The Paradox of Graded Justification

Artūrs Logins

Departament of Philosophy, University of Zurich, Switzerland

Email: Arturs.Logins@uzh.ch

[THIS IS THE PENULTIMATE VERSION, PLEASE CITE THE PUBLISHED VERSION]

ABSTRACT

According to a widely held view epistemic justification is a normative notion. According to another widely held assumption, epistemic justification comes in degrees. Given that gradability requires a context-sensitivity that normativity seems to lack, these two assumptions stand in tension. Giving up the assumption of gradability of justification represents a lesser theoretical cost.

1. Introduction

Two popular assumptions in contemporary epistemology seem to stand in tension. The first is the idea that epistemic justification is a normative notion. Epistemic justification is about assessing beliefs (and other doxastic states). Justification-talk seems to involve expressions of evaluation. Assessment seems to be a mark of the normative.

The second assumption is the idea that epistemic justification comes in degrees. Simply put, it is a common assumption that we can compare beliefs (or other doxastic states) justification-wise. It seems natural to think that my belief that the moon landing was real is as justified as my belief that 9/11 was not an inside job. And my belief that Earth is spherical seems to be more justified than my belief that the impact of a comet caused the extinction of the dinosaurs. Certainly, it doesn't appear far-fetched to think that someone with the common knowledge we share today is more justified in believing that World War II ended in 1945 than, say, in believing that lizard people rule the world. It then seems natural to think that epistemic justification admits of degrees. And it appears that this assumption is common in contemporary epistemology.

Taken together these two widely held assumptions lead to some unforeseen troubles. It appears that while both are independently reasonable, they do stand in tension and cannot both be upheld at the same time. Or so I will argue in what follows.

A more specific aim of the present paper will be to show that it is the assumption of gradability of justification that has to go. Endorsing this conclusion will appear to be the most promising way to deal with what I will call the paradox of graded justification and can be stated as follows:

The paradox of graded justification

- (1) (Epistemic) justification is an absolute property (e.g., similar to being straight, being empty, being flat, being pure).
- (2) (Epistemic) justification comes in degrees (e.g., the belief that Earth is spherical is more justified than the belief that the impact of a comet caused the extinction of the dinosaurs).
- (3) Absolute properties don't come in degrees (e.g., empty or not?)1

Premises (1)–(3) lead to a contradiction. But all have initial plausibility. Thus, we seem to have a paradox here.

In what follows, I will first sharpen up the paradox. I will provide some clarifications and specifications of the three premises. I will provide some theoretical motivation for holding each of these. And I will also introduce some constraints on potential solutions (or dissolutions) of the paradox. Second, I will consider the prospects of giving up premise (1). I will show what this would entail, and I will assess the consequences of giving up (1). In particular, and crucially for our overall dialectic, it would seem that giving up (1) leads to giving up the widespread idea that justification is normative. Third, I will consider the prospects of giving up premise (3). Again, we will work out what such a move would entail, and we will assess the consequences of denying (3). Fourth, we will introduce our favourite solution, namely, giving up premise (2). We will provide considerations in favour of this move. Finally, we will conclude by outlining some further consequences of our preferred solution.

2. Sharpening the paradox

In this section, I provide additional clarifications of the three premises of the paradox of graded justification. I will also introduce some *prima facie* motivation for all three of them. In doing so, we will also motivate the idea that the inconsistent triad (1)–(3) is indeed a paradox, since not only does it lead to a contradiction, but all three premises appear to be independently plausible.

2.1 Premise (1): Epistemic justification is an absolute property

In what sense is (epistemic) justification an absolute property? What do we mean by 'absolute property'? We can first approach these questions by providing a minimal characterization of absolute properties. Roughly, being justified is absolute in the same sense as being empty, being full, being flat, being straight, being transparent, being opaque, and being pure are absolute. These latter are all absolute properties, in an intuitive sense. Emptiness is uncompromising – only when a container has

¹ This paradox is strongly inspired by what Heather Burnett has recently called the paradox of absolute adjectives (see Burnett 2017: 70 and further references therein for other places where similar puzzles are discussed). We consider Burnett's paradox in more detail in section 3 below, but note already that our proposal focuses on justification (the thing) rather than absolute adjectives (language) and therefore is not primarily'about language and is not about absolute properties in general.

nothing in it will we count it as truly empty. Only 24-karat gold is pure gold; anything less than that is not pure gold. And even a bit of moisture on the table surface will make the table wet. Absolute properties are in a sense intolerant. This idea is captured in the following quotation from David Lewis (about certain sorts of absolute properties/terms):

But flat is an *absolute term*: it is inconsistent to say that something is flatter than something that is flat. (Lewis 1979: 353; this suggestion is insightfully explored in Burnett 2017: 71)²

Some absolute terms (and, accordingly, absolute properties) are such that we cannot say in a single utterance without a feeling of contradiction that something has the property and at the same time something else has it more, e.g., being flat.³

What this seems to show is that absolute terms/properties don't rely essentially on comparison classes. This essentially non-relative, intolerant aspect in the context of absolute adjectives is further elaborated by Heather Burnett as follows:

[B]oth AAs [Absolute Adjectives] and NSs [Non-Scalar adjectives] have semantic denotations that are assigned independently of a contextually given CC [comparison class]. That is, in order to know which rooms are empty or which sticks are straight, we don't need to compare them to a certain group of other individuals; we just need to look at their properties. (Burnett 2017: 67)

Contrast absolute terms/properties, that is, properties that don't rely on comparisons to relative terms/properties, that is, properties that do rely on comparisons (this contrast is a central element of Kennedy and McNally's 2005 and Kennedy's 2007 framework for adjectives). Compare, for instance, being tall (for people's height) to being empty. There is no feeling of contradiction in asserting in the same utterance that, say, Agatha is tall and at the same time Anastasia is taller. In the case of being tall, there is radical context variation/relativity. What is tall depends intrinsically on what the relevant comparison classes are. Both Agatha and Anastasia count as tall compared to other kids of their age, but of course, they won't count as tall compared to professional basketball players. There is no such radical context variation in the case of being empty, being flat, and so on. Normally, a box cannot be empty for, say, a chocolate box and at the same time be not empty for a biscuit box. It doesn't make much sense to first compare the box to some comparison classes in order to determine whether it is empty or not. One

² Interestingly this observation appears in Lewis's discussion of Peter Unger's argument for scepticism (note also that Unger 1975 appears to be the original source of the term 'absolute' term/adjective). See Lewis's quotation in its wider context: "Peter Unger has argued that hardly anything is flat. Take something that you claim is flat; he will find something else and get you to agree that it is even flatter. You think the pavement is flat – but how can you deny that your desk is flatter? But flat is an absolute term: it is inconsistent to say that something is flatter than something that is flat. Having agreed that your desk is flatter than the pavement, you must concede that the pavement is not flat after all. [...] In parallel fashion Unger observes, I think correctly, that 'certain' is an absolute term; from this he argues that hardly ever is anyone certain of anything" (Lewis 1979: 353–354).

³ Not all absolute properties are like that. Of course, we can say that the table surface is wet, and the wall is wetter (we cannot reasonably insist, however, that the wall is wetter than the table and at the same time/in the same utterance deny that the table is wet). What's important for us is that being justified is absolute. Whether it is like being flat or like being wet is not a question that we aim to adjudicate here. See [reference omitted] for a suggestion that it appears to be sometimes like flat and sometimes like wet; see also Hawthorne and Logins (2020) for related considerations.

has to look straight into the box to make the right judgment. For tallness, on the other hand, it doesn't make much sense to 'just see' the object/individual in order to determine whether they/it are/is tall or not. We first need to know to which class we are comparing the individual.

Thus, we can tentatively say that a mark of absolute properties (e.g., being empty, flat, transparent) is that, contrary to relative properties (e.g., being tall), their possession is not essentially depending on comparisons to some relevant classes of individuals/objects.

Now, justification behaves more like being empty/pure than like being tall in this respect. It would seem that a belief cannot be epistemically justified for an N-class of beliefs and at the same time not epistemically justified for an N'-class of beliefs. It doesn't make much sense to look for some comparison classes before trying to determine whether a belief is justified or not. There doesn't seem to be any essential context relativity involved in the case of justification.

It is important to note that we are not assuming from the outset that absolute =def. not-gradable. The paradox of graded justification need not ascribe plain absurdity to epistemologists who think that justification comes in degrees. The paradox is constituted by the triad (1)–(3), and one can avoid endorsing a contradiction by rejecting any of these three premises. Therefore, it remains a conceptual possibility at this point that absolute properties are gradable. This possibility amounts to a rejection of premise (3). But let us not anticipate things. We will return to this option below.

It is crucial for our purposes at this point to note that premise (1) seems to have initial plausibility, and that orthodoxy in current epistemology entails (1). Let us look at some popular views that entail or presuppose (1).

It is a common assumption nowadays that epistemic justification is epistemic permissibility. According to many, being justified to F just is being permitted to F. Applied to the case of epistemically justified beliefs, it amounts to maintaining that an epistemically justified belief just is an epistemically permitted belief. See Goldman (1986), Pollock and Cruz (1999), Steup (2001), Nelson (2010), and Wedgwood (2013) for a non-exhaustive list of epistemologists who think we can characterize/define justification in terms of permissibility. Now, assuming that permissibility is an absolute property, the view that justification is permissibility entails that justification is absolute too. That permissibility is absolute seems prima facie very plausible. To know whether it is permitted to park a car on this side of the road, you will directly look for parking signs and/or road markings. You will not first compare the act of parking your car here to some class of possible courses of action, say, acts of parking cars on streets or in garages, to determine whether your act of parking your car here is permissible. And to say that while it is morally permitted to put a lifejacket on oneself before helping others, it is more permitted to save others at the expense of running a risk for oneself, sounds like a creative use of language. The same goes for claims about a belief being permitted but another belief being more permitted. In short, it would seem that the property of being permitted has some essential marks of absolute properties. It doesn't depend essentially on comparison classes and is not intrinsically context relative (which doesn't mean that we cannot make sense of some variations, or rather, contextually informed precisifications of exactly which F-ings are permitted).

Another popular view that entails that justification is absolute is the view that justification is appropriateness or compliance with a norm/with an ought (cf. Williamson 2011, 2014, Simion, Kelp, and Ghijsen 2016, Littlejohn 2013, among others). On this view, a justified belief ultimately just is a belief that complies with the relevant belief norm. What exactly the relevant norm is, of course, is a matter of discussion (e.g., maybe it is the norm that one ought to believe only propositions that one knows, cf. Williamson 2000, 2011, Engel 2004, or maybe it is the norm that one ought to believe only propositions that are true, cf. Boghossian 2003, Wedgwood 2002, Shah and Velleman 2005; I take it that this debate doesn't really matter for our purposes). Assuming that compliance is an absolute property, the view that justification is compliance with the relevant norm entails that justification is absolute. And, again, it appears rather plausible to think that compliance is absolute. To see whether you have complied with the parking regulations, you will check the relevant signs and the position of your car, and you will not first look for comparison classes of acts of parking cars in parking lots or in garages. It makes little sense to say that one has complied with the moral ought of not killing a person, but that, at the same time, one could comply more with the norm of not killing the person. Similarly, it is not easy to see how it would be possible that one both complies with the relevant belief norm by believing that p and that one could comply more with that norm applied to believing that p. Compliance seems to be an absolute property. If so, and if justification is compliance with a norm, then justification is absolute too.

Thus the first constraint for any potential solution to the paradox seems to emerge. Any solution to the paradox has to be in a position to account (one way or another) for the initial plausibility of views that identify/characterize justification as permissibility or compliance with a norm, as well as for the initial plausibility of the idea that being justified doesn't depend essentially on comparisons.

2.2 Premise (2): Epistemic justification comes in degrees

The strongest case for the claim that (epistemic) justification comes in degrees appears to come from some observations about our language use.

It is standard to take (i) felicity of comparative structures, (ii) felicity of 'How-F?' questions, and (iii) the possibility of combination with degree modifiers to be indicators of the gradability of an adjective (cf. Kennedy 2007, Burnett 2017, Lassiter 2017: 7–18, Wellwood 2019: 13–15, among many others). And from what it appears, 'justified' satisfies all three of these criteria.

As we have recently observed in Hawthorne and Logins (2020), 'justified' (i) admits of comparative constructions (e.g., "Over goes Rakitic and the red card that follows is as justified as the challenge was futile and brainless"; see Hawthorne and Logins 2020: 2 for the source of the quotation), (ii) can be felicitously used in 'How-F?' questions (e.g., "But how justified were fan fears that Leeds would be starting the new season with a new man in the dugout for the eighth year in a row?"; see Hawthorne and Logins 2020: 2 for the source of the quotation), and (iii) combines well with degree modifiers (e.g., "Defence analyst David Perry of the Canadian Global Affairs Institute said those concerns are completely justified [...]"; see Hawthorne and Logins 2020: 2 for the source of the quotation). See also Siscoe (forthcoming) for somewhat different observations but, still, the same view that 'justified'

seems to belong to the general category of gradable adjectives. Staffel (2019: 162–165, as well as Siscoe forthcoming) proposes further observations in favour of the gradability of 'rational'.

Before considering the importance of these observations, note that at this point, one need not be committed to any specific ontology of degrees. Merely observing that 'justified' seems to behave like a gradable adjective in some respects need not commit one to a specific way of understanding what degrees are. But note, however, that not all ontologies of degrees are on a par. There are two major competing frameworks for degrees. On one view, degrees exist. They are *abstracta*, and more specifically, representations of measurement. On this 'realist' view, they are independent of objects. And scales on this view are understood as ordered sets of degrees (cf. Von Stechow 1984, Kennedy 2007, among many others).

According to a different view, which we may call 'anti-realist', degrees are not independent of objects/individuals but are best understood as equivalence classes (cf. Cresswell 1976, Klein 1991, Bale 2011), where equivalence classes are '[s]ets of objects all of which bear the "exactly as *P* as" relation to each other, for the relevant property *P*' (Lassiter 2017: 16).

The two approaches differ radically with respect to their ontological commitments. The antirealist approach is less ontologically committed and thus more parsimonious, a quality that many philosophers would find *ceteris paribus* very attractive (cf. Lassiter 2017: 16).

Taking the linguistic observations at face value leads us to consider another constraint for any potential solution to the paradox of graded justification. Namely, the second constraint is that a viable solution to the paradox has to account (one way or another) for the apparent felicity of gradable uses of 'justified' and our commonsense judgments about graded justification.

Before moving on to the third and final premise of the paradox, let me stress the theoretical significance of this premise. It is common for epistemologists to assume that justification comes in degrees. And it is common to rely on the gradability of the justification assumption in arguments for substantive theories of justification. It would seem that some famous arguments in the literature take the following general structure: our theory T explains the gradability of justification best; therefore, by abductive inference, theory T holds. Let me illustrate this point using some examples. For instance:

We can and do regard certain beliefs as more justified than others. Furthermore, our intuitions of comparative justifiedness go along with our beliefs about the comparative reliability of the belief-causing processes. [...] Again, the degree of justifiedness seems to be a function of reliability. (Goldman 1979/2008: 338–339)

A seldom-noted fact about Goldman's groundbreaking 1979 article in defence of process reliabilism is that one of the main positive arguments that he provides there in favour of his version of reliabilism relies crucially on the assumption that justification (or justifiedness in his terminology) comes in degrees. The idea is that since reliability comes in degrees, process reliabilism is best positioned to explain the gradability of justification. Thus, the positive case for process reliabilism relies on the assumption that justification comes in degrees.

Consider another example from a more formal approach to justification:

So far we have only talked about the existence or nonexistence of justification [= "is some evidence for" / confirmation]. But sometimes there is more rather than less justification, so we need the concept of *degree of justification*, that is, E justified H to degree c. It then follows that one piece of evidence may confirm a hypothesis more than another piece of evidence does. So we need the concept of *comparative justification*, that is, E1 justifies H more than E2 justifies H. The Standard Bayesian Response [to the Paradox of Ravens] makes this comparative claim: a black raven justifies H more than a white sneaker justifies H (where H = All unobserved ravens are black[...]). (Bradley 2015: 123)

Bayesians often motivate their approach by putting forward problems and puzzles that their framework can easily deal with. One can read, for instance, Howson and Urbach (2006) as a book-length argument of just this sort. What the above quotation says can then be understood as a more specific instance of such an argument. One could read it as follows: the Bayesian approach, given its technical apparatus applied to degrees of justification, explains comparative justification perfectly since by proving it, it can solve the so-called Paradox of Ravens.⁴ Of course, it is then essential for the Bayesian project that justification does come in degrees.

A similar argumentative pattern can be identified in defences of what Martin Smith calls 'the risk minimization conception of justification' (which is prevalent in contemporary epistemology and which he opposes to his preferred normalcy view):

As well as appraising the justification that one has for believing a single proposition, it is often natural to compare the justification one has for believing distinct propositions [...]. The risk minimisation conception [...] includes a straightforward account of justification comparisons: one has more justification for believing P than Q iff one's evidence makes P more probable than Q. (Smith 2016: 92)

Roughly, according to the risk minimization conception, justification doesn't require certainty, but something like high enough probability. And according to Smith, it is a very popular, indeed orthodox, view in contemporary epistemology. Assuming that justification comes in degrees, the risk minimization conception (and its appeal to probability) can be seen as having an extra argument in its favour as long as it is in the best position to explain the apparent gradability.

extension, our general point still holds and the above quotation illustrates it, as long as S having confirmation for a proposition entails S being epistemically justified in believing that proposition; that is, if it is true that confirmation entails justification. And this doesn't seem to be a particularly controversial assumption.

⁴ At this point, one might wonder whether justification as understood by Bayesians, and as referred to in the above quotation, can be identified with epistemic justification as we understand it in the present debate. Namely, one might think that what Bayesians are after is a theory of scientific reasoning and scientific confirmation and this is not the same as justified belief. For one might insist that not all justified beliefs need to be scientifically confirmed. However, for the present purposes we can set aside the question of the exact nature of the relation between epistemic justification of beliefs and confirmation. Even if the two come apart conceptually or don't have the same extension, our general point still holds and the above quotation illustrates it as long as S baying confirmation for a

Now, the interesting point for us is that, if at the end of the day justification turns out not to come in degrees, then theories that predict the gradability of justification are mistaken. Or, at minimum, if justification is not gradable, then all these (and still other) inference to the best explanation arguments that appeal to degrees of justification are undermined and, hence, fail to provide support for the corresponding substantive theories of justification (see the quotations above). But let's not anticipate our more thorough discussion of the gradability of justification yet.

2.3 Premise (3): Absolute properties don't come in degrees

The strongest case for premise (3), that is, for the claim that absolute properties don't come in degrees, seems to rely on the following reasoning. Absolute properties don't depend essentially on comparisons (being empty doesn't depend on comparison classes, while being tall does). And correspondingly, true ascriptions of absolute adjectives don't depend essentially on contextually determined standards of application for these predicates. However, gradability is strictly and essentially an affair of comparisons and of contextually set standards. Thus, absolutes are not graded. Let me elaborate slightly on this line of thought.

If a property is genuinely gradable, then essentially it allows for comparisons. If P comes in degrees, then it has to be the case that either x can have more of P than y, or x can have less of P than y. Now, suppose for ease of argument (in line with the anti-realist conception of degrees) that degrees are sets of individuals that bear the 'exactly as P as' relation to each other for some property P (cf. Lassiter 2017: 16). So, a given degree of tallness is a set of individuals who are all exactly as tall as each other. But there are various degrees of tallness. So, on this account, each different degree of tallness is a different set of individuals who are all exactly as tall as each other. We don't get gradability of being tall, without having various sets of individuals, various equivalence classes that obey an ordering. Compare this to non-gradable properties. There are no varying degrees of being, say, nuclear or pregnant, since there are no orderings of various equivalence classes of ways of being nuclear or being pregnant. If there are no sets that could be ordered (in a well-defined way) with respect to equivalence classes of some property, then we won't get gradability. Comparisons correspond to the relevant orderings. Thus, it would seem we cannot get gradability without involving comparison classes (comparisons).

A slightly different way to argue for the idea that gradability is essentially linked to comparison classes is that this way of understanding gradability fits well with a promising semantics of graded adjectives (cf. Burnett 2017), namely, the Delineation approach. Consider comparatives and absolutes in this framework:

Informally, in this framework, *John is taller than Mary* is true just in case there is some CC [Comparison Class] with respect to which John counts as tall and Mary counts as not tall. (Burnett 2017: 62)

And:

[I]n order to know which rooms are empty or which sticks are straight, we don't need to compare them to a certain group of other individuals; we just need to look at their properties. [...] [W]hat it means to be non-context sensitive is to have your denotation be invariant across classes. (Burnett 2017: 67)

That absolute adjectives don't come in degrees has also been suggested elsewhere, in the semantics literature. That is, the point doesn't seem proper only to proponents of the specific Delineation approach developed by Burnett. For instance:

As a matter of fact, we know perfectly well which property the adjective *empty* expresses. It is the property (for a container) of not containing anything, of being devoid of contents. This is how we define *empty*. Note that this is an absolute property, a property which a container has or does not have. Either it contains something, or it does not contain anything. So the property which the adjective expresses and which determines its extension is not a property that admits of degrees. How, then, can we explain the gradability of the adjective? (Recanati 2010: 117; cf. Burnett 2017: 70–71)

3. Prospects of giving up the premise (1): Can there be justification without absoluteness?

Giving up premise (1) effectively avoids a contradiction and prevents arriving at an absurd conclusion. If (1) is false, we need not worry any more, the paradox would be solved. However, there are several worries with this move. For one thing, there is the bedfellows problem. If justification is not absolute (in a sense specified above), then what about being empty, straight, flat, pure, and so on? Denying absoluteness for justification only and not for its bedfellows without providing any further independent theoretical motivation for it appears very much like an *ad hoc* move in the present dialectical context. Denying the absolute character of being empty, being straight, being wet, being pure, being impure, and so on, is impalpable. Thus, if one is going to deny premise (1), one has to deal with the 'justified' bedfellows problem, but it is not clear how one could successfully deal with it.

Another problem with denying premise (1) is that if justification is not absolute, then we need an extra explanation of the apparent non-relative aspect of being justified. As we observed above, being justified is clearly different from things like being tall or being rich. Being tall is intrinsically relative – it depends on contextually set standards and comparison classes. If premise (1) is false, and justification is not absolute, it is relative. But if it is relative, how can we make sense of the fact that it appears so different from standard relative properties/terms like being tall/'tall', or being rich/'rich'? Thus, if one denies premise (1), one has to deal with the apparent non-relativity of being justified, but it is not clear how exactly one could successfully deal with this while insisting that justification lacks absoluteness.

A third, and perhaps philosophically more interesting and more substantial, problem with denying (1) is that endorsing this move would make it difficult to make sense of the normative aspect of justification. The rest of the present section explores this worry in more detail.

3.1 Can justification be normative if not absolute?

The worry with normativity is that most, if not all, of the standard ways of understanding (epistemic) normativity seem to involve absolute properties/concepts. But if epistemic justification is not absolute, then it raises a challenge. Is it even normative, and if it is, how should we understand its normativity?

As already seen above (see section 2.1), an increasingly popular way to understand (epistemic) justification is to characterize/define it in terms of permissibility (Wolterstorff 1983, Goldman 1986, Pollock and Cruz 1999, Steup 2001, Nelson 2010, Wedgwood 2013). It is a broadly deontic conception since it is connected to an ought, which is commonly seen as a dual of permissible. Whether there are any positive epistemic duties/oughts, as opposed to the claim that there are only negative epistemic duties (cf. Nelson 2010), is a further debate within this approach. Alternative, broadly deontic views would also include the view that justification just is compliance with a norm (Simion, Kelp, and Ghijsen 2016, Williamson 2011, 2014, among others), and the view that justification is faultlessness (Beddor 2017). Importantly for us, a crucial feature of all these broadly deontic properties is that they are clearly absolute. It cannot be the case that something is more permitted/has more compliance/is more faultless than something that is permitted/has complied/is faultless. To say that my belief that Earth is spherical is more permitted than my permitted belief that it is raining outside sounds funny if not outright absurd.

If understanding justification in terms of permissibility or compliance with norms stands in tension with denying its absoluteness, then opponents of premise (1) might want to explore other ways of accounting for the normative aspects of justification. And as it happens, there are alternative proposals about justification that do (aim to) preserve its normative aspect, broadly understood.

A natural place to begin looking for alternative views would be to consider axiological properties or values rather than deontic properties. That is, once one gives up the idea of understanding justification by appeal to permissibility or compliance with a norm (or faultlessness), it is only natural to investigate the idea of understanding justification in terms of values/axiological properties. For in general, it is common in normative philosophy to oppose values to norms/oughts (and then, sometimes, to try to reduce one to the other). If theories that tie justification to deontic properties are given up, then one might want to explore theories of justification that tie it to axiological properties (that is, values). (In doing so, one might assume that normativity is to be understood broadly enough, perhaps, as an umbrella term/property, of which deontic and axiological properties are two more specific instances.)

Now, the most prominent version of an approach that attempts to explain justification in terms of axiological properties is epistemic consequentialism. And it can be summed up, roughly, as follows:

EC: A belief/doxastic state is justified iff it is conducive to an increase in epistemic value and a decrease in epistemic disvalue (see Ahlstrom-Vij and Dunn 2014, Goldman 1999, Singer 2018, Joyce 1998, 2009, Pettigrew 2013, 2015).

A common claim that epistemic consequentialists typically make is that beliefs (and/or degrees of belief) can be ranked in terms of whether they are at least as good (on some accounts as strictly better) in bringing about an increase in the relevant value as their alternatives are. There is some variation on how

to conceive epistemic value. Should we think of it in terms of, say, true beliefs, knowledge, or perhaps understanding? But this specification need not worry us here. There is also a debate on how to understand the conduciveness element in EC. Is it to be understood in terms of causing, maximizing over time, or indirect promoting? But, again, nothing in our present discussion seems to depend on these further specifications (see Dunn 2015 for an overview of epistemic consequentialism). The important thing for us is the assumption that epistemic consequentialists typically take on board the possibility of ranking beliefs (and/or other doxastic states) with respect to their capacity to bring about an increase in epistemic value.

The consequentialist assumption about the possibility of ranking doxastic states is important for our present discussion since it might provide hope that a non-absolute but still normative (broadly understood) account of justification could be worked out, for one might think that ranking could be something relative and not absolute. One might think that ranking something as (strictly) better (with respect to some aspect) than something else could depend essentially on context-sensitive factors, interests, and so on. So, one might hope that a consequentialist theory of justification could provide theoretically sufficient grounds for maintaining the normative aspect of justification while giving up the idea that justification is an absolute property.

However, it would seem that appealing to epistemic consequentialism and the contextsensitivity of (some) rankings can bring only false hope to opponents of the absoluteness of justification. Indeed, it is not clear why the context-sensitivity of rankings would help here. The crucial worry with it is that the property of being ranked at least as high as is an absolute property. It just cannot be the case that X is ranked at least as high as Y, and Z is even more ranked at least as high as Y. It doesn't make sense. It cannot be the case that a belief that p be ranked at least as high as a belief that q and a belief that r be even more ranked at least as high as the belief that q. Thus, if being ranked at least as high as is an absolute property, but epistemic justification (and its gradability) is understood by appeal to the property of being ranked at least as high as (with respect to conduciveness to an increase in epistemic value), then it would seem that epistemic consequentialism is committed to the idea that epistemic justification is absolute. Compare it to the property of being pure. Let's assume that detection of 1.0 uSv/h indicates danger with respect to radioactivity. That is, let's say that being ranked at least as high as 1.0 uSv/h makes an area impure radioactivity-wise. And one can now rank areas with this measure in hand. Some areas will have higher levels of radiation (detection of more than 1.0 uSv/h), others will have lower. Given this, we can perfectly appeal to these measures to rank areas in terms of their radioactive purity/impurity. However, that doesn't make the property of being pure or impure a nonabsolute one. That we can rank areas as at least as high as 1.0 uSv/h, or lower, and that we establish that to be pure, an area has to be strictly lower than 1.0 uSv/h, doesn't make purity (radioactivity-wise) essentially dependent on a context and comparison classes. To know whether an area is pure, one has to look at what one's Geiger counter indicates. One need not look at what levels of radioactivity other areas have. Similarly with epistemic justification, even if epistemic consequentialism is correct and justification is to be understood in terms of conductivity to an increase of epistemic value (and a decrease of epistemic disvalue), it is still difficult to see why justification would be relative and not absolute. Once doxastic states are ranked with respect to their conductivity to increase in epistemic value (and decrease in disvalue), there is no place for relativity, intrinsic context dependence, and

essential reliance on comparisons in determining whether a belief is at least as highly ranked as another belief. Once the measures are fixed in accordance with consequentialist rules, all we have to do is look at scores that a doxastic state obtains. We don't need to compare it to the scores obtained by other doxastic states to see whether it is conducive to an increase in epistemic value and a decrease in disvalue. Thus, it seems that while it might have appeared to be an initially promising strategy for an opponent of premise (1), endorsing epistemic consequentialism doesn't vindicate the relativity of justification.⁵

It appears then that neither characterizing justification as permissibility or compliance with a norm/ought nor describing justification by appeal to an increase of epistemic value (and a decrease of disvalue) is compatible with denying its absolute character. Are there any alternatives left for an opponent of premise (1) to explore in an attempt to show that their position can be compatible with the justification being normative? Well, some options might remain. The next natural option to explore would be to claim that justification is to be understood in terms of aretaic properties, that is, virtue-related properties.

Consider, for example, a promising recent version of virtue reliabilism (of a knowledge-first flavour) (cf. Kelp 2018). According to knowledge-first virtue reliabilism about epistemic justification:

KFVR-JB. One justifiably believes that p if and only if one competently believes that p. (Kelp 2018: 69)

A roughly competent belief is to be understood as being "formed by an exercise of an ability to know" the relevant propositions (cf. Kelp 2018: 69). Of course, many other versions of virtue/competence-centred accounts of justification (and knowledge) exist in the literature. It is not our aim to survey them all. The critical point for us is that whatever the details of the view, it is not clear how a focus on virtues could help an opponent of (1).

The worry with the virtue-centric approaches is that a notion central to those approaches, namely, the property of *being competent/virtuous in believing that p*, is itself an absolute property. To take the exact formulation from above (from KFVR-JB), that one competently believes that p is not something relative. If S's belief that p is competent, then S's belief/one's belief that p cannot be more competent. And, again, once we are clear about how exactly competence is to be specified, to determine whether a belief is competent (whether one competently believes that p), we only need to check whether one competently believes (whatever this amounts to exactly) and we need not compare the belief to other beliefs in order to determine whether it is competent. Thus, it seems that appeal to virtues won't give opponents of premise (1) what they are looking for. One cannot vindicate the relativity of justification by appeal to virtue-centred approaches about justification.

Are there still other options for an opponent of (1) to explore? One might think so. One could appeal to normative reasons in defining justification. So, what about reasons-first approaches (cf. Scanlon 1998, Parfit 2011, Schroeder 2007, forthcoming, among others)? One might think that

⁵ Interestingly, some epistemic consequentialists seem to recognize the worry (or a close one, at any rate) and only apply their view to deontic properties, e.g., the right; see Singer (2018).

normative reasons (reasons to believe) come with different weights. And one might hope that somehow an appeal to these weights and weighing of reasons one against another could bring in the desired relative aspect for justification.

But, again, it doesn't look like a promising strategy, after all. A reasonable way of spelling out the reasons-first view of justification would be to characterize a justified belief that p as a belief that is (appropriately) based on sufficient reasons to believe that p. But again, the property of having sufficient reasons to believe that p doesn't seem to be a relative property. It doesn't depend essentially on comparison classes and contexts. If S's belief is sufficiently supported (by reasons), then it cannot be even more sufficiently supported (by reasons). It doesn't make much sense. So, again, once reasons are fixed (that is, once we know what reasons are) and their weight is understood, it would seem that to determine whether a belief is justified, all that one has to do within the reasons-first approach would consist in checking what reasons one has for believing the relevant proposition and not to compare the sufficiency of reasons to other alternative situations of the sufficiency of reasons. Again, the situation is closer to that of being empty than to that of being tall. To see whether a box is empty, one has to look inside and need not compare it to other boxes, whereas to see whether x is tall, one has to compare x to other individuals of the relevant class of comparison. In order to see whether there are sufficient reasons for one to believe that p, one has to consider reasons (relevant facts) and their weight. One need not consider some comparison class for the relevant reasons/facts. Thus, it seems that the reasons-first approach to the normativity of justification is not a promising strategy for an opponent of premise (1).6

At this point one might be concerned whether we are not committing a problematic confusion in our way of thinking about gradability of justification of beliefs. One might think that there are two distinct questions that we tend to confuse in our discussion. On the one hand there is the question of whether epistemic justification is gradable. On the other hand there is the question of whether justified belief is gradable (i.e., whether being epistemically justified is gradable). One might worry that we present considerations that speak in favour of one of these as considerations that speak in favour of the other one. But these considerations need not be interchangeable. This then is problematic insofar as justified belief is a belief that is justified enough for P (where P is a given property). And while justified enough for P is not something that comes in degrees, justification might still come in degrees. According to this worry, it is our confusion of justified belief (which is admittedly absolute and non-gradable) with (epistemic) justification (arguably, gradable) that leads us to think that the non-gradability of justified belief raises serious worries for the existing accounts of epistemic justification. Authors might have failed to be careful enough in their discussions, but the philosophically interesting issue, according to

⁶ At this point one might also want to explore evidence-based views of justification as a possible alternative (cf. Conee and Feldman 2008). However, while evidentialism might constitute a possible substantial theory of justification, it doesn't seem to be a promising way to define/characterize justification in its capacity as a normative term/property. For one thing, evidence is not normative. See Whiting (2018) for a defence of the idea that evidence is not normative; see, however, Hofmann (2020) for a dissenting view. For another thing, the same problem as that with reasons-first approach also seems to arise for evidence-centred views. Having sufficient evidential support (/being overall supported by one's evidence) isn't something relative.

⁷ Thanks to an anonymous reviewer for bringing this worry to my attention!

this worry, is that of epistemic justification, which can be thought of as a scale with a threshold. (One might think that, for instance, when Goldman wrote about degrees of justification, he probably had in mind the link between degrees of reliability and degrees of justification, and would agree that justified belief is non-gradable.) It might be of interest for precision and clarity to point out that we should be more careful and precise that justified belief is not gradable, but we should also admit, according to this objection, that there is no fundamental worry in our proposal for any theory that insists that (epistemic) justification comes in degrees. For the considerations that we have proposed in favour of the absoluteness and non-gradability of justified belief don't support the claim that epistemic justification doesn't come in degrees. The suggestion is that justified belief is to be understood as justified enough for P, and this is similar to being tall enough for P, which doesn't come in degrees, while justification and being tall come in degrees. In short, our considerations, according to this worry, haven't shown that justification doesn't come in degrees, but only that being justified (understood as being justified enough for P) is absolute and, arguably, doesn't come in degrees.

To this worry I would like to respond by providing several independent considerations. First, I would like to suggest that if the worry is on the right track and there's a widespread confusion of mixing together two distinct properties, that of epistemic justification and that of being justified (understood as being justified enough for P), then it is already an interesting result. If the confusion is common, it is worthwhile to explore it. After all, the devil is in the details, and the fact that philosophers are widely confused in lumping these two distinct properties together might point to some philosophically interesting consequences. At any rate, precision is also a value in itself; it allows us to pose better and less confused questions.

That being said, I am not entirely convinced that there is really a confusion here. After all, admitting that we are confused here would also require that we provide an error theory of why so many philosophers are commonly confused in using the two expressions 'justification' and 'being justified' as pointing to the same property and don't realize that they actually point to clearly distinct properties. On the face of it, the above quote from Goldman's famous article does use the two expressions as pointing to the same property. Such a massive confusion asks for a substantial error theory. And I don't see what such an error theory could plausibly appeal to.

More substantially, we can observe that 'justification' is a noun phrase and, like 'tallness', can be used to pick out a property and allow us to talk about that property in a more abstract and general way. It seems to be a common feature of nouns that they pick out a property and ease our more general or abstract discussions about the property in question. Exactly as the noun of 'redness' allows us to talk in a more abstract or general way about being red, that is, about some/all things that are red, the nouns 'justification' and 'tallness' seem to enable us to talk in a more abstract and general way about the property of individuals/entities that are justified or tall. So, on a plausible understanding the property picked up by 'justification' is that of being justified, similarly as in the case of 'tallness' the property referred to by 'tallness' is the property of being tall. In a way, we use these noun constructions to ease the discourse, and it need not imply any reification of being tall or being justified. It is a way to refer to the property that justified/tall individuals/entities have. Crucially then, 'justification' and 'justified' don't

need to pick out distinct properties, exactly as 'tallness' and 'being tall' need not pick out different properties, or 'redness' and 'being red' need not pick out different properties.

Finally, the assimilation of 'justified' to 'justified enough' can also be questioned. Linguists distinguish clearly between a positive form of an adjective (e.g., 'tall', 'open') that expresses a measure function, from adjectives in marked forms (e.g., 'taller than', 'too open', 'open enough'), where degree modifiers are applied to adjectives to form an adjective group (by providing an argument for the relevant measure function) that then combines with the subject. These are two clearly distinct categories that have different syntactic roles. The fact that there is this substantive distinction poses a unique challenge to semanticists in terms of providing a unified treatment of adjectives in positive and marked forms. A classic treatment of this can be found in Kennedy (2007). For instance:

- (3) a. Gradable adjectives map their arguments onto abstract representations of measurement, or DEGREES
- b. A set of degrees totally ordered with respect to some DIMENSION (height, cost, etc.) constitutes a SCALE.
- (...) I will follow Bartsch and Vennemann (1972, 1973), and Kennedy (1999) and analyze gradable adjectives as measure functions (type <e, d>).(..) The adjective expensive, for example, is a function from the subset of the domain of individuals that have some cost value to (positive) degrees of cost. Measure functions are converted into properties of individuals by degree morphology, which in English includes (at least) the comparative morphemes (more, less, as), intensifiers (very, quite, rather, etc.), the sufficiency morphemes (too, enough, so), the question word how, and so forth. Degree morphemes serve two semantic functions: they introduce an individual argument for the measure function denoted by the adjective, and they impose some requirement on the degree derived by applying the adjective to its argument, typically by relating it to another degree. (...)

Turning now to the unmarked, positive form, it is a bit paradoxical that the most morphosyntactically simple form of a gradable predicate turns out to be the hardest to adequately characterize in terms of a compositional semantic analysis. (...) The positive form does not have any overt degree morphology (...). (Kennedy 2007: 4–7)

Thus, it would seem, we should not assimilate 'justified' to 'justified enough' straightforwardly. Yet, the intuition that there is *some* sort of similarity between 'justified' in the positive form and 'justified enough' in the marked form that uses the sufficiency degree modifier is telling. This might even constitute another piece of data in favour of the absolute interpretation of 'justified' in the sense that 'justified' in its positive form leads naturally to an interpretation on which the entity/individual being justified has already (by default) reached sufficiency for one to be justified. Somehow, the sufficiency endpoint seems to be already encoded in the meaning of the positive form 'justified'. This contrasts radically with what we observe in the case of relative adjectives, such as 'tall'. There is no natural tendency to assimilate 'tall' to 'tall enough', for we need more contextual details (we need to know the relevant comparison class) to know what exactly one is tall enough for. Yet, 'justified enough' seems to

be more easily interchangeable with the mere 'justified' in the positive form, suggesting that the positive form already presupposes that the sufficiency endpoint on the relevant background scale for justification has been reached (assuming that the individual is justified).

In short, in reply to the above worry I would like to suggest that while it is an interesting and insightful observation that 'a belief being justified' seems to have (often) a natural paraphrase in terms of 'a belief being justified enough', the suggestion that 'justification' and 'being justified' pick out distinct properties doesn't seem to be theoretically well motivated. The former, a noun phrase, seems to have the function of enabling us to discuss more abstractly and more generally the property of being justified that is referred to by the adjective 'justified'. In this sense, 'justification' and 'justified' seem to stand in the same sort of relation as 'redness' and 'red', for instance.

In this section, we have explored all the most popular and at least initially somewhat plausible ways of understanding justification as a normative property. And none of these seems to be compatible with giving up premise (1). On all of these views, justification seems to be an absolute property. Thus, it seems safe to say that unless one is prepared to give up the normativity of justification, one cannot plausibly reject premise (1). We conclude then that, given the standard assumptions in contemporary epistemology, rejecting premise (1) in an attempt to solve the paradox of graded justification is not a viable option.

4. Prospects of giving up the premise (3): Can there be absolute properties with degrees?

Another attempt to solve the paradox would consist of giving up premise (3). On the face of it, denying premise (3) presents a possible way to effectively solve the paradox. That is, again, this option would avoid a contradiction and would avoid an absurd conclusion. However, to be a fully successful, genuine solution, the option of rejecting premise (3) has to be theoretically well-motivated. The rest of this section explores exactly this. Is there a way to provide independent and theoretically satisfactory motivation for giving up premise (3)? In order to examine the prospects of giving up premise (3), we will consider a possible proposal that relies on the recent developments in the semantics of gradable adjectives. The rest of this section explores the details of such a proposal and looks into what appear to be the main worries with it.

The denial of premise (3) amounts to the claim that (at least some) absolute properties do come in degrees. How could we understand such a proposal? One way to make sense of it would be to appeal to very exciting and insightful developments in recent linguistic treatments of graded adjectives. In this context, a major insight comes from a framework developed by Chris Kennedy and Louise McNally (Kennedy 1999, Kennedy and McNally 2005, Kennedy 2007).

Following Kennedy and McNally, we can distinguish *relative* and *absolute* adjectives within the general category of gradable adjectives (similar enough distinctions can also be found elsewhere; however, for reasons of simplicity of presentation, we focus here on the widely popular framework from Kennedy and McNally). The adjectives *tall*, *rich*, and *empty* are examples of gradable adjectives, whereas *hexagonal*, *pregnant*, and *atomic* are examples of non-gradable adjectives. Among the gradable

adjectives, *tall, rich*, and *small* are paradigmatic examples of relative adjectives, whereas *empty, full, wet, dry, pure*, and *impure* are typical examples of absolute gradable adjectives. Interestingly, and insightfully, Kennedy and McNally distinguish further maximal, minimal, and maximal-minimal adjectives within the general category of absolute adjectives. That is, among absolute adjectives, there are three further subclasses. Absolute maximal adjectives, e.g., *dry, pure*, require an instantiation of the maximum amount of the relevant property by the relevant individual/item for the application of the adjective to be truthful (e.g., a surface is dry only when it has a maximal amount of dryness – when there is no moisture on it). Absolute minimal adjectives, e.g., *wet, impure*, require only a minimal amount of the relevant property (e.g., for a surface to be wet, it suffices for there to be any minimal amount of moisture on it). Absolute maximal-minimal adjectives, e.g., *full, empty, opaque, transparent,* might require either a maximal or minimal amount, depending on further contextually and pragmatically determined factors (see the principle of *Interpretive Economy* in Kennedy 2007: 36).

These differences among gradable adjectives are theorized within Kennedy and McNally's approach by appeal to possible variations of formal features of the underlying scales that correspond to different adjectives (cf. Glanzberg 2014: 274 for a helpful summary). Very roughly, the idea is that any scale has to be of one of the following sorts: a completely open scale (both upper and lower ends of the scale are open), a scale closed only at the upper end, a scale closed only at the lower end, or a totally closed scale (closed at both upper and lower ends). Applying this observation about the topological properties of scales to the semantics of adjectives provides an insight on how to explain the four categories of gradable adjectives. Relative adjectives have totally open underlying scales of the relevant property. The fact that on these scales, there are no upper or lower fixed elements explains the radical, intrinsic context-dependence of relative adjectives. Since no endpoint is fixed for a truthful application of a relative adjective, the threshold needs to be fixed given contextual specifications. So, for instance, only once it's fixed that we are talking about five-year-old children or professional basketball players, or what have you, can the phrase 'x is tall' be evaluated as true or false. The context sets the relevant threshold for tallness, and once the threshold is fixed (which might again shift later, once the context changes), the application of 'tall' will be true or false.

In the case of absolute adjectives, Kennedy and McNally's framework specifies that the endpoints are already fixed independently of any context. This is because absolute adjectives have partially or totally closed underlying scales of the relevant properties. For instance, *dry* has a scale (of dryness) that is closed at the upper end. The endpoint on that scale is fixed. And this explains why for something to be dry, it has to have a maximal amount of dryness. It must hit the upper endpoint on the scale of dryness. *Dry* is an absolute adjective that requires a maximal amount of dryness since the formal structure of the underlying scale is such that it is closed at the upper end. Similarly, the scales corresponding to absolute minimal adjectives are closed at the lower end. Their formal structure explains why any truthful application of these requires that there is at least some minimal amount of the relevant property. *Wet* is an absolute adjective that requires only a minimal amount of wetness to have a true application to an item since the formal structure of its underlying scale is such that it is closed at the lower end. Minimal-maximal adjectives, *opaque*, *transparent*, can either have a scale that's closed at the upper end or a scale that is closed at the lower end. Further pragmatic and contextual factors determine which one is triggered in which context.

The tenets of this approach are summed up by McNally as follows:

Kennedy and McNally argue that the standard value for the truthful applicability of a gradable predicate is, like the scale itself, also subject to linguistically relevant parameterization: Specifically, it can be relative, i.e. determined contextually (typically with respect to a comparison class), or absolute, i.e. fixed at a particular value. (McNally 2011: 3)

A number of tests can be used to classify a given gradable adjective in one of these categories, such as entailments with antonyms, degree modifiers, varying patterns of entailments, acceptability of 'is P for an X' constructions, possibility of natural precisifications and others. For reasons of space these are not reproduced here; the reader is encouraged to refer to Kennedy (2007), Kennedy and McNally (2005), Burnett (2017: 38ff) and further references therein for further details on tests that make it possible to distinguish various categories of gradable adjectives.

Now, one might think that to motivate the rejection of premise (3) of the paradox of graded justification, one might rely on the insights from Kennedy and McNally's framework. One might claim then that being empty is certainly absolute, but may still come in degrees. Just like the adjective 'being empty' is theorized as gradable in the specific sense introduced above, of picking out the upper-end element on the scale of emptiness, being empty (that is, the property) is absolute and requires the maximal amount of emptiness to obtain. Now, if this line of thought is generalized to all absolute properties, it would also apply in the case of being justified. One might argue that 'justified' is an absolute gradable adjective (see Hawthorne and Logins 2020 for more arguments for this line of thought; see also Siscoe forthcoming for a different take; see Staffel 2019: 162–165 for similar considerations and classification of 'rational' as an absolute gradable adjective). But if 'justified' is an absolute gradable adjective, and we are allowed to conclude that the property of being justified is absolute and yet comes in degrees (for instance, in the sense of requiring a maximal amount of justifiedness), premise (3) should be rejected. If this line of thought is accepted, then it would seem we have independently motivated grounds for rejecting (3).

According to the above line of thought then, the source of the paradox of graded justification is that epistemologists speak about graded justification by assuming an open scale model for 'justified'. But, the thought goes, 'justified' is best understood on the basis of a (partially) closed model.

Despite appearing as an initially promising and intriguing strategy, such a move raises some serious worries (we are not questioning the classification of 'justified' within the absolute category of gradable adjectives). Consider the following triad, 'The paradox of absolute adjectives', summed up helpfully by Heather Burnett:

⁸ Note that this doesn't entail that there cannot be any variation in amount of justifiedness in different situations when the property of being justified obtains. Kennedy's framework (cf. Kennedy 2007) allows for specifications for applicability of absolute adjectives. There can be some contextual changes – a bowl I place on the table for breakfast will count as empty, in the ordinary context, but, arguably, not in a scientific context where an absolute vacuum is required for a container to be empty. But this doesn't affect the central point that these variations are not essential to the meaning of 'empty'; its meaning is not intrinsically tied to standards set by a context; the standards are intrinsic – part of the meaning of 'empty'.

The paradox of absolute adjectives

- [a.] If gradability is derived from CC [Comparison Class]-based context sensitivity, and
- [b.] AAs [Absolute Adjectives] are not context sensitive, then
- [c.] How can they be gradable? (Burnett 2017: 70, original emphasis)

The worry here is that while rejecting premise (3) can provide a way to avoid arriving at a contradiction in the case of the paradox of graded justification, it is not clear that the appeal to Kennedy and McNally's framework can avoid a closely related paradox of absolute adjectives.

On a promising and elaborate semantic theory of gradability, it is closely tied to comparison classes. Indeed, Burnett (2017) provides a strong book-length case for the claim that gradability in adjectives is grounded in context sensitivity, understood in terms of comparison classes. But, as Burnett shows, absolute adjectives are not context-sensitive in this sense. Recall that on the Delineation approach that she proposes:

John is taller than Mary is true just in case there is some CC [Comparison Class] with respect to which John counts as tall and Mary counts as not tall. (Burnett 2017: 62)

Whereas:

[I]n order to know which rooms are empty or which sticks are straight, we don't need to compare them to a certain group of other individuals; we just need to look at their properties. [...] [W]hat it means to be non-context sensitive is to have your denotation be invariant across classes. (Burnett 2017: 67)

So, if gradability is derived from comparison class-based context sensitivity, and absolute adjectives are not context-sensitive in this sense, then absolute adjectives (including 'justified') cannot be gradable. Endorsing gradability for absolute adjectives then would lead to a contradiction. Thus, we seem to be in a paradoxical situation; since absolute adjectives do *appear* to be gradable, they seem to pass the relevant tests for gradability.

Burnett's own solution to the paradox of absolute adjectives is to insist that the apparently graded uses of absolute adjectives are 'coerced' uses. In other words, graded uses of absolute adjectives are instances of loose talk. Only on a stretched meaning are absolute adjectives graded. Strictly speaking, absolute adjectives are not graded. Burnett's solution 'lies in the appropriate analysis of the existential context sensitivity property that holds of these predicates [i.e., absolute adjectives] [...] and its relation to the phenomenon of potential vagueness' (Burnett 2017: 72). And the appropriate analysis here is one where 'tolerant and strict structures [are built] on top of classical semantic structures' (Burnett 2017: 72). Burnett refers 'to the (classical) denotations that we assigned to (non-)scalar predicates as their semantic denotations and refer[s] jointly to the secondary tolerant and strict denotations as pragmatic denotations' (Burnett 2017: 73, original emphasis). According to Burnett these 'tolerant and strict denotations are construed from classical (i.e. basic semantic) denotations in conjunction with context sensitive indifference (~ relations)' (Burnett 2017: 73, footnote 12).

Applying this approach to absolute adjectives and to 'justified' in particular leads to the following thought, introduced in Hawthorne and Logins (2020):

On this picture the basic semantics of 'justified' does not associate it with a scale but we can nevertheless "coerce" a scalar use out of it by using general linguistic mechanisms that allow us to generate a scalar use out of expressions that do not get associated with a scale via our foundational understanding of them. (Hawthorne and Logins 2020: 9)

Once we accept the above picture, i.e., Burnett's proposal, and don't think that absolute adjectives have a graded *semantic* (or literal) meaning (but only a coerced graded meaning), it is not easy to see how there could be absolute properties that come in degrees. Why would coerced uses correspond to the relevant properties of individuals?

Note further that in some conversational contexts, we can coerce even clearly and evidently non-gradable adjectives into graded uses. Consider 'hexagonal', which is clearly non-gradable. To take a version of a well-known example, consider a context where one is comparing the shapes of various countries, and in such a situation, one could say, 'France is more hexagonal than Poland is'. This sentence doesn't sound infelicitous. We can make sense of it. But of course, being hexagonal doesn't come in degrees. There is no background scale of hexagonality. It is not the case that things might possess the property of being hexagonal to various degrees. The graded use of 'hexagonal' is coerced and dependent on pragmatic aspects. This coerced use doesn't correspond to the properties of individuals/items. Similarly, if the graded use of 'justified' is coerced (and not literal), then it need not be the case that the property of being justified comes in degrees. The coerced use of 'justified' need not tell us anything substantial about the property of being justified, exactly as the coerced use of 'hexagonal' doesn't tell us anything about the property of being hexagonal.

One might wonder, however, whether gradability and comparison classes are as closely tied as we have suggested above. One source of doubt about the tightness of that connection and about the general explanation of gradability in terms of context sensitivity comes from an observation about adjectives such as wet and similar. That is, from considerations about the minimal absolute adjectives that require a minimal amount of the relevant property for a truthful application of the predicate. The problem is that wet and similar adjectives are clearly absolute. Whether a t-shirt is wet doesn't seem to depend on comparison classes. To determine whether the t-shirt is wet one only needs to look at it and to determine whether there's any amount of moisture in it. One need not look at any other t-shirts or clothes and compare their levels of moisture, for that matter. And yet, being wet seems clearly to come in degrees. My black t-shirt might well be wetter than this other white t-shirt. Perhaps the former fell into the bathtub, while the latter just got some drops of water on it. It is not easy to see why we should think that such gradable uses of wet are coerced uses and should amount to loose talk without being strictly speaking true. Now, the worry for our proposal here is that, as we observed earlier (see footnote 3), we haven't taken a strict stance on which more specific group of absolute adjectives justified belongs to. In particular, we haven't ruled out the option that justified is a minimal (or a minimal-maximal) sort of absolute adjective, just like wet and bent are. But if justified is a minimal (or a minimal-maximal) adjective, then we still have to provide a more convincing argument that, contrary to appearances, the gradability of justified is coerced. On the face of it, the gradability of minimals and minimal-maximals

doesn't seem to be due to mere coercion. In short, we need to do more to show that absoluteness and gradability can *never* go hand in hand and that *justified* cannot be both an absolute and a gradable adjective.⁹

To this worry I would like to respond by putting forward two brief lines of thought. Before going into these, let me observe that a proper treatment of this issue might lead us into a full-blown assessment of the whole Delineation Semantics project, which would lead us to assess our most basic ontological commitments. After all, the most appealing selling point (at least for me) of Delineation Semantics is its minimal ontological commitment. It only needs to appeal to the domain of individuals and it treats properties as functions from individuals to truth values while degrees are reduced to equivalence classes. There is no need to postulate any sort of mysterious abstract objects, such as irreducible degrees that would be included in the domain alongside concrete individuals. This simplicity should be seen as an attractive theoretical feature that cannot be taken out of the ultimate equation in our assessment of Delineation Semantics and the fruitfulness of its treatment of adjectives.

Now, more concretely, and first of all, I would like to acknowledge that the above worry seems to be a substantial one and, as far as I can see, does have problematic consequences for the Delineation Semantic framework and in particular for Burnett's treatment of gradable adjectives. It does seem very plausible to maintain that minimal adjectives still have gradability in their semantic, literal meaning that is not coerced by contextual, pragmatic, extra-linguistic elements. After all, the wetness of a shirt or of a surface can be measured directly. We don't need to know any further contextual details to assess whether a shirt or a surface is wet to a given degree, or whether it is wetter than some other shirt or surface. Yet Burnett seems to be committed to the view that in their semantic, literal meanings, absolute adjectives only have trivial scales. On this view strictly speaking something either is wet or is not, and that's all there is in the semantic meaning of a given application of the adjective wet. Recall also that in the Delineation framework comparative statements with respect to x being more P than y are literally true insofar as there are some context-given comparison classes in which x is P but y isn't in the whole domain of individuals. Of course, this is possible only for relative adjectives, since their semantic meanings, on Burnett's framework, are gappy (similar to indexicals), they are universally contextsensitive and depend in this way on extra-linguistic elements for filling in the details. So, given the radical context sensitivity of relative adjectives (e.g., tall) it is possible for them to satisfy the criteria for truthful comparative statements. Yet, absolutes are not so radically context-sensitive, and their literal meaning doesn't depend on comparison classes. Whether an absolute adjective P can be truthfully (literally) applied to an individual depends on whether the individual is in the semantic denotation of P, which is given by the whole domain and is not radically dependent on comparison classes. So, it looks like the apparent gradability of absolutes, wherever it comes from, cannot come from their semantic meaning. It has to depend on pragmatic factors, and be given by pragmatic denotation. This works well with maximal absolutes such as empty, and straight, since we see how, e.g., loose talk (that gives pragmatic denotation) could explain gradable uses of empty or straight – we are just not precise enough when we use these adjectives in gradable ways. Yet, it is very difficult to see how this could explain the

⁹ Thanks to an anonymous reviewer for this journal for raising this worry. Thanks also to Sven Rosenkranz for raising a similar concern.

apparent gradability of minimal absolutes, e.g., wet, bent. For we may be very precise when we say that one shirt has exactly five drops of water more on it than another shirt and thus is wetter than this other shirt. It would seem that gradable uses with minimals need not involve any loose talk. Thus, on the face of it, I tend to think that the above considerations seems to have some problematic implications for the general Delineation approach, and Burnett's proposal in particular.

However, even if we accept that minimal absolutes have literal gradable meaning, our problems are not gone yet. If justified is indeed (sometimes) like wet and we take the above considerations at face value and question the suggestion that no absolute adjective is strictly speaking gradable, then we might have here a solution to our initial version of the paradox of graded justification, since we might have a good theoretical motivation for rejecting premise (3) of the initial triad ("Absolute properties don't come in degrees (e.g., empty or not?)"). After all, if some absolute adjectives can be gradable, why shouldn't the corresponding properties also come in degrees? (This presupposes that it is the strict use of terms that fix semantic, literal meaning, and that this strict, semantic meaning picks out the relevant property; see our discussion in section 5 below for more on this assumption.) But note that the solution would not be entirely satisfactory, however, since another very close triad would constitute a more specific, updated version of the paradox of graded justification. Namely, we would still have a version of the paradox that focuses only on maximal uses of justified. Even if we admit that justified has minimal uses, it seems difficult to deny that it also has maximal uses. At best, it could be classified as a maximalminimal adjective. But then the paradox could still be constructed specifically for the properties corresponding to maximal uses of justified. Here is one version of the paradox that would still be problematic, assuming that justified has maximal uses:

- (1*) (Epistemic) justification as picked out by maximal uses of *(epistemically) justified* is an absolute maximal property (e.g., similar to being straight, being empty, being flat, being pure, or perhaps being open, being full).
- (2*) (Epistemic) justification always (either as a maximal or minimal property) comes in degrees (e.g., the belief that Earth is spherical is more justified than the belief that the impact of a comet caused the extinction of the dinosaurs).
 - (3*) Absolute maximal properties don't come in degrees (e.g., empty or not?).

Premises (1*)–(3*) lead to a contradiction. That *justified* has maximal uses seems to be rather safe to assume. After all, as mentioned above, *justified* does seem to pass standard tests for absolute maximals: it combines well with degree modifiers that are dedicated to maximals (e.g., *absolutely justified*, *maximally justified*), it seems to validate patterns of inference proper to maximals, and it performs badly in constructions such as 'X is P, but could be more P' ('Peter is justified to think that the gas price will increase but he could be more justified' sounds a bit odd); see Hawthorne and Logins (2021) for more on the claim that *justified* is a maximal. If so, then it would seem that admitting that *justified* also has minimal uses in some contexts and that minimal absolutes have, *pace* Burnett, their gradability semantically encoded (and that the corresponding properties picked out by minimal uses of adjectives come in degrees) doesn't help us to solve a slightly updated version of the paradox of graded justification. We still have to motivate the rejection of one of (1*)–(3*). And as our discussion above

shows, Burnett's approach might well give us the right tools to reject the claim that absolute maximal properties come in degrees. After all, graded uses of maximal adjectives do appear to depend on specific contexts and pragmatically given elements. The appeal to loose talk still seems like an appropriate move in the context of the updated version of the paradox of graded justification. The only worry that we might still have here would be the absence of a simple and unified treatment of both versions of the paradox. Distinguishing strictly between maximal and minimal properties of epistemic justification would also invite further questions about which of these two is the focus of epistemologic theories. Do theorists focus on different properties and when they seem to disagree about it, actually tend to speak past each other? These are uneasy questions that further error theories. Thus, I would tentatively suggest that a solution to our paradoxes of graded justification that doesn't lead us to postulate two distinct properties of epistemic justification would be *ceteris paribus* preferable (of course the more general constraint of parsimony also invites us to explore other options before turning to postulating two distinct properties of justification).

My second line of response stands in tension with what I just admitted one paragraph earlier (and is also a way of exploring a solution to the paradox that doesn't commit us to there being two distinct properties of epistemic justification). It is to explore a potential way of mitigating the apparently strong intuition that the gradability of minimal absolutes is different from the gradability of maximal absolutes, where the former seems to be more strongly encoded in the literal, semantic meaning of the relevant adjectives (the proposal here is inspired in particular by Burnett 2017: chapter 4). The suggestion is that in order to assess whether an adjective is gradable or not, we should look both on its affirmative sense and on its negation. So, we should look not merely at the behaviour of empty and wet alone, but should rather consider the behaviour of the pairs empty-not-empty and wet-not-wet. And this is where it becomes puzzling. On a natural assumption the semantic denotation of an absolute adjective P should straightforwardly correspond to the anti-extension of P. For instance, if empty strictly speaking denotes things that have absolutely nothing in them, then not-empty should strictly speaking denote things that have something in them, even if only a slightest bit of something. Now, we know that empty can be coerced into loose talk and has a pragmatic denotation on which some things that are not absolutely empty can be felicitously called *empty*. Yet no empty thing can ever be seen as not-empty. That is, admittedly not-empty will in no context pick out things that are actually absolutely empty. But how then should we think about the extension of *not-empty* in the very contexts where we use *empty* loosely and count things that are strictly speaking not empty as empty? The pragmatic mechanism that is at play here, according to Burnett, is that of the cognitive indifference (similarity/approximation) relation. This relation governs our cognitive/epistemic judgements and assessments of judgements of indifference (or similarity/approximation) among things and, crucially, this relation is not symmetric, according to Burnett. According to this proposal the cognitive indifference relation allows that x is indifferent to y with respect to some property P, but y is not indifferent to y with respect to P. The suggestion that this surprising feature is indeed in play in our cognitive or epistemic assessments is supported by empirical results in psychology that have been observed in the study of our judgements of prototypicality (cf. Tversky 1977 for instance; see further references in Burnett 2017: 75–76). According to these studies, in some situations we tend to judge elements that are less prototypical of a category as similar/indifferent to the prototypical members of the category, and yet we don't tend to judge the

prototypical elements of the category as similar to the non-prototypical ones. In short, in some contexts our judgements of similarity/indifference appear to be directional. One of the examples that Burnett reports is that subjects in the above-mentioned studies tend to accept the judgement that North Korea is similar to China and yet feel less comfortable with the judgement that China is similar to North Korea (cf. Tversky 1977). Burnett's ingenious suggestion here is that this same mechanism is at play in the cognitive indifference relation applied to absolute adjectives. Empty things are more prototypical – in some contexts we will readily judge things that are not empty as similar/indifferent to empty ones, and yet we will not judge (in any context) not-empty things as similar to empty ones. This is why we will never accept an empty thing as a not-empty one. Appealing to the failure of similarity in the cognitive indifference relation allows us to explain how there is a difference in the apparent gradability of empty and not-empty while maintaining the plausible idea that empty is gradable in the pragmatic denotation sense (when we talk loosely). Now, the crucial point for us is that, on Burnett's proposal, the exact reverse also goes for the minimal absolute adjectives, such as wet, and the parallel pairs, e.g., wet-notwet. The cognitive indifference relation at play here is also not symmetric, yet the prototypical element at play here is *not-wet*. That is, while in some contexts we will judge wet things as similar to not-wet things, we will never judge a not-wet thing as a wet one. The example of the former that Burnett provides is that of a not-bone-dry towel. When looking for a towel after a shower to dry oneself, one will readily accept as not-wet a towel that might have a few drops of water on it, yet we will never accept that a towel that is bone dry could count as a wet one.

The quick detour into Burnett's sophisticated and insightful proposal about the role of the cognitive indifference relation leads us naturally to a suggestion that, contrary to appearances, even the graded uses of wet have to be interpreted pragmatically. If we don't do that, we will have a hard time explaining how gradable uses of not-wet are possible. The appeal to the cognitive indifference relation that is not symmetric and relies on our judgements of prototypicality provides an independently well motivated premise for the argument that apparently graded uses of wet and similar adjectives are to be understood in a pragmatic sense after all. In this they are just like the non-graded uses of not-empty. It would seem we need to appeal to this sort of pragmatic element (that is, the non-symmetric cognitive indifference relation) to make sense of the puzzling behaviour of the pairs empty-not-empty and wetnot-wet. This might then help a proponent of the Delineation approach to account for the intuition that gradability in minimals seems to be encoded semantically. If this suggestion is on the right track, then we could extrapolate from these results and conclude that in the case of justification, we need to accept that justification itself doesn't come in degrees even if 'justified' behaves (sometimes) like a minimal absolute adjective. To claim the contrary would lead us to a puzzle of how to understand the apparently radical non-gradability of the absence of justification (as picked out by not-justified), which corresponds admittedly to the positive property of being at fault (see also Beddor 2017). Exploring this proposal in further detail, however, has to be left for another occasion. As I mentioned at the beginning of my response to the above worry, to explore fully our theoretical options here would lead us to the monumental task of inquiring into some of our most basic ontological commitments and preferences (that is, do we want to include degrees in our ontology as non-reducible abstract objects?).

If what I have proposed in this section is on the right track, then we should give pause to the claim that absolute properties and absolute justification in particular come truly and strictly speaking in degrees.

5. A tentative proposal: justification doesn't come in degrees

The aim of the present section is to elaborate the suggestion that the best option in the face of the paradox of graded justification is to give up premise (2).

The suggestion that justification doesn't come in degrees is to be understood as the claim that strictly speaking belief (or other doxastic states) is never more justified than another belief (or state), where the difference in beliefs might be due to the content of beliefs, the subject who believes, or the time of believing (this is not meant to be an exhaustive list of potential differences). So, for instance, my belief that the Earth is spherical is not more justified than my belief that dinosaurs became extinct because of a meteorite. Although we can express these sorts of comparisons and they don't appear infelicitous, strictly speaking, they are not true.

A proper defence of this proposal has to provide something more elaborate in terms of an error theory of why comparative (and other, typically gradability-involving) constructions with 'justified' appear felicitous, contrary to what the present proposal suggests. The initial data seemed to suggest 'justified' is gradable. Recall that one of the constraints (e.g., the second one; see section 2.2 above) that we fixed for any viable solution to the paradox is that the solution is able to validate or explain away, that is, account one way or another for, the appearance of felicitous uses of graded uses of 'justified'.

Here then is the beginning of an error theory. In explaining why graded uses of 'justified' can be felicitous, we can take inspiration from the view developed by Heather Burnett once more and appeal to the 'loose talk' error theory. According to this error theory, coerced graded uses of absolute gradable adjectives are on a par with coerced graded uses of clearly non-gradable (non-scalar) adjectives, such as 'hexagonal' or 'prime'. Thus, on this view, graded uses of 'justified' (e.g., 'my belief that Earth is circular is more justified than my belief that a meteorite caused dinosaurs to become extinct') are of the same kind as the graded uses of 'hexagonal' or 'prime', e.g., 'France is more hexagonal than Belgium' or '5 is more prime than 6' (compare to Burnett 2017: 70). That is, they appear to make sense; they appear felicitous. But this apparent felicity is explained in terms of loose talk – they appear felicitous only as long as they are loose, and the sentences are not understood strictly speaking.

Here is a further error theoretic element that might provide an additional explanation of the apparent felicity of graded uses of 'justified': getting in the comparative mindset when we are thinking about the closeness of a belief to being justified, while focusing on how probable/reliable/intuitive the relevant target proposition or the belief-producing mechanism is, contributes to the apparent felicity of graded uses of 'justified'. That is, we can easily get into a state of mind where we can compare propositions or belief-producing mechanisms. The standards of comparison can vary. For instance, we can compare propositions with respect to how probable they are, and we can compare belief-producing mechanisms with respect to how reliable they are in terms of leading to true beliefs. And these are

genuinely comparable aspects. Some propositions are more probable than others. Some belief-producing mechanisms are more reliable than others and so on. Now, once we are in this mindset where we have a focus on these comparative aspects of propositions, processes, and so on, we are naturally led to interpret a belief's (or doxastic state's) closeness to being justified as actually having a given degree of justification. That is, beliefs (or states) can vary in terms of how close they are to being justified. This variation in closeness to justification is seen as a variation in degrees of justification. The focus on degrees of probabilities, reliability, and so on, fosters this interpretation of closeness to justification as actual degrees of justification. Of course, this is only the beginning of a viable error theory. A fuller development of the error theory will be left for another occasion. But, if the above line of thought is on the right track, we can see how we could explain why justification appears to but doesn't really come in degrees.

Furthermore, interestingly, if the present proposal is correct and the above error theory doesn't completely miss the point, then the following predictions should bear out. If our proposal and error theory are right, then we would see no unified approach in the literature on what the underlying property is in justification. Why is that? Because there should be many ways of concocting scales for measuring closeness to justification on the spot. That is, if our error theory is right, it should be the case that there is variation in how epistemologists theorize the alleged gradability of justification. If one is in a state of mind with a focus on probabilities, then it is only expected that one will theorize gradability of justification by appeal to degrees of probabilities. If one focuses on reliability, then the supposed underlying scale will be different. And this is what we see in epistemology. The prediction seems to bear out. We have a variety of different theories of degrees of justification. This observation might be taken as an additional point in favour of the above proposal.¹⁰

At this point a reader could worry whether the option of giving up premise (2) of the paradox and accepting that strictly speaking justification doesn't come in degrees is as dramatic as our proposal seems to suggest. That is, one might wonder whether the fact that strictly speaking justification doesn't come in degrees puts any pressure on existing theories of justification that treat it as if it came in degrees. For one could admit that even if strictly speaking justification doesn't come in degrees and appears to be gradable only in the sense of loose talk, our loose talk may still matter. That is, if we can build a precise theory of graded justification, it may well be a theory of justification in a loose sense, but

-

¹⁰ In Hawthorne and Logins (2020) we make a somewhat similar but slightly different point. We classify graded uses of 'justified' within the 'scale-derivative' and 'unruly' category. Where this unruliness can be characterized as follows: "For many coerced uses there is no particularly salient natural scale to turn to when contriving a scale, and whichever scale is constructed on the fly will be highly sensitive to the goals and background of the conversation in which the coerced use is in play" (Hawthorne and Logins, 2020: 12). And we then observe that: "if graded uses are unruly in the way described, then those philosophers who (by our lights incorrectly) proceed as if justification simpliciter were somehow to be analysed in terms of a scale of justification would be all over the place in regards to which kind of analysis they settle upon. And this predicted divergence is arguably precisely what we find" (*ibid*.: 13). Our present proposal is congenial to this suggestion, but it is also somewhat different. On the present proposal the divergence in epistemic theories is predicted by the present error theory; namely, that the divergence comes from the possibility of epistemologists focusing on different graded aspects, e.g. probability, reliability, combined with the idea that closeness to justification does come in degrees. Nonetheless, the two approaches can be perfectly combined, or so it seems.

an epistemically interesting theory nonetheless. The idea here is that the phenomenon of coercion (when absolute adjectives are coerced into graded uses) shows that we can construct scales that we apply to 'justified' when we use it in a loose sense. What this shows, one might worry, is that even though 'justified' (or 'justification' for that matter) doesn't refer to a property that comes in degrees when used in a strict sense, there might be a property (and a corresponding scale) that comes in degrees and is picked up by 'justified' ('justification') when used in a loose sense. And if our use of 'justified' in the loose sense is precise and consistent enough, the property it picks up might play epistemologically interesting and relevant roles, theoretical roles that epistemologists are ultimately interested in. In this sense 'justified' might be similar to 'pregnant', for instance. When a speaker is not asserting anything literally true by 'Mary is more pregnant than Jane is', we may still clearly understand what is communicated in a context, say, that Mary has been pregnant for a longer time than Jane has been. The scale of temporal length of pregnancy is the relevant background scale here that is made salient by the loose talk that is at play here in coercing 'pregnant' into a gradable use. The suggestion then is that the loose talk of degrees of justification might still be meaningful and theoretically usable, indeed insightful for the subject matter that is of interest for epistemologists, even if strictly speaking justification doesn't come in degrees. Thus, according to this worry, denying premise (2) might well be the right option, but the consequences of this move are less relevant for epistemological theorizing than our discussion seems to suggest.¹¹

To this worry I would like to respond that strictly speaking I don't see a worry here but rather a suggestion about one possible way of making our proposal more precise. I would agree that the consequences of denying premise (2) are dramatic only to those approaches that would stick to the claim that their theorizing about justification as coming in degrees is to be taken in the strict, literal sense. Theorists who would agree to relax this requirement and would admit that their theorizing about degrees of justification should be understood in a loose, more approximate sense, as not being strictly speaking about the property of justification but as about something close enough to it (similarly to the case of the loose talk about degrees of pregnancy being about the measure of duration of one's pregnancy), should not be much concerned about the fact that strictly speaking justification doesn't come in degrees. Yet, I also think that this mere fact, the fact that only theorizing in a loose sense about degrees of justification wouldn't be problematic, is already an interesting consequence in itself. What is communicated in the loose sense is notoriously dependent on the context, on pragmatic considerations, one's preferences, interests, and aims. If it makes sense to talk about degrees of pregnancy, it's because of our aims and preferences of using the degree-of-pregnancy talk as a proxy for the duration-ofpregnancy talk. We might be motivated by principles of conversational economy, more fluent and quick discussion. We might assume in our communication a shared background, that our interlocutors are also interested in talking about the duration of pregnancy and so on. The meaningfulness of the loose talk about degrees of pregnancy depends on these assumptions, aims, and preferences in a context of communication. Now, if our theorizing about degrees of justification is to be understood in a loose talk sense, then it too makes sense only given some further assumptions, given our aims, preferences, and possible further pragmatic considerations. This need not be problematic per se. We only need to be

¹¹ Many thanks to an anonymous reviewer for this journal for bringing this point to my attention.

lucid about the potentially surprising consequences here. Once we accept that the gradability talk of justification is loose, we should also admit that differences in theories of gradable justification might be due to pragmatic considerations, preferences, and aims. Perhaps the aims and interests of those who put forward reliabilist theories of graded justification are tied to exploring epistemic roles of reliable belief-forming mechanisms and their talk of gradability of justification reflects their background assumptions and focus on degrees of reliability. And those who focus on probabilistic theories of graded justification have different aims, theoretical aims, and background interests. Presumably, the aims of the proponents of probabilistic theories would be tied to exploring the epistemic import of degrees of probability of given propositions. But if so, we should also realize that disagreements between different theorists might be merely apparent, or at least, due to differences in their background assumptions, aims, and interests. This, of course, opens up new perspectives for thinking about theories of justification but it also brings in new challenges. One such challenge would be how to explain that epistemologists who participate in debates about justification (and assume it is gradable) don't seem to be merely talking past each other, but seem to exhibit genuine theoretical disagreement about how exactly we should theorize epistemic justification. Unfortunately, figuring out how exactly the view that admits that the talk of graded justification makes sense insofar as it is to be understood in a loose sense can meet the challenge of explaining the apparently genuine disagreement among theorists of justification that is not merely based on background assumptions, aims, and theoretical interests has to be left for another occasion. I would merely like to put forward the more general, nonspecific conclusion that theories that assume that epistemic justification comes in degrees strictly speaking have problematic consequences and should be revised given our results about the paradox of graded justification. The topic of whether and how loose talk of graded justification can be theoretically insightful is left for another occasion.

6. Conclusion

Solving a paradox typically involves making concessions. By its nature, a paradox is an inference such that independently plausible premises lead to a contradiction. Thus, avoiding a contradiction in paradoxes in non-ad-hoc ways typically involves tradeoffs.

The situation is no different in the case of the paradox of graded justification. Any viable solution to the paradox will involve some tradeoffs. In what precedes, we have specified the paradox, clarified its premises, and then considered the main options that we have in the face of it. The paradox arises from a combination of some otherwise plausible assumptions, namely, (1) that epistemic justification is an absolute property (like being straight, being empty); (2) that epistemic justification comes in degrees (e.g., the belief that Earth is spherical seems to be more justified than the belief that the impact of a comet caused the extinction of the dinosaurs); and (3) that absolute properties don't come in degrees (e.g., something is empty or not, and it feels strange to think of degrees of emptiness). We argued that rejecting premises (1) and (3) leads to implausible results. Our preferred solution then is to deny (2). Of course, denying the idea that justification comes in degrees might still appear counterintuitive for many, especially for proponents of epistemological theories that rely essentially on the idea that justification is gradable. But such theoretical proclivity on its own need not be a reason to maintain

premise (2). Moreover, we have proposed the beginning of an error theory for why it might seem that justification comes in degrees but ultimately does not. Crucial to our proposal is the claim that graded uses of 'justified' make sense as long as we are talking loosely. In a sense, to say that my belief that Earth is spherical is more justified than my belief that a comet caused the extinction of the dinosaurs is similar to saying that France is more hexagonal than Belgium. We can surely make sense of such talk. We seem to understand what a speaker might be up to when she makes such comparisons. However, such comparisons make sense only if they are not taken strictly speaking or literally. If the above proposal is on the right track, then we have here a solution to the paradox of graded justification. This solution has further repercussions and consequences; namely, that views that tie justification straightforwardly and directly to probabilities, reliability, coherence, or similar clearly graded categories cannot be taken to be literally true. Thinking about these graded elements cannot be transposed directly and literally into a theory of the property of epistemic justification. Beyond this implication, the present proposal is compatible with a wide range of popular views about justification.

Bibliography

Ahlstrom-Vij, Kristoffer and Dunn, Jeff (2014). A defence of epistemic consequentialism. *Philosophical Quarterly*, 64 (257):541–551.

Bale, Alan Clinton (2011). Scales and comparison classes. Natural Language Semantics 19 (2):169–190.

Beddor, Bob (2017). Justification as faultlessness. Philosophical Studies 174 (4):901–926.

Bird, Alexander (2007). Justified judging. Philosophy and Phenomenological Research 74 (1):81–110.

Boghossian, Paul A. (2003). The normativity of content. Philosophical Issues 13 (1):31-45.

Bradley, Darren (2015). A Critical Introduction to Formal Epistemology. London: Bloomsbury.

Burnett, Heather (2017). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford: Oxford University Press.

Conee, Earl and Feldman, Richard (2008). Evidence. In Quentin Smith (ed.), *Epistemology: New Essays*. Oxford University Press.

Cresswell, Max J. (1976) The semantics of degree. In Montague Grammar, pp. 261–292. Academic Press.

Dunn, Jeff (2015). Epistemic consequentialism. In Internet Encyclopedia of Philosophy.

Engel, Pascal (2004). Truth and the aim of belief. In Donald Gillies (ed.), *Laws and Models in Science*, pp. 77–97. London: King's College Publications..

¹² Many thanks to the audience of a seminar at the eidos research group at the University of Geneva, Claudio Calosi, Fabrice Correia, Robert Michels, Giovanni Merlo, Sven Rosenkranz, and an anonymous reviewer for this journal for discussion and comments on earlier versions of this article. The research work that led to this article was supported by the Swiss National Science Foundation Grant 'Graded Epistemology' Number 186137.

Glanzberg, M. (2014). Explanation and partiality in semantic theory. In A. Burgess and B. Sherman (eds.), *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press.

Goldman, Alvin I. (1979). What is justified belief. In George Pappas (ed.), *Justification and Knowledge*, pp. 1–25. Boston: D. Reidel.

Goldman, Alvin I. (1986). Epistemology and Cognition. Harvard University Press.

Goldman, Alvin I. (1999). Knowledge in a Social World. Oxford University Press.

Hawthorne, John and Logins, Artūrs (2020). Graded epistemic justification. *Philosophical Studies*, 1–14.

Howson, Colin & Urbach, Peter (2006). *Scientific reasoning: the Bayesian approach*. Open Court Publishing.

Hofmann, Frank (2020). Is evidence normative? Philosophia, 1-18.

Joyce, James (2009). Accuracy and coherence: prospects for an alethic epistemology of partial belief. In Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of Belief*, pp. 263–297. Synthese.

Joyce, James M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65 (4):575–603.

Kelp, Christoph (2018). Good Thinking. A Knowledge First Virtue Epistemology. London: Routledge.

Kennedy, Christopher (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison.* New York: Garland.

Kennedy, Christopher (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30 (1):1–4.

Kennedy, Christopher and McNally, Louise (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 82 (2):345–381.

Klein, Ewn (1991). Comparatives. In Arnim von Stechowand Dieter Wunderlich (eds.), *Semantik: Ein internationales handbuch der zeitgenssisöchen forschung*, pp. 673–691. Berlin: Walter de Gruyter.

Lassiter, Daniel (2017). *Graded Modality: Qualitative and Quantitative Perspectives*. Oxford: Oxford University Press.

Lewis, David (1979). Scorekeeping in a language game. Journal of Philosophical Logic 8 (1):339–359.

Littlejohn, Clayton (2013). XV—The Russellian Retreat. *Proceedings of the Aristotelian Society*, 113 (3pt3):293–320.

McNally, Louise (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In R. Nouwen, R. van Rooij, U. Sauerland and H. C. Schmitz (eds.), *Vagueness in Communication. VIC 2009. Lecture Notes in Computer Science*, vol. 6517. Berlin and Heidelberg: Springer.

Nelson, Mark T. (2010). We have no positive epistemic duties. Mind, 119 (473):83-102.

Parfit, Derek (2011). On What Matters: Two-Volume Set. Oxford University Press.

Pettigrew, Richard (2013). Epistemic utility and norms for credences. *Philosophy Compass*, 8 (10):897–908.

Pettigrew, Richard (2015). Accuracy and the credence-belief connection. *Philosophers' Imprint*, 15:1–20.

Pollock, John L. and Cruz, Joseph (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield Publishers.

Recanati, François (2010). Truth-Conditional Pragmatics. Oxford University Press.

Scanlon, Thomas (1998). What We Owe to Each Other. Belknap Press of Harvard University Press.

Schroeder, Mark (2007). Slaves of the Passions. Oxford University Press.

Schroeder, Mark (forthcoming). Reasons First. Oxford University Press.

Shah, Nishi and Velleman, J. David (2005). Doxastic deliberation. Philosophical Review, 114 (4):497-534.

Simion, Mona, Kelp, Christoph, and Ghijsen, Harmen (2016). Norms of belief. *Philosophical Issues*, 26 (1):374–392.

Singer, Daniel J. (2018). How to be an epistemic consequentialist. *Philosophical Quarterly*, 68 (272):580–602.

Siscoe, Wes (forthcoming). Belief, rational and justified. *Mind*.

Smith, Martin (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press. Sorensen, Roy A. (1991). Rationality as an Absolute Concept: Roy A. Sorensen. *Philosophy* 66 (258):473-486.

Staffel, Julia (2019). Unsettled Thoughts: A Theory of Degrees of Rationality. Oxford University Press.

Stechow, Arnim von (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3(1-2): 1-77.

Steup, Matthias (2001). Epistemic duty, evidence, and internality. In Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.

Sutton, Jonathan (2007). Without Justification. MIT Press.

Tversky, Amos (1977). Features of similarity. Psychological Review, 84:327-352

Unger, Peter (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.

Whiting, Daniel (2018). Right in some respects: reasons as evidence. *Philosophical Studies*, 175 (9):2191–2208.

Wedgwood, Ralph (2002). The aim of belief. Philosophical Perspectives, 36 (s16):267–297.

Wedgwood, Ralph (2013). Doxastic correctness. *Aristotelian Society Supplementary Volume*, 87 (1):217–234.

Wellwood, Alexis (2019). The meaning of *more*. Oxford University Press.

Williamson, Timothy (2000). Knowledge and its Limits. Oxford University Press.

Williamson, Timothy (2011). Knowledge first epistemology. In Sven Bernecker and Duncan Pritchard (eds.), *The Routledge Companion to Epistemology*, pp. 208–218. Routledge.

Williamson, Timothy (2014). Very improbable knowing. Erkenntnis, 79 (5):971–999.

Williamson, Timothy (forthcoming). Justifications, excuses, and sceptical scenarios. In Fabian Dorsch and Julien Dutant (eds.), *The New Evil Demon*. Oxford: Oxford University Press.

Wolterstorff, Nicholas (1983). Can belief in God be rational if it has no foundations? In Alvin Plantinga and Nicholas Wolterstorff (eds.), *Faith and Rationality: Reason and Belief in God*, pp. 135–186. Notre Dame, IN: University of Notre Dame Press.