Rational Supererogation and Epistemic Permissivism

Robert Weston Siscoe
Florida State University
wsiscoe@fsu.edu

Abstract

A number of authors have defended permissivism by appealing to rational supererogation, the thought that some doxastic states might be rationally permissible even though there are other, more rational beliefs available. If this is correct, then there are situations that allow for multiple rational doxastic responses, even if some of those responses are rationally suboptimal. In this paper, I will argue that this is the wrong approach to defending permissivism – there are no doxastic states that are rationally supererogatory. By the lights of contemporary linguistics, 'rational' is an absolute gradable adjective, and as such, can only be applied to things that satisfy the top of the scale of rationality. For this reason, it is not possible to believe what is rational while also failing to believe what is rationally optimal.

Keywords: Epistemic Supererogation, Rationality, Uniqueness, Permissivism, Gradable Adjectives

Word Count: 7,975

Introduction

A number of authors have noted that, despite there being a number of striking analogies between the moral and epistemic domains, there has been very little discussion of whether there are cases of epistemic supererogation:

Jennifer Nado – Ethicists standardly recognize several deontological statuses an action may have: it may be obligatory or forbidden; it may be permissible; it may be supererogatory [...] Strangely enough, however, the category of supererogation is rarely discussed within epistemology.[1]

Trevor Hedberg – Supererogatory acts, those which are praiseworthy but not obligatory, have become a significant topic in contemporary moral philosophy [...] However, despite the similarities between ethics and epistemology, epistemic supererogation has received very little attention.[2]

Han Li – Moral theorists, of course, have long written about [...] "supererogatory" actions. Yet despite the relatively large and varied literature on moral supererogation, practically nothing has been written about its epistemic counterpart.[3]

As Nado, Hedberg, and Li all point out that, even though moral and epistemic theorizing have a number of commonalities, this has not resulted in many investigations of epistemic supererogation. This is surprising given all of the commonalities between the two normative spheres – there are both actions that we are required to perform as well as propositions that we are required to believe, as well as actions and beliefs that are merely permitted but not required.

Despite the promising parallels between moral and epistemic normativity, however, further exploration of the analogy between moral and epistemic supererogation raises a number of questions. An issue that immediately arises is that simply saying that a doxastic state is epistemically supererogatory is not yet precise enough – to say that a belief is supererogatory in an epistemic sense does not isolate the way in which it goes above and beyond the call of duty. This is because 'epistemic' is descriptive rather than evaluative. Actions can be evaluated as moral or immoral, but it is forced at best to evaluate a belief by calling it 'epistemic.' So even though it is relatively simple to move from speaking about moral actions to speaking of morally supererogatory actions, there is an open question with epistemically supererogatory beliefs about the sort of evaluation that is at play.

In what respect are epistemically supererogatory beliefs exemplary? Doxastic attitudes can be evaluated as wise, rational, intelligent, and reasonable, so

---

[1]See Nado (2019), p. 121.
[2]See Hedberg (2014), p. 3621.
[3]See Li (2018), p. 351.

there are a number of dimensions along which beliefs might be epistemically supererogatory. While accounts of epistemic supererogation might be worth exploring for all of these spheres of evaluation, in this paper our focus will be on theories of **Rational Supererogation**, accounts that take it that doxastic states can be both rationally permissible and rationally better than other permissible alternatives:

Rational Supererogation

Advocates of Rational Supererogation think that rational doxastic states are not always the states that are most rational. Instead, rationality sometimes permits believing what is rationally suboptimal. This view is taken up by theorists like Han Li who says that "there are beliefs that are rationally permissible to hold even if they are not maximally rational."[4] Thus, according to proponents of Rational Supererogation, rationality allows for supererogatory doxastic states, states that are both rationally permitted and more rational than other permissible options.[5]

Rational Supererogation is not without consequence. Most notably, it has been wielded to defend permissivism in the uniqueness/permissivism debate. According to Christensen's (2007) canonical account of uniqueness, "there is a unique maximally epistemically rational response to any given evidential situation."[6] If Christensen is correct, this is often taken to show that there is only one doxastic attitude that an agent can adopt and remain rational, a rejection of permissivist views on which multiple doxastic attitudes can be rational on the same evidence. Yet if **Rational Supererogation** is true, then it is possible both that Christensen has things right about maximal rationality and that there are other rationally permissible attitudes besides the maximally rational one that believers can adopt.

In this paper, I will argue that **Rational Supererogation** is false, preventing this line of defense for the permissivist. 'Rational' is an absolute gradable adjective, and as such, always picks out the uppermost point on its underlying scale. The rational doxastic state, then, is the maximally rational doxastic state, refuting **Rational Supererogation**.[7] Here's how we'll proceed. Section

---

[4]See Li (2019), p. 171.

[5]Along with Li, Enoch (2010) and Jackson (2019) also explore accounts of epistemic rational supererogation. Because it is possible to distinguish between epistemic and practical rationality, theories of rational supererogation can extend to both normative domains. Due to limitations of space, this paper will only consider epistemic cases of rational supererogation, though for an application of how the ideas discussed in this paper might apply to the practical case, see Siscoe (Forthcoming). For supporters of rationally supererogation in the practical domain, see Benn and Bales (2019), Kawall (2003), McElwee (2017), and Slote (1989).

[6]See Christensen (2007), p. 190 and 192.

[7]'Rational' can modify a number of different objects beyond just doxastic states – we also often speak of rational actions and rational persons. For the majority of this paper, I will confine my attention to the behavior of 'rational' as it applies to doxastic states, though it may be that 'rational,' when applied to other sorts of entities, is not an absolute gradable adjective, a possibility I consider in Section 4.

1 begins by further detailing **Rational Supererogation** proposals, showing how these accounts have been used to defend permissivism. I will then turn to the distinction between relative and absolute gradable adjectives in Section 2, arguing in Section 3 that 'rational' is an absolute gradable adjective. In Section 4, I respond to the objection that 'rational' can sometimes behave like a relative gradable adjective by drawing a distinction between the capacity sense and the sanctioning sense of 'rational.' I then conclude by considering how permissivists can move forward given the impossibility of **Rational Supererogation**.

1   RATIONAL SUPEREROGATION AND EPISTEMIC PERMISSIVISM

Han Li gives a range of cases that he claims are instances of **Rational Supererogation**. In particular, Li explores cases in which agents exercise extraordinary rational capacities. In the Sherlock Holmes stories written by Sir Arthur Conan Doyle, Sherlock Holmes and Doctor Watson form a crime-solving duo. Sherlock Holmes is a brilliant detective, often seeing solutions that would not have occurred to the casual observer. Even though Doctor Watson is intelligent, he cannot reach Holmes's genius solutions. Li argues that cases like these have all the elements of **Rational Supererogation**:

> Holmes and Watson share the same evidence — there is nothing relevant that Holmes knows which Watson doesn't. Assume also that the shared evidence really does support the solution Holmes eventually came to. Holmes is a peculiarly brilliant epistemic agent — most of us would not solve the case along with Holmes, putting us in Dr. Watson's shoes. Clearly this fact makes us epistemically worse than Holmes, at least in this particular situation. But [...] is our doxastic response to this evidence actually irrational?[8]

Let's assume, for the sake of argument, that Sherlock and Watson have yet to discuss the case, and so Watson does not have the benefit of Sherlock's testimony about his ingenious solution. At this point, Holmes and Watson share all of the same evidence, and yet Holmes adopts a response that is based on rational capacities that are superior to Watson's. At the same time though, it seems wrong to say that Watson is irrational. Just because Watson doesn't display the extraordinary acumen of Sherlock, this does not make him and his doxastic response irrational. If all this is right though, then when Holmes solves a case that stumps Watson, Holmes exhibits **Rational Supererogation**. Watson forms a rationally permissible doxastic response even though there was a more rational response available, making the conclusion that Holmes reached supererogatory.

Li appeals to cases like these to defend intrapersonal permissivism, the thought that there are multiple doxastic states that are rational for an agent on a single body of evidence. Even if we suppose that Sherlock reached the most rational

---

[8]See Li (2018), p. 350.

solution given the evidence, we can still count Watson's belief as rational as well, leading Li to summarize his permissivist take of the situation as follows:

> Even if there is always one unique belief that is maximally rational to hold in any given situation, perhaps agents are not irrational if they fail to hold this belief. Sometimes, there are beliefs that are rationally permissible to hold even if they are not maximally rational. Maybe rationality sometimes cuts us slack, letting it go even if we don't always do the optimal thing. For rationality, sometimes good enough is good enough."[9]

On Li's view, Sherlock adopts the maximally rational doxastic state while Watson adopts a less rational, but still rationally permissible, doxastic response. Li takes this to demonstrates that, even if Christensen is right about maximal rationality, rationality still allows beliefs that are less than fully rational to count as rational. This confirms the permissivist thought that, given a single body of evidence, there may be multiple doxastic states that are rational to adopt.[10]

Li is not the only one to appeal to epistemic supererogation in the uniqueness versus permissivism debate. Like Li, David Enoch and Elizabeth Jackson both suggest that epistemic supererogation can help make room for permissivism. Enoch proposes that the permissivist might be able to appeal to epistemic supererogation, suggesting a view on which "there is one degree of belief that is maximally rational, but some other degrees of belief — though less than maximally rational — are still rationally permissible."[11] Jackson argues that critical reflection on one's evidence is not always required. There are many everyday propositions that, if we reflected on our evidence, we would come to rationally believe, but such reflection is supererogatory. So "if critical reflection on everyday beliefs is epistemically good but not required, this supports the idea that this kind of reflection is supererogatory."[12] It could be, then, that some beliefs that are arrived at by critical reflection are rationally supererogatory, while beliefs that neglect this sort of reflection are nevertheless rationally permissible. This would confirm the permissivist view that there are multiple doxastic states that are rational on a single body of evidence.

One thing to note is that one of these proposals is described in terms of epistemically supererogatory *states* and the other in terms of epistemically supererogatory *actions*. Enoch suggests that there may be some degrees of belief that are rationally supererogatory, while Jackson suggests that the action of critical reflection is supererogatory. In **Rational Supererogation**, I speak in terms of states rather than acts. This is because it is controversial whether there are any epistemic norms or reasons that apply to actions – a number of epistemol-

---

[9] See Li (2019), p. 171.

[10] Because this paper is focused on **Rational Supererogation**, I do not consider whether or not rational ties might be helpful for the permissivist.

[11] See Enoch (2010), p. 957.

[12] See Jackson (2019), p. 6. See also Tidman (1996) and Hedberg (2014) for similar examples.

ogists take it that actions and belief-like states are governed by different sorts of norms.[13] Now it may be that much of what is said about epistemically supererogatory actions can be described entirely in terms of supererogatory states. For example, if we are considering someone who hasn't undergone critical reflection, Jackson claims that they are "perfectly rational in having certain beliefs, but their situation is such that there are other beliefs, inconsistent with the first set, that 'go above and beyond' what rationality requires."[14] These comments suggest that we can understand critical reflection as leading to a supererogatory state. The same can be said of Li's original case — it could be that Sherlock reached an epistemically supererogatory state via critical reflection. Thus, even if it is inaccurate to describe these situations in terms of performing a supererogatory epistemic action, we can still interpret them as claiming that the protagonist adopted a supererogatory epistemic state.[15]

2 Relative and Absolute Gradable Adjectives

On **Rational Supererogation**, being rational only requires being rational enough. Thus, it is possible to believe rationally even while not believing what is most rational. For the remainder of this paper, I will argue that **Rational Supererogation** is mistaken – the only rational doxastic states are those that are at the top of the scale of rationality. In the next two sections, I will make the case for this position by showing that 'rational' is an absolute gradable adjective.[16] The meanings of absolute gradable adjectives are fixed on the tops of their scales, and so if 'rational' is an absolute gradable adjective, then only fully rational doxastic states qualify as rational.

According to contemporary linguistics, gradable adjectives are divided into relative and absolute.[17] Both relative gradable adjectives (RAs), including 'large,' 'long,' and 'tall,' and absolute gradable adjectives (AAs), words like 'closed,' 'straight,' and 'pure,' employ an underlying degree scale, allowing them to be used comparatively.[18]

---

[13] For proponents of this view, see Berker (2018), Feldman (2000), Harman (2004), Kelly (2002) and (2003), and Whiting (2014).

[14] See Jackson (2019), p. 7.

[15] There are other discussions in the literature of epistemic supererogation beyond what we have discussed here, though the majority are primarily interested in accounts of supererogatory epistemic actions; see Bernstein (1986), Hedberg (2014), McElwee (2016), Nado (2019), and Tidman (1996).

[16] The view that 'rational' is an absolute gradable adjective is anticipated by Sorensen's (1991) claim that rationality is an absolute concept.

[17] Those who distinguish between relative and absolute gradable adjectives include Kennedy (2007), Kennedy and McNally (2005), Rotstein and Winter (2004), and Rusiecki (1985).

[18] Due to its ability to capture the distinction between RAs and AAs (Kennedy, 2007), the scale approach taken by the Cresswell (1977), Heim (2000), Kennedy (2007), and von Stechow (1984) has been the most influential. There is also a view, however, on which the basis for comparisons between gradable adjectives are quantifications over possible precisifications of that adjective's extension. For more on this view, see Fine (1975), Kamp (1975), Klein (1980), Larson (1988), and Pinkal (1995).

Relative Gradable Adjective	Comparative
tall	The adult is taller than the child
long	The road is longer than the trail
large	The fox is larger than the mouse
Absolute Gradable Adjective	Comparative
straight	The line is straighter than the stick
pure	The gold is purer than the silver
closed	The first door is more closed than the second

If one rope has a greater degree of length than another, the first rope is longer than the second. Likewise, if a piece of gold has fewer impurities than a piece of silver, then the gold is purer than the silver.[19]

$$[\![\text{er}|\text{more}]\!]([\![\text{pure}]\!])([\![\text{than the silver}]\!])$$
$$=\big[\lambda g\lambda y\lambda x.g(x)\succ g(y)\big](\text{pure})(\text{the silver})$$
$$=\big[\lambda y\lambda x.\text{pure}(x)\succ\text{pure}(y)\big](\text{the silver})$$
$$=\lambda x.\text{pure}(x)\succ\text{pure}(\text{the silver})$$

Despite this similarity, however, RAs and AAs behave quite differently when it comes to their positive, unmarked forms. Consider, for example, the unmarked form of 'tall' in (1):

(1) The basketball player is tall

Even if it is true that a basketball player is taller than a child, this by itself does not make (1) true. In order to determine the truth of (1), we need to know more information about the context. Are we trying to determine whether the basketball player is tall for a basketball player, or just for a person more generally? Depending on the contextual standard at play, (1) may take different truth values. Thus, in appraising the truth or falsity of (1), both the height of the player and the standard of comparison are relevant. If the player is clearly taller than the contextual threshold at issue, then (1) is true, and if the player is clearly shorter, then (1) is false. We thus have it that even the unmarked, *pos* constructions of gradable adjectives have a place for comparison in their semantics:

(2) $[\![[_{\text{DegP}}\, pos\, [_{\text{AP}}\, \text{tall}]]\!] = \lambda x.\text{tall}(x) \succeq c(\text{tall})$

With the unmarked form of gradable adjectives, an object is judged by how it measures up to the function c, a function that chooses a standard of comparison such that the objects that satisfy it "stand out" along the underlying dimension.[20] Thus, a statement like (1) is true only if the player stands out in terms

---

[19]Following Kennedy and McNally (2005), p. 369, I treat the semantics of 'purer' and 'more pure' identically.

[20]Kennedy (2007), Kennedy and McNally (2005), and Rotstein and Winter (2004) all advocate for the view on which the extension of the *pos* form of gradable adjectives includes

of their height relative to the standard of comparison picked out by our function c.

As we have seen, alterations in the domain will effect our c function, changing what it takes in order to stand out along a particular dimension. This domain-shifting can occur even in the case of one-off comparisons. Because RAs have an adjustable contextual threshold, they can be used to differentiate between objects that fall in the middle of the underlying degree scale.[21] For example, suppose that you were presented with Pipes 1 and 2 and given the following command:

(3) Point to the long one

The fact that RAs make use of fairly accommodating contextual thresholds can be seen with the felicity of (3). Even though Pipes 1 and 2 fall relatively close to one another on the underlying scale, we can still differentiate between them using the unmarked 'long'. The thresholds of RAs, however, are not endlessly accommodating. Take two very long roads, the Pan American Highway, coming in at over 18,000 miles, and Highway 1 in Australia, a length of over 9,000 miles. Even though the Pan American Highway is much longer than Highway 1, (3) is now infelicitous. Though the contextual threshold can draw distinctions between objects in the middle of the scale, like Pipes 1 and 2, the threshold in question cannot be used in such a way that 9,000 miles fails to be long. This is because a 9,000 mile road can never fail to stand out on the underlying scale of length, preventing a comparison between the Pan American Highway and Highway 1 using the *pos* form of 'long.'

Even though shifting the domain can effect the standards of application for RAs, the same does not occur with AAs. Whereas RAs permit one-off comparisons between two objects in the middle of the scale, the same type of command is anomalous with AAs:

(4) #Point to the closed one
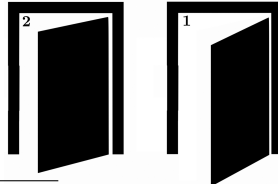
those items that "stand out" relative to the contextual threshold. The view that gradable adjectives enlist a standard of comparison, however, has a much longer history – see Barker (2002), Bartsch and Vennemann (1972), Bierwisch (1989), Cresswell (1977), Fine (1975), Kamp (1975), Klein (1980), Lewis (1970), Pinkal (1995), Sapir (1944), von Stechow (1984), and Wheeler (1972).

[21]See Kennedy (2007), Kyburg and Morreau (2000), Sedivy et al. (1999), and Syrett et al. (2006 and 2010).

Even though Door 2 is more closed than Door 1, (4) is infelicitous. This is because the AA 'closed' does not operate using the same sort of contextual threshold as RAs, preventing (4) from drawing a distinction between the more closed Door 2 and the less closed Door 1. The behavior of 'closed' also departs from that of RAs at the extremes of the underlying scale. If Door 2 is completely shut and Door 1 is slightly ajar, (4) would be felicitous because the uppermost degree is what stands out on the underlying scale. RAs and AAs thus differ in the following way – RAs can be used to distinguish between objects in the middle of a scale but not at the scale's extreme, while uses of AAs cannot distinguish between objects in the middle of the scale, but can at the end of the scale.

What explains the divergence between RAs and AAs in the "point to" test? AAs can be used to distinguish between objects at the scale's extreme because AAs have closed scales. If a door is completely sealed shut, it is not possible for it to be closed any further. Likewise, if a piece of gold is one hundred pure, then it is not possible for it to become any more pure. RAs, on the other hand, have open scales. It is always possible to be a bit longer or a bit taller, at least conceptually speaking, as the scales of tallness and longness are not closed. One way to appreciate this distinction is to see how RAs and AAs pair differently with maximal degree adverbs, terms like 'totally', 'perfectly', and 'completely'. When it comes to AAs, these terms can be used to pick out the top of the underlying scale:

    (5) The door is completely closed
    (6) The line is perfectly straight
    (7) The gold is totally pure

The same is not true, however, with RAs:

    (8) #The fox is completely large
    (9) #The river is perfectly long
    (10) #The skyscraper is totally tall

Because maximal degrees readings of RAs are not possible, (8)–(10) are infelicitous. AAs, on the other hand, allow for maximal degree interpretations, making (5)-(7) acceptable.[22]

Since absolute gradable adjectives have top-closed scales, they are not context sensitive in the same way as RAs. RAs employ a contextually-determined threshold because they lack a semantically-given default value to fix their meaning, but AAs already come with obvious default values, their maximal points.[23]

---

[22]There is reason for caution when using this diagnostic, as words like 'totally' and 'completely' can, in certain contexts, mean something very close to 'very'. For more, see Kennedy (2007), pp. 34-35, and Kennedy and McNally (2005), p. 354, and Section 3 of this paper.

[23]Kennedy (2007) calls the principle that enlists the top of the scale in a term's meaning the Principle of Interpretive Economy. For more on Interpretive Economy and its relationship to relative and absolute gradable adjectives, see Franke (2012), Frazier, Clifton, and Stolterfoht (2008), Rett (2014), and Sassoon (2012).

These maximal points stand out relative to the other points on their underlying scales, enriching our standard of comparison function c with the fixed, top point on the scale. For this reason, domain-shifting can effect the application conditions of RAs, whereas the application conditions of AAs are always centered on their uppermost points.

One way to see the semantic impact of this uppermost point is through instances of natural precisification. RAs are often vague, and their meanings can only be made more precise by introducing stipulative definitions. AAs, on the other hand, allow precisification in the midst of a natural language context because they are often used imprecisely.[24] Even though it may be permissible to describe a door as closed when it is still open an inch, nothing but being completely closed will do when the door in question is the entrance to an airplane that is about to take off. It is thus possible to rule out barely open doors with a statement like (11):

(11) The airplane door needs to be closed, but right now it is still open an inch.

Because what matters is that the door is completely closed, a sentence like (11) raises the standard of precision, naturally precisifying the use of 'closed.' Likewise, even though Britannia silver only permits impurities up to four percent, a manufacturer attempting to make fine silver may rule out loose uses of 'pure' with (12):

(12) The silver for the bullion bars needs to be pure, but this is Brittania silver, so it will not do.

With (12), the manufacturer makes it clear that imprecise uses of 'pure' are unacceptable. If before 'pure' was being used in a loose sense, that is now no longer possible. AAs thus allow eliminating loose speech via natural precisification, confirming that they can be used imprecisely.

Vague RAs, on the other hand, do not allow for natural precisification. Suppose that a zoologist is looking for a record-length snake (the longest snake on record is a thirty-three foot python). Even with this goal, it is not possible to precisify 'long' to only pick out record-setting snakes:

(13) #The snake needs to be long, but this one is only thirty-one feet so it will not do.

With (11), we made it clear that a barely open door does not fall in the extension of 'closed,' but (13) cannot do the same with our snake and the extension of

---

[24]Technically speaking, vagueness and imprecision are properties, not of particular gradable adjectives, but of particular *uses* of gradable adjectives. As shorthand, I will speak of vague and imprecise terms for terms that can give rise to vagueness and imprecision in particular contexts, but vagueness and imprecision should always be thought of as a property of specific uses of gradable adjectives.

'long.' This is because 'closed' has a precise point that constrains its meaning, making it possible to naturally precisify it up to its maximal point, while no such thing is possible with the RA 'long.' We can thus sort vague RAs from imprecise AAs terms by testing for natural precisification. Because AAs employ an uppermost point on their scale, it is possible to eliminate loose talk from consideration.

3   'Rational' as an Absolute Gradable Adjective

As we can now see, whether there is conceptual space for a distinction between rationality and maximal rationality will depend on whether 'rational' is a RA or an AA. If 'rational' is an absolute gradable adjective, then supporters of **Rational Supererogation** are mistaken to think that it is not possible to do what is rational without doing what is most rational. AAs pick out the uppermost point on their underlying scales, making 'rational' apply only to actions that are fully rational. In this section, I will argue that 'rational' is an absolute gradable adjective, laying the groundwork for discussions of **Rational Supererogation** and the permissivism/uniqueness debate in Sections 4 and 5.

Let's start with the "point to" test as it relates to the following example, **Math Test**. As was the case with Pipes 1 and 2 and Doors 1 and 2, the case is designed to see whether it is possible to distinguish between two subjects using the unmarked *pos* form of 'rational':

Math Test
Tyler and John are both overconfident in their performance on the most recent calculus test, assigning too high a credence to getting an A or better on the test. They have both received Cs on all of their past tests, and they studied the exact same amount for the test they just completed. Tyler is very confident that he got an A+, the best grade in the class, and John is very confident that he got an A.

When it comes to rationality, Tyler and John are both somewhat lacking. Given their previous math test performances, they are both overconfident in how they did on the most recent. They are, however, not completely irrational. Even though it might be difficult, it is not impossible to move from receiving a C on past tests to receiving an A on the most recent test. Nevertheless, (14) is infelicitous, making it appear that 'rational' behaves more like an AA than a RA:
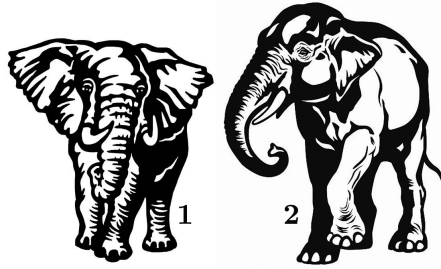
(14) #Point to the one who has a rational degree of confidence

The same result occurs with belief. If Tyler believes that he got an A+ and John believes that he got an A, then (15) is anomalous:

(15) #Point to the one who believed rationally

Unlike 'long', 'rational' cannot be used to differentiate between two things in the middle of the scale, suggesting that 'rational' has a maximum-degree standard. Tyler and John are both irrational to some degree, making their doxastic states irrational.

For further confirmation that the "point to" test paints 'rational' as an AA, the following study was conducted with twenty native English speakers via Amazon Mechanical Turk. First came the control cases. Starting out with the RAs 'long' and 'large', participants were shown Pipes 1 and 2 as well as Elephants 1 and 2. Along with Pipes 1 and 2, participants were asked to respond to this question: "If I asked you to point at the long one, which would you point to?" Next, study subjects were shown Elephants 1 and 2 and asked to respond to the following prompt: "If I asked you to point at the large one, which would you point to?" With both 'long' and 'large', the majority of participants indicated that they would point to the object that had more of that property, the longer Pipe 2 and the larger Elephant 2.

After responding to the two cases testing RAs, participants then viewed two cases that included AAs. After viewing Doors 1 and 2, subjects were asked, "If I asked you to point at the closed one, which would you point to?," and after viewing Lines 1 and 2, subjects were asked "If I asked you to point at the straight one, which would you point to?" Whereas the majority of participants selected the longer pipe or the larger elephant, in this case the majority of subjects said that neither door was closed and that neither line was straight. After establishing the control cases, the study then asked participants about 'rational'. Subjects read the **Math Test**, and then were asked to respond to the following prompt: "If I asked you to point to the one who has a rational degree of confidence, who would you point to?" As with 'closed' and 'straight', the majority of respondents thought that neither could be dubbed as 'rational', and only 25%of participants that that it was possible to differentiate between Tyler and John using 'rational', showing that 'rational' behaves as an AA in the "point to" test:[25]

---
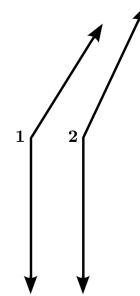
[25]Complete raw survey data for all of the empirical work contained in this paper can be found in the Appendix.
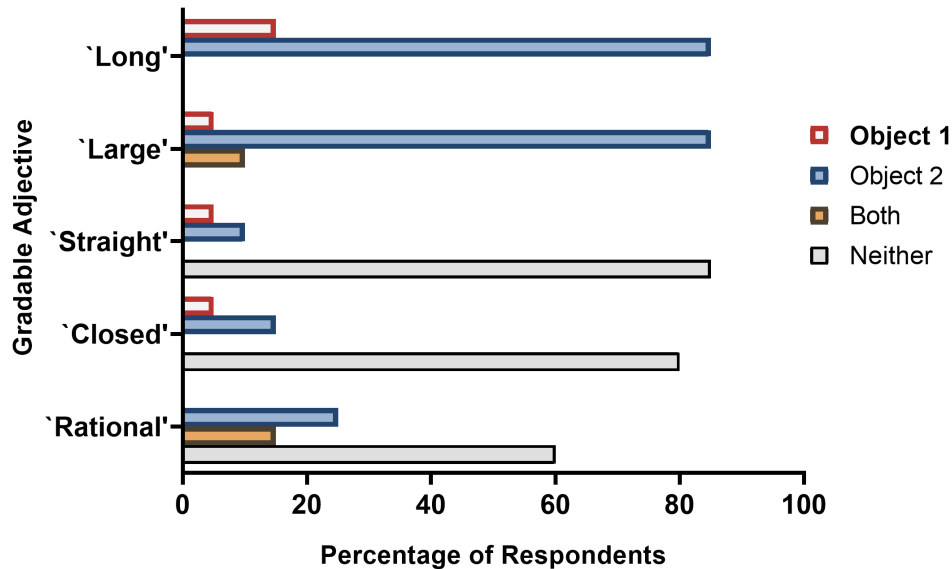
Even though 'rational' cannot be used to differentiate between subjects in the middle of its scale, like other AAs, it can still do so at the top of its scale. If John is appropriately confident that he will get a C and Tyler is too confident that he will get a B, then (14) is felicitous, demonstrating that 'rational' can be used to distinguish between those who are fully rational and those who are irrational to a slight degree. This is possible because 'rational' has a top-closed scale, a thought confirmed by the fact that 'rational' allows modification by maximal degree adverbs. We have seen from both defenders of uniqueness and permissivism that 'rational' accepts modification by a host of maximality adverbs, including 'fully,' 'perfectly,' 'maximally,' and 'completely.' This test is not fail proof. In some instances, words like 'totally' and 'completely' are synonymous with 'very,' making modification by these terms an imperfect diagnostic for the presence of a top-closed scale. Fortunately, this obstacle can be overcome by considering the entailments of maximal degree constructions. If particular uses of 'totally' and 'completely' are synonymous with 'very,' then it should still be possible to possess more of the property in question, whereas if they indicate a top-closed scale, no greater degree of the property can be possessed.

When this entailment data is taken into account, 'rational' tests as having a top-closed scale. Let's say that John has done very well in his math class so far – he only needs a score of 20% on the last test to pass the class. As in our other example, he has received a C on all the other tests and has studied just as much for this one. Because he is almost certain to pass the class, it is completely rational for him to be very confident that he will pass. This complete rationality does not allow that he could do something even more rational, as the following description is infelicitous:

12

422 (16) #John is completely rational to be very confident that he will pass
423 the class, but his confidence is still somewhat irrational

424 As we can see in (16), when 'rational' is part of a maximal degree construction,
425 this entails that it is not possible to possess any more rationality, providing
426 strong reason to think that 'rational' is an absolute gradable adjective.

427
428

Adjective (Type)	Maximality Modifier	Entailment Test
jealous (RA)	$\not\Rightarrow$	I am completely jealous of your house, but my wife is even more jealous
intrigued (RA)	$\not\Rightarrow$	Jerry is totally intrigued by the game, but Marsha is watching even more intently
straight (AA)	$\Rightarrow$	#The road is perfectly straight, but it is slightly curved
flat (AA)	$\Rightarrow$	#The table is completely flat, but it still has a few bumps
rational (AA)	$\Rightarrow$	#John is completely rational to be very confident that he will pass the class, but his confidence is still somewhat irrational

429

430 Because 'rational' is an AA, it is also subject to natural precisification. Let's
431 say that Tyler needs a B on the final calculus test in order to pass the class. He
432 has reflected a bit on his past test performance (all Cs) and is becoming more
433 realistic. Instead of thinking that he will get an A+ on the final test, he has
434 moderated his expectations and is now very confident that he will score at least
435 a B despite the fact that he has not studied any more for the upcoming test.
436 Frustrated that Tyler will not study any more, Tyler's calculus tutor tries to
437 convince him to continue studying with (17):

438 (17) Tyler, we need to be rational about this. Right now, you are still
439 too confident.

440 Even though Tyler is more rational than he was before, the tutor can use (17)
441 to say that Tyler is not yet thinking rationally. He is still somewhat irrational,
442 which allows the tutor to prevent an imprecise reading of 'rational' through nat-
443 ural precisification. 'Rational' thus behaves like an absolute gradable adjective
444 in a range of contexts. It can be modified by maximality adverbs, can be used
445 to differentiate between objects in the middle of its underlying scale but not at
446 the scale's extreme, and can be precisified in a natural language context.

4 Rational Persons, Rational States

At this point, it seems fairly clear that 'rational' is an absolute gradable adjective. What can this tell us about **Rational Supererogation**? Because 'rational' is an AA, it picks out the default value at the top of its scale. Thus, in order for a doxastic state to be rational, it must be fully rational, showing that **Rational Supererogation** is mistaken. Contrary to Li's defense of permissivism, it is not possible to believe what is rational without believing what is maximally rational.

The defender of **Rational Supererogation** might respond that, even if this paper has shown that 'rational' is an absolute gradable adjective, **Rational Supererogation** concerns actions that are rationally permissible, not actions that are rational or irrational simpliciter. Fortunately for our purposes, there is a plausible link between the rational and the rationally permissible. Consider instances of moral supererogation. Even though jumping on the grenade to save his comrades is morally supererogatory, the soldier does not act immorally if he flees. Fleeing is morally permissible, and thus not positively immoral, even though it is not the best moral action that the soldier could have chosen. We thus get the result that morally permissible actions are not at the same time immoral actions. If the arguments of this paper are correct, then rationality differs from morality in that any actions below the top of the scale of rationality are positively irrational, making them also rationality impermissible. With the Holmes and Watson case, the claim is not that Watson believed something rationally permissible and yet irrational – rather, the claim is that it is possible to believe rationally while not at the same time believing what is most rational. If 'rational' is an AA, however, this is not the case. Believing rationally requires believing what is rationally optimal.

Another way in which supporters of **Rational Supererogation** might respond is by arguing that there are cases in which 'rational' behaves like a RA instead of an AA. A study similar to the previous one was conducted with twenty-one native English speakers. The same four control cases were used, but this time, the final scenario was changed from **Math Test** to **Famous Detective**:
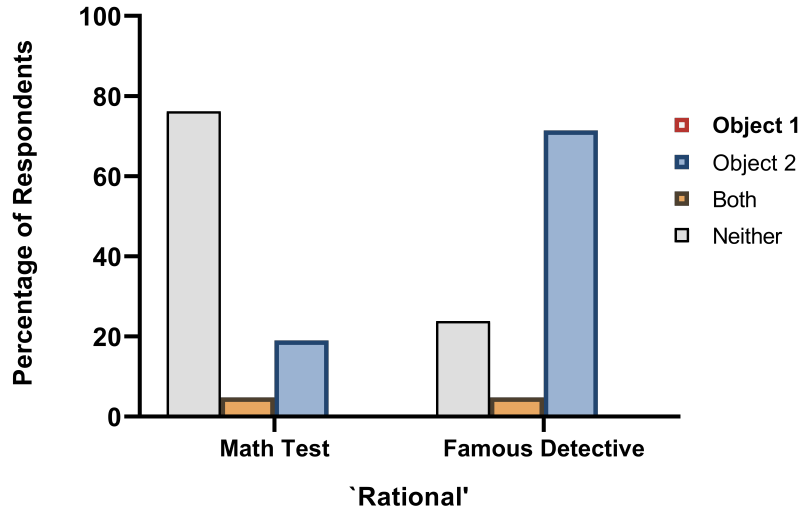
> **Famous Detective**
> In the Sherlock Holmes stories written by Sir Arthur Conan Doyle, Sherlock Holmes and Doctor Watson form a crime-solving duo. Sherlock is a master detective, displaying above average intelligence and reasoning, often solving cases with just a few clues. While still able to help with the cases, Doctor Watson is not as sharp as Sherlock, often making mistakes and reaching conclusions more slowly.

After being presented with the above sketch of the characters of Sherlock Holmes and Doctor Watson, study participants were given the following prompt: "If I asked you to point at the rational one, who would you point to?"[26] With

---

[26]This vignette was also used in a study conducted for Siscoe (Forthcoming).

Math Test, the majority of test subjects thought that it was not possible to differentiate between the credences of Tyler and John using the *pos* form of 'rational.' These results flipped, however, when comparing Sherlock Holmes and Doctor Watson:

In **Math Test**, the majority of participants said that neither Tyler's nor John's credences were rational, whereas in **Famous Detective**, the majority said that they would point to Sherlock. Thus, 'rational' went from behaving like an AA in response to **Math Test** to resembling a RA in response to **Famous Detective**. Do these results create trouble for the thesis that 'rational' is an absolute gradable adjective?[27]

The first thing to say in response to this concern is that the contrasting cases come about when 'rational' modifies distinct sorts of objects. In all of the evidence adduced in Section 3, 'rational' was being used to describe doxastic states. However, in **Famous Detective**, 'rational' is being used to describe persons. Because the behavior of 'rational' changes depending on the type of object it modifies, this opens up the possibility that 'rational' is ambiguous between two different gradable adjectives, one that acts as an AA, and one that acts as a RA.

The most straightforward evidence that 'rational' is ambiguous comes from the *Oxford English Dictionary*. The *OED* contains four senses of 'rational', but we will limit our discussion to just two of these. The first sense, what we will call the *capacity* sense of 'rational,' applies to things that are "endowed with reason" or "have the faculty of reasoning." The second sense, which we will call the *sanctioning* sense of 'rational,' applies to attitudes that are "in accordance with reason."[28] Along with being gradable, these two senses of 'rational' usually

---

[27]All survey data from both studies can be found in the Appendix.
[28]See the *Oxford English Dictionary*, 3rd Edition (2008).

apply to different sorts of things. Generally speaking, the sanctioning sense of 'rational' is used to pick out particular attitudes, while the capacity sense of 'rational' is used to describe a person or persons.

One test - complexion is fair, his business dealings are fair
These two senses of 'rational' can be further teased apart by standard tests for ambiguity.[29] The multiple meanings of ambiguous words can be used to create zeugmas, sentences that only make sense when an ambiguous word receives two different interpretations. Consider, for instance, the two different senses of 'fair':

(18) His complexion is fair
(19) His business dealings are fair

The first use of 'fair' refers to a particular skin tone, while the second picks out business transactions that are just or legitimate. If we try and use both of these sentences, however, with only one occurrence of 'fair', we get a zeugmatic sentence:

(20) His complexion and his business dealings are both fair

The only way to interpret (20) so that it is accurate is to employ two different senses of 'fair', making the sentence a zeugma.[30] We can also create zeugmas using our two senses of 'rational'. Despite the fact that Sherlock is brighter than Watson, Watson is not completely hopeless. Sometimes, Watson even solves a case on his own. Suppose that, in the current case, the chauffeur committed the crime. If Watson has examined the evidence and is appropriately confident that the chauffeur is the perpetrator, then (21) and (22) are both true:

(21) Sherlock is rational
(22) Watson's credence is rational

When we combine these, however, with only a single use of 'rational,' (23) is zeugmatic:

---

[29]Even though I have described 'rational' as ambiguous, it might be more apt to call 'rational' polysemous. No completely precise definitions exist for differentiating between case of ambiguity and polysemy, see Sennet (2016a) and (2016b) and Viebahn (2018), but both are used to describe terms that have various meanings. Ambiguity is often used to describe a term whose meanings are not very closely related ('bank' can mean a financial institution or the side of a river) while polysemy applies when a term has two closely related meanings ('fish' can refer to a particular animal or the method of catching that animal). Because ambiguity and polysemy are discovered using the same kinds of tests, I will not make this distinction in the body of the text, although I do note wherever a polysemy theory of 'rational' would importantly differ from a case of homonymy.

[30]This test is often called the identity or conjunction reduction test, see Asher (2011), Bach (1998), Cruse (1982), (1986), and (2011), Falkum and Vicente (2015), Gillon (2004), Lakoff (1970), Sennet (2016a) and (2016b), and Zwicky and Sadock (1975). The conjunction reduction test has been employed for a wide range of philosophical applications, including debates over ambiguity theories of definite descriptions, Koralus (2013) and Sennet (2002), and know how, Stanley (2005). The conjunction reduction test can be susceptible to false negatives, see Geeraerts (1993), Moldovan (2019), and Viebahn (2018), but (23) provides positive evidence that 'rational' can be used to create zeugmas.

(23) Sherlock and Watson's credence are both rational

Sherlock is rational in that he is endowed with reason, while Watson's confidence is rational because it is formed on the basis of good reasoning. Because we need both of these senses of 'rational' in order to make sense of (23), (23) confirms that the capacity sense and sanctioning senses of 'rational' both show up in tests for ambiguity.

We have seen a number of reasons to think that the sanctioning sense of 'rational' is an AA. The reason that 'rational' does not always behave as an AA is because there are multiple senses of 'rational'. It may be that the capacity sense of 'rational' is a RA, a fact that is relevant for thinking about the rationality of persons, but because this paper is interested in the rationality of doxastic states, I will leave the capacity sense of 'rational' for future work. The survey results for **Famous Detective** thus do not bear on our general thesis, that, because 'rational' is an absolute gradable adjective, **Rational Supererogation** is not possible.

5   Rational Capacities and Epistemic Permissivism

If the argument of this paper is correct, then **Rational Supererogation** is not the right strategy for defending epistemic permissivism. The sanctioning sense of 'rational' is expressed using an absolute gradable adjective, leaving no room for **Rational Supererogation**. However, a number of questions remain. First of all, what is the connection between rational persons and rational doxastic states? Do the rational faculties of a particular person impact which beliefs are rational for them? And what should we say about Li's original case? Is it a counterexample to uniqueness?

Even though each of these questions deserve a paper-length treatment in their own right, a few things can be said. To begin with, it is helpful to consider how more advanced rational capacities effect the rationality of a particular belief. Having more developed rational capacities has several benefits – agents who are more rational can better survey possible hypotheses that explain their evidence and make more connections between their evidence and those hypotheses. Because Sherlock is more rational than Watson, he can dream up solutions that Watson could never have considered. Furthermore, Sherlock is better able to draw connections between his total evidence and the explanations it supports. He can hold more information in mind at once, granting a deeper appreciation of what the evidence supports than what Watson can attain.

These differences in rational capacities in turn affect the sanctioning rationality of Sherlock's and Watson's beliefs. Even though Sherlock may, through his superior rational capacities, reach a solution that enjoys more evidential support, that does not mean that Watson's belief would have been rational had

17

he adopted Sherlock's hypothesis. If Watson does not consider Sherlock's solution or fails to see how the evidence best supports that hypothesis, it would be positively irrational for him to choose that conclusion. Does this, then, make Watson's belief irrational when he reaches a different verdict?

There are a couple ways to go here. Even though I have argued that the only rational belief is the maximally rational belief, what I have said is silent on what factors determine whether or not a belief is maximally rational for an individual. If the maximally rational belief for Watson is to eventually reach Sherlock's conclusion, then anything else he believes will be irrational. If this account is correct, that only Sherlock's solution to the case is maximally rational, then uniqueness is likely the correct account of rationality. Some permissivists, however, have claimed that it is possible to believe rationally without having superior rational capacities. Igor Douven has defended the view that sanctioning rationality does not require such brilliant flashes of insight,[31] while Abelard Podgorski has argued that sanctioning rationality does not require taking all of one's total evidence into account.[32] On these views, Watson's inability to reach Sherlock's solution does not make the conclusion he does reach automatically irrational. Instead, given that he cannot make the same connections or dream up the same solutions as Sherlock, the maximally rational belief for him is not the same as the maximally rational conclusion that Sherlock can reach. In these cases, Watson's belief can still believe what is maximally rational because, given his rational capacities, there's not a more rational belief that he could adopt. If this latter view is right, with what is maximally rational for Watson differing from what is maximally rational for Sherlock, then Li's case might be a successful counterexample to uniqueness after all. All I have said, then, leaves open other possible defenses of epistemic permissivism.

Importantly though, even if Li's case is a persuasive instance of permissivism, this does not also make it a successful defense of **Rational Supererogation**. In cases of moral supererogation, there are multiple options that are all morally permissible. The soldier does not do something immoral if he flees even though he does something morally better by jumping on the grenade. This is not so for Watson - there is no situation in which he has multiple, rationally permissible options. If he does not have Sherlock's flash of insight, then it is only rational for him to adopt his original conclusion – it would be irrational for him to reach Sherlock's solution without realizing how his evidence supports that hypothesis. On the other hand, if he does somehow manage to share Sherlock's stroke of genius, then his original conclusion is no longer rational. Once he understands that Sherlock's hypothesis is better supported by the evidence, he is rationally required to adopt that solution. Thus, there is no situation in which there are multiple options that are rationally permissible, preventing either from being genuine instances of **Rational Supererogation**.

---

[31]See Douven (2009), pp. 315-352.
[32]See Podgorski (2016), pp. 1934-1936.

Appendix

Survey Data*

Study 1

Image/ Scenario	Prompt	Percentage of Respondents (20 Total)			
		Object 1	Object 2	Both	Neither
Pipes	If I asked you to point at the long one, which would you point to?	15	85	0	0
Elephants	If I asked you to point at the large one, which would you point to?	5	85	10	0
Lines	If I asked you to point at the straight one, which would you point to?	5	10	0	85
Doors	If I asked you to point at the closed one, which would you point to?	5	15	0	80
Math Test	If I asked you to point to the one who has a rational degree of confidence, who would you point to?	0	25	15	60

Study 2

Image/ Scenario	Prompt	Percentage of Respondents (21 Total)			
		Object 1	Object 2	Both	Neither
Pipes	If I asked you to point at the long one, which would you point to?	4.8	90.5	4.8	0
Elephants	If I asked you to point at the large one, which would you point to?	0	95.2	4.8	0
Lines	If I asked you to point at the straight one, which would you point to?	0	14.3	9.5	76.2
Doors	If I asked you to point at the closed one, which would you point to?	0	19.0	4.8	76.2
Famous Detective	If I asked you to point at the rational one, who would you point to?	4.8	71.4	23.8	0

*All survey participants were native English speakers and high school graduates

Bibliography

Asher, Nicholas. 2011. *Lexical Meaning in Context: a Web of Words*. Cambridge University Press.

Bach, Kent. 1998. "Ambiguity." In the *Routledge Encyclopedia of Philosophy*. Taylor and Francis.

Barker, Chris. 2002. "The Dynamics of Vagueness." *Linguistics and Philosophy*, 25, no. 1: pp. 1-36.

Bartsch, Renate and Theo Venneman. 1972. *The Grammar of Relative Adjectives and Comparison. Linguistische Berichte* 20: pp. 19-32.

Berker, Selim. 2018. "A Combinatorial Argument against Practical Reasons for Belief." *Analytic Philosophy* 59: pp. 427-470.

Bernstein, Mark. 1986. "Moral and Epistemic Saints." *Metaphilosophy* 17: pp. 102-108.

Bierwisch, Manfred. 1989. "The Semantics of Gradation." In *Dimensional Adjectives*, ed. Manfred Bierwisch and Ewald Lang. : pp. 71-261. Berlin: Springer-Verlag.

Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116: pp. 187-217.

Cresswell, Max. 1977. "The Semantics of Degree." In *Montague Grammar*, ed. Barbara Partee: pp. 261-292. New York: Academic Press.

Cruse, D. A. 1982. "On Lexical Ambiguity." *Nottingham Linguistic Circular* 11: pp. 65-80.

Cruse, D.A. 1986. *Lexical Semantics*. Cambridge University Press.

Cruse, D. A. 2011. *Meaning in Language: An Introduction to Semantics and Pragmatics*. Oxford University Press.

Falkum, Ingrid and Agustin Vicente. 2015. "Polysemy: Current Perspectives and Approaches." *Lingua* 157: pp. 1-16.

Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60: pp. 667-695. Press.

Fine, Kit. 1975. "Vagueness, Truth, and Logic." *Synthese* 30: pp. 265-300.

Franke, Michael. 2012. "On Scales, Salience and Referential Language Use." *Logic, Language and Meaning*: pp. 311-320.

Frazier, Lyn, Charles Clifton Jr., and Britta Stolterfoht. 2008. "Scale Structure: Processing Minimum Standard and Maximum Standard Scalar Adjectives." *Cognition* 106: pp. 299-324.

Geeraerts, Dirk. 1993. "Vagueness's Puzzles, Polysemy's Vagaries." *Cognitive Linguistics* 4: pp. 223-272.

Gillon, Brendan. 2004. "Ambiguity, Indeterminacy, Deixis, and Vagueness: Evidence and Theory." In *Semantics: A Reader* Edited by Steven Davis and Brendan Gillon. Oxford University Press.

Harman, Gilbert. 2004. "Practical Aspects of Theoretical Reasoning." In *The Oxford Handbook of Rationality*. Edited by Alfred Mele and Piers Rawling. Oxford University Press.

Hedberg, Trevor. 2014. "Epistemic Supererogation and its Implications." *Synthese* 191: pp. 3621-3637.

Heim, Irene. 2000. "Degree Operators and Scope." In *Proceedings of the Tenth Semantics and Linguistic Theory Conference*. Ed. by B. Jackson and T. Matthews: pp. 40-64. Ithaca, New York: CLC Publishing.

Jackson, Elizabeth. 2019. "A Defense of Intrapersonal Belief Permissivism." *Episteme*: pp. 1-15.

Kamp, J.W.A. 1975. "Two Theories of Adjectives." In *Formal Semantics of Natural Language*, ed. Edward Keenan: pp. 123-155. Cambridge University Press.

Kelly, Thomas. 2002. "The Rationality of Belief and Some Other Propositional Attitudes." *Philosophical Studies* 110, pp. 163-196.

Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66: pp. 612-640.

Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30, no. 1: pp. 1-45.

Kennedy, Christopher and Louise McNally. 2005. "Scale Structure and the Semantic Typology of Gradable Predicates." *Language* 81, no. 2: pp. 345-381.

Klein, Ewan. 1980. "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4: pp. 1-45.

Koralus, Philipp. 2013. "Descriptions, Ambiguity, and Representationalist Theories of Interpretation." *Philosophical Studies* 162: pp. 275-290.

Kyburg, Alice and Michael Morreau. 2000. "Fitting Words: Vague Language in Context." *Linguistics and Philosophy* 23: pp. 577-597.

Lakoff, George. 1970. "A Note on Vagueness and Ambiguity," *Linguistic Inquiry* 1: pp. 357-359.

Larson, Richard. 1988. "Scope and Comparatives." *Linguistics and Philosophy* 11: pp. 1-26.

Lewis, David. 1970 "General Semantics." *Synthese* 22: pp. 18-67.

Moldovan, Andrei. 2019. "Descriptions and Tests for Polysemy." *Axiomathes.* https://link.springer.com/article/10.1007/s10516-019-09445-y

McElwee, Brian. 2016. "Supererogation Across Normative Domains." *Journal of Philosophy* 95: pp. 505-516.

Nado, Jennifer. 2019. "Who Wants to Know?" In *Oxford Studies in Epistemology*, Vol. 6. Edited by Tamar Gendler and John Hawthorne. Oxford University Press.

Pinkal, Manfred. 1995. *Logic and Lexicon.* Dordrecht: Kluwer.

Rett, Jessica. 2014. *The Semantics of Evaluativity.* Oxford University Press.

Rotstein, Carmen and Yoad Winter. 2004. "Total Adjectives vs. Partial Adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* 12: pp. 259-288.

Rusiecki, Jan 1985. *On Adjectives and Comparisons in English.* New York: Longman Linguistics Library.

Sapir, Edward. 1944. "Grading: A Study in Semantics." *Philosophy of Science*, 11: pp. 93-116.

Sassoon, Galit. 2012. "A Slightly Modified Economy Principle: Stable Properties have Non-Stable Standards." *Proceedings of the Israel Association of Theoretical Linguistics* 27: pp. 163-181.

Sedivy, J., Tanenhaus, M., Chambers, C. and Carlson, G. 1999. "Achieving Incremental Semantic Interpretation through Contextual Representations." *Cognition* 71: pp. 109-147.

Sennet, Adam. 2002. "An Ambiguity Test for Definite Descriptions." *Philosophical Studies* 111: pp. 81-95.

Sennet, Adam. 2016a. "Ambiguity." In *The Stanford Encyclopedia of Philosophy.* Edited by Edward N. Zalta: https://plato.stanford.edu/entries/ambiguity/.

Sennet, Adam. 2016b. "Polysemy." In *Oxford Handbooks Online: Philosophy.* Edited by Sanford Goldberg: http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935314.001.0001/oxfordhb-9780199935314-e-32.

Siscoe, Robert Weston. Forthcoming. "Being Rational Enough: Maximizing, Satisficing, and Degrees of Rationality." *Australasian Journal of Philosophy.*

Sorensen, Roy. 1991. "Rationality as an Absolute Concept." *Philosophy* 66: pp. 473-486.

Stanley, Jason. 2005. "Hornsby on the Phenomenology of Speech." *Proceedings of the Aristotelian Society* 79: pp. 107-145.

Stechow, Arnim von. 1984. "Comparing Semantic Theories of Comparison." *Journal of Semantics* 3: pp. 1-77.

Syrett, K., Kennedy C. and Lidz, J. 2010. "Meaning and Context in Children's Understanding of Gradable Adjectives." *Journal of Semantics* 27: pp. 1-35.

Tidman, Paul. 1996. "Critical Reflection: An Alleged Epistemic Duty." *Analysis* 56: pp. 268-276.

Titelbaum, Michael and Matthew Kopec. 2018. "When Rational Reasoners Reason Differently." In *Reasoning: Essays on Theoretical and Practical Thinking.* Edited by Magdalena Balcerak Jackson and Brendan Balcerak Jackson. Oxford University Press.

Viebahn, Emanuel. 2018. "Ambiguity and Zeugma." *Pacific Philosophical Quarterly* 99: pp. 749-762.

Wheeler, Samuel. 1972. "Attributives and Their Modifiers." *Nous* 6, no. 4: pp. 310-334.

Whiting, Daniel. 2014. "Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action." In *Epistemic Norms: New Essays on Action, Belief, and Assertion.* Edited by Clayton Littlejohn and John Turri. Oxford University Press.

Zwicky, Arnold and Sadock, Jerrold. 1975. "Ambiguity Tests and How to Fail Them." *Syntax and Semantics* 4: pp. 1-36.