


ARTICLE

Agent Reliabilism and Inferential Knowledge from Gettiered Belief

K. Merrick Olivier 

School of Divinity, History and Philosophy, University of Aberdeen, Aberdeen, UK
Email: olivier.kevin@abdn.ac.uk

(Received 15 July 2019; revised 21 December 2019; accepted 28 February 2020;
first published online 24 June 2020)

Abstract

Epistemologists have generally accepted that competently deduced, known conclusions must issue from known premises, as the principle of Counter-Closure demands; however, some have recently challenged the notion, arguing that knowledge may be inferred from non-knowledge. In this paper, I focus on the yet unexamined topic of inferential knowledge from Gettiered belief with regard to Greco's virtue-epistemic framework, which he refers to as 'agent reliabilism'. I argue that agent reliabilism allows for instances of Counter-Closure violation. In presenting my argument, I construct and provide an analysis of a case that challenges Counter-Closure, and defend the case against various possible objections, including some that feature in mainstream Counter-Closure literature. I contend that, if my analysis is accepted, Greco is left with two options: either he rejects Counter-Closure, or he must substantially revise his view in such a way as to preserve Counter-Closure in light of cases of inferential knowledge from Gettiered belief.

Keywords: Greco; virtue epistemology; counter-closure; agent reliabilism; Gettier

Epistemologists have generally accepted that competently deduced, known conclusions must result from known premises, as the principle of Counter-Closure demands; however, some have recently challenged the notion, arguing that knowledge may be inferred from non-knowledge.¹ In this paper, I focus on the topic of inferential knowledge from Gettiered belief with regard to Greco's virtue-epistemic framework.² To date, no one has examined the relationship between Counter-Closure and virtue epistemology; although I only focus on agent reliabilism in this paper – a version of virtue reliabilism championed by John Greco – if my argument is correct, it will have important implications

¹Epistemologists critical of Counter-Closure have argued that knowledge may be inferred from an unjustified belief (Murphy 2015), a justified false belief (Warfield 2005; Klein 2008; Fitelson 2010; Arnold 2011; Luzzi 2014), an assumption that falls short of belief (Murphy 2013), or a Gettiered belief (Luzzi 2010).

²Virtue epistemologies are generally regarded as fitting into one of two categories: virtue reliabilism or virtue responsibilism. Greco's version of virtue epistemology, which he refers to as agent reliabilism, is a form of virtue reliabilism. Considering the dissimilarity between how virtue epistemologists conceive of intellectual virtues, it would be best to examine agent reliabilism with regard to Counter-Closure separately from other virtue-epistemic theories. As such, I only focus on agent reliabilism with regard to inferential knowledge from Gettiered belief in this paper.

for all virtue-epistemic frameworks which characterize the knowledge-relevant virtues, and knowledge in general, in a similar manner as Greco. I preface my argument by briefly outlining the tenets of agent reliabilism.³ I then examine various Gettier-type cases in light of agent reliabilism, which Greco argues is able to accommodate our intuitions regarding such cases. Finally, I construct and provide an analysis of a case that challenges Counter-Closure. If my analysis of the case is accepted, then Greco is left with two options: either he rejects Counter-Closure, or he must substantially revise his view in such a way as to preserve the principle in light of cases of inferential knowledge from Gettiered belief.

1. Agent reliabilism and Gettier cases⁴

I will begin by outlining the main tenets of agent reliabilism. As with other virtue-epistemic frameworks, agent reliabilism refers to an approach to epistemology according to which intellectual agents are regarded as the fundamental source of epistemic normativity, and intellectual virtues constitute the primary focus of epistemic evaluation. What distinguishes agent reliabilism from other virtue-epistemic theories is how knowledge is defined, as well as what constitutes an intellectual virtue. Knowledge is defined, according to Greco's agent reliabilism, as follows:

S knows that p iff:

1. S's believing that p is produced by a knowledge-relevant ability (i.e., an intellectual ability *of the relevant sort* required for knowledge);
2. S is in conditions relevant for the exercise of that ability; and
3. S has a true belief *because* S's belief is produced by a knowledge-relevant ability, while in relevant conditions. (Greco 2016: 57)

Greco defines knowledge-relevant abilities as "disposition[s] to believe relevant truths, in relevant circumstances, relative to some environment, with a sufficient degree of reliability" (Greco 2010: 18).⁵ The parameters relevant to these abilities are defined according to pertinent informational needs; although parameters may often be set to values that are normal for the exercise of a relevant ability, they may be altered if a task requires the exercise of an ability to obtain information under abnormal conditions. If, for example, an individual was tasked with separating balls by colour under normal conditions (e.g., being fully sober and in a well-lit room with nothing obscuring their view) and came to form true beliefs regarding the colour of each ball, then they may be credited with having the pertinent knowledge-relevant ability. If, however, the individual was asked to perform the same task but with the lights dimmed to such an extent that their ability to distinguish colour is poor and unreliable, they would not be credited with having the knowledge-relevant ability to separate balls by

³While there are other epistemologists who identify as virtue reliabilists, I choose to focus mainly on Greco's account of agent reliabilism in this paper, as I find his reliabilist epistemic framework to be one of the most comprehensive and spelled-out. That being said, if my argument against Counter-Closure with regard to agent reliabilism is accepted, then I believe it would have implications for other virtue reliabilist theories which characterize inferential reasoning as an intellectually virtuous process; however, these views fall outside of the scope of this paper.

⁴The term 'Gettier case' broadly refers to a wide range of thought experiments – first proposed by Edmund Gettier (1963) – in which a subject possesses a justified true belief that, intuitively, does not constitute knowledge.

⁵The term "knowledge-relevant ability" will be used interchangeably with "intellectual virtue" throughout this paper.

colour under the circumstances, even if they happened to form true beliefs regarding the colour of some of the balls. In this case, the separation task would require the ability to reliably distinguish between colours under (abnormal) extremely poor lighting conditions; as such, only individuals with this superior capacity may be credited with having the knowledge-relevant ability.

With knowledge generally characterized as true belief “produced by the right sort of ability in the right sort of way” (Greco 2012: 19), Greco claims that agent reliabilism is able to explain why subjects in a wide range of Gettier cases do not possess knowledge. In such cases, the subject is traditionally regarded as possessing a justified true belief in a proposition that – as a result of the influence of epistemic luck – they do not know. Greco’s response to most Gettier cases is as follows:

In cases of knowledge, the fact that S believes from an intellectual ability *explains why* S has a true belief. In Gettier cases, S believes from an ability and S has a true belief, but the fact that S believes from an ability does not explain why S has a true belief. (Greco 2009: 19)

This position, according to Greco, is motivated by the pragmatics of causal explanation language. As information-sharing agents, we value explanatory salience, which is “partially a function of our interests and purposes”, as well as “partially a function of what is normal and usual” (Greco 2009: 20). When it comes to forest fires, for example, human error (e.g., not completely stamping out a campfire) is often cited as a cause in order to assign blame, even though there may have been numerous factors contributing to the fire, such as the dryness of the forest and propitious weather conditions. Likewise, in situations where an occurrence is the result of an abnormal influence, explanatory salience is granted to the abnormality. For example, imagine that a skilled basketball player is about to take a shot, but right before they release the ball, their ankle gives out, causing them to release the ball in such a way that it is guaranteed to miss the basket by a significant distance; however, in a stroke of luck, the ball ends up hitting a bird flying by and ricochets into the basket. In this case, although the aim of the basketball player’s performance was achieved (i.e., putting the ball through the basket), the player is not credited with the success of the shot, since explanatory salience is granted to the abnormality (i.e., the ball ricocheting off a bird), even though multiple factors contributed to the success, such as the player releasing the ball just as the bird flies overhead. Concisely put, in normal cases of knowledge, “intellectual abilities have a default salience in explanations of true belief”, whereas in a wide range of Gettier cases the “default salience is trumped by something abnormal in the way that S gets a true belief” (Greco 2009: 20). For those Gettier cases in which intellectual abilities do not have a default salience in explanations of true belief, Greco’s third condition for knowledge is not met: the subject does not possess a true belief *because* it is produced by a knowledge-relevant ability; rather, the subject has a true belief as a result of the influence of epistemic luck.

To exemplify how common Gettier cases may be diagnosed according to Greco’s characterization of knowledge, let us consider the following popular case:

Chisholm’s Sheep: A subject notices what appears to be a sheep standing in a field, as a result of which they form the belief that there is a sheep in the field. As it turns out, what they take to be a sheep is really a dog; however, it is true that there is a sheep in the field, as one is standing right behind the dog, outside of the subject’s line of sight.⁶

⁶This is a slightly altered version of Chisholm’s (1989) Sheep case.

In this case, the subject is commonly regarded as being justified in believing that there is a sheep in the field, which happens to be true, yet they arguably do not know that there is a sheep in the field. While the subject's belief might be true, agent responsibilists would argue that they lack knowledge of the proposition since their belief isn't true *as a result of* its being produced by a knowledge-relevant ability; rather, explanatory salience for the truth of their belief is granted to there unexpectedly being a sheep in the field somewhere outside of their line of sight. Although the subject possesses a true belief formed on the basis of a knowledge-relevant ability (i.e., perception), the fact that they believe on the basis of such an ability does not explain why they possess true beliefs; as such, the subject is not regarded as knowing the proposition in which they believe.

While appealing to explanatory salience allows for the diagnosis of a wide range of Gettier cases, this approach cannot account for a subject's lack of knowledge in Gettier cases of the Fake Barn variety.⁷

Fake Barns: Henry is driving through the countryside and perceives a barn in clear view. On the basis of his perception, he forms the belief that the object he sees is a barn. Unbeknownst to Henry, he is driving through Fake Barn County, which is littered with barn façades that are indistinguishable from real barns from the perspective of those driving on the road; however, in a stroke of luck, Henry happens to be looking at the only real barn around.⁸

Contrary to the previously examined cases, the problem presented by Goldman's Fake Barn case cannot be understood as a "deviant causal chain" between ability and true belief" (Greco 2009: 21). That being said, Greco argues that Henry's lack of knowledge is the result of his belief not being founded on a knowledge-relevant ability, since such abilities are defined relative to specific circumstances and environments. Just as a funambulist who possesses the ability to walk a tightrope under normal circumstances would lack that ability while in hurricane-force winds, one's ability to form true beliefs on the basis of perceptual evidence is limited to conditions and environments propitious for the exercise of one's knowledge-relevant faculties. Whereas Henry may have the ability to reliably discern real barns from mere façades relative to normal environments and circumstances, he lacks that ability while driving through Fake Barn County. Perhaps if the circumstances were such that Henry had to circle around each barn façade as he drove, he would reliably be able to make the distinction; however, considering that the conditions of the case are such that real barns and façades appear identical from Henry's vantage point, he cannot be credited with having the knowledge-relevant ability to discern real barns from façades in this environment and under these circumstances. In this case (as well as those similarly structured), Greco's second condition for knowledge is not met: the subject is *not* in conditions propitious for the exercise of what, under normal circumstances, would be considered their knowledge-relevant ability.

2. Inferential knowledge from Gettiered belief

Having provided a brief overview of the relevant tenets of agent reliabilism, as well as how Greco explains various Gettier-type cases, I will now present and discuss a case that

⁷While epistemologists generally agree that knowledge is absent in Fake Barns cases, a minority, including some experimental philosophy theorists, reject this position and argue that knowledge is present (Hetherington 1999; Sosa 2007; Turri 2012).

⁸This is a slightly altered version of Goldman's (1976) Barn Façade case, which he attributes to Carl Ginet.

challenges Greco's ability to preserve the principle of Counter-Closure, which – motivated by the widely accepted belief that the epistemic status of a conclusion can be no better than that of the essential premise(s) from which it was inferred – states that “knowledge-yielding competent deductive inference must issue from known premises” (Luzzi 2010: 673). The principle may be expressed as follows:

Counter-Closure: If (i) S knows that Q, and (ii) S's belief that Q is based exclusively on competent deduction from a set of essential premises including P, then (iii) S knows that P.

The principle of Counter-Closure has long enjoyed widespread acceptance among epistemologists, bolstered by the view that deductive inference can only transmit knowledge, not create it. This point is evidenced by the fact that the principle of Counter-Closure itself, and views that entail it, have been endorsed by numerous prominent epistemologists and often accepted as orthodoxy.⁹ Note, however, that not everyone thinks that Counter-Closure is true. Since Warfield's (2005) publication on knowledge from falsehood, there has been substantial debate regarding the truth of the principle, with numerous examples of alleged inferential knowledge from non-knowledge being proffered, including cases of knowledge from Gettiered belief. To date, this debate has led a minority of epistemologists to conclude that Counter-Closure is false. If my argument is correct, then Greco must abandon the principle that the majority of epistemologists hold, or substantially revise his view in order to preserve Counter-Closure.

I should note here that the aim of this paper is purely to demonstrate the incompatibility of agent reliabilism (in its current form) and Counter-Closure; as such, I will not provide detailed argument regarding which of the two aforementioned responses to my case Greco should choose, if he accepts my analysis. Given that virtue epistemology, in general, is gaining in popularity, the compatibility of prominent virtue-epistemic frameworks with entrenched epistemic views is worth analysing in and of itself. For proponents of Counter-Closure, finding that the principle is incompatible with their favoured epistemic framework leads one to re-evaluate one's epistemic commitments and abandon either Counter-Closure or their favoured framework, or alter one's framework in such a way that it is compatible with the principle. Whichever option one chooses, it is apparent that acknowledging such an incompatibility leads one to further develop one's overall epistemic view.

With regard to agent reliabilism specifically, if Greco were to break with the majority of epistemologists and simply abandon Counter-Closure he would have to accept that knowledge may be gained via inference from premises which are unknown, whether due to their being false, or due to their not being the product of a knowledge-relevant ability. This would further require Greco to acknowledge that unknown beliefs may be reliable with regard to inferring certain known propositions, a point which I will later address in my discussion of subjective justification. To clarify, Greco would not have to hold that inferential knowledge may be acquired via inference from *any* unknown premise-belief; rather, he would have to accept that competent inference from *certain* unknown premise-beliefs can reliably produce true beliefs, some of which, if the agent is conscientious, are known. In this paper, I do not attempt to discern all the conditions under which an unknown premise-belief might serve as the basis for one's knowledge conferring inference. I do, however, suggest that for premise-beliefs of the

⁹For examples of prominent epistemologists relying on the truth of Counter-Closure in their writings, see Armstrong (1973: 189–99), Nozick (1981: 231), Stanley (2005: 89–90) and Audi (2010: 184–91).

Fake Barns variety, an inference based on such a belief can be reliable, and therefore knowledge-conferring, if the conclusion is entailed by the Gettiered belief and certain other key features are in place.

In the remainder of this section, I construct a case in which a subject infers a proposition from a Gettiered belief. I then argue that, according to the tenets of agent reliabilism, the subject knows their conclusion, even though their premise-belief is unknown (because Gettiered). In the subsequent sections, I defend my argument against possible objections.

Consider the following case:

Structures: James is driving from County A to County B. He knows that there are no obvious border signs separating County A from County B; the only noticeable difference between the counties is the presence of barns and barn-shaped structures in County B. James knows that in County A all barns are absolutely forbidden. In fact, in County A anything even resembling a barn is forbidden; as a preventative measure, the county has developed a surveillance system that has been proven to be completely effective, which automatically detects anything resembling a barn and instantly destroys it. In addition to knowing barns are forbidden in County A, James knows that there is an abundance of barns in County B, though not within 5 miles of its border with County A. While driving between counties, James perceives what appears to be a barn a short distance away (<5 miles), on the basis of which he forms the belief that (*p*) the object ahead is a barn. Remembering that barns are absolutely forbidden and instantly destroyed in County A, as well as the fact that barns in County B must be at least 5 miles away from the county line, James deduces from *p* that (*q*) he is in County B. Unbeknownst to James, he happens to be driving through a part of County B that is littered with barn façades; he does not know that this area is littered with barn façades, and is not blameworthy for this lack of knowledge, as the façades unexpectedly popped up the night before, with no publicly available information regarding their presence. Although James is driving through an area littered with barn façades, it turns out he is looking at the only real barn around.

In this case, James's belief that (*q*) he is in County B is true and was competently deduced from his belief that the object he perceives ahead is a barn. In order for this to qualify as an instance of inferential knowledge from non-knowledge, James must have come to believe *q* solely on the basis of competent deduction from a set of essential premises including *p*; he must know that *q*; and he must not know that *p*. Considering the fact that the observation of barns (and things resembling barns) is the only reliable way to recognize when one has left County A and entered County B, it seems uncontroversial to accept that James came to believe *q* solely on the basis of competent deduction from a set of essential premises including *p*; indeed, the case is designed in such a way that James could only determine that he has entered County B via deductive reasoning.

According to Greco's conditions for knowledge, in order for James to know that (*p*) the object ahead is a barn, his belief that *p* must be produced by a relevant intellectual ability, for which he must be in the appropriate conditions for its exercise, and his belief must be true *because* it was produced by the ability within relevant conditions. Since James based his belief that *p* on his perception of a barn, the knowledge-relevant ability of concern in this instance is the ability to identify real barns. As was the case with Henry in Fake Barns, Greco should accept that James possesses the disposition to reliably discern real barns from fake ones relative to normal environments, but deny that James possesses the same ability given the environment he finds himself in. Since James

is driving through an area of County B that is littered with barn façades, he would have to be able to reliably tell the difference between real and fake barns while driving in order for him to be credited with possessing a true belief regarding a barn on the basis of the stated knowledge-relevant ability; however, as he is not able to reliably distinguish real barns from fake ones relative to his environment and the prevailing circumstances, it must be concluded that James lacks the pertinent knowledge-relevant ability. Consequently, since James's belief that p was not produced by a knowledge-relevant ability, it follows that he does not know that p . This conclusion is unsurprising and is in line with Greco's analysis of Fake Barns, which is wholly appropriate given the similarities between the cases.

While Greco must commit to the view that James does not know that p , by his own theoretical commitments, he must also accept that James knows that (q) he is in County B. In order to effectively demonstrate that James's belief that q meets Greco's conditions for knowledge, I will first explain how his belief that q was produced by a knowledge-relevant ability, and show that he was appropriately situated for the exercise of that ability; in doing so, I will demonstrate that James's belief that q is both subjectively and objectively justified. On occasion, Greco's account of knowledge explicitly features these conditions; however, as he believes that "intellectual virtues involve a motivation to believe the truth", and that virtuously held beliefs are objectively justified, he often opts to discuss 'knowledge' simply in terms of virtuous belief (Greco 2002: 311). Once I've established that James's belief that q was virtuously produced, I will demonstrate that his belief meets Greco's third condition for knowledge – that is, that his belief that q is true *because* it was virtuously produced. Finally, I will argue that, since James knows that q and came to believe that q via deductive inference from his belief that p , which is unknown, Structures undermines Counter-Closure.

With regard to identifying the salient knowledge-relevant ability exercised in instances of deductive reasoning, it should be noted that Greco argues that cases of inferential knowledge involve "dispositions to form one's beliefs on the basis of reliable evidence" (Greco 2010: 45); subjects exercising such dispositions exemplify a "sensitivity to the reliability of [their] evidence," which amounts to what Greco refers to as subjective justification (Greco 1999: 289). 'Reliability', as used in the last quotation, denotes the likelihood that a belief is true given the evidence upon which the belief is based (Greco 1999: 289). Since, plausibly, the premises of one's inference serve as the evidence for the inference, for a subject to have inferential knowledge of a conclusion (according to agent reliabilism), they must, in part, be sensitive to the reliability (or truth) of their premise(s), which is to say that their belief in the conclusion must be subjectively justified (Greco 1999: 289). According to Greco (1999: 289), a belief is subjectively justified for a subject if and only if it is "grounded in the cognitive dispositions that [the subject] manifests when [they are] thinking conscientiously". It is important to note that conscientious thought does not require that a subject think "with an explicitly voiced purpose of finding the truth" (Greco 1999: 289); rather, Greco considers conscientiousness to be a feature of one's default mode of thought, manifested in trying to form one's beliefs accurately.

Since James came to believe that q on the basis of deductive inference from p , the knowledge-relevant intellectual ability in this case is, conceivably, the ability to form reliable beliefs on the basis of reliable evidence or, more specifically, the ability to form true beliefs on the basis of deduction from true premises, which is assumed to be a reliable disposition under normal circumstances (e.g., not under the influence of mind-altering substances). Framing the knowledge-relevant ability involved in deductive reasoning in terms of true belief is in keeping with Greco's conception of 'reliable evidence'; however, some may argue that this is too strong, since one's premise

might constitute reliable evidence for one's conclusion even if it is not true. For example, one might argue that the premise (p) it is raining outside is reliable evidence for the conclusion (q) if I leave the house without an umbrella I'll get wet, even under some circumstances where p is false, such as if the water were really falling from some atmospheric sprinkler system, rather than from clouds. In this example, even though the premise is false, in all nearby possible worlds where the subject comes to believe that q via the same method, q is true. I will address this point in greater detail in a subsequent section where I discuss conscientious inference.

In exercising his knowledge-relevant ability to obtain true beliefs via inference from true premises, James carried out the following deduction in Structures: (1) (p) the object ahead is a barn; (2) if the object ahead is a barn, then I am in County B; (3) (q) therefore, I am in County B.¹⁰ Considering that James's argument is both valid and sound, his capacity for deductive reasoning from true evidence is evident; as such, Greco's first condition for knowledge is satisfied. Furthermore, since James is not under the influence of any mind-altering substances when carrying out his inference from p to q , nor is there anything in his environment that could hinder his ability to reason properly, it follows that he possesses this knowledge-relevant ability relative to the environment he is in and the circumstances under which he carries out the deduction. Being appropriately situated for the propitious exercise of his knowledge-relevant ability fulfils Greco's second condition on knowledge with respect to q .

If one accepts that James's belief that q is the product of an intellectually virtuous process, then his belief that q is, according to Greco's epistemic framework, objectively justified. This is, as previously stated, one condition for knowledge that Greco mentions on occasion, though it is often collapsed into his virtue-centric account of knowledge, since a belief's being objectively justified is, according to Greco, entailed by its being virtuously produced. In addition to q being objectively justified, in order to establish that James knows that q , he must also be subjectively justified in believing that q . A subject's belief is subjectively justified if it is grounded in cognitive dispositions that they manifest when thinking conscientiously; this involves, according to Greco, a sensitivity to the reliability of one's evidence, and does not necessitate reflection on the part of the subject. Since Greco uses the term 'conscientious' to refer to individuals' default mode of thought – which involves trying to form one's beliefs accurately – absent any indication that James is engaging in cognitively dubious thought, we must assume that he is thinking conscientiously. Although James's belief that p is unknown, having taken all of the evidence available to him into account, he has no reason to consider the possibility that he might be looking at a barn façade. Furthermore, James understands that the truth of p entails that he is in County B, and he is also subjectively justified in believing that p , since his belief that p , while Gettiered, was formed in a conscientious manner. To this point, it is important to note that James does not simply believe that *if* p is true then q is true; rather, he (justifiably) believes that p is true, and bases his belief that q on inference from p .

Some might argue that one cannot be subjectively justified in believing a proposition inferred from a Gettiered premise because such inferences can easily lead to unreliable beliefs. While it may be the case that some inferences based on Gettiered beliefs are unreliable, this isn't necessarily true for all such inferences. For any single belief there is a set of propositions entailed by that belief; for example, the belief that 'I am writing on a computer' entails that I am not writing on a typewriter, that I am utilizing

¹⁰Some may argue that there are further premises – more or less tacit – necessary in order for James to competently infer that q . Even if that were the case, it would not make a difference to my argument, since an unknown proposition (i.e., p) would still play an essential role in the deductive inference.

a digital machine, etc. When it comes to inference based on a Gettiered belief of the Fake Barns variety, such an inference is reliable with regard to a set of propositions that are entailed by both the subject's Gettiered belief and the proposition that the conscientious subject would believe if they were told of the Gettierizing circumstances surrounding their original belief. In Structures, if James were told that the area he is in is littered with Barn façades and that there is only one barn present, he would have (being a conscientious individual) adjusted his belief that p in such a way as to account for this fact; for example, he may instead have believed that (p') either the object ahead is a barn or a barn façade. Clearly, there are several propositions entailed by p that are also entailed by p' , such as: the object ahead is man-made; the object ahead is shaped like a barn; and, the object ahead is a structure. If James inferred from p a proposition that isn't also entailed by p' , it is easy to see how his belief could have been false had he not luckily been looking at the only real barn around; however, since James's belief that q is entailed by both p and p' , his inference from p is reliable with respect to q , even though p is Gettiered.

The abovementioned considerations regarding James's inference from p to q support the notion that James acted conscientiously in coming to believe that q ; as such, James's belief that q is arguably not only objectively justified, but also subjectively justified. The only thing left to establish in order to ascribe knowledge of q to James is whether or not his belief is true *because* it was produced by the knowledge-relevant ability in question. According to Greco, "a success is attributable to S's ability just in case S's ability contributes to that success in the right way, where 'in the right way' means 'in a way that would regularly serve relevant purposes'" (Greco 2012: 14). With regard to deductive reasoning from true evidence, the relevant purpose of the ability is to validly deduce conclusions from sets of true premises. Considering that James believes q solely on the basis of competent deduction from a set of essential, true premises including p , it stands to reason that his belief is true *because* it was produced by his knowledge-relevant ability, rather than because of some non-relevant ability. Since James came to believe q via deductive reasoning from true premises, explanatory salience for his true belief is granted to his deductive reasoning capabilities; in other words, James possesses the true belief that q *because* it was produced by his relevant intellectual ability.

I have argued that James's belief that q meets Greco's agent-reliabilistic requirements for knowledge. It should be noted that his belief also meets some of the constraints that have been previously proposed as explanations for the lack of knowledge in various Gettier cases. For example, James's belief that q was not inferred from any falsehoods. As was demonstrated, while James does not know that p , since he is not appropriately situated for the exercise of his knowledge-relevant ability (i.e., the ability to identify real barns), the proposition is nonetheless true. Additionally, James's supporting belief that the presence of barns is indicative of being in County B is true. Given that there are no falsehoods involved in James's inferring that q , his belief that q meets the requirements proposed by the 'no false lemmas' condition for knowledge. Moreover, James's belief that q meets the constraints proposed by certain modal conditions for knowledge, such as safety and sensitivity. Generally stated, a subject's belief in a proposition is sensitive if, in the closest possible world where the proposition is false, the subject does not believe it;¹¹ accordingly, James's belief that q is sensitive since, in the closest possible world where he is not in County B, he would not have perceived a barn and therefore would not have inferred that q on the basis of his belief that the object ahead is a barn. Additionally, the safety condition for knowledge states that, in all nearby possible worlds where a subject believes a proposition on the same basis as the one used in

¹¹This is a version of Nozick's (1981) sensitivity condition.

the actual world, the proposition would not be false;¹² on this account, James's belief that q is safe, since there is no nearby possible world where he believes q on the basis of deducing it from p , and where q is false.

Although there are many variations of the sensitivity and safety conditions for knowledge, Greco notes that only a certain version of safety applies with respect to agent reliabilism, which he refers to as "(Ability + ProperConditions)-relative Safety"; according to this version of the condition, a belief is safe just in case "in close worlds where S believes p from ability A , and in conditions proper for the exercise of A , p is true" (Greco 2016: 54). This characterization of safety is motivated by the virtue-theoretic approach to knowledge, as intellectual virtues are regarded as dispositions to reliably believe truths when under conditions propitious for the exercise of the knowledge-relevant abilities. Even on this account of safety, James's belief that (q) he is in County B is safe, since in all close worlds where James believes that q on the basis of competent deductive inference from premises that constitute reliable evidence for q (which is a knowledge-relevant ability) while under proper conditions for the exercise of his reasoning faculties, q is true. After all, given that James knows that there are no barns or barn-like structures in County A, seeing what is, in fact, a barn is enough for James to know that he is in County B.

To summarize, Greco must concede that agent reliabilism allows for instances of Counter-Closure violation. As was demonstrated, Greco must accept that James does not know that p , even though p is true, since he lacks the knowledge-relevant ability to discern real barns from fake ones within the environment he finds himself in – that is, an environment in which real barns and barn façades are both present. Nonetheless, as I have argued, Greco must accept that James knows that q , as his belief that q meets Greco's conditions for knowledge – that is, James's true belief that q derived from a knowledge-relevant ability (i.e., his ability to form true beliefs on the basis of deduction from true premises), it is subjectively justified for him, and his belief that q is true *because* it was produced by the ability in question. If Greco grants that James knows that q , as well as that James inferred q from a set of true essential premises including the unknown belief that p , then he must reject the principle of Counter-Closure.

3. Virtuous inference objection

Some may argue that the pertinent purpose of the knowledge-relevant ability appealed to in the above thought experiment is not simply to validly deduce conclusions from sets of *true* premises, but rather to validly deduce conclusions from sets of *known* premises. Those who adopt this stance may contend that a subject's true belief derived via deductive reasoning from an unknown premise cannot count as knowledge, since the subject's ability to reason deductively did not contribute to their obtaining true belief *in the right way*; this would entail that the subject did not infer their conclusion virtuously, even though the same deductive inference would have been virtuous had they known their premise. Along these lines, Greco might regard deductive reasoning as a complex process involving multiple abilities, one of such being the ability to start from an appropriate premise – that is, the ability to judge whether one's premise is known and forbear from attempting to obtain knowledge via inference from an unknown premise. To this point, we would have to distinguish between simply reasoning virtuously, which is successfully completing the last step of an inference (i.e., drawing the conclusion), and forming one's inferential belief on the basis of a more complex virtue, which involves not only this last step, but also prior steps, including

¹²This is a version of Sosa's (1999) safety condition.

appropriately selecting one's premise(s). If we apply this standard of knowledge-relevant deductive inference to Structures, James would not be credited with knowing that q , as his belief that q was not formed on the basis of this complex virtue; while he may have competently deduced q from p , he does not know that p , and therefore did not act virtuously when carrying out his inference.

Although some may find this line of argument to be promising, it is implausible that knowledge-relevant deductive reasoning is intellectually virtuous only if one reasons from a known premise; indeed, there are countless real-world examples where one intuitively obtains knowledge via reasoning competently from an unknown premise. Take, for instance, an engineer, living before the time Einstein published his work on General Relativity, calculating the momentum of two objects travelling at the same speed; in carrying out their calculation, the engineer uses the formula $p = mv$ to determine that the object with the greater mass has the greater momentum. From this point forward, the engineer believes that momentum is a function of mass and velocity; on this basis, they quickly infer that, of objects travelling at the same constant velocity, the object with the greatest mass has the greatest momentum. This inference, I would argue, is virtuously carried out. Assuming that the engineer was rightly situated in order to effectively exercise their reasoning faculties, they inferred a true conclusion from a premise that is widely accepted as true in the scientific community, and the fact that they came to believe the conclusion on the basis of their reasoning ability explains why they have a true belief. However, once Einstein published his thesis on General Relativity, the scientific community changed the basic principle of momentum; the new formula multiplies an object's velocity by its invariant mass and the Lorentz factor, and is depicted as $p = \gamma m_0 v$. Since the engineer used the (now outdated) Newtonian formula for momentum as the basis of their inference, they apparently inferred a true conclusion from a false premise.

Should we conclude that the engineer's previous inference was non-virtuous, or even vicious, since they reasoned from a falsehood? I believe we should not, as they based their inference on a formula that was widely accepted, on the basis of good evidence, by the scientific community. If there is a complex ability involved in obtaining inferential knowledge, which includes the ability to virtuously choose which premise to reason from, it would appear that the engineer exercised such an ability; being a conscientious scientist, they reasoned on the basis of the best available evidence, which led to a true conclusion, even though their premise was false. To hold that virtuous inference requires that one knows one's premise would entail that all scientists calculating objects' momentum (and similar properties) prior to Einstein introducing his theory of General Relativity did not form their inferential beliefs virtuously, even if they reasoned competently from well-supported evidence. Rather than requiring that inferential reasoning involves, among other things, the ability to determine whether one knows a given premise in order to virtuously carry out an inference, it is more appropriate to require merely that one act conscientiously when reasoning inferentially – that one take into account all available evidence responsibly in determining what to believe as a premise. This stance is in line with Greco, who considers conscientiousness on behalf of a subject to be necessary in order for them to be subjectively justified in a belief.

Given this standard, we judge the pre-Relativism engineer as making a virtuous inference, even though they did so from a false premise; likewise, we should consider James's inference as virtuous, as he did act conscientiously in inferring that q from what appeared to him to be a barn. After all, the conditions were propitious for the reliable exercise of James's perceptual faculties (the only exception being with regard to differentiating between real barns and barn façades), he wasn't misjudging anything, and he took into account all of the evidence at his disposal.

Greco might object to the above characterization of virtuous inference, arguing that such inferences do not require that one's premise-beliefs be known, but rather, along the lines of Goldman's (1979) definition of conditional justification, that they were reliably formed, which is to say that they were produced by a knowledge-relevant ability.¹³ In response to this objection, I offer two points. First and foremost, even if Greco were to conceive of virtuous inference in this way, he would not be able to preserve Counter-Closure, as reliably formed premises are not necessarily known premises. If all that is required of one's premises in order to obtain inferential knowledge therefrom is that they be reliably formed, then it is possible to infer a proposition that meets all of Greco's requirements for knowledge from a premise-belief that is reliably formed yet unknown.

Take, for example, the abovementioned Chisholm's Sheep case, where the subject's belief that there is a sheep in the field is true and the product of a knowledge-relevant ability. In this case, their belief is reliably formed, yet they do not know that there is a sheep in the field because their belief is not true *as a result of* its being produced by a knowledge-relevant ability; rather, the salient reason for their belief's being true is that there is a sheep in the field behind the dog which they misidentified as a sheep. From their reliable, yet unknown, belief, imagine that the subject competently infers that there is at least one animal in the field. While their premise-belief is unknown, the inferred conclusion meets all of Greco's requirements for knowledge: it is true that there is at least one animal in the field; their belief was produced by a knowledge-relevant ability; and, their belief is true *as a result of* its being produced by the knowledge-relevant ability. Furthermore, even though the subject misidentified the dog as a sheep, the fact that they properly identified the object as an animal lends credence to the notion that the subject is sensitive to the reliability of their premise, at least with regard to propositions concerning the presence of animals in the field. If my analysis is accepted, then this case demonstrates that agent reliabilism allows for violations of Counter-Closure, even under the assumption that inferential knowledge requires that one's premise-belief(s) be reliably formed.

As a second point, I would argue that requiring that one's premise-belief(s) be reliably formed if one is to obtain inferential knowledge is, in any case, too strong a condition, as it leads to counterintuitive conclusions. While inference on the basis of an unreliably formed premise may lead one to falsehood, it may also lead one to reliably obtain a wide range of true beliefs. With regard to Structures, for example, although James's belief that he is looking at a barn is unreliably formed, given that he is a conscientious individual, he may reliably infer true beliefs from that premise, such as the belief that he is looking at a man-made object. Even though his premise-belief is not the product of a knowledge-relevant ability, it seems highly intuitive that he would know he is looking at a man-made object if he inferred as much from what he believes to be a barn. If this were the case, then James would be sensitive to the fact that what he perceives to be a barn is a reliable indication that he is looking at a man-made object.

Examples such as these abound, which bolsters the case for allowing that knowledge may be inferred from unreliably formed premises. This is not to suggest that unreliably formed premises can serve as the basis for acquiring inferential knowledge of any competently deduced proposition; however, it does suggest that, relative to one's evidence for an unreliably formed premise-belief, one may be able to obtain inferential knowledge of propositions in a certain domain from that premise. More work would have to be done to determine exactly which sets of propositions one can know via inference from an unreliably formed premise-belief. That being said, I suggest that a good starting

¹³I am grateful to an anonymous referee for pointing out this possible objection.

point would be to compare the evidence one has for an inferred conclusion with what they would believe, as a conscientious individual, if they were made aware of the fact that their premise was unreliably formed. If one would still hold their inferred belief knowing that their premise was unreliably formed, which would entail that they are sensitive to the fact that their evidence is a reliable indicator of the truth of their conclusion, and that belief further meets Greco's constraints on knowledge, then there is good reason to accept that the subject knows their conclusion.

4. Proxy-premise objection

In this section I am going to examine one potential objection to my argument that already features in Counter-Closure literature. This objection is often raised in response to alleged cases of knowledge from falsehood, and involves arguing that knowledge of one's conclusion in such cases is not based on inference from a falsehood, but rather on some true proposition in the neighbourhood of the falsehood (i.e., the proxy-premise).

Greco may attempt to preserve Counter-Closure by accepting that James knows that q and lacks knowledge of p , yet maintain that Structures does not exemplify an instance of Counter-Closure violation on the grounds that p is not essential to James's deduction. In support of this position, Greco might argue that James would still know that q if p were simply removed from his set of beliefs. In order for this to be the case, there must be some other known premise in the proximity of p that plays the epistemizing role in James's deductive inference. This move, which some refer to as the 'proxy-premise strategy', is widely debated in Counter-Closure literature, and was originally offered in response to Warfield's (2005) argument in favour of deductive knowledge from falsehood.¹⁴ With regard to cases in which a subject allegedly infers a known proposition q from a false proposition p , Ball and Blome-Tillman (2014) argue that the subject's belief that q is not based on p , but rather on some tacit (subconscious) belief which is known by the subject and is entailed by p . Similarly, Montminy (2014) argues that in such cases, if the subject knows that q , then their knowledge is due to some true proposition that is 'in the neighbourhood' of p , which the subject dispositionally believes and knows. While James does not come to believe that q via inference from any falsehoods, one may nonetheless attempt to employ the proxy-premise strategy to argue that Counter-Closure is not violated in Structures, or, for that matter, in any other case in which a known proposition is allegedly inferred from a Gettiered belief; in doing so, one must identify a known premise 'in the neighbourhood' of p which is entailed by p and can epistemize q , as well as provide sufficient reason to support the claim that the proxy-premise, rather than the Gettiered proposition, is essential to the subject's deductive inference.

With regard to James's belief that (p) the object ahead is a barn, one may suggest a suitable proxy-premise by invoking the following disjunction: (p^*) the object ahead is a barn or something that resembles a barn. As can be seen, p^* is clearly entailed by p , and, while James may not explicitly believe that p^* , it is reasonable to assume that he at least tacitly believes (and knows) that p^* . Furthermore, as q is entailed by the truth of either disjunct, q could serve as the conclusion of a competent deduction from p^* . There are undoubtedly many other suitable proxy-premises which can epistemize q (e.g., 'the object ahead is a man-made construction', 'the object ahead is a barn-shaped building', etc.); however, examining Structures in light of p^* is sufficient to demonstrate the inadequacy of the proxy-premise strategy with regard to preserving Counter-Closure in alleged cases of inferential knowledge from Gettiered belief.

¹⁴See Ball and Blome-Tillman (2014), Montminy (2014) and Schnee (2015).

The problem with the view that p^* , rather than p , plays the essential role in James's deductive inference is that if p^* is believed simply *because* it is entailed by p , then p would still play an essential role in James's coming to believe that q ,¹⁵ if this were the case, then James's inference would still violate Counter-Closure, as one of the essential premises in the deduction would be unknown. While it is true that p would be essential to James's deduction if p^* is believed *because* it is entailed by p , one might argue that p^* is actually believed because it is entailed by some other true proposition which is known, and which grounds both p^* and p .¹⁶ By claiming that James's belief that p^* is based on some known proposition other than p , one is able to avoid the objection that p is essential to James's deduction. According to this strategy, if James's knowledge that q is not based on inference from a Gettiered belief, then there must be some known proposition that evidentially supports both his belief that (p) the object ahead is a barn, as well as his tacit belief that (p^*) the object ahead is a barn or something that resembles a barn. A suitable supporting proposition may be that (t) the object ahead is a barn-shaped structure. Although James lacks the necessary knowledge-relevant ability to discern real barns from fake barns under the conditions specified in Structures, it would appear that he at least possesses the knowledge-relevant ability to recognize objects shaped as barns; if we grant as much, then it follows that James knows that t , since the proposition is true in such a way that manifests a knowledge-relevant ability.

While James's belief that t may be sufficient to demonstrate that p does not play an essential role in his deducing q from p^* , appealing to tacitly known proxy-premises when objecting to potential cases of inferential knowledge from non-knowledge prompts an important question – namely, can tacit beliefs play an active role in inferential reasoning? If Ball and Blome-Tillman (2014) are correct, then tacit beliefs may epistemize inferred conclusions; however, there is reason to doubt this position. In response to arguments in favour of a proxy-premise strategy, Buford and Cloos (2017) assert that inferentially active propositions satisfy the *taking condition*, according to which an act of inference “necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact” (Buford and Cloos 2017: 3); this condition is motivated by the notion that inferential reasoning is a person-level, agential process. If the *taking condition* is true, then it follows that tacit (subconscious) beliefs cannot play an essential role in inferential reasoning and are, therefore, inferentially inert, as these beliefs are not held in such a way that they can be “readily accessed or deployed in inference” (Buford and Cloos 2017: 5n6). Accordingly, even if James tacitly believes that p^* , the proposition is inferentially inert and cannot epistemize q , since he does not take p^* to support his belief that q , but rather believes that q specifically because he believes that p ; as such, p satisfies the *taking condition* and may therefore be regarded as an inferentially active proposition with regard to James's inference.¹⁷

¹⁵This objection is similar to that offered by Warfield (2005: 410) concerning the notion that, in alleged cases of inferential knowledge from falsehood, there is some true proposition in the vicinity of the subject's explicitly believed false proposition which functions as the basis for their knowledge of the inferred conclusion. He states that, in such cases, if the true proposition is believed *because* it is entailed by the false proposition, then the falsehood still plays an important epistemizing role.

¹⁶This explanation of why p is not essential to James's deduction is similar to that offered by Ball and Blome-Tillman (2014: 558) in response to Warfield's objection. They argue that their explanation of apparent cases of knowledge from falsehood avoids the problem raised by Warfield, as the proxy-premise in each case is not believed *because* it is entailed by the falsehood, but rather because it is entailed by some other true proposition which grounds the subject's explicit belief in the premise, as well as their belief in the proxy-premise.

¹⁷Some may consider the *taking condition*, as Buford and Cloos (2017) formulate it, to be too strong; they might argue, as Ball and Blome-Tillman (2014) do, that there are instances of subconscious inferential

Some may object to the *taking condition*, arguing that tacit beliefs may indeed play an active role in inferential reasoning; as Ball and Blome-Tillman note, much of our actual cognitive processes are automatic and subconscious (Ball and Blome-Tillman 2014: 558). But, even if tacit beliefs were able to play active roles in inferential reasoning, it does not necessarily entail the *impossibility* of the subject's coming to believe a proposition solely based on inferential reasoning from a Gettiered premise; in order to demonstrate this point, all one must do is show that there is a possible case in which a subject infers this way. If a subject is sincerely convinced of the truth of a proposition that happens to be Gettiered (on the basis of which they carry out a competent deduction), on what grounds should we regard the proposition as simply causally relevant, while ascribing epistemic relevance to some proxy-premise? As Luzzi (2014: 265) notes, "in the absence of a general account of the distinction between epistemically relevant and merely causally relevant premises", it appears question-begging to insist that in alleged cases of inferential knowledge from non-knowledge, explicitly believed premises are merely causally relevant, whereas in cases of inferential knowledge from known premises, the premises are considered both causally and epistemically relevant. Indeed, without a robust account of why one should consider proxy-premises to be epistemically relevant with regard to certain cases of deductive inference and not others, the proxy-premise strategy appears to simply be an ad hoc attempt to preserve Counter-Closure.

5. Conclusion

To summarize, the Structures case outlines a scenario in which the subject, James, infers the true proposition that (q) he is in County B from the unknown proposition that (p) the object ahead is a barn. From an agent reliabilistic perspective, although p is true, the proposition is nonetheless unknown, since James lacks the knowledge-relevant ability to discern real barns from fake ones under the conditions present while driving through an area littered with barn façades. Despite the fact that p is unknown, q is known, since James arrives at the true conclusion on the basis of competently deducing q from p in conditions propitious for the exercise of that knowledge-relevant ability, which entails that his belief that q is both objectively and subjectively justified. If these points are conceded, then Structures shows that Greco's agent reliabilism – one of the most prominent epistemic frameworks within the field of virtue epistemology – is incompatible with the widely endorsed principle of Counter-Closure. In light of this finding, Greco is faced with two options: either he must accept that knowledge may be obtained via inference from non-knowledge and abandon Counter-Closure, or he must revise his virtue-epistemic framework in such a way as to be compatible with the principle. While I am neutral with regard to whether this counts against Greco or Counter-Closure, my analysis of Structures nonetheless illuminates an interesting point that warrants taking note of.¹⁸

reasoning. While this may be the case, there is still reason to believe that the *taking condition* applies to instances of inferential reasoning where there is at least one explicitly believed premise. Indeed, when presented with cases in which a subject competently deduces a conclusion from a known premise, it is generally accepted that the subject took that premise to support their conclusion, rather than some proxy-premise.

¹⁸I would like to thank Federico Luzzi and Luca Moretti for the detailed feedback they provided on earlier drafts of this paper. I would also like to thank the anonymous reviewer for *Episteme* for their insightful and helpful comments.

References

- Armstrong D.** (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Arnold A.** (2011). 'Some Evidence is False.' *Australasian Journal of Philosophy* **91**(1), 165–72.
- Audi R.** (2010). *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. New York, NY: Routledge.
- Ball B. and Blome-Tillman M.** (2014). 'Counter-Closure and Knowledge Despite Falsehood.' *Philosophical Quarterly* **64**, 552–68.
- Buford C. and Cloos C.M.** (2017). 'A Dilemma for the Knowledge Despite Falsehood Strategy.' *Episteme* **15**(2), 166–82.
- Chisholm R.** (1989). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- Fitelson B.** (2010). 'Strengthening the Case for Knowledge from Falsehood.' *Analysis* **70**(4), 666–9.
- Gettier E.** (1963). 'Is Justified True Belief Knowledge?' *Analysis* **23**(6), 121–3.
- Goldman A.** (1976). 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy* **73**(20), 771–91.
- Goldman A.** (1979). 'What is Justified Belief?' In G. Pappas (ed.), *Justification and Knowledge*, pp. 1–23. Dordrecht: D. Reidel.
- Greco J.** (1999). 'Agent Reliabilism.' *Philosophical Perspectives* **13**, 273–96.
- Greco J.** (2002). 'Virtues in Epistemology.' In P. Moser (ed.), *The Oxford Handbook of Epistemology*, pp. 287–315. New York, NY: Oxford University Press.
- Greco J.** (2009). 'Knowledge and Success from Ability.' *Philosophical Studies* **142**(1), 17–26.
- Greco J.** (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Greco J.** (2012). 'A (Different) Virtue Epistemology.' *Philosophy and Phenomenological Research* **85**(1), 1–26.
- Greco J.** (2016). 'Knowledge, Virtue and Safety.' In M.A.F. Vargas (ed.), *Performance Epistemology: Foundations and Applications*, pp. 51–61. New York, NY: Oxford University Press.
- Hetherington S.** (1999). 'Knowing Failably.' *Journal of Philosophy* **96**(11), 565–87.
- Klein P.** (2008). 'Useful False Beliefs.' In Q. Smith (ed.), *Epistemology: New Essays*, pp. 25–62. Oxford: Oxford University Press.
- Luzzi F.** (2010). 'Counter-Closure.' *Australasian Journal of Philosophy* **88**(4), 673–83.
- Luzzi F.** (2014). 'What Does Knowledge-Yielding Deduction Require of its Premises?' *Episteme* **11**(3), 261–75.
- Montminy M.** (2014). 'Knowledge Despite Falsehood.' *Canadian Journal of Philosophy* **44**(3–4), 463–75.
- Murphy P.** (2013). 'Another Blow to Knowledge from Knowledge.' *Logos & Episteme* **4**(3), 311–17.
- Murphy P.** (2015). 'Justified Belief from Unjustified Belief.' *Pacific Philosophical Quarterly* **98**(4), 602–17.
- Nozick R.** (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Schnee I.** (2015). 'There is no Knowledge from Falsehood.' *Episteme* **12**(1), 53–74.
- Sosa E.** (1999). 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* **33**(13), 141–53.
- Sosa E.** (2007). *Apt Belief and Reflective Knowledge*. Oxford: Clarendon Press.
- Stanley J.** (2005). *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Turri J.** (2012). 'Is Knowledge Justified True Belief?' *Synthese* **184**(3), 247–59.
- Warfield T.** (2005). 'Knowledge from Falsehood.' *Philosophical Perspectives* **19**, 405–16.

K. Merrick Olivier is a PhD student at the University of Aberdeen. He is the recipient of the Elphinstone Scholarship for the study of Knowledge from Non-Knowledge and is conducting research under the supervision of Dr Federico Luzzi and Dr Luca Moretti. His research is primarily focused on whether prominent virtue-epistemic theories can preserve Counter-Closure.