INTELLECTUAL VIRTUES AND BIASED UNDERSTANDING

Andrei Ionuţ Mărăşoiu

University of Bucharest

andrei.marasoiu@icub.unibuc.ro

ABSTRACT: Biases affect much of our epistemic lives. Do they affect how we understand things? For Zagzebski (2001), we only understand something when we manifest intellectual virtues or skills. Relying on how widespread biases are, Carter and Pritchard (2016) raise a skeptical objection to understanding so conceived. It runs as follows: most of us seem to understand many things. We genuinely understand only when we manifest intellectual virtues or skills, and are cognitively responsible for so doing. Yet much of what we seem to understand consists in conceptions whose formation could have easily been due to biases instead, and the work of biases is opaque to reflection. If conceptions constituting how we understand things could have easily been due to biases, then we are not cognitively responsible for them because we cannot reflectively appraise what we understand. So, we are mistaken in thinking we genuinely understand most of the time. I will defend the grounding of understanding in intellectual virtues and skills from Carter and Pritchard's objection. We are cognitively responsible for understanding when we manifest our expertise. We can do so, I will argue, without being required to reflectively appraise what we understand.

I. UNDERSTANDING

We use the word 'understanding' in many ways.[1] Honorific uses aside, we can illustrate several varieties of understanding. When we understand what someone says, perhaps we only know what they mean. When we understand what we see, perhaps that only comes down to segmenting the visual scene by applying concepts. When we understand what a friend is going through, perhaps that only calls for empathy. When we understand how to ride a bike, perhaps that only involves the ability to do it and familiarity with having done it in the past. Linguistic, perceptual, emotional, and practical understanding doubtless exist.

However, I will focus on intellectual understanding, following Zagzebski (2001), who is the main proponent of a virtue epistemology of understanding. I do so in order to consider and reject a skeptical challenge to virtue epistemologies of understanding. The challenge, raised by Carter and Pritchard (2016), relies on the fact that biases are ubiquitous and pervasive in thinking. In the rest of this section and the next, I sketch the relevant details of Zagzebski's virtue epistemology of understanding. In Section III-V, I formulate Carter and Pritchard's skeptical challenge starting from the pervasiveness of biases and the risks they pose to thinking. Sections VI-VII articulate my response to the skeptical challenge on behalf of a virtue epistemology of understanding, which appeals to a notion of responsibility that doesn't presuppose that we are able to reflectively appraise what we understand.

What is understanding?

Understanding is a cognitive state that arises from *technê* ... The person who

has mastered a *technê* understands the nature of the product of the *technê*

and is able to explain it. (Zagzebski 2001: 240)

For Zagzebski, to understand something manifests skill. One's understanding is "the

product of the *technê*" – and one has "to be able to explain it." For instance, what is it to

understand the position on a chessboard? Most players can, on inspection, identify the state

of the game that they are playing and forming conceptions about, e.g., what weaknesses

the opponent's position has, and how those can be exploited. Conceptions of this sort,

which constitute how players understand a position, are available for them to use in

explaining that position.[2] This conception of understanding is in the background of

Zagzebski's virtue epistemology, to which I now turn.

II. VIRTUES AND SKILLS

Zagzebski writes: "Understanding is an epistemic state that arises from *technê*" (240).

Technai are virtues or skills. Which ones are at play in understanding varies from case to

case. Understanding some phenomena may require open-mindedness to see alternatives,

and understanding other phenomena may require a vivid imagination.

Zagzebski's virtue epistemology is best construed as addressing four aspects of understanding: justification, formation, excellence, and responsibility. The conception by which we understand something is epistemically justified because it was formed in the right way: its formation, in the initial episode of coming to understand, must have primarily manifested intellectual virtues or skills appropriate to the context of inquiry. The manifestation of virtues and skills explains what is excellent in cognition, aiming at optimally meeting epistemic norms such as accuracy, coherence, or explanatoriness.

Other epistemologies of understanding can also account for the formation, justification, and excellence of the conception by which we understand something. Consider reliabilism. As Zagzebski (2001: 238) argues, virtues and skills are reliable: those who understand possess varying degrees of expertise, and an expert is "the person who is the most reliable source of knowledge" in her field of inquiry. The reliability of virtues and skills explains why the conceptions they help form are mostly true. Or consider coherentism. "Understanding involves seeing how the parts of that body of knowledge fit together" (Zagzebski 2001: 243). Presumably, it is the conceptions of experts that are most coherent in their domains of expertise.

What's novel in Zagzebski's virtue epistemology is a focus on epistemic agents. In understanding, we don't merely house special states of mind. Rather, we are responsible for how we understand things; we author our understanding inasmuch as *our* virtues and skills explain how we came to conceive what we understand.

With this sketch of a virtue epistemology of understanding in the background, we can now consider Carter and Pritchard's skeptical challenge to it.

III. BIASES

Carter and Pritchard's skeptical challenge relies on the ubiquity and importance of biases in everyday thought. To set the stage for their challenge, this section asks what biases are. Carter and Pritchard say:

> A bias, in the most general sense, is a disposition, implicit or explicit, to reach a particular kind of conclusion or outcome – in the kind of case we're interested in, the outcomes will be representational. Call these *cognitive* biases. (Carter and Pritchard 2016: 273)

This description doesn't lead far. Any disposition is "implicit or explicit." And, we can likewise assume that any cognitive disposition is manifested in "a particular kind of conclusion or outcome." So far, all we have to go on is that cognitive biases are dispositions to produce presumably cognitive representations. This doesn't say why biases threaten understanding.

To properly appreciate the skeptical challenge biases occasion, it pays to improve on Carter and Pritchard's characterization of biases. Two options seem open for what might be wrong with biases. Either the representation produced is defective: biases help form conceptions which are systematically *erroneous.* Or the way in which representations are

produced is defective: biases taint the *etiology* of the conceptions they help form. Let's

take them in turn to see if fleshing them out can occasion a skeptical challenge.[3]

Kahneman (2002: 3-4) voices the often-met view that biases typically produce

inaccurate representations: "Systematic errors are known as biases, and they recur

predictably in particular circumstances." However, the defective-representation construal

faces a problem: it's unclear why biases would systematically depart from the truth.

Perhaps the appearance that biases produce inaccurate representations is largely due to

unrealistic benchmarks for what should count as sufficiently accurate human reasoning.[4]

Deliberation might sometimes take too long, whereas we often have to reach decisions

quickly, and in conditions of uncertainty. Fast and frugal heuristics (cf. Chase, Hertwig and

Gigerenzer 1998) might be more effective in advancing our adaptive goals, being accurate

enough most of the time, in typical circumstances. Seen this way, biases wouldn't so much

fuel skepticism as indicate context-sensitive trade-offs between accurate goal-achievement

and limited resources within overall successful cognitive performances.

A skeptical challenge might more readily[5] be mounted based on how biases feature

in the *etiology* of our conceptions. In virtue epistemology, conceptions constituting our

understanding are justified by having been formed in the right way – primarily, through

manifesting our intellectual virtues or skills. Why would biases playing a role in how our

understanding is formed be epistemically vicious? Morewedge and Kahneman say:

> It is often useful to think of judgments as a weighted combination of items of
>
> information... In this scheme, judgment biases can always be described as an

6

overweighting of some aspects of the information and underweighting or

neglect of others, relative to a criterion of accuracy or logical consistency...

In this fashion, the principles of associative activation help explain biases of

judgment. (Morewedge and Kahneman 2010: 435)

They see the most common cognitive biases as side-effects of associative semantic

memory. Abstracting away from the details of their hypothesis, the threat biases pose

seems to be the following. How activation spreads within a semantic network is not up to

the agent. It is an informational and biological *datum*. So, even when conceptions which

partly result from biases are accurate, it's far from clear that we are responsible for what

we believe: because we didn't have a say about how those conceptions were formed, and

can't effectively control them. This worry about how we control our own thoughts – how

they are, in a sense, *ours* – seems principled. The worry applies most poignantly to those

thoughts we value most, the ones by which we understand the world around us. I'll now

argue it is this kind of worry that motivates Carter and Pritchard's skepticism.

IV. CARTER AND PRITCHARD'S CHALLENGE

In this section, I sketch Carter and Pritchard's bias-driven skepticism.[6] They write:

On the supposition that the individual's cognitive success is not primarily explained by her exercise of cognitive ability, but also partly down to an unconscious bias, is one able to enjoy the kind of cognitive ownership of this fact that is characteristic of understanding? We suggest not. (287)

Grant that understanding manifests virtues and skills. However, most of us are frequently subject to biases, enough so that we can't sort out biases or virtues-plus-skills as primary causal factors in how we have come to understand something. Then, even if our conceptions were in fact primarily due to virtues and skills, they could have easily been primarily due to biases instead.[7] Why is this possibility problematic for understanding? Because biases, being sub-personal cognitive routines that operate relatively automatically and implicitly, undermine our cognitive responsibility – "the kind of cognitive ownership ... that is characteristic of understanding."[8] As Carter and Pritchard put it:

Like cognitive achievement, understanding requires a particular kind of cognitive ownership, that isn't essential for merely knowing. But, unlike cognitive achievement, that cognitive ownership also essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities. (Carter and Pritchard 2016: 288)

Here, understanding is likened to cognitive achievement. "Achievement," however, may stand for any of the following: (i) epistemic success achieved, i.e., individuated in terms of how it was acquired or produced; (ii) success creditable to the epistemic agent; (iii) success due to the epistemic agent's virtues or skills; (iv) success that the epistemic agent is responsible for. These four conditions seem, pending further argument, to be logically independent from each other. Comparing understanding to cognitive achievement seems to involve all four conditions being met in the case of understanding as well.

Carter and Prichard claim (in the excerpt just given) that cognitive responsibility for understanding implies that one should have *a basis, in reflection*, for excluding "live-error" possibilities. What are some live-error possibilities being excluded? If the conception constituting how we understand something could have *easily* been due to biases rather than our intellectual virtues or skills, then the possibility of an epistemically vicious etiology for what we seem to understand is *live*. The problem is that, since biases often operate implicitly, we typically *don't* have a basis in reflection for what we seem to understand whenever what we seem to understand could have easily been due to biases rather than our intellectual virtues or skills. By their reflective opacity yet cognitive efficacy, biases undermine our cognitive responsibility. Since understanding implies cognitive responsibility, we often don't *genuinely* understand what we *seem to*.

Although they only aim to *motivate* a skeptical worry, I believe Carter and Pritchard's challenge is best represented in argument-form. Carter and Pritchard don't offer this argument word for word. My own reply to their skeptical challenge depends on how faithful to their intentions this argument is:[9]

(1) All instances of genuine understanding are primarily due to the epistemic agent's intellectual virtues and skills.

(2) If an instance of understanding is primarily due to an epistemic agent's intellectual virtues and skills, then that agent is cognitively responsible for it.

(3) If an epistemic agent is cognitively responsible for understanding, then that agent must be able to appraise what they understand in reflection.

(4) Most instances of seeming to understand could have easily been due to biases (though they weren't).

(5) If an instance of seeming to understand could have easily been due to biases (though it wasn't), then the epistemic agent couldn't appraise what they understand in reflection.

(6) So, most instances of seeming to understand are not instances of genuine understanding.

The premises each do important work in supporting the conclusion. (1) is key to how virtue epistemology construes understanding: as a state of mind individuated in terms of its acquisition – an acquisition the epistemic agent can be credited for. I defend understanding, construed as based on virtues and skills, against skepticism; so I grant (1).

Including (2) purchases Carter and Pritchard a narrower target. If one had a reliabilist approach to intellectual virtues or skills (Greco 1999), instances of

understanding that are luckily free from bias could be recognized to exist by denying (2).

Carter and Pritchard's challenge only properly applies to responsibilist virtue

epistemologies of understanding,[10] e.g., Zagzebski's (2001), not to reliabilist ones. Given

the dialectic, we should grant (2).[11]

(4) is an empirical premise to the effect that, given how widespread biases are, it is

always possible that conceptions formed well, and constituting our understanding of

something, might have easily been produced primarily by our biases rather than our

intellectual virtues and skills. I, of course, grant that biases are widespread, and that they

often insinuate themselves in how we think.

(5) concerns how biases interact with reflection. No one can reflectively appraise

what one understands when how one conceives of what is understood is reflectively

opaque. As biases typically operate implicitly, their absence is equally unavailable to

reflective scrutiny.[12] Since Carter and Pritchard's challenge trades on what typically

happens when biases luckily fail to be primary causal factors in how we come to

understand something, we should grant (5).

I argue against skepticism about understanding. So I deny conclusion (6) that most

instances in which we seem to understand something aren't instances in which we

genuinely understand. The premise I deny is (3). Before arguing, in Section VII, against

(3), that one can be cognitively responsible without reflectively appraising one's

understanding, I'll clarify the central terms involved: the appeal to reflection (Section V),

and our responsibility for manifesting virtues or skills in understanding (Section VI).

V. REFLECTIVE APPRAISAL

In Section IV, I called the requirement Carter and Pritchard impose on the kind of cognitive responsibility we have for understanding 'reflective appraisal.' Their exact words refer to a "cognitive ownership" that "essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities" (2016: 288). Notice two key phrases: "reflective grip" and "rational basis."

What do they mean? One has a *reflective grip* on a matter when one can reflect on it, and one's grip consists in one's ability to reach some conclusion about that issue in reflection. In particular, when it comes to the justification of conceptions that constitute how we understand something (e.g., the position on a chessboard), one can reach, in reflection, some conclusion – however tentative and qualified – about whether one is right to have the conceptions in question or not; whether they are justified by their own lights. One has a *rational basis* for holding the conceptions one does when, on balance, the reasons favoring one's conceptions outweigh the reasons against them.[13] The requirements of having a reflective grip on what one understands and having a rational basis for it are logically independent.[14]

On Carter and Prichard's view, we count as understanding something only if the conceptions that constitute our understanding can sustain reflective scrutiny *because* they have a rational basis available to reflection. This adds the reflective-grip to the rational-basis requirement.[15]

What is the relation between responsibilist virtue epistemology and a reflective appraisal requirement for understanding? For Carter and Pritchard's challenge to hit its target,[16] it has to be the case that reflective appraisal, as a requirement on cognitive responsibility, somehow flows from the fact that one comes to understand by manifesting intellectual virtues or skills. For all "neo-Aristotelian" or responsibilist virtue epistemologies, writes Pritchard:

> epistemic virtues are reliable cognitive traits which also demand a certain
> level of reflective responsibility on the part of the agent. (Pritchard 2005:
> 194-195)

"Reflective responsibility" meshes well with Zagzebski's view, on which "it is impossible to understand without understanding that one understands" (Zagzebski 2001: 246).[17] And the reflective dimension of understanding matters for cognitive responsibility: "One of the central features of agency is self-reflectiveness" (Zagzebski 2001b: 152). So Carter and Pritchard's argument seems to *succeed* against Zagzebski, for she accepts (3). Yet, as Pritchard (2005: 195) notes, Zagzebski's is only "a version" of virtue epistemology. In Section VI, I will argue that a responsibilist virtue epistemology of understanding is coherent even in the absence of a reflective appraisal requirement for cognitive responsibility. In Sections VI-VII, I will argue that experts *don't* always have a reflective grip on what they understand. Hence neither having a reflective grip on what is understood nor being able to reflectively appraise what is understood can be required for

understanding – even when one's cognitive behavior is responsible. It follows that premise (3) in Carter and Pritchard's argument is false.

VI. ACTS OF VIRTUE AND SKILL

In this section, I sketch a notion of cognitive responsibility for the 'irreflective' manifestation of our intellectual virtues and skills, *contra* (3).

Zagzebski's virtue epistemology of understanding presupposes that one is cognitively responsible for how one understands something (e.g., the position on a chessboard). Both I and Carter and Pritchard share that assumption. That raises the question of what cognitive responsibility for understanding is.[18] Watson writes:

> Because many of these appraisals concern the agent's excellences and faults – or virtues and vices – as manifested in thought and action, I shall say that such judgments are made from the *aretaic perspective*. (Watson 1996: 231)

The sage is responsible for their deed *because* they have the virtues and skills that deed manifests. Those virtues and skills[19] are part of *who one is*, and their development is part of one's learning history. Since character is usually seen as comprised of virtues, I'll speak of expertise in a domain as a catch-all for one's skills in that domain. Suppose you understand the moral horrendousness of a situation, or that the eradication of poverty isn't

feasible. To be *responsible* for how you understand that is to be responsible for the virtues

and skills manifested in the formation of the conception by which you understand what

you do. To be responsible for those virtues and skills is for them to be a part of your

character and expertise, explaining why you act by them. In chess, Tal's sacrifices are his

partly because he is inventive; Botvinnik's positional game is his partly because he is

prudent. And how each understands the position they see – even when playing each other –

is informed by what possibilities that position affords: of sacrifice, for Tal, or of stepwise

advance, for Botvinnik.

　　　　Among other things, this means that responsibility for manifesting virtues doesn't

presuppose reflective appraisal. As Hookway writes:

　　　　Virtues, we have said, enable us to respond to reasons: they provide a

　　　　sensitivity to rational requirements in particular cases, and they are usually

　　　　motivating. Although they can guide behavior through regulating

　　　　deliberation and inquiry, their operation is not transparent to consciousness

　　　　or open to reflective self-control. (Hookway 2000: 154)

While "reflective self-control" may exist, Hookway suggests it needn't[20] be operative

when we manifest virtues or skills. If Hookway is right, to tie acts of virtue to reflective

appraisal would put the cart before the horse, as virtues themselves may be "regulating

deliberation." Rather, what seems to be at play is a virtue of reflectiveness (Goldie 2008),

which prompts reflection *only* when the situation calls for it.[21] Suppose we are nonetheless

15

able to ask ourselves in reflection if the acts committed were appropriate even when the *occurrence* of reflection isn't called for. In reply, we may point out that the ability to reflect plays no obvious role in explaining *why* the act committed was one of virtue if reflection isn't called on to act.

Notwithstanding, one might think that if our mental acts are *acts* properly speaking, then requirements of free will should apply to them too. In particular, the presumptive requirement that agents should *know* what they are doing might be construed as one's ability to represent to oneself what one is doing, were one to reflect on the matter. If cognitive responsibility for a thought entails that thought being thought freely, perhaps the thinker should then be able to recognize themselves in reflection as thinkers. In reply, notice that it is doubtful – or at least we seem to lack a compelling reason for thinking – that *aretaic* cognitive responsibility is best construed as involving an appeal to reflection. Thinking carefully, paying attention to how one understands something, knowing when to trust one's abilities and when to hesitate – these all illustrate thinking freely, yet none of them seems to presuppose the *necessary* exercise of discursive reflective capacities – even if and when reflection does occur.

Here, I don't develop a full concept of responsibility for 'irreflective' acts of virtue (in coming to understand included).[22] If such a concept failed to be coherent, that would be a larger problem for neo-Aristotelian virtue epistemologies in general, and not just for questions about Carter and Pritchard's skeptical challenge to understanding. Rather, the current point is only that we may coherently deny (3), which requires reflective appraisal in order for one to qualify as being responsible for the formation of the conception by

16

which one understands.[23] In the next section, I'll describe a set of cases best construed as counterexamples to (3) if the coherence of 'irreflective' cognitive responsibility is granted.

VII. AN EXAMPLE

I deny (3) by providing examples in which one is cognitively responsible for understanding but doesn't reflectively appraise what one understands. The examples are drawn from some of the chess games played by Mikhail Tal, the Estonian 1960 world champion. One doesn't always reflectively appraise what one understands because reflective appraisal presupposes having a reflective grip on the rational basis for the conceptions by which one understands. Expert cognitive behavior (e.g., understanding chess positions) presupposes having a rational basis for conceptions one holds as an expert. However, expertise doesn't also require having a reflective grip over what one understands. So (3) is false; the chess examples will illustrate this.

Tal was one of the most tactical thinkers in 20[th] century chess. He was world champion for a year, 1960-61, until his predecessor Botvinnik won the rematch. I'm especially interested in Tal's early successes, 1957-1960, chronicled by Liepnieks (1961). Tal understood the positions he played, and was cognitively responsible in so playing. His inventiveness, surprising sacrifices, and eye for middle-game complications were well-known. Many of his spectacular moves withstand criticism. These facts indicate that he had a rational basis for the conceptions that constituted his understanding of the positions on the board. Reporting on his live performance, Tal says:

17

Calculation is only one side of it. In chess no less important is intuition, inspiration, or, if you prefer, mood. I for example cannot always explain why in one position this move is good, and in another bad. In my games I have sometimes found a combination intuitively simply feeling that it must be there. Yet I was not able to translate my thought processes into normal human language.[24]

The games where Tal couldn't put into words the intuitive way he understood the positions on the board are evidence for thinking Tal didn't have a reflective grip on those positions. (Tal seems to be suggesting that any explanations he might provide would often count as retrospective rationalizations of what he there and then only *saw*.) Given the discussion in Section VI, such cases may plausibly be construed as counterexamples to premise (3) of Carter and Pritchard's skeptical argument.

We can use Tal's avowals to set up a telling instance of Carter and Pritchard's challenge. Consider *the availability heuristic* (Tversky and Kahneman 1973),[25] by which what comes to mind first is taken to be what is most relevant in context – its springing to mind first being an extra reason for thinking it is most relevant in context. Novices who don't stop and think before moving use this heuristic, with amusing consequences. But experts often use the heuristic as well. They simply *see* the position for what it is, and move accordingly. They make what looks like a right move given their familiarity with positions of that kind.

Given the ubiquity of applying the availability heuristic (with varying results), it is fair to say that although most of Tal's moves were due to his chess expertise (including appropriate uses of the availability heuristic), they could have easily been due to choosing a move because it happened to occur to him first, manifesting not his chess skill but, say, an unscrutinized overconfidence in his possession of that skill; a bias. After all, Tal was playing in the *same* way when his combinations and sacrifices were successful and when they were criticizable.[26] Through tongue-in-cheek remarks like "Some sacrifices are right; the rest are mine," Tal seems to acknowledge this.

We can apply Carter and Pritchard's challenge by saying that Tal couldn't have been responsible for understanding positions on the chessboard that could have easily had a *vicious* rather than virtuous epistemic etiology. Despite how spectacular his play was, a question should then linger as to whether Tal always genuinely understood those positions.

In reply to Carter and Pritchard's challenge, successful and unsuccessful applications of the availability heuristic jointly support the view that not only biased thinking, but expert cognitive behavior too, is often intuitive and opaque to reflection. Intuitive expert performance such as understanding the position on a chessboard may be cognitively responsible even when one cannot appraise what one understands in reflection.

VIII. CONCLUSION

On virtue-based approaches to understanding (e.g., Zagzebski 2001), genuinely understanding something is primarily due to the epistemic agent's skills and virtues. Against such views, Carter and Pritchard (2016) raised the prospects of a skeptical position driven by the ubiquity and importance of biases in our epistemic lives: conceptions constituting how we understand something (e.g., the position on a chessboard) should, if arrived at primarily by intellectual virtues or skills, be accessible to reflection. This is because one is cognitively responsible for manifesting intellectual virtues and skills. But if the conceptions constituting our understanding could have easily been due to biases instead of virtues and skills, then a crucial aspect of their formation is reflectively opaque, and so we wouldn't count as fully responsible for understanding things the way we do. Whenever it could have easily been primarily due to biases instead of virtues and skills, our seeming understanding would not be genuine. As the threat of biases looms large for any degree of expertise, skepticism about understanding would be rampant.

I replied to Carter and Pritchard's skeptical challenge by denying a premise in the argument I construed them as advancing. Epistemic agents may be cognitively responsible for understanding something even when their understanding (e.g., of chess positions) isn't appraised in reflection. This is typical, I have suggested, for experts: possibly acceding to reflection is not a requirement for responsible expert cognitive performance.[27]

REFERENCES

Alfano, Mark. 2012. "Expanding the Situationist Challenge to Responsibilist Virtue Epistemology." *Philosophical Quarterly* 62: 223-249.

Antony, Louise M. 2016. "Bias: Friend or Foe? Reflections on Saulish Skepticism." In *Implicit Bias and Philosophy*, Vol.1 *Metaphysics and Epistemology,* eds. Michael Brownstein and Jennifer Saul. Oxford: Oxford University Press, 157-188.

Baehr, Jason. 2006. "Character, reliability and virtue epistemology." *Philosophical Quarterly* 56: 193-212.

Blackburn, Simon. 2001. "Reason, Virtue, and Knowledge." In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility,* eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press, 15-29.

Cargile, James. 1970. "A Note on Iterated Knowings." *Analysis* 30: 151-155.

Carter, J. Adam and Duncan Pritchard. 2016. "Cognitive Bias, Scepticism and Understanding." In *Explaining Understanding,* eds. Stephen R. Grimm, Christoph Baumberger and Simone Ammon. London: Routledge, 272-292.

Chase, Valerie M., Ralph Hertwig and Gert Gigerenzer. 1998. "Visions of rationality." *Trends in Cognitive Sciences* 2: 206-214.

Cianciolo, Anna T. and Robert J. Sternberg. 2018. "Practical Intelligence and Tacit Knowledge: An Ecological View of Expertise." In *The Cambridge Handbook of Expertise and Expert Performance*, eds. K. Anders Ericsson, Robert R. Hoffman, Aaron Kozbelt and A. Mark Williams. Cambridge: Cambridge University Press, 770-792.

Dancy, Jonathan. 2006. "Reasons, Relevance and Salience: A Response to Hookway."

    Philosophical Studies 130: 71-79.

de Regt, Henk W. 2004. "Discussion note: Making sense of understanding." *Philosophy of*

    Science 71: 98-109.

Dreyfus, Hubert L. 2005. "Overcoming the Myth of the Mental: How Philosophers Can

    Profit from the Phenomenology of Everyday Expertise." *Proceedings and*

    Addresses of the American Philosophical Association 79: 47–65.

Fennell, John and Roland Baddeley. 2012. "Uncertainty Plus Prior Equals Rational Bias:

    An Intuitive Bayesian Probability Weighting Function." *Psychological Review* 119:

    878-887.

Franklin, Rosalind L. 1981. "Knowledge, belief and understanding." *Philosophical*

    Quarterly 31: 193-208.

Fridland, Ellen. 2014. "They've lost control: reflections on skill." *Synthese* 191: 2729–

    2750.

Gobet, Fernand and Neil Charness. 2018. "Expertise in Chess." In *The Cambridge*

    Handbook of Expertise and Expert Performance, eds. K. Anders Ericsson, Robert

    R. Hoffman, Aaron Kozbelt and A. Mark Williams. Cambridge: Cambridge

    University Press, 597-615.

Goldie, Peter. 2008. "Misleading emotions." In *Epistemology and Emotions,* eds. Georg

    Brun, Ulvi Dogluoglu and Dominique Kuenzle. Burlington: Ashgate, 149-165.

Greco, John. 1999. "Agent reliabilism." *Philosophical Perspectives* 13: 273-296.

———. 2007. "The nature of ability and the purpose of knowledge." *Philosophical Issues* 17: 57-69.

Hookway, Christopher. 2000. "Regulating Inquiry: Virtue, Doubt, and Sentiment." In *Knowledge, Belief, and Character: Readings in Virtue Epistemology,* ed. Guy Axtell. Lanham: Rowman & Littlefield, 149-162.

Korsgaard, Christine. 1996. "Reflective Endorsement." In *Sources of Normativity.* Cambridge: Cambridge University Press, 49-89.

Kahneman, Daniel. 2002. "Introduction. Answering an Easier Question. The Illusion of Understanding." In *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Levinson, Robert 1994. "Experience-Based Creativity." In *Artificial Intelligence and Creativity: An Interdisciplinary Approach,* ed. Terry Dartnall. Dordrecht: Springer, 161-179.

Liepnieks, Alexander. 1961. *The Chess Psychologist: World Champion Tal*. North American Latvian Chess Association. Ann Arbor: University of Michigan Press.

Lombrozo, Tania. 2016. "Explanatory Preferences Shape Learning and Inference." *Trends in Cognitive Sciences* 20: 748-759.

McDowell, John. 1979. "Virtue and Reason." *The Monist* 62: 331-350.

Morewedge, Carey K. and Daniel Kahneman. 2010. "Associative processes in intuitive judgment." *Trends in Cognitive Sciences* 14: 435-440.

Owens, David. 2000. "Rationality and Reflection." In *Reason Without Freedom: The Problem of Epistemic Normativity*. Oxford: Oxford University Press.

Pritchard, Duncan. 2005. "Cognitive Responsibility and the Epistemic Virtues." In

    Epistemic Luck. Oxford: Oxford University Press, 181-201.

———. 2014. "A Modal Account of Luck." *Metaphilosophy* 45: 594-619.

Saul, Jennifer. 2013. "Scepticism and Implicit Bias." *Disputatio* 37: 243-264.

Trout, Joseph D. 2002. "Scientific explanation and the sense of understanding."

    Philosophy of Science 69: 212-233.

Tversky, Amos and Daniel Kahneman. 1973. "Availability: A Heuristic for Judging

    Frequency and Probability." *Cognitive Psychology* 5: 207-232.

Watson, Gary. 1996. "Two Faces of Responsibility." *Philosophical Topics* 24: 227-248.

Wilson, Timothy D., David B. Centerbar and Nancy Brekke. 2002. "Mental contamination

    and the debiasing problem." In *Heuristics and biases: The psychology of intuitive*

    judgment, eds. Thomas Gilovich, Dale Griffin and Daniel Kahneman. New York:

    Cambridge University Press, 185-200.

Zagzebski, Linda. 1996. "Using moral theory in epistemology." In *Virtues of the Mind: An*

    Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge.

    Cambridge: Cambridge University Press, 1-28.

———. 2001. "Recovering Understanding." In *Knowledge, Truth, and Duty: Essays on*

    Epistemic Justification, Responsibility, and Virtue, ed. Matthias Steup. Oxford:

    Oxford University Press, 235-249.

———. 2001b. "Must Knowers be Agents?" In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility,* eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press, 142-157.

ENDNOTES

[1] The perceptual, practical, emotional, linguistic and intellectual uses of the word 'understanding' seem so different that I doubt they depict the same state of mind. Franklin 1981 argues that there is a cluster of uses, corresponding to what I call in the text 'intellectual understanding', that depict a state of mind with positive epistemic status, similar to but distinct from knowledge and belief.

[2] Zagzebski 2001: 243 has a richer conception of understanding: "Understanding involves seeing how the parts of that body of knowledge fit together, where the fitting together is not itself propositional in form." The resulting conception isn't clear, for three reasons. First, Zagzebski agrees that, although skill-based understanding isn't *propositional* on the face of it, that doesn't imply that one cannot propositionally specify the conceptions by which one understands; so one shouldn't be sanguine about non-propositionality. Second, it's unclear whether the requirement of *coherence* (or fitting-together) applies to "parts of that body of knowledge" or to our conscious experience of it. Third, the quote suggests that only parts of a "body of knowledge" may cohere to form an understanding. Yet Zagzebski 2001: 243-244 also writes that "understanding does not always build on a base of knowledge". It is far from clear that cognitive skills may

genuinely be exercised only in the formation – or sustenance – of accurate conceptions, as opposed to inaccurate, or perhaps even to accurate but misleading, ones.

3 Carter and Pritchard 2016: 273 follow Saul 2013 in characterizing any bias whose outcome is representational as cognitive. This effaces the distinction between biases that do and those that don't often trigger *affect.* The former include biases in categorizing black men as more dangerous, women as less competent, or the poor as lazier; Saul, Carter and Pritchard focus on these. In contrast, I discuss biases such as overconfidence, which, even if they may not readily evoke as strong an affective response, seem closer to the epistemology of understanding. I am grateful to a reviewer for questions here.

4 To remedy this, one may, for instance, follow Fennell and Baddeley 2012 in mixing prior probabilities of ignorance – e.g., 0.5 – and prior probabilities learned through experience. The mix would be weighted by how often situations encountered were familiar to the subject, needing priors of experience, and how often situations encountered were unfamiliar, needing priors of ignorance.

5 Trout 2002: 223 ff. argues that biases may make the beliefs underwriting our seeming to understand actually false. Trout identifies the hindsight and overconfidence biases as two relevant sources of error. However, I'm not quite sure what Trout's challenge is. Take hindsight, for instance. Doubtless, we often give way to thoughts like "I knew it all along." But it's quite unclear how thoughts of this kind can partly constitute an *understanding* of something, which is what is at issue. For a more substantive reply to Trout, cf. de Regt 2004. By framing their skeptical challenge as they do, Carter and Pritchard seem to be granting (at least for the purpose of argument) that most beliefs

constituting understanding are accurate, focusing on whether we are cognitively responsible for their formation. Similarly, we may raise the question of how understanding fares with respect to experimental studies showing either that we often exhibit "illusory knowledge," cf. Koriat 2000, or that we have lasting "explanatory preferences," cf. Lombrozo 2016, for, say, the simplicity of hypotheses by which we understand the world around us, irrespective of whether simpler hypotheses happen to be closer to the truth or not. In reply, it is worth noting that none of these experimentally evinced phenomena fully overlaps with understanding as construed by virtue epistemology, providing a principled reason why psychological studies and the virtue epistemology of understanding haven't made contact, viz., that they seem to target related yet distinct mental phenomena. I'm very grateful to an anonymous reviewer for making this point.

[6] Saul 2013 supports bias-driven skeptical conclusions about knowledge. Carter and Pritchard 2016: 288, however, think that virtue epistemology doesn't apply to knowledge. To give just one example of why not: We often gain knowledge by testimony. Even when listeners are discerning, knowing who to trust, knowledge by testimony isn't primarily due to their virtues. However, Carter and Pritchard think virtue epistemology applies to understanding. So they correspondingly consider skeptical theses based on how biases would undermine understanding. Carter and Pritchard also consider Alfano's 2012 situationist challenge to virtue epistemology as raising a similar epistemic challenge to that of Saul 2013. However, to the extent that biases are dispositional, they can be characterized across a range of different situations, so

situation specificity raises different concerns than biases. Alfano's challenge goes beyond the scope of this paper; Greco 2007 outlines a reply to the challenge.

7 It would be unrealistic to focus on those rare cases where experts *lack* bias entirely. However, it is reasonable to assume that many of the conceptions by which we understand aren't *primarily* due to bias. When could forming an understanding of something have *easily* been due to biases instead of intellectual virtues or cognitive skills? The possibility of bias is *live* if the probability of a bias-inclusive etiology of our conceptions *isn't* low. Yet the probability of a bias-inclusive etiology of the conceptions by which we understand couldn't be *too* high either. For otherwise their virtuous etiology would be unexplainable, sheer luck. How high the probability of a biased etiology of the conceptions by which we understand, then, has to be a context-sensitive matter: neither too low (so as to make it negligible) nor too high (so as to undermine the actual operation of our virtues and skills). I'm grateful to Paul Humphreys for questions about this. Also see Pritchard 2014.

8 What is it to have *ownership* over one's understanding? We are obviously bearers of our states of mind. The question of ownership has to turn on whether we're also cognitively responsible for them, as authors.

9 Carter and Pritchard don't anywhere claim that skepticism about understanding is *true*. Instead, they only wish to articulate a skeptical threat, and motivate it by appeal to the pervasiveness of biases. However, it is hard to see precisely what the skeptical challenge or threat might be unless it were a challenge that the *truth* of skepticism hasn't been ruled out. I'm grateful to a reviewer for emphasizing that Carter and

Pritchard's aim might be more modest than the argument I attribute to them might suggest.

[10] Pritchard 2005 explicitly contrasts credit due to an epistemic agent for conceptions formed by manifesting reliable competencies with the cognitive ownership characteristic of manifesting intellectual virtues we are responsible for. He writes that "the cognitive achievement is entirely at a sub-personal level, and in this sense the agent proper is not cognitively responsible for her reliably formed true beliefs at all... agent reliabilist accounts of knowledge might ensure that agents are able to take a very minimal form of cognitive responsibility for their beliefs," viz., credit (2005: 190). Contrast this with Carter and Pritchard's 2016: 288 view that "understanding requires a particular kind of cognitive ownership" that includes but goes beyond credit because it "essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities." Call this requirement "reflective appraisal;" I'll return to it in Section V.

[11] This is not to say that the distinction between purely reliable competencies and virtues we are responsible for isn't controversial; Baehr 2006 suggests the two may often overlap. While the mental items discussed may be the same, responsibilism and reliabilism still differ on what is explanatorily prior: the fact that we are responsible for the intellectual virtues in question; or the fact that we thus get at the truth. The equation "Truth is what the sage aims to believe" may be read differently according to whether one emphasizes the left-to-right or the right-to-left order of explanation; cf. Blackburn 2001.

As Wilson et al. 2002 note, conscious de-biasing in reflection is difficult but not impossible. It is, then, open to a virtue epistemologist to say that we might form some conceptions early on, under the influence of bias, and only *later* those conceptions may come to constitute a genuine understanding – later, if and when de-biasing and reflective appraisal become possible. One should note that reflective access and reflective effectiveness differ greatly; one may be capacious enough as to become aware of a bias one has, while unable to change the pattern of vicious thinking however much one tried. Alternative strategies like changing the environment or focusing on areas where the bias isn't operative may work better, cf. Antony 2016. I'm grateful to an anonymous reviewer for the point.

13 One could sharpen this requirement further in several ways, adding that the epistemic agent should be able to become aware *of* one or more reasons favoring what they believe, or that such reasons should be *salient* in the context of inquiry. A good menu of options of extra requirements is Owens 2000. However, believing what, on balance, one's reasons favor, suffices for having a rational basis.

14 On the one hand, what one's conceptions are rationally based on may be unavailable to the epistemic agent's reflection. A rational basis for what one believes may be ascertained by one's responsiveness to reasons, or rational sensibility. A rational sensibility which doesn't always involve reflection seems to underlie the very manifestation of virtues and skills. On the other hand, having a reflective grip on one's conceptions need not be rationally based. We may, in reflection, identify different reasons, or weigh them differently, than what genuinely supports or undermines our

conceptions, cf. Dancy 2006. And our reflection may also assess not whether our conceptions are based in reasons, but whether they promote our projects or are otherwise useful.

[15] Carter and Pritchard are concerned with understanding's having a rational basis "to exclude live-error possibilities," where a biased etiology counts as the relevant possibility envisaged. Having "a rational basis to exclude live-error possibilities" sounds ambiguous (to me, at least). For it is one thing to have reasons in favor of the conception by which I understand the position on a chessboard, reasons strong enough that they – in fact, and my awareness of them regardless – exclude even live-error possibilities. And it is another thing to have a second-order reason – if we were to use Dancy's 2006 terminology –, ascertained in reflection, to think that first-order reasons in favor of our conceptions outweigh or override first-order reasons against our conceptions. A virtuous etiology is a reason in favor of our conceptions; the non-negligible risk of bias interference is a reason against them. This richer construal of the reflective requirement Carter and Pritchard impose on cognitive reponsibility for understanding includes, and is perhaps stronger, than both the reflective grip and the rational-basis requirement. So, if what I call "reflective appraisal" in the text is too cognitively demanding to fit all cases of understanding, as Section VI-VII below argue, then the criticism of over-demandingness will apply to Carter and Pritchard's stated requirement as well.

[16] Parenthetically, notice that if reflective appraisal were *independent* from the requirement that your conception be *formed* by manifesting intellectual virtues and

skills, then Carter and Pritchard's objection would seem to beg the question against a responsibilist virtue epistemology. For a supposedly crucial epistemic aspect of what makes us responsible for understanding – the amenability of our conceptions to reflective scrutiny – would *not* be accounted for in virtue-theoretic terms.

[17] Taking Zagzebski at face value about this is unpromising. Compare Cargile's 1970 critical remarks on knowing that one knows, which could easily be adapted for understanding that one understands.

[18] One unsuitable view is doxastic voluntarism. Surely we don't come to understand things the way we do because *we wish it*. From a virtue-theoretical perspective, to think that cognitive responsibility for understanding has to amount to doxastic voluntarism is to conflate aretaic and act-relative perspectives in attributing responsibility. For how these perspectives differ, see Watson 1996 and Zagzebski 1996.

[19] Watson 1996: 244 argues against the generalization of this notion of responsibility from virtues to skills. To reply: virtues do differ from skills in important respects. But that falls short of showing that the notion of responsibility applicable to acts of virtue doesn't apply to acts of skill. If Hans writes a complex logical proof with ease and skill at calculation, we may properly say he's responsible for it: not for his scribbles on paper but for the proof – precisely because of his logical skill. The example doesn't settle the issue; but a rationale is needed for restricting aretaic evaluations from acts of skill. I should note that I haven't seen aretaic evaluations for responsibility explicitly applied to experts so far in the literature. However, the possibility of such evaluations is

guaranteed by the fact that, unlike virtue ethics, virtue epistemologies often treat virtues and skills jointly.

[20] Indeed, part of the reliability of one's acquired expertise presupposes it is tacit knowledge, operating seamlessly and for which the interference of reflection may be disruptive, cf. Cianciolo et al. 2018.

[21] One might construe "reflective appraisal" along the lines of Korsgaard's 1996 notion of "reflective endorsement." To be cognitively responsible for how one understands things presupposes a set of norms governing one's understanding (e.g., accuracy, coherence, etc.), and to be guided by such norms involves recognizing their normative force. Such a recognition, on views similar to Korsgaard's, presupposes acknowledging those norms to be *good* and *ours* were we to reflect on the matter. In doing so, we would "reflectively endorse" the thinking performed in light of, and guided by, those norms. If reflective appraisal as a requirement on understanding were construed as reflective endorsement à la Korsgaard, one might then argue that our very recognition of intellectual virtues or skills (whichever those may be) as epistemically *good* – i.e., as conducive to accuracy, warrant, coherence, etc. – presupposes recognizing the normative force of what is epistemically good, a recognition only achieved in reflective endorsement. The topic is large, but consider sketches of two replies. First, if successful, this line of reasoning would commit *all* virtue epistemology to some form of a reflective endorsement requirement, and that is quite implausible. Second, notice how far this defence of the requirement of reflective appraisal for understanding has strayed from Carter and Pritchard's initial view. Reflectively appraising reasons for how one

understands something, e.g. reasons that rule out likely sources of error in how one conceives of what one understands, *doesn't* entail any possible awareness of the epistemic norms governing one's cognitive activity. I am grateful to an anonymous reviewer here.

[22] Nor do I propose a positive account of control, especially as it applies to one's skilled cognitive activity. For principled difficulties facing such accounts of control, see Fridland 2014.

[23] One might think, *pro* both Zagzebski and Carter and Pritchard, that we can only understand something when we are able to *explain* it, and that explanations require the ability to reflectively appraise what is explained. However, just as in Section V we distinguished a reflective grip on what is understood from having a rational basis for holding one's conception, here we should distinguish the reasons one invokes in explaining what one understands from the verbal explanatory end-product. One may be sensitive to reasons in the way needed to satisfy the rational-basis requirement on understanding without thereby being required to be able to have a reflective grip on what one understands. Since what is at issue is virtue-based understanding, the point made here is essentially McDowell's 1979 point that one's virtuous sensibility to reasons in favor of acting as one does need not always (even possibly) be put into words. To say that verbal codifiability into rules for acting is *no constraint* on what gets to count as virtuous acting is less committal than Dreyfus' 2005 view that chess-masters and seasoned drivers *don't* follow any rules.

[24] Cited in Levinson 1994.