

FAKE BARNS AND OUR EPISTEMOLOGICAL THEORIZING

ÁNGEL GARCÍA RODRÍGUEZ
Universidad de Murcia
Departamento de Filosofía
agarcia@um.es

SUMMARY: Pure virtue epistemology faces the fake barn challenge. This paper explains how it can be met. Thus, it is argued that the thought experiment contains a hidden ambiguity concerning the visual ability typically ascribed to, or denied, fake barn subjects. Disambiguation shows fake barn subjects to have limited knowledge of the target proposition (e.g. that there is a barn ahead). This accords with a pure virtue-theoretic conception of knowledge that predicts and explains all the intuitions elicited by the thought experiment. As a result, the relationship between knowledge, luck and ability is illuminated, and our epistemological theorizing improved.

KEY WORDS: ability, achievement, intuitions, knowledge, luck

RESUMEN: La epistemología de virtudes pura se enfrenta al reto de los falsos graneros. Este artículo explica cómo superarlo. Así, se constata que el experimento mental es ambiguo respecto de la habilidad visual típicamente atribuida, o negada, a los sujetos del ejemplo. La desambiguación muestra que dichos sujetos tienen conocimiento limitado de la proposición relevante (*v.g.*, que hay un granero enfrente). Esto casa con la epistemología de virtudes pura, que predice y explica todas las intuiciones producidas por el ejemplo. Con ello se esclarece la relación entre las nociones de conocimiento, suerte y habilidad, y se mejora nuestra teorización epistemológica.

PALABRAS CLAVE: conocimiento, habilidad, intuiciones, logro, suerte

1. *Fake Barn Cases*

Contemporary virtue epistemology places the notions of cognitive achievement and ability at the centre of our theorizing about knowledge. For knowledge is considered to be a form of achievement, and the latter requires appropriate connection to ability. More formally,

VIRTUE EPISTEMOLOGY (VE): S knows that p only if S 's true belief that p is a cognitive achievement.

COGNITIVE ACHIEVEMENT: S 's true belief that p is a cognitive achievement iff it is appropriately connected to a given ability A , i.e. if S 's true belief that p manifests A . (For short: S 's belief that p is true through ability.)

VE imposes a necessary condition on knowledge. In addition, some defenders of VE endorse a sufficient condition, to the effect that one

knows that p if one's true belief that p is a cognitive achievement (see Sosa 2007, 2011; Greco 2010). The resulting view has been called a *pure* (or robust) version of VE.¹ More formally,

PURE VIRTUE EPISTEMOLOGY (PVE): S knows that p iff S has a true belief that p through ability.²

PVE has emerged as a strong contender in the contemporary debate as to the nature of knowledge, because *prima facie* it meets the challenges raised by traditional Gettier cases as well as other prominent puzzles, such as Norman the clairvoyant, Truetemp, blindsighters and chicken sexers. Thus, if the intuition is that these subjects lack knowledge (as often is the case), PVE explains it because the relevant beliefs are not true through ability.³ Let us concede that PVE succeeds here, and focus instead on the challenge posed by fake barn cases (Goldman 1976); i.e. the challenge to accommodate the intuition that knowledge is sensitive to the nature of the modal environment.

A cursory glance at the literature shows that opinion is divided, as some authors claim that PVE has the necessary resources to meet the fake barn challenge, and others deny it. This is due to the fact that fake barn cases elicit different, and apparently contradictory intuitions, about whether the true beliefs of fake barn subjects count both as knowledge and as cognitive achievements. In turn, these conflicting intuitions are underwritten by standing assumptions about the nature of (knowledge-undermining) luck and its relation to cognitive ability. Therefore, an assessment of the prospects of PVE to meet the fake barn challenge must proceed *via* a detailed tour and analysis of the intuitions elicited, as well as the underlying assumptions.

¹ By contrast, an *impure* version of VE only endorses the necessary condition on knowledge. Pritchard (2009, 2010a) coined the label *robust* VE, to oppose it to his own anti-luck VE. Kelp (2013) has distinguished between pure and impure VE, including Pritchard's anti-luck VE as an example of the latter.

² There is no agreement among proponents of PVE as to the exact nature of the link between ability and true belief. Sometimes, a "quantitative" reading is assumed, for it is required that true belief *sufficiently* depend on ability (Carter 2016). A "qualitative" suggestion is that true belief depend on ability *in the right way*, current proposals favouring either pragmatic (Greco 2010, 2012) or metaphysical (Sosa 2011) considerations. Turri (2011) has helpfully argued that a charitable development of the latter proposals should include the idea that true belief *manifests* ability. "Through" naturally conveys such an idea, and should be read that way here.

³ For a recent defence, see Sosa 2015. Miracchi (2015) disagrees about the prospects of PVE vis-à-vis traditional Gettier cases. But see fn. 19.

Against this backdrop, the paper has three aims: to clarify the fake barn puzzle; to show that PVE has the resources to meet the fake barn challenge; and to enhance our epistemological theorizing, by elucidating the relations between knowledge, luck and ability. The overall result is a new defence of PVE. Central to this is the unveiling of an ambiguity in the fake barn thought experiment, concerning the visual ability normally ascribed to, or denied, subjects. On this basis, section 6 clarifies the example with pure virtue-theoretic tools; section 7 shows that this allows PVE to predict and explain all the intuitions elicited by fake barn cases; and section 8 draws some general conclusions about our epistemological theorizing, in particular about the relations between knowledge, luck and ability. Prior to this, the methodology of the paper is set out (Section 2), and all the intuitions elicited by the thought experiment both charted (Section 3) and connected to the ambiguity already mentioned (Sections 4 and 5).

2. Methodology

In contemporary philosophy, intuitions about fake barn cases are generally regarded as tests for our theories of knowledge. The idea is that, independently of one's theorizing, fake barn cases elicit intuitions to the effect that the protagonist of the example knows, or does not know, a target proposition (e.g. that there is a barn ahead), which must then be predicted and explained by an adequate theory of knowledge. The fact that those intuitions are independent of one's theorizing means that, to the extent that a preferred theory of knowledge fails to predict and explain them, the theory gives counter-intuitive results, and must be discarded, or at least amended.

Let us assume this testing view of intuitions,⁴ and ask ourselves: which intuitions about fake barn cases must be accommodated by our theorizing? As already mentioned, fake barn cases elicit conflicting epistemic intuitions, some epistemologists finding it intuitive to say that fake barn subjects know the target proposition (Hetherington 1999; Lycan 2006; Turri 2011, 2012); others (perhaps most) finding it intuitive to say that they do not know it (Goldman 1976; Pritchard

⁴The empirical thesis that intuitions about fake barn cases, as generally used by epistemologists, play a testing role for our philosophical theorizing must be distinguished from the evaluative claim that it is right for epistemologists to award intuitions such a testing role. For the purposes of developing the argument of this paper, it is the former empirical thesis that must be assumed. For discussion, see Cappelen 2012.

2009, 2010a; Greco 2010, 2012; Millar 2010); and others still finding it intuitive to say both (Sosa 2011, 2015; Navarro 2015).⁵ Now, should our epistemological theorizing be tested against the intuition that fake barn subjects know, the intuition that they do not know, or both?⁶

In addition, fake barn cases elicit conflicting intuitions regarding whether the true beliefs of fake barn subjects count as cognitive achievements. Some epistemologists find it intuitive to say that they do (Pritchard 2009, 2010a); others find it intuitive to say that they do not (Greco 2010; Millar 2010; Jarvis 2013; Littlejohn 2014); and others still find both reactions intuitive (Sosa 2011, 2015; Navarro 2015). Given the centrality of the notion of achievement for VE, these intuitions are particularly relevant when considering the fake barn challenge to PVE. But, should all of them be accommodated by our epistemological theorizing?

If epistemologists stick to their own (epistemic and achievement) intuitions, and test theories of knowledge accordingly, unacceptably varying results will follow, a theory turning out to be satisfactory or not, depending on the intuitions used to test it. This is unacceptable, for surely the nature of knowledge is not relative to what individual epistemologists consider intuitive to say about particular examples. This result could be avoided if all epistemologists shared the same intuitions about fake barn cases, but this is clearly not so. Therefore, the question still stands —which intuitions from the array actually avowed by epistemologists should prevail as the test for our theorizing about knowledge?

One methodological approach is to stage a battle of the intuitions, to see which intuitions emerge on top. The rationale for this may be that, as not all intuitions can be expected to track the truth (for either fake barn subjects know the target proposition or not, and

⁵ Greco is often read as claiming that fake barn subjects lack knowledge, but he also says that the truth-value of a knowledge claim depends on the nature of the practical environment, and that there are possible practical environments where fake barn subjects have knowledge. So, there may be some reason for categorizing Greco as having mixed epistemic intuitions about fake barn cases.

⁶ Experimental work has shown that both knowledge and no-knowledge intuitions are also avowed by lay people. Thus, as part of an experiment into the effect of age on epistemic intuitions in fake barn cases, Colaço et al. (2014) collected the participants' responses in a seven-item scale ranging from 0 (absence of knowledge) to 6 (knowledge). As their data demonstrate, there are responses at all the points in the scale. Similarly, in an experiment testing the effect on knowledge attributions of a salient but failed threat to the truth of a belief obtained perceptually (as in fake barn cases), Turri et al. (2015) collected data showing that both knowledge and no-knowledge intuitions were reported.

either their true beliefs are achievements or not), only intuitions that track the truth should be awarded a testing role. Those that fail to track the truth are faulty intuitions, and there is no pressure to make room for them in our epistemological theorizing. Instead, they should be diagnosed or explained away.

This approach forecloses the possibility that fake barn cases contain an overlooked complexity to account for all the intuitions reported by epistemologists, concerning both knowledge and achievement. Therefore, so as not to foreclose this possibility, the following methodological recipe will be adhered to in this paper: if at all possible, our theorizing about knowledge should accommodate *all* normal intuitions elicited by fake barn cases.⁷ The recipe fits the testing view of intuitions, for a theory of knowledge that manages to accommodate all normal intuitions elicited by fake barn cases is in good shape. Such a theory will have a story to tell as to why different epistemologists report different intuitions, and will not have counter-intuitive results.

Adopting this methodological recipe will prove fruitful. For, as will be seen below, the fake barn thought experiment contains a hidden ambiguity regarding the typical ability ascribed to, or denied, fake barn subjects. The ambiguity will help clarify the example, and will also help show that all the intuitions reported by epistemologists can be predicted and explained by PVE. Therefore, the disclosure of the ambiguity will be crucial for the purposes of this paper. But before that, all the reported intuitions must be clearly in view.

3. *Intuitions and Options*

For ease of presentation, the intuitions elicited by fake barn cases can be organized in a chart with a knowledge axis and an achievement axis (see table on page 34). Four possible options are available, though only three have been endorsed in the VE literature.⁸

⁷ For “normal”, read the intuitions reported by honest people, who fully understand the example, are not blinded by theory, and so on. This clause will apply throughout our discussion.

⁸ In line with the methodology explained earlier, the table summarizes the data for our theorizing. Contemporary epistemologists have elaborated on those intuitions in ingenious ways, and some of those elaborations will be discussed below.

INTUITIONS	Achievement	No-Achievement
No-knowledge	OPTION 1 Pritchard	OPTION 2 Greco Millar Sosa Jarvis Littlejohn Navarro Broncano-Berrocal
Knowledge	OPTION 3 Sosa Turri Navarro	OPTION 4

Table. Intuitions and options about fake barn cases endorsed by contemporary epistemologists

OPTION 1 (ACHIEVEMENT WITHOUT KNOWLEDGE). It is intuitive to say (i) that subjects in fake barn cases do not know the target proposition, although (ii) their true belief (regarding the target proposition) counts as an achievement (Pritchard 2009, 2010a).

OPTION 2 (NEITHER ACHIEVEMENT NOR KNOWLEDGE). It is intuitive to say both (i) that subjects in fake barn cases do not know the target proposition, and (ii) that their true belief (regarding the target proposition) does not count as an achievement (Greco 2010; Millar 2010;⁹ Sosa 2011, 2015; Jarvis 2013; Littlejohn 2014; Navarro 2015; Broncano-Berrocal 2017).¹⁰

OPTION 3 (ACHIEVEMENT AND KNOWLEDGE). It is intuitive to say both (i) that subjects in fake barn cases know the target proposition, and (ii) that their true belief (regarding the target proposition) counts as an achievement (Sosa 2011, 2015; Turri 2011, 2012; Navarro 2015).¹¹

⁹ Millar does not use the label “virtue epistemology” to describe his work, but there are significant affinities, so he is included here.

¹⁰ Some authors claim that fake barn subjects lack the required ability (Greco 2010; Millar 2010); others that their true belief fails to manifest it (Jarvis 2013; Littlejohn 2014). Both alternatives will be accommodated in the discussion below. Broncano-Berrocal claims that the target belief is a partial achievement, whereas knowledge requires complete achievement.

¹¹ Sosa and Navarro are included in options 2 and 3. According to Sosa (2015), fake barn subjects possess animal knowledge through the exercise of the first-order

OPTION 4 (KNOWLEDGE WITHOUT ACHIEVEMENT). No one has found it intuitive to say both (i) that subjects in fake barn cases know the target proposition, and (ii) that their true belief (regarding the target proposition) does not count as an achievement. (This is unsurprising, for the analysis of the prospects of PVE vis-à-vis fake barn cases takes as common ground that knowledge is a form of achievement.)¹²

It is important to note that each of options 2 and 3 allows for positive verdicts on the prospects of PVE as a theory of knowledge. These are the details. PVE is the thesis that cognitive achievement amounts to knowledge. Now, according to the intuitions avowed in option 2, fake barn subjects have true beliefs regarding the target proposition that do not qualify as either knowledge or (complete) achievements, for either the subjects lack, or their beliefs fail to manifest, the pertinent ability. Therefore, fake barn cases fall squarely outside the definition provided by PVE, without contravening it. As for the intuitions avowed in option 3, fake barn subjects have true beliefs regarding the target proposition that count both as achievements and knowledge; therefore, fake barn cases exactly match the definition provided by PVE. Either way, the result is that PVE does not give counter-intuitive results when tested against fake barn cases, and is vindicated as a theory of knowledge.

It is equally important to note that these verdicts are incomplete, for they are based on some, not all, of the intuitions reported, and so are unlikely to settle the issue for good. For one thing, in line with the testing view of intuitions, the verdicts reached may be contested by those who report different intuitions about fake barn cases: namely, that the target true belief counts as a (complete) achievement (*contra* option 2), or that it does not count as knowledge (*contra* option 3). For another, the positive verdicts have been reached on the basis of apparently contradictory intuitions about fake barn cases, which in

ability to form true beliefs regarding the target proposition; but they lack knowledge full well, because they are not guided by the second-order ability to judge that the situation is right to form true beliefs through the first-order ability. According to Navarro (2015), there are coarse- and fine-grained descriptions of fake barn subjects' cognitive action and target belief. He concludes that their coarse-grained belief that there is a barn ahead is both an achievement and a case of knowledge, whereas their fine-grained belief that there is a barn ahead in fake barn country is neither.

¹²Some contemporary epistemologists, like Hetherington and Lycan, have defended the intuition that fake barn subjects know, whilst remaining silent on the issue of achievement. To preserve the focus of the paper, i.e. the prospects of PVE vis-à-vis fake barn cases, their views will not feature here.

turn stands in the way of our obtaining clarity about the puzzle: in other words, is PVE upheld because fake barn cases involve achievement, hence knowledge (option 3); or because they involve neither (option 2)?

To make headway, in line with the methodology set out earlier, it will now be shown that the apparently contradictory intuitions elicited by fake barn cases, regarding both knowledge and achievement, are in fact compatible. This is so because the target belief of fake barn subjects is based on a visual ability whose exact nature has been left ambiguous. Epistemologists have been unaware of this, but the ambiguity has been operating quietly behind the scenes, bringing about conflicting intuitions. Therefore, it needs to be unveiled, both to illuminate the discrepancies about the reported intuitions and to clarify the puzzle itself. The task is undertaken over the next few sections.

4. *Achievement and No-Achievement Intuitions*

Given the intricacies of post-Gettier epistemology, it is probably to be expected that fake barn cases should elicit conflicting epistemic intuitions from epistemologists. But, is it also to be expected that fake barn cases should elicit conflicting intuitions about whether fake barn subjects' performances count as achievements? Perhaps, given that the nature of achievements, and cognitive achievements in particular, is an issue for debate. Thus, it has been claimed that not every success through ability counts as an achievement, for the task might not be (sufficiently) challenging, or valuable. But regardless of how this dispute is settled (see Pritchard 2010a and Navarro 2015, for opposing views), fake barn cases make prominent a different dispute—namely, does the target belief of fake barn subjects *manifest* a visual ability, or not? As noted above, contemporary epistemologists disagree about this, but why?

Here is one suggestion: epistemologists disagree, because they have different visual abilities in mind. The suggestion may not seem very promising, because in the standard example considered so far—that is, the target true belief that there is a barn ahead—only one ability seems relevant for determining achievement—namely, the ability visually to distinguish barns from non-barns. But the crux of the suggestion being made now is that there is a hidden ambiguity in the category *non-barn*.

To see why, consider that one can visually distinguish barns from cows, tractors, or silos (to stay with Goldman's original example);

and one can visually distinguish barns from cows, tractors, silos *and barn façades*.¹³ Furthermore, subjects may possess, or their beliefs manifest, the former ability, but not the latter. So, a subject's true belief that there is a barn ahead may count as an achievement, because it is appropriately connected to the former ability. Yet, it may not count as an achievement, because it is not appropriately connected to the latter ability. This makes the following explanation plausible: epistemologists reporting a no-achievement intuition have the more inclusive ability in mind; whereas epistemologists reporting an achievement intuition have the less inclusive ability in mind.¹⁴

This distinction between visual abilities fits Goldman's original presentation of the thought experiment, and helps make sense of the achievement vs. no-achievement intuitions reported by epistemologists. But it may be suspected to be an *ad hoc* manoeuvre to save PVE from the fake barn challenge. To see that it is not, consider that a perceptual ability to discriminate Fs is the ability to tell Fs from a class of non-Fs, which in turn requires (explicit or implicit) individuation. Furthermore, different individuations of the class of non-Fs result in different abilities. The following examples will help.

VELÁZQUEZ: Diego has just completed a History of Art undergraduate degree, and seeks admission to a postgraduate course on the great Spanish painters of the past, co-organized by the Prado Museum and a nearby university. The number of applications exceeds the number of places available; so, an entry exam is set for applicants to show that they can visually tell Velázquez's works from those by other great masters, like Goya and Murillo. Diego has studied Velázquez's style and technique closely, comparing it with Goya's and Murillo's. On the established date he passes the exam with flying colours, and is offered a place on the course.

FAKE VELÁZQUEZ: Diego has just completed a History of Art undergraduate degree, has studied Velázquez's style and tech-

¹³ Both visual abilities must be read as containing a *ceteris paribus* clause, to the effect that the lighting conditions are good, the perceiver is fit (e.g. psychologically and neurologically), and so on.

¹⁴ It may be suggested that, as well as visual abilities, subjects need other abilities to acquire true beliefs on the basis of visual experience. This may well be so, depending on one's conception of the nature of visual experience and belief —namely, whether one thinks that seeing is believing (at least in default cases), or whether one takes believing to include some form of inference from one's visual experiences (also in default cases). Hence, readers may fill in all references to the visual beliefs of subjects as required by their preferred view of visual experience and belief.

nique, and is able on that basis visually to tell his works apart from those by other Spanish great masters. A vacancy arises at the Prado Museum for a curator specializing in Velázquez, and Diego applies. At the job interview, he is asked visually to distinguish Velázquez's paintings from paintings by minor artists from Velázquez's own workshop in Seville. Naturally, the style and technique of the latter was very similar to Velázquez's own, for they strived to imitate the master himself, managing it to a large extent. Unfortunately, Diego is unable to carry out the task, and is not offered the job.

Now, the reason why Diego is offered a place on the course in the first story, but is not offered the job in the second, is that while his performance manifests the visual ability required for admission to the course, it does not manifest the visual ability required for the job. The former performance is an achievement, because it is appropriately connected to the ability visually to distinguish Velázquez's paintings from non-Velázquez paintings, i.e. paintings by other great masters, like Goya or Murillo. However, the latter performance is not an achievement, because it is not appropriately connected to the ability visually to distinguish Velázquez's paintings from non-Velázquez paintings, including those by minor artists from Velázquez's own workshop. So, Diego's performance is assessed differently in each story (achievement vs. no-achievement), but there is no contradiction or conflict in doing so, as two different abilities are at stake. For on the one hand, there is the visual ability to distinguish Velázquez's paintings from Goya's and Murillo's; and on the other hand, there is the visual ability to distinguish Velázquez's paintings from Goya's, Murillo's and some minor artists' paintings. In general, then, the ability visually to distinguish Fs from non-Fs involves different perceptual-discriminative tasks, depending on how non-Fs are individuated. The point that fake barn cases comprise two different visual abilities is, quite simply, an instance of this general phenomenon.

To be sure, one of the abilities encompasses the other. But this does not undermine the reality and independence of the less inclusive ability. The stories make this clear. For, the fact that explains the different outcomes of each story is that Diego's performance manifests the less inclusive ability, even though it fails to manifest the other, more inclusive ability. Alternatively put, the explanatory fact is that Diego has a limited ability, i.e. one to make certain visual

discriminations, but not others. The disambiguation of the category *non-barn* helps us see this.

To conclude, the suggestion that two different visual abilities underlie fake barn cases is hereby proposed as a tool to obtain clarity about the thought experiment. Moreover, it is a suggestion that a complete verdict on the prospects of PVE to meet the fake barn challenge should employ. As seen so far, the distinction helps us make sense of the achievement and no-achievement intuitions reported by epistemologists. But does it also help us make sense of the epistemic intuitions elicited? It may seem unlikely, for the no-knowledge intuition reported by epistemologists is traced to the presence of luck, and it may not seem obvious that luck is connected to ability. We turn to this now.

5. *Epistemic Intuitions*

Contemporary epistemologists who endorse the no-knowledge intuition about fake barn cases trace it to the presence of luck. The basic thought is that, given the overwhelming number of barn façades in the surrounding (modal) environment, fake barn subjects might easily have had a false target belief, which in turn shows that their actually true belief is only luckily true. This suggests a view of luck, according to which if an event is lucky, then it obtains in the actual world, although it fails to obtain in (most, if not all) nearby possible worlds. However, there are two ways to cash this out. For, it could be argued that what makes the target belief of fake barn cases only luckily true is the overwhelming presence of barn façades in the surrounding (modal) environment. Alternatively, it could be argued that what makes the target belief of fake barn cases only luckily true is not the overwhelming presence of façades *per se*, but rather the fact that the target belief is not an achievement, due to the overwhelming presence of façades in the surrounding (modal) environment.

What makes all the difference here is a plausible principle subscribed to by some contemporary epistemologists —namely, the modal robustness of abilities (see Greco 2010; Millar 2010; Turri 2011; Littlejohn 2014; Sosa 2015). It can be formulated as follows: if an actual exercise of an ability is an achievement, it would not easily fail to be an achievement in nearby possible worlds; and, if exercises of an ability in nearby possible worlds failed to count as achievements, so does an actual exercise of the same ability.

The principle is apparently exposed to counter-examples. For instance, talented golfers sometimes hit a hole-in-one. This is an event

unlikely to obtain in nearby possible worlds, yet it is a performance by a talented golfer, hence an achievement. But it is not plausible to claim that everything that talented golfers do (while playing golf) counts as an achievement. Sometimes, talented golfers hit lucky shots. Hitting a hole-in-one is one of those lucky shots: it is an unlikely success, so the shot is not successful through ability. In support, consider that the ability to place the ball *near* the hole with a single shot must not be confused with the ability to hit a hole-in-one. Talented golfers have the former ability (reliably), so when a shot ends up near the hole through the exercise of their ability, it counts as an achievement. But, while possessing this ability, talented golfers may lack the ability to hit a hole-in-one. So when a shot is a hole-in-one, it does not count as an achievement, but as a lucky shot.

The modal robustness of abilities sheds light on fake barn cases, as follows. Fake barn subjects' performances in nearby possible worlds are not achievements: for, they fail to produce true beliefs appropriately connected to ability (after all, subjects cannot visually tell barns from façades). Given the modal robustness of abilities, their actual performance is not an achievement, either. Despite not being an achievement, the target belief is actually true. Therefore, it is a lucky true belief. Thus, what accounts for the actual truth of the target belief is luck, rather than its appropriate connection to ability.¹⁵

The upshot is that the “environmental luck” (Pritchard 2005) of fake barn cases can be understood in two different ways, one independent from, and one dependent on, achievement. Now, to link this to the previous sections of the paper, it is important to appreciate that

¹⁵ A consequence is that an actual performance is an achievement just in case its success is a matter of ability, not luck. However, a number of epistemologists now endorse the view that actual success is a matter of both ability and luck (e.g. McKinnon 2013; Carter 2016). Does this pose a problem here? There are two ways to take their view. According to the first, actual success is lucky, in the sense that exercising an ability does not absolutely guarantee success, even though it makes it likely; hence, success is only one of two possible outcomes. According to the second, ability and luck are two separate positive properties of actual success. In the former sense of luck, saying that an actual performance is lucky means that nothing went wrong that could have gone wrong, which is another way of saying that success happened. The latter sense of luck means something different: that if ability is exercised *and* success is the actual outcome, ability alone does not explain it, but rather ability plus luck. There is room for the former sense of luck in the point made in the main text (for, the modal robustness of abilities only requires likely success across possible worlds), but not for the latter sense of luck. However, it is unclear that luck has been proven to be a positive property of actual success, by those who view actual success as a matter of both ability and luck; perhaps all they have in fact shown is that actual success is lucky in the former sense.

these two notions of environmental luck rest on the different visual abilities already found to be doing work in the achievement vs. no-achievement intuitions elicited by fake barn cases. Thus, it could be argued that, regardless of the ability actually exercised by fake barn subjects (in particular, the ability to distinguish barns from non-barns, façades excluded), their target belief is luckily true, given the overwhelming presence of barn façades in the modal environment. Hence, luck in fake barn cases is independent of (actual) achievement. But, it could also be argued that fake barn subjects lack, or that their beliefs fail to manifest, the visual ability to tell barns from non-barns, façades included, which together with the modal robustness of abilities, latches the lucky nature of the actual target belief to the absence of, or the failure to manifest, that ability. Hence, luck is now dependent on achievement, or more precisely its absence. The significance of this for our understanding of the fake barn puzzle, and the intuitions it elicits, is that the existence of two different visual abilities helps make sense, not only of the achievement vs. no-achievement intuitions, but also of the place of luck in different reconstructions of the no-knowledge intuition.

Moreover, the distinction between visual abilities also helps make sense of the knowledge intuition elicited by fake barn cases. To see how, let us begin by noting that fake barn subjects' beliefs manifest the visual ability to tell barns from cows, tractors or silos (i.e. the ability to distinguish barns from non-barns, façades excluded). Furthermore, their beliefs would not easily fail to manifest this ability in nearby possible worlds, even though there they would fail to manifest a different visual ability —namely, the ability to tell barns from cows, tractors, silos and barn façades. So, the target belief that there is a barn ahead is an achievement relative to the ability to distinguish barns from non-barns, façades excluded. In turn, this means that the target belief is not only actually true, but also not lucky, despite the overwhelming number of façades in the surrounding environment. For, exercising the ability to tell barns from non-barns, façades excluded, produces achievements across nearby possible worlds. Therefore, there is no reason to return a no-knowledge verdict about fake barn cases. Not due to luck anyway, for the target belief is true through the exercise of a safe ability. Combined with the further thought that cognitive achievement is necessary and sufficient for knowledge, the above means that there is reason to return a knowledge verdict.

This verdict rests partly on the assumption that safety, as a condition for knowledge, is secured by the modal robustness of abili-

ties. However, some authors have denied that safety is necessary for knowledge (e.g. Comesaña 2005; Baumann 2008). Their case relies on counter-examples involving look-alikes in nearby possible worlds, as in fake barn cases. Thus, in Baumann's Wild West story of bank robbers Gottit and Nogood, a passer-by forms the true belief that Gottit just robbed the bank upon seeing his mask slip while getting away. In nearby possible worlds, the belief is false though, because the bank robber is in fact Nogood, wearing a Gottit mask underneath the mask that slipped. Baumann endorses the intuition that the passer-by actually knows that Gottit just robbed the bank, but given the proximity of these other possible worlds the true perceptual belief is not safe. Therefore, Baumann concludes that knowledge does not entail safety. In reply, it can be noted that the passer-by's true belief manifests the perceptual-discriminative ability to distinguish Gottit from other townsfolk, such as the bank clerk, the sheriff, and so on. This is an ability that the passer-by retains in nearby possible worlds, including those in which the passer-by cannot distinguish real Gottit from the Gottit-masked look-alike. Therefore, the passer-by's true perceptual belief has a safe, though limited, ability as its basis, and *pace* Baumann, qualifies as (limited) safe knowledge.¹⁶ (On limited knowledge, see below.)

Now, where does all this leave us regarding both the clarification of the puzzle and the prospects of PVE to meet the fake barn challenge?

6. *The Puzzle Clarified*

Section 4 has shown fake barn subjects' performance to manifest the less inclusive visual ability to tell barns from non-barns, façades excluded, but not the more inclusive ability to tell barns from non-barns, façades included. Section 5 has shown that, given the modal robustness of abilities, manifestation of the less inclusive ability leads to a robustly true belief, hence knowledge; and that failure to manifest the more inclusive ability leads to lucky true belief, hence absence of knowledge. This helps clarify the puzzle as follows.

Fake barn subjects have a limited visual ability; that is, an ability to make certain visual discriminations, but not others. For that reason, fake barn subjects have limited perceptual-discriminative knowledge of the target proposition. To explain: on the one hand, fake barn subjects can visually distinguish barns from cows, tractors, silos

¹⁶ Similarly for Broncano-Berrocal's (2017) dachshund example: the subject's true belief has a safe ability as its basis. Therefore, his safety dilemma for PVE does not arise.

and so on, and this ability is modally robust; therefore, fake barn subjects know that there is a barn ahead. On the other hand, fake barn subjects cannot visually distinguish barns from barn façades, either actually or counterfactually; hence, their knowledge is limited. But, how exactly is this idea of limited knowledge to be understood; is it perhaps a contrastivist idea?

The point that fake barn subjects have limited perceptual-discriminative knowledge of the target proposition may be expressed as follows: they know that there is a barn ahead, rather than a cow, a tractor and so on; yet they do not know that there is a barn ahead, rather than a barn façade. This could be taken in at least two ways. According to epistemic contrastivism (Karjalainen and Morton 2003; Schaffer 2005), knowledge is a three-place relation between a subject, a target proposition and a contrast class of propositions. Therefore, taken contrastively, the claim is that fake barn subjects know a target proposition rather than a disjunction of other propositions; and similarly for what they do not know. According to an alternative reading, the claim is that fake barn subjects non-contrastively know a target proposition, which in turn entails making certain visual discriminations, but not others. On the view that concepts are discriminative abilities (see Geach 1971, §5) and that propositional thought involves the exercise of concepts (see Evans 1982, pp. 100 ff.), this alternative reading emphasizes that (non-contrastively) knowing a certain proposition is a matter of making some correct discriminations (but not others) through ability. Here, knowledge remains a two-place relation. For current purposes, it is not necessary to choose between these readings, for both capture the crucial point about the story—namely, that fake barn subjects' performance manifests a limited visual ability, hence that they have limited knowledge.

There is nothing remarkable about the fact that fake barn subjects have a limited visual ability, for all (human) abilities are limited. But remarking on it helps with the puzzling nature of the example. Thus, it is not immediately obvious how the nature of the (modal) environment connects with the truism that (perceptual-discriminative) knowledge is limited. But the disclosure of the two visual abilities present in fake barn cases, together with the modal robustness of abilities, makes the connection apparent. In turn, the sense of puzzlement is eased, and the example clarified. For, talk of limited (perceptual-discriminative) knowledge is a way explicitly to acknowledge the limits of the epistemic situation of fake barn subjects: they know the target proposition (on account of the modally robust visual ability manifested in their true target belief), albeit in a limited way

(on account of the visual discriminations they cannot make, actually and counterfactually, given the visual ability they lack).

PVE theorists assume that there is a fact of the matter as to the epistemic status of the target belief of fake barn subjects. The preceding analysis has shown what the fact of the matter is —namely, that fake barn subjects' perceptual-discriminative knowledge is limited, on account of the (limited) visual ability manifested by their target belief. Here is an analogy. Mary can reliably score two-point shots at basketball, but she cannot reliably score three-point shots. Her throwing ability is limited. We capture this by saying that she has a limited throwing ability, for she can reliably score some shots, but not others. This is the fact of the matter, in Mary's case. Remember also Diego, who can reliably tell Velázquez's from Goya's and Murillo's paintings, but not from the paintings by minor artists from Velázquez's own workshop. The fact of the matter is that Diego's perceptual-discriminative knowledge is limited. Similarly, fake barn subjects can reliably (actually and counterfactually) tell barns from non-barns (other than façades), but not from façades. The fact of the matter is that they have limited perceptual-discriminative knowledge.

It is important to note that the analysis of fake barn cases provided here does not give rise to a generality problem. The generality problem for reliabilism arises when the token process leading to a particular belief can be subsumed under different types, with dissimilar success ratios. Now, in the analysis of fake barn cases provided earlier two different visual abilities have been invoked, one of which is more inclusive than the other; so one might have the impression that there are two possible types for a single, actual, token process. But this misplaces the complexity of fake barn cases, which does not lie in the existence of two different types (read, visual abilities) to subsume the token process leading to the target true belief, but rather in the existence of a single (limited) visual ability to make some discriminations, but not others. Crucially then, no generality problem arises.

Greco (2010) has claimed that PVE faces a version of the generality problem for reliabilism, regarding the individuation of abilities (what they are abilities to do), the conditions for their exercise, and the range of nearby possible worlds. But here the discriminative boundaries of the visual ability of fake barn subjects have been specified through disambiguation, the conditions for their exercise are normal, and the range of nearby possible worlds is determined by the presence of barns and non-barns (including barn façades). These

descriptions follow from the familiar understanding of the example in the literature, and the (alleged) relevance of alternative (or bizarre) descriptions can be ruled out on the basis of the shared understanding of the example. True, *ceteris paribus* and “and so on” clauses have been employed, but these are common currency in philosophical theorizing, and can be filled in further, if necessary; therefore, they are not an impediment to the plausibility of the analysis provided.

7. *PVE Vindicated*

The view that fake barn cases involve limited knowledge of the target proposition, on account of the limits in the visual discriminations fake barn subjects can make, is consistent with PVE. For it is based on the virtue-theoretic notions of ability and cognitive achievement, modally construed, together with a distinction between visual abilities. For this reason, PVE can predict and explain all the intuitions reported about fake barn cases. Here are the details.

According to the intuitions reported in option 3, the target belief of fake barn subjects counts as both an achievement and knowledge. According to the intuitions reported in option 2, it counts as neither a (complete) achievement nor knowledge. If fake barn subjects have limited knowledge of the target proposition, on account of the limits in their visual ability, these results are to be expected. For, evaluated with one visual ability in mind (i.e. the ability to tell barns from non-barns, façades excluded), the target true belief manifests a modally robust ability; so it is a case of knowledge. And, evaluated with a different visual ability in mind (i.e. the ability to tell barns from non-barns, façades included), the target belief does not manifest a modally robust ability; so it is not a case of knowledge. But it is only one (limited) visual ability that is possessed and exercised by fake barn subjects. Therefore, their true target belief is a case of limited knowledge, though apparently conflicting intuitions are reported when the performance of fake barn subjects is evaluated with the two visual abilities mentioned in mind.

What about option 1? According to it, fake barn cases involve achievement without knowledge. The underlying thought is that, though the target belief of fake barn subjects actually manifests their visual ability, the overwhelming number of façades in the surrounding (modal) environment turns the true belief lucky, and therefore deprives subjects of knowledge. As seen, this means that luck is thought of as independent of achievement. In turn, this has led

Pritchard to claim that an extra anti-luck condition, i.e. safety, is needed for knowledge.

But as also seen, abilities can be plausibly claimed to be modally robust; therefore, cognitive achievements do not need to be supplemented with an extra safety condition to amount to knowledge. In that case, fake barn subjects can be said to have some perceptual-discriminative knowledge, despite the overwhelming number of façades present in the surrounding environment. For their true target belief manifests the modally robust ability visually to tell barns from other objects, barn façades excluded. Therefore, the most that could be concluded from the existence of an overwhelming number of façades in the (modal) environment is that fake barn subjects lack some perceptual-discriminative knowledge. This is so because, *pace* Pritchard, the discriminative knowledge fake barn subjects lack is related to their actually, and counterfactually, lacking or failing to manifest a particular visual ability —namely, the ability visually to tell barns from other objects, barn façades included. Again, this is consistent with fake barn subjects having knowledge of the target proposition, albeit limited, on account of the limits in the visual ability they exercise.

Lately, Pritchard (2016) has claimed that the problem underlying fake barn cases is not luck *per se*, but risk; and that in defending the view that fake barn subjects know the target proposition, one is committed to the counter-intuitive compatibility between knowledge and the high risk of having a false belief. The risk is high, because of the proximity of a possible world where the belief in the target proposition turns out false, despite being formed on the same basis as in the actual world. But the basis on which the target true belief is actually formed is the exercise of a visual ability to distinguish between barns and non-barns, façades excluded, and as this is a modally robust basis, the risk of forming a false belief is not high at all. Therefore, *pace* Pritchard, the intuition that fake barn subjects have limited knowledge does not lead to high-risk beliefs. So, shifting the focus from luck to risk does not pose a problem for the analysis provided in this paper.

This pretty much completes two of the tasks undertaken in the paper —namely, to show that PVE has the resources to clarify the fake barn puzzle, and to meet the ensuing challenge. To sum up, the puzzle has been shown to contain an overlooked complexity traceable to the limited nature of the visual ability underlying the target belief of fake barn subjects. Overlooked, because of an ambiguity in the

characterization of the ability. After disambiguation, the limits are made apparent, and the epistemic fact of the matter can be clearly appreciated. As the example has been clarified through the use of virtue-theoretic tools alone, PVE can be shown to have the resources to predict, and explain, all the intuitions reported by contemporary epistemologists. Hence, it has been vindicated against the fake barn challenge.

8. *Our Epistemological Theorizing*

At this point, one may both accept that the analysis developed here clarifies the fake barn puzzle and shows how PVE can meet the challenge, and still worry that this result may not help our epistemological theorizing very much. After all, it can be argued, the fake barn case is only one example, and diagnosing what is puzzling about it need not automatically provide us with a better understanding of what knowledge is. Even showing that PVE is not troubled by fake barn cases may not drastically enhance that understanding, for there are plenty of puzzling cases in post-Gettier epistemology, and something other than a piecemeal approach to the analysis of such cases is required if our epistemological theorizing is to be improved. Rather, what is required is a more general story than this paper has told, so far.

In reply to this worry, it can be noted that the diagnostic tools used to clarify the fake barn puzzle can also be applied to other prominent cases in post-Gettier epistemology; in particular, epistemic Twin-Earth cases and traditional Gettier cases. For that reason, a case can be made for the claim that the virtue-theoretic tools used earlier can enhance our epistemological theorizing, by illuminating the relationship between knowledge, luck and ability.

To begin with, consider epistemic Twin-Earth cases, as described by Kallestrup and Pritchard (2011). On Earth, *S* forms a true belief that some demonstratively grasped liquid is water, through the visual ability to tell it apart from other liquids. On Twin-Earth, twin-*S* forms the same true belief on the same basis, but unlike on Earth, there is a nearby perceptual possibility to the effect that the liquid in question is twin-water. Can PVE predict and explain the intuitions avowed about this case —namely, that *S* knows, whereas twin-*S* does not? And can it do so, while acknowledging that *S* and twin-*S* are physical duplicates, and that abilities supervene on physical properties? The answer is that it can, thanks to the familiar distinction

between visual abilities. Thus on the one hand, twin-*S* cannot visually distinguish water from twin-water, and in that sense lacks some perceptual-discriminative knowledge. But on the other hand, *S* and twin-*S* alike can visually tell water apart from other liquids, twin-water excluded. Hence, twin-*S* has some perceptual-discriminative knowledge, for the same reason as *S* —namely, their target belief is true through the exercise of a modally robust visual ability to distinguish water from other liquids, twin-water excluded.¹⁷ The supervenience claim is not violated, because the epistemic difference between *S* and twin-*S*, though predicted and explained, is not traceable to the abilities shared by both. In fact, as *S* and twin-*S* share a modally robust ability a verdict of limited knowledge can be granted in both cases, though only in Twin-Earth is the limit made explicit. Therefore, PVE has the tools needed both to predict and explain the no-knowledge intuition reported by Kallestrup and Pritchard, and to clarify the puzzling nature of epistemic Twin-Earth cases, in line with the previous analysis.

Next, consider some traditional Gettier cases; for instance, Chisholm's (1989) sheep-in-the-field example. Here, through vision *S* forms a true belief that there is a sheep in a field. However, the belief turns out to be only luckily true, for *S* is in fact looking at a sheep-looking object, not a sheep, and there just happens to be a sheep elsewhere in the field, though out of view. The unanimous intuition elicited by this example is that *S* lacks knowledge, on account of the lucky nature of the belief. For, it is thought, in most, if not all, nearby possible worlds where *S* formed a belief on the same (visual) basis, there would not be an out-of-view sheep in the field; so, *S*'s belief would be false. The interesting point to note now is that this no-knowledge intuition can be predicted and explained with the help of the virtue-theoretic tools used earlier in the paper. In particular, the example allows for a distinction between two different visual abilities —namely, the ability to distinguish sheep from many other objects, sheep look-alikes included, and the ability to distinguish sheep from many other objects, sheep look-alikes excluded. There is no reason to think that *S* lacks, or that *S*'s belief fails to manifest, the latter visual ability, even though the former ability is either lacking, or not manifested. Furthermore, that *S* lacks, or that *S*'s belief fails to manifest, the ability to distinguish sheep from many

¹⁷ Kallestrup and Pritchard (2011, p. 343) write: “the mere presence of twin-water on Twin-Earth does not prevent twin-*S* from being able to reliably tell water apart from petrol or beer”. But they fail to notice that this helps PVE.

other objects, sheep look-alikes included, explains the no-knowledge intuition. For, given that abilities are modally robust, if S 's true actual belief manifests the ability, then S 's performance would not easily fail to be an achievement in nearby possible worlds. But, as noted, S 's performance in most, if not all, nearby possible worlds would likely be false (hence, not an achievement), so S 's actual performance also fails to be an achievement. So, S 's epistemic luck, and lack of knowledge, can be traced to the absence of achievement, i.e. the failure to possess or manifest a particular ability.¹⁸

So, a common pattern is beginning to emerge throughout the examples; for what accounts for epistemic luck, and therefore absence of knowledge, is absence of achievement, i.e. absence of true belief through appropriate connection to a (modally robust) visual ability. Similarly, what accounts for (limited) knowledge in fake barn cases (and epistemic Twin-Earth cases), and would account for (equally limited) knowledge in traditional Gettier cases, if suitably modified, is presence of achievement. Think, for instance, of the subject in the sheep-in-the-field example looking at a sheep, and on the basis of a visual ability to distinguish sheep from many other objects, sheep look-alikes excluded, forming the true belief that there is a sheep (not a dog, a farmer, and so on) in the field ahead. The subject's actual performance delivers a true belief, and given the modal robustness of abilities, it would not easily fail to do so in nearby possible worlds (even if there were sheep look-alikes in the surrounding modal environment, and the subject could not visually tell them apart from sheep).

Succinctly put, the emerging pattern is that absence of achievement entails luck, hence absence of knowledge; and that presence of achievement entails no luck, hence knowledge. In other words, given the modal robustness of ability, achievement is necessary and sufficient for knowledge, as PVE claims.¹⁹ Importantly though, this vindication of PVE as a theory of knowledge is tied not to the

¹⁸ The same analysis accounts for the no-knowledge intuition normally elicited by Lehrer and Paxson's (1969) Tom Grabit case and by Zagzebski's (1994) husband-in-the-room example.

¹⁹ According to Miracchi's (2015) variation on Chisholm's example, in most (if not all) nearby possible worlds, there is (systematically) an out-of-view sheep elsewhere in the field. This also fits the pattern noted above. The belief does not count as knowledge, because the subject actually (and counterfactually) lacks, or fails to manifest, the ability to distinguish sheep from many other objects, sheep look-alikes included; so her belief is not an achievement. Thus, absence of knowledge is still traced to absence of achievement; it is lucky that an out-of-view sheep is (systematically) present in most (if not all) nearby possible worlds.

analysis of fake barn cases alone, but also of other puzzling cases in contemporary epistemology. For, the same tools have been applied consistently to all the considered examples. So, those tools have proven their worth, by providing us with a better understanding of the relations between knowledge, luck and ability. In that respect, our epistemological theorizing has been enhanced.

Two final objections before the end. The first is that in the resulting picture knowledge boils down to discrimination, which is not plausible (for discussion, see Pritchard 2010b). Thus, if fake barn subjects know that there is a barn ahead (as claimed earlier), surely they know that there is not a barn façade ahead. Yet they cannot discriminate between barns and façades (as also claimed above); so the latter piece of knowledge cannot be a matter of discrimination. The reply is that, although fake barn subjects lack the required perceptual ability, they may still know that there is not a barn façade ahead, if they competently deduce so from their (limited) perceptual-discriminative knowledge that there is a barn ahead, together with the entailment relation between both propositions. Competently deducing a conclusion from its premisses is an ability, indeed one that can be exercised to produce cognitive achievements. Therefore, although not all knowledge boils down to perceptual discrimination, as the existence of inferential knowledge shows, the emerging picture is not impugned, for non-perceptual knowledge can still be a matter of cognitive achievement, albeit of a different, non-perceptual variety.

In support, consider that the same pure virtue-theoretic tools employed above can help with non-perceptual Gettier cases. Thus, in Gettier's (1963) coins-in-the-pocket or in Lehrer's (1965) Nogot examples, the relevant belief is actually true but is not an achievement, so it is only luckily true; hence, not an instance of knowledge. It is not an achievement, because it is not true through appropriate connection to a modally robust inferential ability. (In nearby possible worlds, the same inferential basis would likely not lead to true beliefs.) This is the familiar pattern connecting achievement and knowledge, despite the non-perceptual nature of the examples.²⁰

The final objection is that the resulting theory of knowledge is too permissive, for in some sense of ability, an ability is possessed and exercised whenever one gets things right; so in some sense,

²⁰ Introducing a Miracchi-style, systematic element does not alter the pattern: absence of knowledge is still a matter of absence of achievement, for truth (though preserved across nearby possible worlds courtesy of the systematic element) is not a matter of appropriate connection to a modally robust inferential ability.

knowledge rather than lucky belief is always the outcome. But the theory is not too permissive. For one thing, exercising an ability is not the same as manifesting it; therefore, even if it is true that in some sense an ability is exercised whenever one gets things right, this is not the same as getting things right through the exercise of the ability. For another, not only does the resulting theory of knowledge predict that one can get things right by luck, but it also explains why —namely, because the performance is not successful through a modally robust ability.

To sum up and conclude, the analysis provided here not only clarifies the fake barn puzzle, and shows how PVE can meet the ensuing challenge, but also significantly enhances our understanding of the relations between knowledge, luck and ability, and in doing so, contributes to our epistemological theorizing.²¹

REFERENCES

- Baumann, P., 2008, “Is Knowledge Safe?”, *American Philosophical Quarterly*, vol. 45, no. 1, pp. 19–30.
- Broncano-Berrocal, F., 2017, “A Robust Enough Virtue Epistemology”, *Synthese*, vol. 194, no. 6, pp. 2147–2174.
- Cappelen, H., 2012, *Philosophy without Intuitions*, Oxford University Press, Oxford.
- Carter, J.A., 2016, “Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution”, *Pacific Philosophical Quarterly*, vol. 97, no. 1, pp. 140–155.
- Chisholm, R., 1989, *Theory of Knowledge*, 3rd ed., Prentice-Hall, Englewood Cliffs.
- Colaço, D., W. Buckwalter, S. Stich and E. Machery, 2014, “Epistemic Intuitions in Fake-Barn Thought Experiments”, *Episteme*, vol. 11, no. 2, pp. 199–212.
- Comesaña, J., 2005, “Unsafe Knowledge”, *Synthese*, vol. 146, no. 3, pp. 395–404.
- Evans, G., 1982, *The Varieties of Reference*, ed. J. McDowell, Clarendon Press, Oxford.
- Geach, P., 1971, *Mental Acts*, 3rd ed., Routledge and Kegan Paul, London (First published, 1957).
- Gettier, E.L., 1963, “Is Justified True Belief Knowledge?”, *Analysis*, vol. 23, no. 6, pp. 121–123.

²¹ Work for this paper was supported by a grant from the Spanish government (ref. FFI2012–38908–C02–02). For encouraging feedback or constructive criticism on earlier drafts, David García Carreres, Tobías Grimaltos, Noreen Mabin and the appointed referees are gratefully acknowledged.

- Goldman, A., 1976, “Discrimination and Perceptual Knowledge”, *The Journal of Philosophy*, vol. 73, no. 20, pp. 771–791.
- Greco, J., 2012, “A (Different) Virtue Epistemology”, *Philosophy and Phenomenological Research*, vol. 85, no. 1, pp. 1–26.
- , 2010, *Achieving Knowledge*, Cambridge University Press, Cambridge.
- Hetherington, S., 1999, “Knowing Failably”, *The Journal of Philosophy*, vol. 96, no. 11, pp. 565–587.
- Jarvis, B., 2013, “Knowledge, Cognitive Achievement and Environmental Luck”, *Pacific Philosophical Quarterly*, vol. 94, no. 4, pp. 529–551.
- Kallestrup, J. and D. Pritchard, 2011, “Virtue Epistemology and Epistemic Twin Earth”, *European Journal of Philosophy*, vol. 22, no. 3, pp. 335–357.
- Karjalainen, A. and A. Morton, 2003, “Contrastive Knowledge”, *Philosophical Explorations*, vol. 6, no. 2, pp. 74–89.
- Kelp, C., 2013, “Knowledge: The Safe-Apt View”, *Australasian Journal of Philosophy*, vol. 91, no. 2, pp. 265–278.
- Lehrer, K., 1965, “Knowledge, Truth and Evidence”, *Analysis*, vol. 25, no. 5, pp. 168–175.
- Lehrer, K. and T. Paxson, 1969, “Knowledge: Undefeated Justified True Belief”, *The Journal of Philosophy*, vol. 66, no. 8, pp. 225–237.
- Littlejohn, C., 2014, “Fake Barns and False Dilemmas”, *Episteme*, vol. 11, no. 4, pp. 369–389.
- Lycan, W., 2006, “On the Gettier Problem problem”, in S. Hetherington (ed.), *Epistemology Futures*, Oxford University Press, Oxford, pp. 148–167.
- McKinnon, R., 2013, “Getting Luck Properly under Control”, *Metaphilosophy*, vol. 44, no. 4, pp. 496–511.
- Millar, A., 2010, “Knowledge and Recognition”, in Pritchard, Millar and Haddock 2010, pp. 91–190.
- Miracchi, L., 2015, “Competence to Know”, *Philosophical Studies*, vol. 172, no. 1, pp. 29–56.
- Navarro, J., 2015, “No Achievement beyond Intention”, *Synthese*, vol. 192, no. 10, pp. 3339–3369.
- Pritchard, D., 2016, “Epistemic Risk”, *The Journal of Philosophy*, vol. 113, no. 11, pp. 550–571.
- , 2010a, “Knowledge and Understanding”, in Pritchard, Millar and Haddock 2010, pp. 1–88.
- , 2010b, “Relevant Alternatives, Perceptual Knowledge and Discrimination”, *Noûs*, vol. 44, no. 2, pp. 245–268.
- , 2009, *Knowledge*, Palgrave Macmillan, Basingstoke.
- , 2005, *Epistemic Luck*, Oxford University Press, Oxford.
- Pritchard, D., A. Millar and A. Haddock, 2010, *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press, Oxford.

- Schaffer, J., 2005, "Contrastive Knowledge", in T.S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 1, Oxford University Press, Oxford, pp. 235–271.
- Sosa, E., 2015, *Judgment and Agency*, Oxford University Press, Oxford.
- , 2011, *Knowing Full Well*, Princeton University Press, Princeton and Oxford.
- , 2007, *A Virtue Epistemology*, Clarendon Press, Oxford.
- Turri, J., 2012, "Is Knowledge Justified True Belief?", *Synthese*, vol. 184, no. 3, pp. 247–259.
- , 2011, "Manifest Failure: The Gettier Problem Solved", *Philosophers' Imprint*, vol. 11, no. 8, pp. 1–11.
- Turri, J., W. Buckwalter and P. Blouw, 2015, "Knowledge and Luck", *Psychonomic Bulletin and Review*, vol. 22, no. 2, pp. 378–390.
- Zagzebski, L., 1994, "The Inescapability of Gettier Problems", *Philosophical Quarterly*, vol. 44, no. 174, pp. 65–73.

Received: May 17, 2016; revised: December 18, 2017; accepted: February 8, 2018.