 1980, p. 176) or reliabilist ideas (Goldman, 1988, p. 44) to be able to deny knowledge in fake barn scenarios.

[12] The need of explanation can be motivated by practical or epistemic considerations. Here we focus on epistemic incentives for explanation. According to Wong and Yudell (2015), a phenomenon stands in need of epistemic explanation to the extent that it does not fit our theory (or map) of the world.

Declarations

Conflict of interest Sven Bernecker has no conflict of interest to disclose.

References

Bird, A. (2000). *Philosophy of science*. University College London Press.

Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind, 118*, 241–294.

Bogardus, T., & Perrin, W. (2022). Knowledge is believing something because it's true. *Episteme, 19*, 178–196.

Cabrera, F. (2020). Does IBE require a 'model' of explanation? *British Journal of Philosophy of Science, 71*, 727–750.

Chisholm, R. (1989). *Theory of knowledge* (3rd ed.). Prentice-Hall.

Cohen, S. (1998). Contextualist solutions to epistemological problems: Skepticism, gettier and the lottery. *Australasian Journal of Philosophy, 76*, 289–306.

Craig, E. (1990). *Knowledge and the state of nature*. Oxford University Press.

DeRose, K. (2009). *The case for contextualism*. Oxford University Press.

Fairweather, A., & Montemayor, C. (2017). Epistemic dexterity: A Ramseyian account of agent-based knowledge. In A. Fairweather & O. Flanagan (Eds.), *Naturalizing epistemic virtue* (pp. 118–142). Cambridge University Press.

Faraci, D. (2019). Groundwork for an explanationist account of epistemic coincidence. *Philosophers' Imprint, 19*, 1–26.

Gardiner, G. (2018). Virtue epistemology and explanatory salience. In H. Battaly (Ed.), *Routledge handbook of virtue epistemology* (pp. 296–308). Routledge.

Goldberg, S. C. (2015). Epistemic entitlement and luck. *Philosophy and Phenomenological Research, 91*, 273–302.

Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy, 64*, 357–372.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy, 73*, 771–791.

Goldman, A. H. (1984). An explanatory analysis of knowledge. *American Philosophical Quarterly, 21*, 101–108.

Goldman, A. H. (1988). *Empirical knowledge*. University of California Press.

Greco, J. (2002). Virtues in epistemology. In P. Moser (Ed.), *Oxford handbook of epistemology* (pp. 287–315). Oxford University Press.

Greco, J. (2003). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue. Perspectives from ethics and epistemology* (pp. 111–134). Clarendon Press.

Greco, J. (2008). What's wrong with contextualism. *Philosophical Studies, 58*, 416–436.

Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.

Greco, J. (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research, 85*, 1–26.

Greco, J. (2017). Contextualism and gettier cases. In J. J. Ichikawa (Ed.), *Routledge handbook of epistemic contextualism* (pp. 190–201). Routledge.

Greco, J. (2019). The gettier problem and intellectual virtue. In S. Hetherington (Ed.), *The gettier problem* (pp. 108–124). Cambridge University Press.

Greco, J. (2020a). Knowledge-producing abilities. In C. Kelp & J. Greco (Eds.), *Virtue theoretic epistemology: News methods and approaches* (pp. 124–146). Cambridge University Press.

Greco, J. (2020b). *The transmission of knowledge*. Cambridge University Press.

Harman, G. (1973). *Thought*. Princeton University Press.

Harman, G. (1980). Reasoning and evidence one does not possess. *Midwest Studies in Philosophy, 5*, 163–182.

Jenkins, C. S. (2006). Knowledge and explanation. *Canadian Journal of Philosophy, 36*, 137–164.

Kallestrup, J., & Pritchard, D. (2011). Virtue epistemology and epistemic twin earth. *European Journal of Philosophy, 22*, 335–357.

Kallestrup, J., & Pritchard, D. (2012). Robust virtue epistemology and epistemic anti-Individualism. *Pacific Philosophical Quarterly, 93*, 84–103.

Kelp, C. (2009). Pritchard on virtue epistemology. *International Journal of Philosophical Studies, 17*, 583–587.

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese, 158*, 345–361.

Lackey, J. (2009). Knowledge and credit. *Philosophical Studies, 142*, 27–42.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy, 74*, 549–567.

Littlejohn, C. (2014). Fake barns and false dilemmas. *Episteme, 11*, 369–389.

Millar, A. (2010). Knowledge and recognition. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The nature and value of knowledge: Three investigations* (pp. 89–188). Oxford University Press.

Nimtz, C. (2012). Knowledge, abilities, and 'because' clauses: A critical appraisal of virtue-theoretic analyses of knowledge. In T. Henning & D. Schweikard (Eds.), *Knowledge, virtue, and action: Putting epistemic virtues to work* (pp. 188–208). Routledge.

Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.

Pritchard, D. (2010a). Cognitive ability and the extended cognition thesis. *Synthese, 175*, 133–151.

Pritchard, D. (2010b). Knowledge and understanding. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The nature and value of knowledge: Three investigations* (pp. 1–88). Oxford University Press.

Rieber, S. (1998). Skepticism and contrastive explanation. *Nous, 32*, 189–204.

Riggs, W. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research, 64*, 79–96.

Sosa, E. (1991). *Knowledge in perspective: Selected essay in epistemology*. Cambridge University Press.

Sosa, E. (2002). Tracking, competence, and knowledge. In P. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 264–286). Oxford University Press.

Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford University Press.

Stanley, J. (2005). *Knowledge and practical interest*. Oxford University Press.

Vaesen, K. (2011). Knowledge without credit, exhibit 4: Extended cognition. *Synthese, 181*, 515–529.

Wong, W.-H., & Yudell, Z. (2015). A normative account of the need for explanation. *Synthese, 192*, 2863–2885.

Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.