

This is a repository copy of *Virtue and Salience*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/90869/>

Version: Accepted Version

Article:

Chappell, Richard Yetter orcid.org/0000-0003-3322-4935 and Yetter-Chappell, Helen orcid.org/0000-0003-3527-5306 (2016) *Virtue and Salience*. *Australasian Journal of Philosophy*. pp. 449-463. ISSN: 1471-6828

<https://doi.org/10.1080/00048402.2015.1115530>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

VIRTUE AND SALIENCE

Richard Yetter Chappell and Helen Yetter-Chappell

Forthcoming in the Australasian Journal of Philosophy

This paper explores two ways that evaluations of an agent's character as virtuous or vicious are properly influenced by what the agent finds salient or attention-grabbing. First, we argue that ignoring salient needs reveals a greater deficit of benevolent motivation in the agent, and hence renders them more blameworthy. We use this fact to help explain our ordinary intuition that failing to give to famine relief (for example) is in some sense *less bad* than failing to help a drowning child right before your eyes, in a way that's compatible with the contention that there's no principled reason to see the one life-saving act as any more or less choiceworthy than the other. Second, we argue that alleged 'virtues of ignorance' (modesty, believing better of friends than the evidence supports, etc.) are better understood as 'virtues of salience'. Rather than placing demands on what we believe, these virtues place demands on what we find salient.

Keywords: virtue, salience, attention, benevolence, modesty, friendship

1. Introduction

This paper explores two ways that evaluations of an agent's character as virtuous or vicious are properly influenced by what the agent finds salient or attention-grabbing. The moral quality of our character is influenced by our patterns of concern: a good person cares deeply about the right things, whereas a bad person does not. What we find salient reveals what we care about. Failing to find the right things salient may reflect poorly on us by *directly* indicating a lack of appropriate concern. Salience can also *indirectly* reveal our moral character, influencing the degree to which a negligent act or omission indicates a problematic lack of moral concern.

What is salience? Salient things are attention-grabbing. A feature of the world is salient to an agent at a time *t* in so far as *that feature exerts an involuntary draw on the agent's attention* at *t*. Just what this means is best illustrated with examples. Imagine a black and white picture of a fruit bowl. The entire picture is black, white, and shades of grey . . . except for a single bright yellow banana. For most of us, the banana will be highly salient: Your attention will be involuntarily drawn to the

banana, not because someone has instructed you to look at the banana or through any other deliberate effort on your part, but simply because the banana itself seems to ‘tug’ at your awareness.¹

Models of salience have been developed by computational neuroscientists to explain how involuntary attention is directed to perceptual features [Koch and Ullman 1985; Itti and Koch 2000]. They argue that involuntary attention is directed by a ‘bottom-up, fast, primitive mechanism that biases the observer towards selecting stimuli based on their *saliency*’ [Itti and Koch 2000: 1490]. The brain does this by creating a ‘saliency map’, combining low-level visual information such as changes in colour, intensity, orientation, and motion into a single global measure of conspicuity.² Those parts of the map which ‘elicit a strong response are thought to *draw visual attention to themselves* and to therefore be experienced as “visually salient”, [whereas] directing attention at any of the other parts is thought to require voluntary “effort”’ (ibid., emphasis added).

While these neuroscientists are focused on forms of *perceptual* salience and attention, we take the basic concept to clearly extend to all forms of attention (which will likewise be guided by some neural underpinnings—or ‘saliency maps’—towards particular thoughts, objects, or features that the agent thereby experiences as salient).³ We will often speak of salience in terms of an agent’s *disposition to attend*, by which we mean specifically this kind of active draw on involuntary attention.

What is salient to an agent thus depends on the psychology of that agent. We could imagine an agent who is naturally inclined to find circular things salient, and who processes colour very differently from us, such that they find the black-and-white oranges in the fruit bowl salient, rather than the bright yellow banana. As these examples suggest, salience is essentially agent-relative. (There’s no basis for attributing objective ‘salience’ to features of the world, independently of what individual – and diverse – psychologies happen to find salient.) But we can often

¹ Note that what is salient to someone at a given moment is what exerts an involuntary draw on their attention at that moment, not merely what *typically* draws their attention.

² Similar models have been proposed for the low-level attention-biasing mechanisms at work in audition [Kayser, Petkov, Lippert, and Logothetis 2005] and in speech perception [Kalinli and Narayanan 2007]. Multi-modal saliency maps have also been proposed [Huber, Khosla, and Dow 2009].

³ We of course take no stand on the empirical question of whether the neural underpinnings of salience are uniform across all possible domains.

predict what typical human agents will find salient. Speech in one's own language is typically more salient than nonsense syllables or speech in an unknown language. One's own child crying might be more salient than the cries of another child. The plight of a refugee might be more salient were you to meet them in person, or if your own family was once in a similar situation, or if they're fleeing from a country with which you have personal ties. In each of these cases, the salient features 'tug' at the agent's awareness—though the natural draw of the salient may be overridden by a deliberate act of will that fixes one's attention elsewhere.

We'll discuss a number of additional examples of salience throughout the paper, as well as some particular features that can plausibly bear on salience, including proximity and in-group/out-group status.

This paper addresses two important relations between salience and character. The first half focuses on the indirect connections between salience and character. We argue for the following comparative claim: All else being equal, negligent inaction is less blameworthy (and less indicative of a problematic callousness in one's character) when the neglected moral considerations are less salient to the agent—in contrast to cases where the agent fails to act because they simply do not care. This can explain—among other things—the intuition that failing to give to famine relief (for example) is *less bad* than failing to help a drowning child right before your eyes, in a way that's compatible with the contention that there's no principled reason to see the one life-saving act as any more or less choiceworthy than the other.

The second half of this paper explores the direct connections between salience and character. There we argue that alleged 'virtues of ignorance' (modesty, believing better of friends than the evidence supports, etc.) are better understood as 'virtues of salience'. Rather than placing demands on what we believe, these virtues place demands on what we find salient.

2. Salience and Quality of Will

2.1 A Puzzle

Singer [1972] invites us to imagine a child drowning in a pond, whom we could save at the cost of ruining our expensive clothes. It seems clear that we ought to save the

child, no matter the (comparatively insignificant) financial cost to ourselves. This motivates the principle that ‘if it is in our power to prevent something very bad from happening, without thereby sacrificing anything morally significant, we ought, morally, to do it’ [Singer 1972: 231]. However, as Singer notes, we are constantly violating this principle by failing to donate as much as we could to effective charities that address preventable harms caused by global poverty: malnutrition, disease, lack of basic medical care, etc. What should we think of this practical inconsistency?

When we reflect on what’s at stake, it’s not implausible to hold that we really should prevent these grave harms rather than buying unnecessary luxuries for ourselves. (Call this claim *Act Evaluation*.) That much of Singer’s argument seems plausible.⁴ But the analogy between the drowning child and the global poor may also be taken to suggest a much more troubling conclusion, via the following argument:

1. It would be morally monstrous to do nothing and let the drowning child die.
2. Saving a distant stranger’s life by donating to an effective charity is relevantly similar to saving a nearby drowning child.
3. So, it would be morally monstrous to let a distant stranger die when we could have saved their life by donating to an effective charity.

Here the conclusion is not just about the status of the *act* of helping—that it merits choosing, or ought to be done—but about the moral *character* of the agent who fails to act as they ought in this case. It is, according to this argument, no minor failure, but one that renders us *morally monstrous* or blameworthy to the highest degree. And this claim—call it *Character Evaluation*—seems entirely incredible. It just isn’t plausible that in failing to save distant strangers we reveal our moral character to be as bad as someone who callously watches a child drown and does nothing about it.

2.2 The Solution

Considerations of salience, together with a *Quality of Will* account of blameworthiness, can help to resolve this puzzle, showing how we can accept Singer’s Act Evaluation without committing ourselves to the implausible Character

⁴ We won’t be defending this claim here, but see, e.g., Unger [1996], Kagan [1991], and Cullity [2006].

Evaluation. According to Quality of Will accounts, an agent is blameworthy to the extent that their actions manifest an insufficient degree of good will towards others [Strawson 1962; Arpaly 2003]. An agent may be understood as having a ‘sufficient degree of good will’ when their desires for others’ welfare are sufficiently weighty, both absolutely and in relation to their self-interested desires. The less that I care about others, and hence the more harms that I’m willing to impose on them, the morally worse my character is. When I perform actions that manifest this lack of concern for others, I am blameworthy in proportion to the moral inadequacy of the desires that I thereby act upon. (Someone who acts with at least some minimal concern for others will be less blameworthy than someone who acts with complete disregard for others’ interests.)

But now notice the following fact about human psychology: Our actions are determined not only by the strengths of our standing desires (both self-interested and altruistic), but also by which desires are emotionally ‘activated’, or occurrently felt, and to what degree. Due to our limited cognitive capacities, we do not—and could not—constantly feel the force of everything that interests us or that we care about. Their full force is typically felt only when *triggered*, perhaps by certain thoughts or the proximity of relevant environmental stimuli.

For example, it’s noteworthy that in the infamous Milgram experiments, participants’ willingness to obediently electrocute a victim was significantly reduced when victims were made more salient via increased proximity. Milgram [1965] ran multiple experimental conditions, varying the ‘immediacy of the victim’ from cases where the subject could hear (but not see) the victim to cases where the subject had to manually place the hand of the protesting victim on the shock plate. ‘The data revealed that obedience was significantly reduced as the victim was rendered more immediate to the subject.’ [ibid.: 62] Milgram’s own explanations of the disparity appeal to the importance of salience in motivation, noting that in closer-proximity cases the victim ‘necessarily intrudes on the subject’s awareness’ [ibid.: 63].

This fact about the variable efficacy of desires indicates that not all failures to help others (even holding fixed the relevant costs and benefits) will necessarily reveal the same insufficiency of good will, and hence the same degree of blameworthiness. In cases where the needs of others are especially *salient* to an

agent, any altruistic desires in that agent can be expected to function at full efficacy. Hence the failure to help others when their needs are salient to the agent signals a much greater deficit of good will: one must be thoroughly lacking in concern for others to not be moved to help even when their needs are most salient. Such egregious deficits of good will yield, on the quality of will account, a greater degree of blameworthiness.

This principle may be helpfully applied to the case of Singer's pond: We feel that a person who could watch a child drown before his eyes must be terribly callous.⁵ The child's need is so obvious, and so emotionally gripping (for anyone with a modicum of good will), that to fail to act in this case reveals a truly monstrous lack of concern for others.⁶ By contrast, good-willed people regularly fail to act in ways that would save the lives of distant strangers. The reduced salience of more distant needs is plausibly at least part of the explanation for this. A child drowning before our eyes shocks us out of complacency, activating whatever altruistic concern we may have, whereas the constant suffering of the global poor is easier to ignore, meaning that inaction doesn't necessarily imply an egregious lack of concern.

In Singer's pond case, salience correlates with physical proximity. Other important factors that can affect salience include social proximity (a far-away friend's marital troubles might well be more salient to you than those of a colleague down the hall) and in-group/out-group status. To borrow an example from Blum [1991: 706]:

Tim, a white male, is waiting for a taxi at a train station. Waiting near him are a black woman and her daughter. A cab comes by, past the woman and her daughter, and stops in front of him. Tim, with relief, gets in the cab.

The degree to which Tim strikes us as culpably complicit in this injustice can depend upon salience in a couple of notable respects. If the taxi driver's racial discrimination

⁵ As Slote [2007: 286] writes: 'In the familiar drowning examples, someone's danger or plight has a salience, conspicuousness, vividness, and *immediacy* ... that engages normal human empathy (and consequently arouses sympathy and concern) in a way that similar dangers we merely know *about* do not.' The role of salience in explaining our intuitions is also emphasized in Unger [1996].

⁶ Likewise in the Milgram experiments, we think worse of participants who would shock victims even in the Touch-Proximity case, compared to those who would only comply in the Remote case. The former group seems to be particularly lacking in concern for the apparent victim, and we take this fact about their egregious lack of concern to reflect poorly on their moral character.

is not salient to Tim, this surely signals a defect of character—a problematic lack of concern—on his part. But what we want to highlight here is a further comparative truth: It intuitively reflects much *worse* on Tim if he notices and *just doesn't care*, than if he fails to notice it altogether (at least if he is disposed to respond appropriately once the issue is drawn to his attention).

Our analysis of the indirect significance of salience secures the common-sense result that Character Evaluation is mistaken: Failure to help the distant needy is typically not as blameworthy (nor as indicative of bad character) as inaction in the case of Singer's pond would be. Moreover, this result is secured on the basis of what, intuitively, seems like the right explanation, namely that *it would take a much worse person* to let a child drown before their eyes, whereas any ordinary non-saint does less than they could (and even, arguably, *should*) for the distant poor.⁷ (Note that this conclusion about character is compatible with Singer's Act Evaluation—that one *should* act to help the distant needy—and even the stronger claim that inaction in either case would be *equally* wrong.)⁸

2.3 Salience vs. Attention

In light of our earlier distinction between the two, you might wonder: is salience (being attention-grabbing) or *actual* attention what's doing the work in our above discussion? Mere attention seems insufficient: Aid fundraisers may momentarily

⁷ In accounting for the falsity of the Character Evaluation, we have assumed that agents find the nearby drowning more salient than far away harms. Since this is true of normal human agents, it seems a fine assumption for explaining our intuitions (which are shaped by consideration of typical humans). But one might wonder about the implications of our view for agents who find radically different things salient. Consider an alien agent who finds what's right before their eyes and what they read of in magazines equally salient. Would they have to be more morally monstrous to ignore the drowning child than the distant needy? We think that once you fully appreciate what this agent finds salient, it's clear that ignoring either need displays equal ill-will. Likewise (in contrast to Kamm [1999: 182-3]), we think that an alien agent who found the needs of distant sufferers more salient than those nearby would seem to need to be more vicious to fail to help them than the nearby needy. (Note that we are not committed to the claim that failing to respond appropriately to salient things is the *only* respect in which salience influences blameworthiness. Beyond the ill-will of failing to respond to especially salient harms, it may be that failing to find certain things salient is itself a sign of ill-will, as illustrated by the case of Tim and the taxi. Perhaps, more generally, there are base levels of salience that are required for having minimal concern for others. The second half of this paper discusses some further respects in which salience can be normatively significant.)

⁸ Cf. Slote [2007: 290], who assumes that these salience considerations are incompatible with Singer's Act Evaluation.

bring to our attention the needs of the distant poor, but even so, it doesn't seem morally monstrous to shrug them off and return to our everyday lives—at least, not in the way that it would be monstrous to shrug off the needs of a child drowning right before our eyes. This makes sense, on our account, if such external promptings tend not to (fully) activate our standing desires in the way that genuinely attention-grabbing stimuli do. Salience, not mere attention, thus plays an essential role in our account.

2.4 Summary

Our analysis shows how facts about what's psychologically *salient* to an agent can indirectly alter our assessment of their moral character, thereby showing that we can accept Singer's Act Evaluation without committing ourselves to the implausible Character Evaluation.⁹ In the next section, we'll go on to explore how virtue might sometimes *require* us to find some things more salient than others. Then, in the final section, we will discuss the relation between these two ideas.

3. Virtues of Salience

Several philosophers [Driver 1989, 1999; Keller 2004; Stroud 2006] have recently defended the surprising thesis that some virtues essentially involve ignorance or epistemic bias. We will focus on two cases in particular: whether the virtue of modesty involves ignorance, and whether friendship demands that we believe better of our friends than the evidence warrants. In both cases, we will argue, these alleged 'virtues of ignorance' are better understood as 'virtues of salience'—placing demands on what we find salient or attention-grabbing. While this frequently affects our initial inclinations to believe, it needn't affect our settled beliefs. We end by considering what other putative virtues might appropriately be understood as virtues of salience. (We'll remain neutral throughout on the larger question of whether these

⁹ It's worth noting that our explanation of the disconnect between Act Evaluation and Character Evaluation could be used by consequentialists to support their position, in so far as it reveals that the consequentialist needn't simplistically read character judgments off of agents' actions.

putative virtues truly are virtues.¹⁰⁾

3.1 Modesty and Ignorance

Driver [1989, 1999] argues that the virtue of modesty consists in a disposition to moderately *underestimate* one's own worth. This explains the 'Moore-paradoxical' infelicity of the assertion:

1. I am modest.

On Driver's account, the infelicity of (1) is to be explained in terms of the general infelicity of claiming that one's own belief is false: If you think it's false, then how can you believe it? Likewise: If you think you're underestimating your self-worth, then isn't that just to say that you actually think your self-worth is rather higher than previously intimated?

Driver notes that 'behavioural' accounts of modesty, either in terms of *understating* one's true worth, or a general reluctance to brag, can also account for the oddity of asserting (1). But they fail for the reason that they cannot distinguish sincere from false modesty. The difference, for Driver, is that the genuinely modest person does not merely behave as though she has less worth, she really believes it.

We may wonder: If modesty really involves ignorance in this way, then how is it a virtue? Driver [1989: 383] initially suggested that it is because modesty-as-ignorance typically arises from 'a reluctance to dwell on one's good qualities' or give much thought to rankings, and it is this disposition, rather than the resulting ignorance as such, that is truly valuable (at least instrumentally, and perhaps intrinsically as well). But in that case, why not take this valuable disposition, rather than the contingently resulting ignorance, to be constitutive of the virtue of modesty?

In later work, Driver [1999: 829] takes the more radical stance:

In writing my original paper I had intended simply to argue against the thesis

¹⁰ It is a separate project – and not one that we undertake here – to give systematic criteria for what a virtuous agent would find salient. Our aim is rather to account for our pretheoretic intuitions about these putative 'virtues of ignorance', and to assess which psychological feature (ignorance, attention, or salience) best explains them.

that no virtue can involve ignorance in any way. ... However, I would like to go further than this and claim that it is the ignorance that we value.

To establish this stronger claim, she considers a case in which ranking dispositions and ignorance come apart. Suppose Albert the scientist puts a great deal of thought into his comparative standing as a physicist. He publicly declares himself to be (as he now believes) the fifth best physicist in the world, though in fact the evidence shows that he is third best. Driver bites the bullet and insists that, since he underestimates himself, Albert is modest—albeit in an ‘anomalous’ fashion, ‘modest *in spite of* his overzealous ranking behaviour’ [ibid.]. Albert’s disposition to underestimate his own worth is itself sufficient to establish this result.

But as we see it, this case is under-described. Imagine several variations of the case. In the first, Albert is obsessed with his own achievements and abilities, which are hyper-salient to him. This hyper-salience is what drives him to compulsively rank himself against others. Nevertheless, he misjudges himself as merely the fifth best physicist in the world. In this case, it seems to us that Albert is not really modest at all, in any normatively significant sense. He is instead (mildly) epistemically irrational, in addition to being immodest. He is, in this way, doubly flawed. It’s the salience of his own achievements that is driving him to rank himself and to proclaim his (perceived) status to others. And despite all his efforts he isn’t even able to accurately assess what it is.

We might imagine other versions of this case, in which Albert is engaged in the ranking process for other reasons and plausibly is modest. Perhaps Albert is just obsessed with ranking things. He ranks historians, philosophers, countries’ GDPs, and so on. It’s this obsession with ranking things, not any perverse salience of his own achievements, that drives him to rank physicists (and himself among them). Or perhaps a wealthy philanthropist will donate one million dollars to Albert’s favourite charity if he gives them a ranking of the worlds’ physicists. Or perhaps, in a state of depression, Albert sets out to rank physicists to see just how terrible he is. In these cases, Albert may well be modest.

What these cases bring out is that, contrary to Driver, underestimating his worth as a physicist is not sufficient for modesty. If his own physicist-virtues are abnormally salient to Albert, but for some other reason he misjudges his worth as a

physicist, this is not indicative of modesty.¹¹

Driver argues that what we value in modesty is ignorance itself (rather than an anti-ranking disposition), by considering a case where the two come apart in the opposite direction: Bob knows, on the basis of reliable testimony, that he is the best, though he hasn't himself engaged in any ranking exercise to confirm this. Driver [1999: 829] objects that '[a]ny professions of inferiority on his account would constitute false modesty', and hence be found objectionably patronizing and condescending by knowledgeable observers. Now, we agree that any such dishonest attempts to placate the presumed jealousy of his audience would constitute condescending rather than modest behaviour. But the mere *possibility* of behaving condescendingly cannot be sufficient to show that an agent is *actually* immodest. (Note that Driver believes Albert, above, to be modest. But he too could behave condescendingly, e.g., by dishonestly reassuring his colleagues that he's 'not even in the top ten'.)

Suppose that Bob never gives a moment's thought to his relative ranking. What he finds salient are, rather, opportunities out there in the world. In so far as he assesses himself at all, he is disposed to do so in non-comparative terms, noticing where he has room for improvement [Brennan 2007]. He doesn't think of himself as better than other people, for he doesn't think in terms of comparative rankings at all. (In this way, he differs from the falsely modest person who merely *pretends* not to think of himself as better than others.) In this case, Bob strikes us as a paradigmatically modest person. This is so even though he could, if asked, retrieve from his dusty memory banks the information that the top-ranked person in the world happens to be . . . him. Whoa!

On our account—call it modesty-as-salience, in contrast to Driver's account of modesty-as-ignorance—the virtue of modesty need not involve any epistemic error or impairment. It merely requires that the agent not be disposed to dwell overmuch on her own excellences or virtues. (See Bommarito [2013] for a similar

¹¹ Bommarito [2013: 106] makes a similar point, noting that what is relevant to immodesty is 'why a person overestimates his or her good quality'. On Bommarito's view, this requires attention to one's own positive features, and not just any attention, but attention 'for the right reasons', e.g., not 'because an attention disorder prevents them from attending to anything for very long or because they are the kind of pessimists who never attend to any good qualities at all' [ibid.: 103]. For a discussion of the differences between our view and Bommarito's, see footnote 12 and section 3.3.3.

account of modesty as a virtue of attention.)¹² She may know the truth of the matter, but it isn't something she cares about.¹³ And so it isn't something she finds salient, or that tends to intrude into her thoughts.

Both Driver's account and ours agree that a modest agent will typically not be aware of her ranking or comparative worth. But we diverge when it comes to explaining *why* this is so. Driver proposes that the modest agent must not *believe* the truth about her self-worth, she must instead underestimate it. We instead propose that, rather than giving the wrong answer, the modest agent simply is not *disposed to attend* to the question.

We can further distinguish these two views by employing a 'Fate of the World' test, which shows that erroneous beliefs are not necessary for modesty. Suppose an evil demon will destroy the world unless Claire offers an accurate assessment of her abilities. Would Claire's answering correctly prove that she lacks modesty? On Driver's view, it would. If modesty requires underestimation, then Modest Claire's best effort at answering the question will yield an incorrect answer. But on our view, this need not be so. Modest Claire is not disposed to dwell on her accomplishments (does not find them salient), so the answer may not immediately spring to mind. She may even be *initially* inclined, just as a matter of first appearances, to underestimate herself, due to giving such little thought to her many achievements.¹⁴ But when the stakes are high, she is able to override her characteristic disposition to refrain from attending to her own importance. She can systematically comb through the evidence in an accurate and dispassionate light. So,

¹² The key difference between our account and Bommarito's is that salience consists in your disposition to find something attention-grabbing, rather than in whether you *actually* end up attending to it. So, for example, if an evil demon threatened to blow up the world unless Dale dwells on his own accomplishments each evening, his subsequent pattern of dwelling would bar him from qualifying as modest, according to Bommarito's modesty-as-inattention view. This strikes us as too quick. We think that Dale may still be modest so long as his accomplishments are not *salient* to him, such that he has to work really hard to keep them at the forefront of his mind, and it feels to him awkward and unnatural to do so.

¹³ Schueler [1997] offers the related proposal that modesty consists in not caring *whether others are impressed* by you. But as Driver [1999] objects, one might disdain others' opinions out of extreme arrogance rather than modesty (though cf. Schueler's [1999] response). Our proposal avoids this worry, since the arrogant person is still interested in ranking, and finds his own status highly salient, it's just that he assumes that others are lower-ranked than he, and disregards their opinions on that basis.

¹⁴ The modest disposition not to dwell on her accomplishments makes it likely that she will not do so, and hence may underestimate herself. But note that neither the underestimation nor (as explained in footnote 12) the inattention that may cause it are essential to her modesty on our view.

upon considering the matter in depth, she is able to give the demon the correct answer. Then, the immediate need having been met, her attention will once again drift away from herself, and back to what she considers to be intrinsically more important matters.

This account of modesty-as-salience has two major advantages over Driver's modesty-as-ignorance view. First, it yields verdicts that are intuitively more plausible in the cases of Albert, Bob, and Claire, discussed above. Second, it seems more appealing on theoretical grounds, in so far as we believe that our account can offer a better explanation of *why* modesty is a virtue.

On our account, the modest agent is one who is not disposed to find their own positive attributes especially salient. These dispositions would naturally result from the agent's having internalized certain important (putative)¹⁵ moral truths—that their own achievements aren't all that important in the grand scheme of things, that each of us is but one person amongst moral equals, and that comparative rankings lack intrinsic importance.¹⁶ In so far as the disposition not to find one's own positive attributes salient stems from such an appreciation of moral truths, it has a kind of intrinsic appropriateness that befits its status as a genuine virtue.¹⁷

By contrast, Driver's account of modesty bestows it with merely instrumental value: '*what* makes modesty as ignorance valuable[?] My claim is that a person who is modest stops problems from arising in social situations' [Driver 1999: 828]. In

¹⁵ Helen is not persuaded that the first two claims are moral truths or that modesty is a virtue. But perhaps you are.

¹⁶ One might wonder whether the non-salience of one's own positive traits is playing any role in modesty's being a virtue. Perhaps it is simply having internalized these moral truths that is virtuous, and not the salience dispositions that stem from this internalization. But arguably the salience dispositions are a constitutive part of *truly* internalizing the egalitarian perspective. We may imagine someone who says, 'Yes, I'm not more important than anyone else from a God's eye perspective. But that's not the perspective *I'm* taking. I'm attending to me and my awesomeness!' But there's a clear sense in which this person hasn't *truly* internalized the egalitarian perspective – their moral belief is too compartmentalized from their personal perspective. Perhaps there is a weaker sense in which they have 'internalized' the moral claim in virtue of readily believing it, but then this very case makes clear that mere 'internalization' *without the right salience dispositions* is not sufficient for modesty.

¹⁷ But there might be other, 'anomalous' explanations for such salience-related dispositions. Perhaps I took a pill that renders me unable to attend to any of my own positive attributes. Our account of the value of modesty would not apply to such a case. We're tempted to conclude that modesty, though *typically* a virtue, is thus capable of being instantiated in a defective form. But if one thinks that modesty is essentially a virtue, one could simply embrace a more restrictive notion of modesty on which the salience dispositions must flow from the appreciation of moral truths to truly count as *modesty* at all. We think this is merely a terminological point, and are happy to accept that there are broader and more restrictive uses of the term.

particular, modesty-as-ignorance may be instrumentally useful in defusing problematic social emotions like jealousy that can arise when faced with superior others. These negative social emotions arise precisely from an exaggerated concern for one's status. Our view of modesty avoids these issues from the start, as the modest agent (i) won't have such concerns, and hence (ii) won't inflate themselves or emphasize their status in a way that is threatening to others. In sum, modesty-as-ignorance results in people putting forward a less threatening presentation of their status, while modesty-as-salience results in people not putting forward status presentations at all. Thus, our view of modesty captures the same instrumental value as Driver's, in addition to opening the possibility of modesty's having intrinsic value.

3.2 Friendship and Epistemic Partiality

Keller [2004: 329] uses the following incident from the sitcom *Friends* to suggest that friendship requires epistemic partiality, or thinking better of our friends than the evidence warrants:

Joey and Chandler are playing a game where the latter gives immediate, unreflective and unfiltered answers to the questions asked by the former. Joey has just landed an acting job in Las Vegas, which he hopes will be his big break. He asks Chandler, 'Is this job going to be my big break?' to which Chandler reflexively answers 'No,' putting their friendship in crisis. Chandler's lack of belief in him causes Joey to feel betrayed, and Chandler to feel guilty, suggesting that their friendship involves normative expectations to think well of each other—even in the absence of evidence warranting such optimism.

Stroud [2006: 508] similarly argues:

[T]he bias of the good friend will normally take the form of casting what she sees or hears in a different light, shading it differently, placing it in a different optic, embedding it in a different overall portrait of her friend. Where our friends are concerned, in short, we become spin doctors.

Both Keller and Stroud conclude that friendship places demands on our *beliefs*. As in the previous section, we want to resist this strong conclusion, and replace it instead with a more subtle demand on what we find *salient*. We agree with Stroud's

characterization of how friendship requires us to *see* things in a more positive light. The good friend will be disposed to attend to their friends' virtues, talents, and successes. Even in light of damning evidence against a friend, the good friend will continue to find her positive attributes disproportionately salient. This bias in salience may—particularly on initial assessment—lead friends to form beliefs about friends that are more rosy than the evidence merits. But, as with modesty, we see such false beliefs as a contingent consequence of good friendship, not as a requirement of friendship in its own right.

Again, we can distinguish the alternative views on offer by means of a 'Fate of the World' test: Suppose that an evil demon will destroy the world unless Chandler answers correctly whether Joey's new job will be his big break. And let's alter the details of the case to specify that Chandler finds Joey's positive attributes more salient than his weaknesses: The dimensions along which he's a good actor are what really stand out to Chandler, those along which he's weaker fade into the background. Joey's past successes are more salient than his past failures. This bias in salience together with Chandler's deep hopes for Joey's success and the natural human inclination towards wishful thinking cause Chandler's unreflective impression to be an optimistic one: 'Yeah! Joey's got a shot.' But then, as the stakes are so high, he pauses and reflects more carefully on the question. He tries to systematically review the evidence, deliberately attending to things that friendship might have made less salient. He forces himself to explicitly attend to the low base rate of success in the acting business. And recognizing that, as a friend, he's naturally going to be a bit biased in Joey's favour, he explicitly adjusts his credence down to correct for this. Weighing it all up carefully, Chandler concludes (against his initial inclination) that the answer is in fact 'No'.

When his process of belief-formation is spelled out in this way, does Chandler still seem like he has been in any way a bad friend? Surely not. The crucial difference between the original case and this one is that while Chandler still comes to a negative conclusion about Joey's career prospects, in our case his initial inclination was more positive. And this seems all that friendship, intuitively, demands. A good friend finds his friend's strengths to be more *salient* than his weaknesses, which naturally leads to an initial inclination towards overestimation. But there is no

requirement that we settle for first appearances. We can (if pressed) correct for our biases, and so reach a more accurate final conclusion, without in any way violating the norms of friendship. The problem with Chandler in the original case is not that he *believed* poorly of Joey, but that it didn't even *appear* to him that Joey would do well.

3.3 Other Virtues of Salience

Driver suggests that there may be a great many virtues that (like modesty on her view) are best thought of as virtues of ignorance. We suspect that she's right that there are many virtues that fall into the same class as modesty and should receive similar treatment, but as with modesty we will argue that these are best thought of as virtues of salience.¹⁸

3.3.1 Blind Charity

A *Blindly Charitable* person is—as Driver [1989: 381] defines it—a person who 'sees the good in [others], but does not see the bad'. The blindly charitable person not only refrains from speaking ill of others, he doesn't *think* ill of others, even when the evidence would warrant it. In this way, Driver argues that blind charity is 'an ignorance about what is bad in other people' [ibid.: 382].

We think that such a tendency to see only the good and not the bad in people is also better understood as a virtue of salience. One does not have to—on reflection—believe that there is no evil in the world, that their neighbours never think unkind thoughts, and so on, to possess this virtue. One could exemplify the virtue of seeing the good, not the bad, and yet recognize on an intellectual level that others sometimes act out of spite, gossip, and don't have their friends' backs.

¹⁸ One might be inclined to think that both Driver's account and our own pick out something of value. One could adopt a hybrid view on which there are two valuable aspects of modesty: ignorance and salience. While nothing we have said strictly rules this out, we're disinclined to this proposal for two reasons. (1) The salience account of these virtues is able to fully account for the intuitive data points that motivated Driver's view, thus undercutting the need to posit value in ignorance. (2) In further cases which differentiate the two accounts (e.g., various 'Fate of the World' cases), the ignorance accounts seem to positively yield the *wrong* results. A hybrid view would inherit this flaw.

What's relevant is that you find the good in those around you much more salient than the bad. Consider Ed: Ed finds it hard to attend to the flaws not only of his friends but of all his fellow humans. When someone takes a parking space that he was clearly waiting and indicating for, he doesn't dwell on the small harm they've done him. What's salient in his mind is the exhausted and frantic look of the driver as she herds her three children across the parking lot, and how she seems patient with her children even when she's so obviously having a difficult day. This is typical of Ed. He always finds the positive traits of others more salient than their negative traits. He views his fellow human beings through rose-tinted lenses.

Ed seems like a paradigm case of the blindly charitable. This seems so even if—were the fate of the world to depend on it—he could systematically comb through the evidence about the characters of his fellow humans, deliberately compensating for his attentional biases, and come to an accurate assessment of the characters of others. As in the case of friendship, it seems that what's valuable is not Ed's ignorance, but his initial inclination to believe the best of others—the fact that he finds the good in others to be more salient than the bad in them.¹⁹

In this way, being blindly charitable is much like being a good friend. But where friendship arguably requires us to find *our friends'* positive traits comparatively salient, the blindly charitable extend this charitable outlook to *all* those they encounter.

3.3.2 *Forgiving and Forgetting*

Another virtue that Driver [1989] describes as a virtue of ignorance is the willingness to 'forgive and forget' rather than holding a grudge for past harms. This too is better captured as a virtue of salience.

Your life-long friend reveals some personal information that you confided in her, ultimately resulting in your losing your job. Because of the value you place on your friendship, you decide to 'forgive and forget'. This does not require you to wipe all traces of the past harm from your memory, or to become *ignorant* of the past

¹⁹ In fact, while Driver [1989: 382] ultimately describes blind charity as a virtue of *ignorance*, that 'requires that one be unknowing about something', she initially describes blind charity as 'a *disposition* not to see the defects, and to focus on the virtues of persons' (Driver [1989: 381], emphasis added).

wrong.

Further, ‘forgiving and forgetting’ doesn’t seem to require that you *actually* fail to attend to the past harm. Imagine that something occurs that forces you to attend to the past harm (despite not finding it particularly attention-grabbing). Perhaps the police have come by asking about the event. The fact that you are, as the police question you, attending to the event, does not mean that you are holding a grudge. Even were you required to continually attend to the event in a drawn out trial, this would not entail that you held a grudge.

So ‘forgiving and forgetting’ is a virtue of salience (rather than of ignorance or attention). This fits naturally with the intuitive idea that if you have truly internalized your forgiveness, the harms done to you in the past will no longer tug at your heartstrings. The memories are silent, even if external circumstances can cause you to attend to them.

3.3.3 Gratitude

We conclude with a virtue that may *not* be so well captured as a virtue of salience, but which clarifies what it is to be a virtue of salience, as opposed to a virtue of attention [Bommarito 2013].

To *feel* gratitude seems to require that we *actually attend* to the value of something someone else has done for us. Someone who found the value of what someone had done for them attention-grabbing (salient), but who through some fluke never actually attends to this, does not feel grateful, though they may be disposed to feel gratitude.

It’s an interesting question just what the relationship is between feeling gratitude and being a grateful person (i.e., possessing the virtue of gratitude).²⁰ If gratitude is a virtue of attention, then you could become a grateful person by setting alarm reminders throughout the day to attend to good things others have done for you. (Perhaps you think that you don’t spend sufficient time attending to the good things others do for you, and decide to remedy this by setting your alarm to ensure you give these goods the attention they deserve.) But the mere frequency of grateful

²⁰ We’re grateful to Eden Lin for prompting us to consider this more carefully.

feelings might seem insufficient for making you a *grateful person* if they require such deliberate efforts. On the other hand, if gratitude is a virtue of salience, then one could be a grateful person despite never actually feeling gratitude (due to a series of flukes that continually draw your actual attention away from what others have done for you). This might seem surprising.

This might lead some to prefer a hybrid account: Perhaps the truly grateful person must *both* find others' assistance to be salient *and* actually end up attending to it (for this reason). We don't have a firm view on this. But whatever the right account of gratitude turns out to be, we think the range of options clearly brings out the difference between virtues of mere attention and virtues of salience proper.

4. Wrapping up

This paper has explored the two-way relation between salience and good character. First, we saw that our evaluation of an agent's moral character needs to take into account what they find salient, since neglecting a salient need reflects a greater deficit of benevolent motivation than does neglecting an objectively similar but much less noticeable need. In this way, facts about salience can serve as an important 'input' to our moral assessments. In the second half of the paper, we saw that the connection also goes the other way: Finding some things more salient than others can be an important 'output', requirement, or downstream consequence of good character. In particular, we argued that the virtue of modesty consists in not finding one's own achievements excessively salient, and that a good friend will find his friends' better qualities to be more salient than an impartial stranger would find them. Our proposals differ from previously floated views, in both cases, because we insist that there's no further requirement for the agent to *believe* in line with initial appearances—he may instead correct for any biases introduced by what he finds more or less salient.

But given that virtue can place constraints on what we find salient, this might lead one to wonder whether being virtuous might also require that, e.g., we find distant harms as salient as nearby ones, find harms to out-groups as salient as harms to in-groups, and so on. Perhaps in the Singer case, the truly virtuous person would

find those starving on the other side of the world just as salient as the child drowning before their eyes. Perhaps it reflects poorly on us that we don't find these things equally salient.

Perhaps.²¹ But this is fully compatible with our contention that—given that we have the salience dispositions that we do—ignoring the drowning child displays a greater deficit of goodwill (and is more blameworthy) than ignoring distant needs. Our aim in the first part of the paper was to establish the comparative claim: neglecting actually salient needs is generally *more* blameworthy than neglecting (actually) non-salient ones. This is compatible with thinking that an ideally virtuous agent would differ from us in what they find salient.²² It's simply to add that, just as what we find salient can reflect well or poorly on us, so it can affect how our (in)actions reflect on our moral character.²³

University of York

References

- Arpaly, Nomy 2003. *Unprincipled Virtue: An Inquiry Into Moral Agency*, New York: Oxford University Press.
- Blum, Lawrence 1991. Moral Perception and Particularity, *Ethics* 101/4: 701–725.
- Bommarito, Nicolas 2013. Modesty as a Virtue of Attention, *Philosophical Review* 122/1: 93–117.
- Brennan, Jason 2007. Modesty Without Illusion, *Philosophy and Phenomenological Research* 75/3: 111–28.
- Cullity, Garrett 2006. *The Moral Demands of Affluence*, Oxford: Clarendon Press.
- Driver, Julia 1989. The Virtues of Ignorance, *Journal of Philosophy* 86/7: 373–84.
- Driver, Julia 1999. Modesty and Ignorance, *Ethics* 109/4: 827–34.
- Huber, D., Khosla, D., and Dow, P. 2009. Fusion of Multi-Sensory Saliency Maps for Automated Perception and Control, *SPIE Proceedings* 7336.

²¹ We indeed find it plausible that we should strive to develop broader sensitivity to others' needs, though we're not sure that it would be ideal for us to find distant needs *equally* salient to nearby ones, for the simple reason that we tend (all else equal) to be in a better position to do something about nearby needs.

²² The cases we considered in section 2 involved patterns of salience that actual people regularly exhibit. By contrast, it is much less clear whether it's humanly possible to be completely impartial in our patterns of salience. We think it's particularly interesting to evaluate dispositions of salience that we can realistically hope to achieve.

²³ Special thanks to Eden Lin for invaluable comments on several drafts of this paper. Thanks also to Elizabeth Harman, Simon Keller, Philip Pettit, Michael Smith, and an anonymous referee for helpful comments and discussion.

- Itti, L. and Koch, C. 2000. A Saliency-Based Search Mechanism for Overt and Covert Shifts of Visual Attention, *Vision Research* 40/11: 1489–1506.
- Kagan, Shelly 1991. *The Limits of Morality*, Oxford: Oxford University Press.
- Kalinli, O., and Narayanan, S. 2007. A Saliency-Based Auditory Attention Model with Applications to Unsupervised Prominent Syllable Detection in Speech. *INTERSPEECH*: 1941–4.
- Kamm, Frances 1999. Famine Ethics: the problem of distance in morality and Singer's ethical theory, in Dale Jamieson (ed.), *Singer and His Critics*. Oxford: Blackwell Publishers.
- Kayser, C., Petkov, C., Lippert, M., and Logothetis, N. 2005. Mechanisms for Allocating Auditory Attention: an Auditory Saliency Map, *Current Biology* 15/21: 1943–7.
- Keller, Simon 2004. Friendship and Belief, *Philosophical Papers* 33/3: 329–51.
- Koch, C. and Ullman, S. 1985. Shifts in selective visual attention: towards the underlying neural circuitry, *Human Neurobiology* 4/4: 219–27.
- Milgram, Stanley 1965. Some Conditions of Obedience and Disobedience to Authority, *Human Relations* 18/1: 57–76.
- Schueler, G.F. 1997. Why Modesty is a Virtue, *Ethics* 107/3: 467–85.
- Schueler, G.F. 1999. Why IS Modesty a Virtue? *Ethics* 109/4: 835–41.
- Slote, Michael 2007. Famine, Affluence, and Virtue, in Rebecca L. Walker and P. J. Ivanhoe (eds.), *Working Virtue: Virtue Ethics and Contemporary Moral Problems*. New York: Oxford University Press.
- Singer, Peter 1972. Famine, Affluence, and Morality, *Philosophy and Public Affairs* 1/3: 229–43.
- Strawson, Peter 1962. Freedom and Resentment, *Proceedings of the British Academy* 48: 1–25.
- Stroud, Sarah 2006. Epistemic Partiality in Friendship. *Ethics* 116/3: 498–524.
- Unger, Peter 1996. *Living High and Letting Die: Our Illusion of Innocence*, New York: Oxford University Press.