

Solving the Moorean Puzzle*

Michael Blome-Tillmann

McGill University

University of Cambridge

Abstract

This article addresses and resolves an epistemological puzzle that has attracted much attention in the recent literature—namely, the puzzle arising from Moorean anti-sceptical reasoning and the phenomenon of transmission failure. The paper argues that an appealing account of Moorean reasoning can be given by distinguishing carefully between two subtly different ways of thinking about justification and evidence. Once the respective distinctions are in place we have a simple and straightforward way to model both the Wrightean position of transmission failure and the Moorean position of dogmatism. The approach developed in this article is, accordingly, ecumenical in that it allows us to embrace two positions that are widely considered to be incompatible. The paper further argues that the *Moorean Puzzle* can be resolved by noting the relevant distinctions and our insensitivity towards them: once we carefully tease apart the different senses of ‘justified’ and ‘evidence’ involved, the bewilderment caused by Moore’s anti-sceptical strategy subsides.

1. The Moorean Puzzle

To begin with, let us consider what I shall call the *Moorean Argument*:¹

The Moorean Argument:

- (i) I have hands.
- (ii) If I have hands, then I am not a handless brain in a vat.
- (iii) Therefore, I am not a handless brain in a vat.

Next, consider a widely accepted version of the closure principle for knowledge:

* For comments on earlier versions of this paper I am indebted to Yuval Avnur, Berit Brogaard, Philip Ebert, Thomas Grundmann, Esa Diaz Leon, Dustin Locke, Luca Moretti, Ram Neta, Jim Pryor, Robert J. Stephens, Martin Smith, Chris Tillman, Ralph Wedgwood, Crispin Wright, Masahiro Yamada, and audiences at the University of Manitoba and at a workshop at the Munich Centre for Mathematical Philosophy. I am especially grateful to two anonymous referees for this journal for providing extensive and valuable comments on an earlier version of the paper. Research for this paper was generously supported by the Social Sciences and Humanities Research Council of Canada and the Fonds Québécois de la Recherche sur la Société et la Culture.

¹ Moore’s (1939) original “proof” derives the existence of an external world from the existence of a hand. The difference to the example used here is irrelevant for present purposes.

Competent Deduction Closure (CDC):

If x knows p and bases her belief that q on a competent deduction from p , then x knows q .²

Given (CDC), it is tempting to think that we can use the logically valid inference from (i) and (ii) to (iii) to obtain knowledge that we are not brains in vats: as long as we know the premises (i) and (ii) of the *Moorean Argument* and as long as we base our belief in the conclusion on our competent derivation from those known premises, we come to know the argument's conclusion. However, many epistemologists will strongly object to this Moorean line of reasoning. That is, they will question its legitimacy by arguing that the Moorean line intuitively does not present a proper basis for one's belief in its conclusion. It is this intuition that constitutes what I shall call the *Moorean Puzzle*: on the one hand, the *Moorean Argument* is a clear case of a deductively valid argument and, as such, should be entirely suitable for coming to know its conclusion. On the other hand, it is undeniable that the argument at least seems unsuitable as a basis for one's belief in its conclusion. How can this puzzle be resolved?

Before looking in more detail at one of the most widely discussed accounts of the intuitions at hand, it is worthwhile noting that there are numerous inferences displaying a phenomenon very similar to that of the *Moorean Argument*. Here are just two familiar examples based on Dretske's (1970) *Zebra Case* and Cohen's (2002, p. 312) *Red Table* example:

Zebra Argument:

- (i) Those animals are zebras.
- (ii) If those animals are zebras then they are not cleverly painted mules.
- (iii) Therefore, those animals are not cleverly painted mules.

Red Table Argument:

- (i) That table is red.
- (ii) If that table is red, then it is not white with red lights shining on it.
- (iii) Therefore, that table is not white with red lights shining on it.

Given the similarities between these two arguments on the one hand and the *Moorean Argument* on the other, it is desirable that an explanation of why the *Moorean Argument* can provide a proper basis for our belief in its conclusion will extend

² Further complications are necessary, such as the condition that x retain her knowledge of the premises throughout the competent deduction of q from p . I shall ignore these subtleties for the sake of simplicity.

naturally to these additional examples. But how are we to account for the puzzle constituted by these examples?

2. Transmission Failure and Entitlement

To achieve a better understanding of the intricacies involved in giving a Moorean explanation of the above examples, let us consider in some detail a position recently defended by Crispin Wright (2002, 2003, 2004, 2008). According to Wright's view, the above inferences are not ways to properly base one's beliefs in their conclusions and thus they are not ways to come to know them. To see why Wright takes this *prima facie* drastic position we need to examine more closely his distinction between the *closure of justification* on the one hand and the *transmission of justification* on the other.

Consider first the notion of closure for justification, which we shall here define analogously to how we defined *Competent Deduction Closure* for knowledge (CDC) in the previous section:

Closure of Justification (CJ):

If x 's belief that p is justified and x bases her belief that q on a competent deduction from p , then x 's belief that q is justified.³

Note that (CJ) leaves open the question: *in virtue of what* is x 's belief that q justified? In particular, note that (CJ) does not claim that, if the conditions in the antecedent of (CJ) are satisfied, then x 's belief that q is justified *in virtue of* the conditions in its antecedent being satisfied. In this respect (CJ) differs from the following principle:

Transmission of Justification (TJ):

If x 's belief that p is justified and x bases her belief that q on a competent deduction from p , then x has thereby acquired (possibly for the first time) a justification for q .

When transmission is satisfied, but not necessarily when closure is, one *gains* a justification for the conclusion by competently performing the inference and forming one's belief that q accordingly. In other words, when justification transmits through a competent deduction of q from p , then one's belief that q is justified in virtue of one's

³ Wright uses the phrase 'has a warrant' rather than 'is justified'. The difference is, however, purely terminological. See, for instance, (Wright 2008, p. 30), where he says: "Let a warrant for a belief be, roughly, an all-things-considered mandate for it: to possess a warrant for p is to be in a state wherein it is, all things considered, epistemically appropriate to believe p ." Now, many will feel the urge to claim that, on this characterization, one possesses a warrant for p iff one knows p : knowledge is the norm of belief. Cp. (Williamson 2000).

competent inference from p : by performing the inference competently one has acquired a justification for its conclusion. (TJ) is accordingly the logically stronger principle of the two, for its consequent entails (CJ)'s consequent, but not *vice versa*.

With the distinction between closure and transmission in place, we can formulate Wright's position with regard to the *Moorean Argument*. According to Wright, the *Moorean Argument* is not—as Dretske (1970) and Nozick (1981) have it—a case of closure failure, but it is, instead, a case of transmission failure. In other words, the *Moorean Argument* cannot, Wright has it, supply us with a justification for its conclusion. In support of this view, Wright points out that one can only justifiably believe premise (i) of the *Moorean Argument* if one already has an *independent* or *antecedent* justification for the argument's conclusion.⁴ More precisely, Wright has it that our evidence for our belief that we have hands—that is, our perceptual experiences as of having hands—can provide a justification for the belief that we have hands only if we are already independently justified that we are not handless brains in vats that are fed misleading experiences as of having hands.⁵

A structurally similar type of dependence can, according to Wright, be observed with respect to a number of other examples, such as the abovementioned *Zebra Argument* and *Red Table Argument*. Intuitively, our evidence for our belief that the animal is a zebra—that is, our perceptual experiences of a black and white striped horse-like animal—can provide a justification for our belief that the animal is a zebra only if we are already independently justified to believe that the animal is not a cleverly disguised mule.⁶ Similarly, our evidence for our belief that a particular table is red—that is, according to Wright, our perceptual experiences as of a red table—can justify our belief that the table is in fact red only on the precondition that we are already independently justified to believe that the table is not white with red light shining on it: if we did not have such an independent justification for our belief that the table is not white with red light shining on it, then our belief that the table is red could not be justified by our perceptual experiences as of a red table.

⁴ (Wright 2002, p. 332).

⁵ See (Wright 2002, pp. 336-338).

⁶ Cp. (Wright 2010, p. 206): “the inference from (i) ‘The animals in that cage are zebras’ to (ii) ‘The animals in that cage are not cleverly disguised mules’, when the warrant for the first is given by casual observation too crude to distinguish the animals in vision from cleverly disguised mules, is arguably a failure of transmission: it is not that one can have a casual observational warrant for (i) but no warrant for (ii), but rather that it is only when (ii) (and a range of kindred propositions) are presupposed that casual observation warrants (i).”

Wright thus claims that in each of the above cases our justification for premise (i) is conditional on our possession of an independent justification for the respective argument's conclusion (iii). If this is so, however, then—Wright argues—transmission fails: by performing the respective inferences competently, we cannot *acquire* a justification for their conclusions and the arguments are therefore unsuitable for the proper basing of our beliefs in their conclusions. In other words, when performing the above inferences competently, we are not justified in believing their conclusions *in virtue of our competent deduction*, but rather in virtue of our antecedent and independent justification for their conclusions. Transmission fails, but closure does not.

Further aspects of Wright's view are worth mentioning at this point. In particular, note that combining the view that transmission fails in the *Moorean Argument* with the negation of scepticism has far-reaching consequences. Assuming, contra the sceptic, that we know ordinary propositions (henceforth '*op*')—and thus that we justifiably believe them—the defender of Wright's view is forced to accept that we have an independent justification for our belief in the negation of the sceptical hypothesis (henceforth '*¬sh*'). In fact, note that there is another important reason for the advocate of transmission failure to accept the view that we have an independent justification for *¬sh*. If transmission fails in the *Moorean Argument*, then, assuming that *Closure of Justification* (CJ) does not fail, it follows straightaway that we must have an independent justification for our belief that *¬sh*. Thus, both Wright's acceptance of (CJ) and his rejection of full-blown scepticism force him to accept that we have an independent justification for *¬sh*.

But on what grounds could we be justified in believing *¬sh*, if not on the basis of a competent derivation from our knowledge that *op*? Note first that it is not implausible to claim, in response to this question, that we have an independent justification for the conclusions of the *Zebra Argument* and the *Red Table Argument*. In fact, with respect to those arguments it can plausibly be argued that we are justified in virtue of possessing certain background knowledge: we know, for instance, that our local zoo would not deceive its visitors by displaying disguised mules in the zebra pen and we also know that the conditions in the furniture store next door are not such that the store manager aims to deceive us about an item's true colour by means of fancy lighting. However, even though these responses may seem appealing, a corresponding response is far less obvious in the case of the *Moorean Argument*. What could

constitute our independent justification or background knowledge for our belief that we are not brains in vats?⁷

Wright aims to answer this question by proposing a view that employs an idea going back to Wittgenstein's *On Certainty*. Ignoring exegetical subtleties,⁸ Wright claims that we have a justification for $\neg sh$ in virtue of being "epistemically entitled" to accept $\neg sh$ —where one can have such an epistemic entitlement for a proposition p without having any evidence whatsoever in support of p .⁹ Moreover, as Wright points out, we can have an epistemic entitlement for a rather restricted subclass of propositions only—namely, for those propositions that he refers to as "cornerstones" or "hinge propositions". Such 'cornerstones' or 'hinge propositions' play, on Wright's account, a special role in our cognitive lives: they are presuppositions of the "cognitive projects" that we, as a matter of fact, engage in, and without which we would be unable to pursue those projects. Here is a slightly more extensive quote from Wright:

Suppose there is a type of rational [justification] which one does not have to *do any specific evidential work* to earn: better, a type of rational [justification] whose possession does not require the existence of evidence—in the broadest sense, encompassing both *a priori* and empirical considerations—for the truth of the [justified] proposition. Call it *entitlement*. If I am entitled to accept p , then my doing so is beyond rational reproach even though I can point to no cognitive accomplishment in my life, whether empirical or *a priori*, inferential or non-inferential, whose upshot could reasonably be contended to be that I had come to know that p , or had succeeded in getting evidence justifying p . [...] Entitlements [are justified] without evidence. (Wright 2004, pp. 174-175; emphasis in original)

Entitlements are thus a type of non-evidential justification—a type of justification that one has, as Pryor (2004, p. 356) puts it aptly, "by default".¹⁰

⁷ Note that, if our evidence for the zebra conclusion and the red table conclusion supervenes on our perceptual experiences, which, by assumption, do not discriminate between the good and the bad case, then we do not have an independent justification in those cases. If, however, we allow for what I have, in the main text, called 'background knowledge' concerning the zoo or the furniture store, why, then, should we not also allow such propositional background knowledge with respect to the *Moorean Argument*? For instance, why should we not be allowed to infer from our knowledge that we are reading a philosophy paper right now that we are not merely brains in vats? For a more detailed argument along these lines see Section 4 of this paper.

⁸ Wright sometimes (see his 2004, p. 206) writes as if he assumed that we know $\neg sh$, while admitting in other places that his account amounts to a "sceptical solution" to sceptical puzzles, by which he means a solution on which "we do indeed have no claim to know, in any sense involving possession of evidence for their likely truth, that certain cornerstones [...] hold good." See (Jenkins 2007) for a careful interpretation and criticism of Wright's view.

⁹ (Wright 2004, p. 175).

¹⁰ (Pryor 2004, p. 356). Pryor (ibid., p. 372, n. 19) also notes that similar views are presented in "Cohen's (1999) and (2000), which claim that certain skeptical hypotheses are a priori irrational, so we're entitled to reject them without evidence."

With the notion of entitlement in place, Wright goes on to claim that while we do not have a justification for the conclusion of the *Moorean Argument* in the same way in which we have a justification for its first premise, our belief in the argument's conclusion has nevertheless a positive epistemic status: we are *epistemically entitled* to believe $\neg sh$, despite the fact that we have no evidence in support of that belief.¹¹ The notion of epistemic entitlement thus opens the door for the claim that the conclusion of the *Moorean Argument* is both known and justifiably believed, while it is nevertheless not the case that we can come to know $\neg sh$ *in virtue of* competently deducing $\neg sh$ from our knowledge that *op*. Our belief in the argument's conclusion, Wright has it, instead derives its positive epistemic status from the fact that we have an epistemic entitlement in support of it.

Let me sum up the presentation of Wright's account so far. As we have seen, Wright disagrees that the *Moorean Argument* can provide us with a justification for its conclusion and would therefore presumably also disagree that the argument can provide us with a way to properly base our belief that we are not brains in vats. On Wright's view, however, our belief in the argument's conclusion is nevertheless justified and plausibly even constitutes knowledge. However, it is crucially not justified, known, or properly based *in virtue of* our competent deduction from our knowledge that we have hands.¹² Rather, its justification derives from an entirely different, non-evidential source: from what Wright calls 'epistemic entitlement'.

While Wright's account may seem appealing at first glance, it is worthwhile noting that it faces some familiar difficulties. Consider first the notion of entitlement. As Carrie Jenkins (2007) and Duncan Pritchard (2007) have pointed out, it is not quite clear whether Wright's notion of entitlement is pragmatic rather than epistemic: if entitlement is, as Wright has it, a kind of non-evidential justification, does that commit him to the view that having an entitlement for *p* amounts eventually to nothing but the possession of a non-epistemic, merely pragmatic justification to believe *p*? Wright's remarks on the notion of entitlement at times suggest that the

¹¹ Wright is, in fact, more guarded, claiming that we only have entitlements to *accept* $\neg sh$, but not to *believe* $\neg sh$. I ignore the issue here, but see (Wright 2004, pp. 175-178) for discussion.

¹² Note again that Wright does not accept that our *belief* that $\neg sh$ is justified. He only claims that we are justified (entitled) to *accept* $\neg sh$. One has to wonder whether such a view can provide a satisfactory response to the sceptic. For Wright's view on this question see his (2004, pp. 175-178).

notion is in fact merely pragmatic.¹³ But if that should be correct, why, then, should we consider Wright’s account a solution to the sceptical puzzle? As Pritchard puts it pointedly, if the notion of entitlement is in fact merely pragmatic, then “[n]on-scepticism is [...] defended on the grounds that it is the practical alternative, but we knew *that* already.”¹⁴

Secondly, note that, despite his attempts to retain closure, Wright’s rejection of *Transmission of Justification* (TJ) also commits him, as Silins (2005, pp. 91-92) has pointed out, to the rejection of *Competent Deduction Closure* (CDC) as discussed in the previous section. Here is (CDC) again, repeated for convenience:

Competent Deduction Closure (CDC):

If x knows p and bases her belief that q on a competent deduction from p , then x knows q .

Remember that, on Wright’s account, an inference from op to $\neg sh$ cannot justify one’s belief that $\neg sh$. Thus, on Wright’s account, one’s belief that $\neg sh$ will be improperly based, if based exclusively on a competent deduction of $\neg sh$ from op . If this is so, however, then (CDC) must fail—despite the fact that we have an entitlement for $\neg sh$: surely, the mere possession of an entitlement cannot turn an improperly based belief into knowledge, assuming that no improperly based belief can ever be knowledge. It is, as a consequence, by no means obvious whether Wright can retain (CDC) while at the same time rejecting (TJ).

Thirdly and finally, note that the view that transmission fails in the mentioned arguments is incompatible with a number of rather plausible assumptions about the nature of both justification and evidence.¹⁵ Consider the following principles:

Knowledge as Evidence ($K \subset E$):

If x knows p , then p is a part of x ’s evidence.¹⁶

Entailment as Evidential Support (EES):

If p entails q , then any body of evidence that comprises p provides evidential support for q .¹⁷

¹³ See, for instance, Wright’s (2004, p. 183) explication of what he calls “strategic entitlement” in terms of game theoretically dominant strategies.

¹⁴ (Pritchard 2007, p. 207; emphasis in original). For an interesting response to Pritchard’s challenge see (Pedersen 2009).

¹⁵ Cp. (Silins 2005, p. 88) and (Klein 1981).

¹⁶ This principle is inspired by (Williamson 2000, ch. 9). Williamson defends the stronger view that $E = K$ (see also below (Section 4)).

It is easy to see why these principles spell trouble for the above view: Wright accepts that we know premise (i) of the *Moorean Argument*, but rejects the view that our evidence supports (iii)—the argument’s conclusion. But if (i) is known and thus part of our evidence, why are we not in a position to acquire a justification for (iii)? After all, if (i) is known and therefore part of our evidence, then our evidence both entails and, given (EES), supports (iii). Clearly, Wright has to reject either $(K \subset E)$ or (EES). However, both of these principles are certainly rather plausible.¹⁸

As we have seen, Wright’s account of the *Moorean Argument* is not entirely unproblematic. However, it is worthwhile emphasizing that, despite the view’s difficulties, it accurately captures our intuition that the arguments at issue differ in an epistemically significant way from more ordinary deductive inferences. More specifically, we ought not to lose sight of the fact that Wright can account for the intuition that the *Moorean Argument* underlies some kind of epistemic circularity and that the argument is, as Pryor (2004, p. 365) has pointed out, unsuitable to “rationally overcome one’s doubt in its conclusion.” We are, as a consequence, facing a puzzle: on the one hand the idea that transmission fails with respect to the arguments under consideration seems attractive and is backed up by a cluster of intuitions that I shall refer to as the *transmission failure intuitions*. On the other hand, the idea that transmission fails in the cases at hand stands not only in direct conflict with exceedingly plausible assumptions about the nature of evidence and justification, but also with the rather intuitive idea that knowledge is closed under competent deduction. In other words, the idea that transmission of justification fails in the cases at hand conflicts with precisely those assumptions and intuitions whose naturalness is the driving force behind the Moorean response to the sceptic.

How are we to resolve this conflict? In what follows, I shall argue that we can account for our transmission failure intuitions within the framework of *Relevant*

¹⁷ A referee for this journal points out that (EES) is implausible, as it entails both that p evidentially supports itself and that the conjunction $(p \wedge q)$ evidentially supports p . I am not too worried about these consequences, but it should be noted that the issue can be averted by reformulating (EES) in terms of an appropriately defined notion of *non-circular evidential support* (for further discussion of circular evidential support, see (Williamson 2000, p. 187)).

¹⁸ See (Williamson 2000, ch. 9) for arguments in support of $(K \subset E)$. Note also that, for the purposes of the argument here, we might retreat to $(K \subset E)$ ’s weaker cousin $(KK \subset E)$ —that is, the principle according to which all propositions that one knows to know are part of one’s evidence. To see why this weaker principle would also spell trouble for Wright, note that, on the assumption that we know that we know that we have hands, the proposition that we have hands would be part of our evidence, and could thus justify or serve as a proper basis for our belief that we are not handless brains in vats.

Alternatives Theory (henceforth ‘RAT’), and that we can do so in a way that allows us to uphold (K \subset E), (EES), (CDC), and Wright’s claims about the failure of transmission in the *Moorean Argument*. Thus, RAT can account for the phenomenon of transmission failure while at the same time accepting that the *Moorean Argument* can provide us with a proper basis for our belief that $\neg sh$: we can come to know $\neg sh$ by competently deducing it from *op*. Once we accept the version of RAT outlined below we can resolve the puzzle of Moorean reasoning in an elegant and straightforward manner.¹⁹

3. RAT and Transmission Failure

In recent years, the discussion of *Relevant Alternatives* theories has become rather stale. However, we shall see here that an RA-framework that makes exceedingly minimalist and—I presume—uncontroversial assumptions about the nature of knowledge, allows us to give an interesting resolution of what I have called the *Moorean Puzzle*. I therefore take the following arguments to not only provide a motivation for the view defended here, but also to provide good reasons for epistemologists to reconsider their sometimes sceptical or even outright hostile attitude towards versions of RA that have clear explanatory virtues.

Let us begin with the following formulation of what I shall call the *Relevant Alternatives Theory* of knowledge:

Relevant Alternatives Theory (RAT):

If *x* knows *p*, then *x*’s evidence eliminates all epistemically relevant $\neg p$ -worlds.

Note that (RAT) formulates merely a *necessary* condition for knowledge—it is not a biconditional, and therefore does not provide us with an analysis or definition of knowledge. Consequently, (RAT) should be significantly less controversial than the stronger, biconditional versions of the view.²⁰ Moreover, note that, in addition to (RAT), I shall, in this article, interpret the notion of evidence employed in the above principle as coinciding with David Lewis’s (1996, p. 553) notion of evidence, which is defined as follows:

¹⁹ I do not mean to suggest that contemporary Mooreans dogmatists have no story to tell about the transmission failure intuitions. Those accounts are just not the topic of this paper.

²⁰ The biconditional cousin of (RAT) would fail because it would not place any constraints on belief formation and sustenance: unreliably based beliefs could count as knowledge.

Lewisian Evidence:

x 's evidence is the sum total of x 's perceptual experiences and memory states.

I take it that (RAT) thus defined and explicated expresses a triviality about the modal structure of knowledge: if one knows p , then one's evidence eliminates all epistemically relevant $\neg p$ -worlds.²¹ However, before putting these two principles to work, some comments and clarifications concerning (RAT) are in order.

To begin with, note that *Lewisian Evidence* is incompatible with *Knowledge as Evidence* ($K \subset E$) as mentioned above. Even though potentially problematic, let me ignore this apparent difficulty for the moment and return to the topic below, where I shall argue that the two principles place constraints on two independent, but philosophically equally legitimate notions of evidence, each of which must play a role in explaining the phenomenon of transmission failure. Next, note that some relevant alternative theorists, such as Dretske's (1970), have held the view that we know op because our evidence eliminates all relevant alternatives to op , while at the same time holding that we do not know $\neg sh$, because our evidence fails to eliminate all relevant alternatives to $\neg sh$. Dretske's view therefore—rather controversially—entails the failure of closure principles such as *Competent Deduction Closure*.

It is important to note that one of the implicit assumptions underlying Dretske's view is that there are some sh -worlds that are relevant alternatives to $\neg sh$ but not to op . This assumption can be rejected. In particular, we can think of epistemic relevance as a property that is not relativized to propositions. Worlds are, on such a view, epistemically relevant *simpliciter*: if a world w is epistemically relevant, it is a relevant alternative to any proposition p that is false in w . Given such a non-relativized, Lewisian (1996) notion of epistemic relevance, Dretske's claim that we do not know $\neg sh$ because our evidence does not eliminate all relevant alternatives to $\neg sh$ no longer follows.²²

²¹ Again, note that this claim is not meant to offer a definition or an analysis of knowledge. Rather, it merely records the obvious fact that if one knows p , one can eliminate certain counterpossibilities to p .

²² The problem with talking about propositions as alternatives is that if we take $\neg op$, for instance, to be a relevant alternative to op , then it becomes unclear which exact $\neg op$ -worlds our evidence has to eliminate: all of them? Surely not, for that would lead to scepticism. Only some of them? Yes, but which ones? Surely, the only clear answer here is that only the epistemically relevant *counterpossibilities* must be eliminated. But that is precisely what I have called a 'Lewisian' conception of relevant alternatives in the main text. Thus, Lewis's formulation of *Relevant Alternatives Theory* is more precise and clearer than Dretske's, for Lewis also provides a clear account of evidence

To see this in more detail, note that if we are to know *op*, then it follows directly from (RAT) that our evidence eliminates all epistemically relevant $\neg op$ -worlds. But if all epistemically relevant $\neg op$ -worlds are eliminated, then, since *sh* entails $\neg op$ (and since epistemic relevance is now assumed to be non-relativized), all epistemically relevant *sh*-worlds must *ipso facto* be eliminated, too.²³ Thus, the relevant alternatives condition in (RAT) doesn't present a threat to closure, as it does on Dretske's version of the theory. To repeat: if we interpret the notion of a relevant alternative along Lewisian lines—that is, as relevance *simpliciter* or as non-relativized to propositions—then our knowing *op* entails that no *sh*-world whatsoever is epistemically relevant, for no *sh*-world whatsoever is eliminated by our evidence: if any *sh*-world were relevant, we wouldn't know *op*. As a consequence, it also follows, firstly, that our evidence eliminates all epistemically relevant $\neg sh$ -worlds and thus, secondly, that (RAT) as formulated above doesn't commit us to closure failure. To the contrary, (RAT) explicates a condition on knowledge that is itself closed under entailment.

Gail Stine (1976) has argued for a similar conclusion along similar lines. Stine proposes a version of RAT while at the same time salvaging closure—and she does so by accepting the view that we can sometimes know a proposition *p* despite the fact that our evidence doesn't eliminate any alternative to *p*. As Stine (1976, p. 258) herself puts it, “if the negation of a proposition is not a relevant alternative, then I know it—obviously, without needing to provide evidence.” I shall agree with Stine in this paper that we know $\neg sh$, despite the fact that—in the terminology of (RAT)—our Lewisian evidence doesn't eliminate any *sh*-world whatsoever. And it should be obvious by now that this position is entirely compatible with (RAT) as formulated at the beginning of this section: since no *sh*-world whatsoever is epistemically relevant, our Lewisian evidence eliminates every epistemically relevant *sh*-world. The condition in (RAT) doesn't stand in our way of knowing $\neg sh$, nor is it incompatible with closure.

Let us move on from the discussion of (RAT) and closure, while bearing in mind that, on the view adopted here, we know *op*, which we in turn saw to entail that no *sh*-world is epistemically relevant. With these fairly innocuous assumptions in place let

and the elimination of a possibility by one's evidence. I ignore the fact that Lewis's version of the view is contextualist.

²³ If *sh* entails $\neg op$, then the *sh*-worlds form a subset of the $\neg op$ -worlds, and thus, the epistemically relevant *sh*-worlds form a subset of the epistemically relevant $\neg op$ -worlds.

us return to the topic of Mooreanism. As I shall argue in what follows, the conceptual resources provided by (RAT) and the assumption that we know *op* allow us to give a surprisingly simple and straightforward way to model the notion of epistemic entitlement as discussed in the previous section. Consider the following definition:

Epistemic Entitlement (ENT):

x has an epistemic entitlement to believe *p* iff all $\neg p$ -worlds are epistemically irrelevant.

Note that according to (ENT), whether one has an epistemic entitlement for *p* depends exclusively on what is and what is not an epistemically relevant alternative to *p*. The proposed notion of entitlement, therefore, has a straightforwardly epistemic dimension, for if one has an entitlement to believe *p*, then one's evidence does not have to eliminate any $\neg p$ -worlds in order for one to know *p*. On the above definition we therefore have a straightforward response to Jenkins and Pritchard's worry, mentioned in Section 2, that the notion of an entitlement is only pragmatic in nature, while not having an interesting or non-trivial epistemic dimension.²⁴

Before going into more detail with respect to (ENT) and its relation to the *Moorean Argument*, however, let us first consider the definition of what I shall call 'relevant alternatives justification' or, for short, 'RA-justification':

RA-Justification (RAJ):

x has an RA-justification for *p* iff:

1. some $\neg p$ -worlds are epistemically relevant and
2. all $\neg p$ -worlds that are epistemically relevant are eliminated by *x*'s evidence.

According to this definition, one only has an RA-justification for *p* if one's evidence does some real epistemic work with respect to *p*: one's evidence must eliminate the counterpossibilities to *p* that are in fact relevant. In this respect RA-justification differs from entitlement, which one can have, as we saw Pryor put it previously, entirely "by default".

Next, note that the above two definitions, (ENT) and (RAJ), allow us to explain all of the Wrightean desiderata about entitlement and transmission failure discussed in the previous section. To begin with, note that it follows from the above definitions that the possession of an epistemic entitlement for *p* and the possession of an RA-

²⁴ For a different approach to entitlement see (Smith 2012). Smith models the notion of entitlement in terms of possible worlds, too, but does so via an analysis of the epistemic support relation in terms of Lewis's notion of a variably strict conditional.

justification for p are mutually exclusive: if one has an entitlement for p , then all $\neg p$ -worlds are epistemically irrelevant, which is incompatible with the possession of an RA-justification. And if one has an RA-justification, then there are some $\neg p$ -worlds that are epistemically relevant, which, in turn, is incompatible with the possession of an epistemic entitlement for p . Moreover, note that knowing p does not entail the possession of RA-justification for p : rather, it only entails the disjunction of RA-justification and entitlement.²⁵ According to RAT, we can know p either in virtue of having an RA-justification for p or in virtue of having an entitlement for p . Finally, note that, given the stipulative nature of the definitions (ENT) and (RAJ), all of the above claims are theorems of RAT: they are entailed by (RAT) in conjunction with (ENT) and (RAJ).

With the above definitions in place let us return to the discussion of the *Moorean Argument*. First, note that (ENT) gives us an elegant and intuitive explanation of Wright’s observation that it is “within our right” to accept $\neg sh$ despite the fact that our evidence does not support $\neg sh$: given the above definitions, we are entitled to accept $\neg sh$ because all sh -worlds are epistemically irrelevant. Given (RAT), there is a clear and precise sense in which we are entitled to accept $\neg sh$, despite the fact that no sh -world whatsoever is eliminated by our evidence: we know $\neg sh$, despite the fact that, given (RAJ), we do not have an RA-justification for $\neg sh$.²⁶ Thus, our above definitions capture exactly Wright’s intuition that we are entitled to accept $\neg sh$ without evidence or justification, and they do so by employing nothing but the familiar and independently motivated conceptual apparatus provided by RAT in conjunction with my definitions of entitlement and RA-justification.

Having dealt with our epistemic stance towards the conclusion of the *Moorean Argument*, let us next turn to the argument’s first premise. To begin with, note that there are some $\neg op$ -worlds that are epistemically relevant—such as, for instance, nearby worlds in which we have lost our hands in a terrible car accident.²⁷ Note

²⁵ Note that an exception will be necessary here for certain disjunctive necessary truths: if one has an RA-justification for p , and q is necessarily true, then one also has an RA-justification, rather than an entitlement, for the disjunctive proposition $(p \vee q)$, despite the fact that there is no world in which $\neg(p \vee q)$, and thus no epistemically relevant world in which $\neg(p \vee q)$.

²⁶ We still need to explain that our belief that $\neg sh$ is properly based. That explanation will be given in the following section.

²⁷ I assume that such worlds are epistemically relevant because they are rather close to our actuality. I shall not engage in an explication of the notion of epistemic relevance in this paper, but rather presuppose an intuitive grasp of the notion. Cp. also Section 5.

further that these epistemically relevant $\neg op$ -worlds are eliminated by our evidence: in the closest worlds in which we have lost our hands in a terrible car accident we have different sensory experiences—namely, sensory experiences as of not having hands. We therefore have, given (RAJ), an RA-justification for the first premise of the *Moorean Argument*. Moreover, note that we do not have, given the definition (ENT), an entitlement for premise (i). This is so because there are, as just mentioned, some $\neg op$ -worlds that are epistemically relevant. Again, the above definitions provide us with a clear and straightforward account of Wright’s claims concerning the epistemic standing of our belief that *op* while employing nothing but the conceptual machinery provided by RAT and our definitions (ENT) and (RAJ).

So far we have seen that we can model Wright’s claims that we have an entitlement for $\neg sh$ but not for *op*, and that we have an evidential justification for *op* but not for $\neg sh$, in a straightforward way within the framework of RAT. This leaves us with the task of accounting for Wright’s further claim that transmission of justification fails in the *Moorean Argument*. Here is how to account for that further claim. First, note that by competently performing the deduction of $\neg sh$ from *op* one cannot acquire an RA-justification for $\neg sh$. This is so because one’s Lewisian evidence—the totality of one’s perceptual states and memory states—does not eliminate any *sh*-world whatsoever, independently of whether or not we have competently derived $\neg sh$ from our RA-justified belief that *op*. Consequently, the epistemic status of being RA-justified doesn’t *transmit*, in the case of *Moorean Argument*, through competent deduction. Moreover, note that RA-justification is not even closed under competent deduction: as can be seen right away, closure for RA-justification fails in the case of the *Moorean Argument* and its cousins. Summing up, the Wrightean intuition that transmission fails in the *Moorean Argument* is rather elegantly explained by RAT and the above definitions.²⁸

Further Wrightean desiderata remain to be explained. Remember that, according to Wright, what justifies our belief in premise (i) of the *Moorean Argument* cannot also justify our belief in its conclusion. The above account has again a straightforward explanation of this desideratum, for what grounds our RA-justification for premise (i)—that is, our evidence or, as Lewis has it, our experiences and memory states—

²⁸ It might be thought that it is implausible that a notion of epistemic justification—namely, that of RA-justification—is not closed under competent deduction. This worry will be addressed below in Section 4.

cannot also provide us with an RA-justification for the argument's conclusion: while our Lewisian evidence eliminates all epistemically relevant $\neg op$ -worlds, it doesn't eliminate any sh -world whatsoever, and so it couldn't provide us with an RA-justification for $\neg sh$, if some $\neg sh$ -worlds were epistemically relevant. Thus, what in fact RA-justifies our belief that op couldn't RA-justify our belief that $\neg sh$. The above definitions again deliver the exact result required by Wright's account.

Next, consider the intuition that our justification for premise (i) of the *Moorean Argument* is conditional or dependent on our possession of an entitlement for the argument's conclusion: as Wright puts it, (i) *presupposes* (iii). Given RAT and the above definitions, this intuition turns out accurate, too, for, according to the definitions (ENT) and (RAJ), we are RA-justified in believing op only if we have an entitlement for $\neg sh$. To see this, note that if we did not have an entitlement for $\neg sh$, then, by definition, some sh -worlds would be epistemically relevant. But since our Lewisian evidence does not, by stipulation, eliminate any sh -world whatsoever, our Lewisian evidence would not, in the envisaged situation, eliminate all epistemically relevant $\neg op$ -worlds: it would not eliminate those epistemically relevant $\neg op$ -worlds that are sh -worlds. As a consequence, if we did not have an entitlement for $\neg sh$, we couldn't have a RA-justification for op : Wright's dependence claim turns out to be true and can be elegantly accounted for within the framework of RAT.

Finally, note that I have so far applied the definitions (ENT) and (RAJ) to the *Moorean Argument* only, but not to the *Zebra Argument* and the *Red Table Argument*. However, analogous considerations apply to the latter cases. To illustrate this briefly for the *Zebra Argument*, note that, by assumption, our evidence—that is, the sum total of our perceptual experiences and memory states—does not eliminate the closest worlds in which the animals in the pen are cleverly painted mules. Nevertheless, we are, according to (ENT), entitled to believe that the animals are not cleverly painted mules: the closest worlds in which the animals are cleverly painted mules are epistemically irrelevant. Furthermore, note that from the fact that we have an entitlement for the conclusion of the *Zebra Argument*, it follows that we do not have an RA-justification for that proposition. But since we have such an RA-justification for the *Zebra Argument's* first premise—our evidence eliminates, for instance, the epistemically relevant possibilities that the animals are elephants or emus—it follows that RA-justification does not transmit through the zebra inference. An analogous

explanation can be given for the *Red Table Argument*, but I shall spare the reader the details.

Before moving on, let me address an aspect of RA-justification that might seem problematic at first sight.²⁹ To see what I have in mind, we need to make a few substantive assumptions about epistemic relevance. Consider Lewis's (1996, pp. 554-555) *Rule of Actuality* according to which the subject's actuality is always epistemically relevant. Note further that, according to the Lewisian definition of evidence, one's actuality is always uneliminated. If one's actuality is always relevant and uneliminated, however, it follows that one's belief that p is RA-justified only if p is true: RA-justification is factive.³⁰ Many theorists, however, will object at this point that this consequence casts doubt on whether RA-justification really is a legitimate notion of epistemic justification, for the idea that one can have epistemic justification for falsehoods is widely regarded as a platitude. In response, it should be noted that RA-justification is nothing but a limiting case of what we may call 'RA*-justification', which is defined as follows:

RA-Justification (RA*J):*

x has an RA*-justification for p iff:

1. some $\neg p$ -worlds are epistemically relevant and
2. some $\neg p$ -worlds that are epistemically relevant are eliminated by x 's evidence.

RA*-justification comes in degrees, depending on how large a subset of the epistemically relevant $\neg p$ -worlds are eliminated by one's evidence: the larger the subset, the greater the degree of one's RA*-justification, with RA-justification as the limiting case in which all relevant $\neg p$ -worlds are eliminated. Moreover, note that RA*-justification is not factive. To illustrate this, imagine that you believe the proposition (p) that Hannah isn't in the office today. You have this belief because you haven't seen Hannah's coat on the coat rack, where it usually is when Hannah is in. In this scenario, your Lewisian evidence eliminates some relevant $\neg p$ -worlds—namely, those in which Hannah is in the office *and* has left her coat out on the coat rack—and you are, as a consequence RA*-justified that p . Imagine further, however, that p is in fact false: Hannah is hiding in her office, and didn't put her coat on the

²⁹ I am indebted to the referees for *Philosophical Studies* for bringing my attention to this issue.

³⁰ If one falsely believes that p , then one's actuality is a $\neg p$ -world. And since one's actuality is always relevant and uneliminated, it follows that, if one falsely believes p , there is one $\neg p$ -world that is relevant and uneliminated: one's actuality. Consequently, one doesn't have RA-justified for p , if p is false.

coat rack today because she intends to skip the department meeting in the afternoon. In this case, your evidence eliminates some relevant $\neg p$ -worlds, but not all of them. More importantly, your Lewisian evidence fails, in the case described, to eliminate actuality, which is a $\neg p$ -world. Thus, in the example just explicated, your false belief that p is RA*-justified and plausibly even RA*-justified to a fairly high degree.

Given that RA-justification as defined previously is nothing but a limiting case of RA*-justification, the fact that RA-justification is factive should be unproblematic. An analogy helps illustrate this point: On certain widely accepted evidentialist accounts of justification, what we may call ‘full propositional justification’ is factive, too, for on a probabilistic interpretation of evidentialism, one has full propositional justification iff $P(p|e) = 1$, and thus only if p is true. Moreover, on probabilistic interpretations of evidentialism full propositional justification is nothing but a limiting case of propositional justification, which doesn’t require factivity and can be had even if $P(p|e) < 1$. Thus, the fact that RA-justification as defined above is factive doesn’t by itself disqualify the notion as a plausible notion of epistemic justification. One might wonder, however, whether only RA-justification and not RA*-justification would fail to transmit through the Moorean Argument. For if RA-justification is nothing but a limiting case of RA*-justification, shouldn’t the latter fail to transmit through the argument, too? Interestingly, RA*-justification fails to transmit through the Moorean Argument as well: since our Lewisian evidence doesn’t eliminate any sh -world whatsoever, we are not RA*-justified in believing $\neg sh$, and an explanation of the abovementioned phenomena concerning transmission failure for RA*-justification can be given that is analogous to the one given for RA-justification in the previous paragraphs.³¹

Let me sum up the discussion so far. As we have seen, (RAT) in conjunction with the definitions (ENT) and (RAJ) provides a simple but powerful and elegant explanation of both our transmission failure intuitions and Wright’s notion of epistemic entitlement. Moreover, the account developed here is formulated entirely within the independently motivated framework of RAT. The explanations outlined

³¹ Note, however, that having an RA*-justification that op doesn’t presuppose or entail having an entitlement that $\neg sh$. That particular Wrightean intuition must be modelled in terms of RA-justification.

here therefore present an interesting and fruitful application of RAT to the *Moorean Puzzle* and the problem of transmission failure.³²

4. RAT and Moorean Dogmatism

It will no doubt seem odd to some theorists that, as I have suggested so far and as Stine (1976, pp. 257-258) has argued in great detail, we can know some propositions without having any evidence whatsoever in their support. On the face of it this is a legitimate and familiar challenge to RAT.³³ But I shall argue in this section that RAT can be interpreted in a way that avoids commitment to the implausible idea that there could be knowledge that p without there being evidence in support of p . To see what I have in mind note that RAT and the account of transmission failure developed in the previous section entail that we can sometimes know a proposition p despite the fact that our *Lewisian* evidence—that is, the totality of our perceptual experiences and memory states—does not eliminate any $\neg p$ -world. It is important to emphasize at this point, however, that Lewis’s notion of evidence as employed in (RAT) is technical and that it is not at all obvious that this technical notion of evidence can be readily equated with the more everyday notion of evidence familiar from scientific and legal discourse and television programmes such as *CSI* and *Law and Order*. In fact, I shall assume in this article that on our more everyday understanding of evidence, evidence is at least sometimes propositional and, moreover, that all propositions that one knows are part of one’s evidence. In other words, I shall assume principle (K \subset E) from Section 2, repeated here for convenience:

Knowledge as Evidence (K \subset E):

If x knows p , then p is a part of x ’s evidence.

The idea that one’s knowledge is part of one’s evidence is inspired by Williamson’s (2000) work, who defends the stronger view that $E = K$:

³² Wright (2004, p. 205) offers a somewhat baroque account of epistemic entitlement, distinguishing in a seemingly *ad hoc* way between what he calls *strategic entitlements*, *entitlements of cognitive project*, *entitlements of rational deliberation*, and *entitlement of substance*. In fact, Wright himself is quite guarded with respect to this account when admitting, for instance, that his “discussion [...] is bound to leave many loose ends” and should be understood as merely an attempt “to outline a *prima facie* case for a number of different possible species of entitlement.” (Wright 2004, p. 175). Wright also (ibid.) points out with respect to his notion of entitlement of substance that he has merely “gestured, in the most promissory and indefinite way, at the possibility of—and need for—[the notion of] entitlement of substance.”

³³ See, for instance, (Cohen 1988, p. 99).

$E = K$:
 p is part of x 's evidence iff x knows p .

I am rather sympathetic to the idea that $E = K$. But given the controversy with which Williamson's view has been met in the literature, I shall steer clear of it in this article, for purely practical reasons: everything that I aim to argue for in what follows is entailed by the considerably weaker and, I take it, exceedingly plausible claim that $K \subset E$.³⁴

Let us thus return to the topic of the *Moorean Argument* and the idea that one can sometimes know a proposition p without having evidence in support of p . Given the assumption that $K \subset E$, we have evidence in support of the propositions that we competently deduce from what we know: if one competently deduces p from a proposition q that one knows, then p is entailed by what one knows, and thus by one's evidence. Applying this to the case of the *Moorean Argument*, it follows that we do have, after all, evidence in support of $\neg sh$ —namely, our knowledge that op .

To see this in more detail, note that with a notion of evidence in hand that is governed by the assumption that $K \subset E$, we can easily honour the intuitive idea that all knowledge is based on evidence. This is so because, given (RAT), our belief that $\neg sh$ turns out to be supported by our evidence, for our evidence comprises op , which is known. And since op entails $\neg sh$, our evidence supports $\neg sh$. Given (RAT) and the idea of knowledge as evidence ($K \subset E$), we can, accordingly, make clear and precise sense of the intuition that, in the case of the *Moorean Argument*, we have evidence in support of what we know. Moreover, note that, given the abovementioned constraint on evidence, we can easily define a notion of epistemic justification according to which we are epistemically justified in believing $\neg sh$. Consider the following evidentialist definition of justification that contrasts sharply with the notion of RA-justification from the previous section:

Propositional E-Justification (PEJ):
 x is evidentially justified to believe p iff x 's evidence makes p sufficiently likely.³⁵

³⁴ It is worthwhile noting that the main criticisms of Williamson's claim that $E = K$ has targeted the left-to-right direction of the biconditional principle $E = K$. See, for instance, (Conee and Feldman 2008) and (Goldman 2009).

³⁵ For a more traditional evidentialist account of epistemic justification see (Conee and Feldman 1985). In this paper, I shall work with (PEJ) for the sake of simplicity.

We are, on this definition, evidentially justified to believe $\neg sh$, for our evidence entails $\neg sh$: our evidence makes $\neg sh$ as likely as a proposition can be.

Next, note that, with (PEJ) in hand, we can also define, in a familiar evidentialist fashion, a notion of *doxastic evidential justification*.³⁶

Doxastic E-Justification (DEJ):

x 's belief that p is evidentially justified iff

1. x 's evidence makes p sufficiently likely and
2. x 's belief that p is properly based on x 's evidence.

With (DEJ) in hand we can now capture a clear and precise sense in which Moorean inferences are transmissive of justification. Consider the following principle:

Transmission of Evidentialist Justification from Knowledge (TEJK):

If x knows p and x comes to believe q on the basis of a competent deduction from p , then x has thereby acquired (possibly for the first time) an evidentially justified belief that q .

Given our (partly) propositional notion of evidence and the above definition of doxastic E-justification, this principle seems rather trivial. To see this in detail, note again that if x knows p , then, by definition, p is part of x 's evidence. Thus, if x bases her belief that q in a proper way on her knowledge that p , then x 's belief that q will be evidentially justified. But surely, if p is part of x 's evidence, p entails q , and x bases her belief that q on a competent deduction from p , then x has *properly* based her belief that q on her evidence for q . Thus, the above transmission principle is rather plausible, given our current conception of evidence as constrained by (K \subset E) and the definition of doxastic evidentialist justification in (DEJ). After all, the intuitive picture underlying (TEJK) is nothing but the simple idea that when we competently deduce a proposition q from our knowledge that p , then it is p that justifies q : p is our *reason* in support of q or the relevant part of our evidence that justifies our belief in the inference's conclusion.³⁷

Let us finally apply the above considerations about evidential justification to the *Moorean Argument*. Firstly, note that since we know op , op is part of our evidence.

³⁶ See (Turri 2010) for an opposing view, on which propositional justification is to be accounted for in terms of doxastic justification.

³⁷ See also (Silins 2005, pp. 87-88) for this view, and cp. (Tucker 2010, p. 507): "Suppose that Harold's belief in P is doxastically justified by his evidence E; he notices that P entails Q; and then he subsequently deduces Q from P. It is natural to identify Harold's reason for accepting Q as P, not E. Since we are supposing that P entails Q, P is presumably a warrant for Q. But if P is Harold's reason for Q and is itself a warrant for Q, it doesn't seem to matter whether the deduction transmits warrant, that is, whether the deduction makes E into a warrant for Q."

Our evidence, therefore, strongly supports—in fact, entails— $\neg sh$. Next, note that our belief that $\neg sh$ is also properly based on our evidence insofar as we have based it on a competent deduction from op . In other words, since basing one's belief on a competent deduction from what one knows amounts to the proper basing of one's belief on one's evidence, we can acquire, possibly for the first time, an evidentially justified belief that $\neg sh$ by basing our belief that $\neg sh$ on a competent deduction from op :

Moorean Transmission of Evidentialist Justification (MTEJ):

If x knows (i) and bases her belief that (iii) on a competent deduction from (i), then x has thereby acquired (possibly for the first time) an evidentially justified belief that (iii).

As we have seen, this principle is fairly trivial, given our acceptance of the idea that $K \subset E$, and the conception of doxastic justification in (DEJ). There is, as a consequence, a surprisingly clear sense in which transmission does *not* fail for competent Moorean derivations.

5. Having One's Cake and Eating It, Too

Given the definitions of evidence and justification from the previous chapters we can now have our cake and eat it, too. In particular, we can combine the view that Moorean inferences are cases of transmission failure—namely, in the sense that they do not transmit RA-justification—with the view that they are cases of successful transmission—namely, in the sense that they transmit E-justification. The fact that the inferences at issue are cases of transmission failure with respect to one sense of 'justified' does not entail that they are also cases of transmission failure with respect to the other sense. The approach sketched in the previous sections is, accordingly, ecumenical in that it allows us to give an account of the intuition of transmission failure while at the same time explaining why Moorean inferences can generate both knowledge and justification. To be clear, transmission fails with respect to the arguments at issue: RA-justification does not transmit. But transmission, in a different sense, doesn't fail: E-justification transmits.

On the account proposed here we can thus resolve the Moorean puzzle by claiming that when some theorists defend the view that justification does not transmit through the *Moorean Argument*, while others take the view that it does, then these theorists

are most charitably interpreted as being both right, but as speaking about different types of epistemic justification.³⁸ Moreover, since each of the above definitions presents an entirely legitimate way of thinking and theorizing about epistemic justification and evidence, we can explain our puzzlement about *Moore's Argument* by simply noting that we have a tendency to equivocate between these subtly different notions: sometimes, when speaking about epistemic justification, we are interested in whether our evidence makes our beliefs sufficiently likely, while at other times we are interested in whether our perceptual evidence and memory states eliminate all relevant counterpossibilities. Distinguishing carefully between these two different ways of thinking about justification provides the key to the solution of the Moorean puzzle.

A final advantage of the account proposed here is worth mentioning. Note that the RAT-based approach from Section 3 is also in a position to explain why, as James Pryor (2004, p. 363) has pointed out, the *Moorean Argument* cannot play a role in rationally overcoming one's doubt in its conclusion.³⁹ In fact, the defender of (RAT) has a rather straightforward explanation of the phenomenon pointed out by Pryor. To see what I have in mind, note that if one doubts a proposition p , then one usually believes that one doesn't know p . As a consequence, doubters will typically accept that there are some relevant alternatives to p that are uneliminated by their evidence.⁴⁰ Those who doubt $\neg sh$ —the conclusion of the *Moorean Argument*—will, therefore, usually accept that not all sh -worlds are epistemically irrelevant. But if some sh -worlds are epistemically relevant, then it follows, from (RAT), that one doesn't know op : no sh -worlds are, after all, eliminated by one's (Lewisian) evidence. Thus, for those subjects who doubt the conclusion of the *Moorean Argument*, the argument itself is based on premises that are not known. It is, therefore, not surprising that those

³⁸ One might insist that the Wrightean intuition is that *one and the same* epistemic property is transmitted across some pieces of deductive reasoning, but not across others. If that is the intuition, then it is, on the account defended here, simply wrong. However, one might wonder whether our intuitions concerning Moorean arguments are in fact as fine-grained as this objection assumes: surely, our pre-theoretical intuitions are that Moorean arguments are peculiar or defective in a way in which more paradigmatic cases of deductive reasoning aren't, or that being warranted in the argument's conclusion presupposes having an antecedent warrant for its premise (i). There is reason to doubt, however, that our pre-theoretic intuitions support the claim that *one and the same epistemic property* is transmitted across some deductive arguments but not others.

³⁹ As Pryor (2004, p. 363) puts it, "anybody who had doubts about [the] conclusion [of Moore's argument] couldn't use the argument to rationally overcome those doubts."

⁴⁰ Note that one can accept that there are some relevant alternatives to p that are uneliminated by one's evidence without explicitly believing any such proposition: the attitude in question amounts to accepting that one doesn't know p because one's evidence isn't good enough.

who sincerely doubt that they are not brains in vats will not be convinced by the *Moorean Argument*.^{41,42}

It might be objected to this argument in particular and to the project defended in this article more generally that the central notion of RAT—the notion of an epistemically relevant alternative—is too vague and unclear to do any real explanatory work. To a large extent I agree with this objection: the notion of relevance at issue is indeed not entirely clear and certainly not explicitly and reductively defined. But why should it be? As Stalnaker points out in a different context,

[s]o long as certain concepts all have some intuitive content, then we can help to explicate them all by relating them to each other. The success of the theory should depend not on whether the concepts can be defined, but on whether or not it provides the machinery [...] to make conceptual distinctions that seem important. With philosophical as well as scientific theories, one may explain one's theoretical concepts, not by defining them, but by using them to account for the phenomena. (Stalnaker 1970, p. 46)

Let me sum up. We have seen in the course of this paper that RAT, if formulated carefully, provides us with a simple and straightforward way to model both the Wrightean position of transmission failure and the Moorean position of dogmatism. Further, I have argued that we are sometimes puzzled by the *Moorean Argument* and its cousins because we are unaware of the subtly different ways in which one can be epistemically justified in believing a proposition. Once we carefully tease apart the different senses of 'justified' and 'evidence' involved, the bewilderment caused by the *Moorean Argument* and its cousins should subside.

References

- Bergmann, M. (2004). Epistemic Circularity: Malignant and Benign. *Philosophy and Phenomenological Research*, 69(3), 709-727.
- Blome-Tillmann, M. (2009). Knowledge and Presuppositions. *Mind*, 118(470), 241-294.
- Cohen, S. (1988). How to Be a Fallibilist. *Philosophical Perspectives*, 2, 91-123.

⁴¹ See Bergman (2004, p. 719), who also argues that a thinker loses her knowledge that p as soon as she starts doubting whether p . However, Bergman does not employ a relevant alternatives conception of knowledge, such as (RAT), in accounting for this datum.

⁴² An alternative explanation relies on a contextualist semantics for 'knows' and the notion of an epistemically relevant alternative, claiming that what is epistemically relevant varies with context. On this view, a speaker's doubting $\neg sh$ in a context C will entail that some sh -worlds are taken seriously in C , which in turn will render those sh -worlds epistemically relevant. Presumably, Lewis (1996) himself would be attracted to such a view, but see also (Blome-Tillmann 2009).

- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- Cohen, S. (2000). Replies. *Philosophical Issues*, 10(1), 132-139.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2), 309-329.
- Conee, E., & Feldman, R. (1985). Evidentialism. *Philosophical Studies*, 48, 15-34.
- Conee, E., & Feldman, R. (2008). Evidence. In Q. Smith (Ed.), *Epistemology - New Essays* (pp. 83-104). Oxford: Oxford University Press.
- Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67, 1007-1023.
- Goldman, A. I. (2009). Williamson on Knowledge and Evidence. In P. Greenough, & D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 73-91). Oxford: OUP.
- Jenkins, C. (2007). Entitlement and rationality. *Synthese*, 157(1), 25-45, doi:10.1007/s11229-006-0012-2.
- Klein, P. D. (1981). *Certainty, a refutation of scepticism*. Minneapolis: University of Minnesota Press.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549-567.
- Moore, G. E. (1939). Proof of an External World. *Proceedings of the British Academy*, 25, 273-300.
- Nozick, R. (1981). *Philosophical Explanations*. Oxford: OUP.
- Pedersen, N. (2009). Entitlement, value and rationality. *Synthese*, 171(3), 443-457, doi:10.1007/s11229-008-9330-x.
- Pritchard, D. (2007). Wittgenstein's On Certainty and Contemporary Anti-scepticism. In M.-S. Danièle, & W. H. Brenner (Eds.), *Readings of Wittgenstein's On Certainty* (pp. 189-224). New York: Palgrave.
- Pryor, J. (2004). What's Wrong With Moore's Argument? *Philosophical Issues*, 14, 349-378.
- Silins, N. (2005). Transmission Failure Failure. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 126(1), 71-102.
- Smith, M. (2012). Entitlement and Evidence. [doi: 10.1080/00048402.2012.732094]. *Australasian Journal of Philosophy*, 1-19, doi:10.1080/00048402.2012.732094.
- Stalnaker, R. (1970). Pragmatics. *Synthese*, 22, 272-289.
- Stine, G. (1976). Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies*, 29, 249-261.
- Tucker, C. (2010). When Transmission Fails. *Philosophical Review*, 119(4), 497-529.
- Turri, J. (2010). On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*, 80(2), 312-326, doi:10.1111/j.1933-1592.2010.00331.x.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: OUP.
- Wright, C. (2002). (Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell*. *Philosophy and Phenomenological Research*, 65(2), 330-348, doi:10.1111/j.1933-1592.2002.tb00205.x.
- Wright, C. (2003). Some Reflections on the Acquisition of Warrant by Inference. In S. Nuccetelli (Ed.), *New Essays on Semantic Externalism and Self-Knowledge* (pp. 57-78). Cambridge, Mass.: MIT Press.
- Wright, C. (2004). Warrant for nothing (and foundations for free?). *Aristotelian Society Supplementary Volume*, 78(1), 167-212, doi:10.1111/j.0309-7013.2004.00121.x.

- Wright, C. (2008). The Perils of Dogmatism. In S. Nuccetelli, & G. Seay (Eds.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (pp. 25-48). Oxford: OUP.
- Wright, C. (2010). Crispin Wright. In J. Dancy, E. Sosa, & M. Steup (Eds.), *A Companion to Epistemology* (pp. 204-210). Oxford: Blackwell.