

Virtue Epistemology and Abilism on Knowledge*

John Turri
john.turri@gmail.com

Abstract: Virtue epistemologists define knowledge as true belief produced by intellectual virtue. In this paper, I review how this definition fails in three important ways. First, it fails as an account of the ordinary knowledge concept, because neither belief nor reliability is essential to knowledge ordinarily understood. Second, it fails as an account of the knowledge relation itself, insofar as that relation is operationalized in the scientific study of cognition. Third, it serves no prescriptive purpose identified up till now. An alternative theory, *abilism*, provides a superior account of knowledge as it is ordinarily and scientifically understood. According to abilism, knowledge is an accurate representation produced by cognitive ability.

Keywords: knowledge; ability; virtue; belief; reliability; philosophical method

A popular view in contemporary Anglophone epistemology is that knowledge is true belief produced by intellectual virtue. Philosophers accepting this are often called “virtue epistemologists.” Virtue epistemologists differ on how to characterize an intellectual virtue. Some theorists focus on refined intellectual character traits, such as open-mindedness or conscientiousness, which the agent cultivates over time through deliberate effort (Code 1987; Montmarquet 1993; Zagzebski 1996). On this approach, knowledge is defined as true belief produced by a refined intellectual

* This is the penultimate version (10 June 2016) of a paper to appear in Heather Battily (Ed.), *The Routledge Handbook of Virtue Epistemology* (Routledge). Please cite the final, published version if possible.

character trait. Other theorists include reliable cognitive faculties, such as excellent vision or memory, among the intellectual virtues (Sosa 1991; Greco 1993; Goldman 1993). On this approach, knowledge is defined as true belief produced by the agent's reliable character traits or reliable cognitive faculties. Here "reliable" means that the trait or faculty produces mostly true beliefs.

An important initial question is what virtue epistemologists are giving theories of. One possibility is that they are offering theories of the ordinary *knowledge concept*, a mental representation shared by members of a community. This aligns with the dominant methodology of contemporary Anglophone epistemology (see Turri 2016c for a review). For example, theorists aim for an account that "matches our everyday practice with the concept of knowledge" (Craig 1990: 3), describes "the concepts and practices of the folk" (Goldman 1993: 272), or reveals "the folk theory of knowledge" (Jackson 1998: 32). In line with that aim, theorists motivate and evaluate proposals "by reflecting on our thinking and practices" regarding knowledge (Greco 2010: 4). The basic idea behind this methodology is that our judgments and behavior provide evidence regarding the content of our concepts. For example, patterns in ordinary knowledge attributions provide evidence regarding the content of our knowledge concept.

Understood as an account of the ordinary knowledge concept, virtue epistemology faces at least two serious problems.

The first serious problem is that knowledge does not require belief, as those categories are ordinarily understood. Virtue epistemologists have offered no serious argument that knowledge requires belief. The field's most widely cited review article covers no argument that knowledge

requires belief (Greco & Turri 2011). One of the field's most influential monographs does not include "belief" in its index (Greco 2010). One influential theorist simply "affirm[s] that knowledge entails belief" without argument (Sosa 2007: 24), and another merely notes that "a rough philosophical consensus has developed" around the idea (Zagzebski 2009: 3). But it is easy to name philosophers who have rejected the idea (e.g. Plato 380 BCE; Radford 1966; Lewis 1996) and no attempt is made to estimate the proportion of philosophers who accept the idea. Aside from that, the content of ordinary concepts is not determined by consensus among professional philosophers and recent findings have shown that alleged philosophical consensus often badly mischaracterizes ordinary concepts, including the knowledge concept (for reviews, see Turri 2016d; Turri 2016c; Buckwalter & Turri in press).

The conceptual relationship between knowledge and belief is a case in point. Researchers have recently discovered that competent speakers are often willing to attribute knowledge without attributing belief (Myers-Schulz & Schwitzgebel 2013; Murray, Sytsma & Livengood 2012). In some cases, the majority of people who attribute knowledge also deny belief. For example, participants in one study read about Karen, a first-year university student taking an introductory science course (Murray, Sytsma & Livengood 2012). Karen's instructor and textbook teach that the earth revolves around the sun. This conflicts with what Karen's religious parents taught her growing up: they taught her that the sun revolves around the earth. On the final exam for the science course, Karen encounters the question, "True or false: the earth revolves around the sun." She answers "true." After reading the case, participants were asked whether Karen knows that the earth revolves around the sun, and whether Karen believes that the earth revolves around the

sun. Of those who attributed knowledge, the vast majority (85%) denied belief. Moreover, in another series of experiments, regression analysis and causal modeling suggested that even when people attributed both knowledge and belief, their knowledge attributions were not based on their belief attributions (Turri & Buckwalter in press; Turri under review; Turri, Buckwalter & Rose under review). Finally, reaction-time studies have shown that people attribute knowledge faster than they attribute belief (Phillips, Knobe & Cushman 2015). If knowledge attributions occur before belief attributions, then the former are not based on the latter.

What explains this conceptual disconnect between knowledge and belief? One hypothesis is that belief ordinarily understood is connected with *feeling* or *emotion* in a way that knowledge is not (Buckwalter, Rose & Turri 2015; Buckwalter & Turri 2016). As William James wrote, by its “inner nature,” belief “is a sort of feeling more allied to the emotions than to anything else” (James 1889: 21). According to Hume (1748/1993: 32), this is how belief is understood “in common life.” By contrast, knowledge might not entail anything about how someone feels. If knowledge does not entail anything about how someone feels but belief does, then knowledge cannot entail belief. Another hypothesis, consistent with the first, is that the belief concept is essentially connected to distinctively human conversational practices of explicitly justifying and excusing behavior, whereas the knowledge concept is essentially connected to the much older animal practice of predicting behavior (Turri under review). On this hypothesis, the two concepts play different roles in our social-cognitive economy and there is no reason to expect entailment.

The second serious problem is that knowledge does not require reliability. Theorists have provided some arguments that knowledge requires reliability. One main passage cited in favor of

the reliability requirement contains a brief explanatory argument (Goldman 1979). The passage begins by claiming that knowledge must be appropriately caused and then asks what kind of causes yield knowledge. Two lists are then produced. One list features processes that intuitively do produce knowledge: perception, introspection, memory, and “good reasoning.” The other list features processes that intuitively do not: “wishful thinking,” “mere hunch,” “guesswork,” and “confused reasoning.” It is noted that members of the first list all seem to be reliable, whereas members of the latter all seem to be unreliable. A possible explanation for the membership of these lists is that knowledge requires reliability. Surprisingly, this has been dubbed the “master argument” lending “powerful prima facie support” to the claim that knowledge requires reliability (Goldman 2012: 4). But it should not persuade us because alternative explanations are not considered. For example, the hypothesis that knowledge requires cognitive ability, reliable or not, can explain the membership of these lists. The list of processes that produce knowledge includes cognitive abilities to detect, discover, and retain truths; the list of processes that do not produce knowledge does not include abilities of detection, discovery, and retention.

A second argument for the reliability requirement is that it explains patterns in ordinary intuitions and judgments about knowledge. Theorists claim that reliabilism is attractive because it “explains a wide range of our intuitions regarding what does and does not count as knowledge” (Greco 2010: 6); it fits with our “ordinary, intuitive judgments” (Dretske 1981: 92ff); it “matches our everyday practice with the concept of knowledge as actually found” (Craig 1990: 3-4); it can explain “our inclinations” and “intuitions” about knowledge (Goldman 1993: 271). This *proto-reliabilist hypothesis* about folk epistemology generates testable predictions. On

one hand, people will not count unreliably produced belief as knowledge. For example, if an ability produces only ten percent or thirty percent true beliefs, then the beliefs it produces will not be judged knowledge. On another hand, vast and explicit differences in reliability will produce large differences in knowledge judgments.

A recent series of studies tested both predictions by eliciting knowledge judgments about simple cases while varying how reliably beliefs were formed (Turri 2016a). Both predictions were falsified. Rates of unreliable knowledge attribution reached 80-90%. People consistently attributed knowledge to reliable and unreliable believers at similar rates. People attributed knowledge despite actively classifying the knower as unreliable, and they overwhelmingly declined to attribute knowledge in closely matched controls. More generally, although people understood and processed explicit information about reliability, they did not seem to consult this information when making knowledge judgments. For example, in one study, participants were divided into two groups (Turri 2016a: Experiment 4). They all read a story about Alvin, who has been out running errands all day. In the morning, Alvin's wife told him that he should stop at the dry cleaners. On his way home, Alvin stops at the dry cleaners. Participants in the unreliable condition were also told that Alvin's memory is very "unreliable"; participants in the reliable condition were told that Alvin's memory is very "reliable." That was the only difference. Participants were then asked whether Alvin knows that he should stop at the dry cleaners. In the unreliable condition, participants categorized Alvin as unreliable and 86% of them attributed knowledge. In the reliable condition, participants categorized Alvin as reliable and 88% of them attributed knowledge. The same pattern occurred when using different dependent measures, differ-

ent narrative contexts, for male and female agents, and for different faculties.

A third and related argument for the reliability requirement is that it explains our intuition that knowledge is better than mere true belief, which has been a central concern in recent Anglophone epistemology (for a review, see Pritchard & Turri 2014). Consider two pleasing art-museum exhibits that, as it turns out, are intrinsically indistinguishable. One was produced by an artist's highly reliable ability to produce pleasing paintings; the other was produced by an artist accidentally spilling some painting supplies. Intuitively, the former is better than the latter. Similarly, consider two true beliefs in the very same proposition. One was produced by reliable eyesight; the other was produced by wishful thinking. Again, intuitively, the former is better than the latter. We seem to value a successful outcome more when it is due to a reliable ability than to luck. Why? One hypothesis is that an agent deserves more credit for an outcome produced by reliable ability than for an outcome produced by luck. But an alternative hypothesis is that an agent deserves more credit for an outcome produced by ability, reliable or not, than for an outcome produced by luck. Recent behavioral experiments support the alternative hypothesis specifically for cognition (Turri 2016a: Experiment 7). When given a choice to credit a correct answer to an agent's ability or to luck, people credited it to luck when it was based on a guess, but they credited it to ability when the agent exercised her ability. Importantly, they credited it to reliable and unreliable abilities alike. And their knowledge judgments were strongly correlated with credit attributions. Crediting the correct answer to ability rather than luck increased the odds of attributing knowledge by a factor of 38 (i.e. 3800%).

In light of current evidence, I conclude that virtue epistemology fails as a theory of the or-

dinary knowledge concept. Knowledge ordinarily understood is not true belief produced by intellectual virtue.

But perhaps capturing the ordinary knowledge concept is not what some virtue epistemologists aim for. Perhaps some virtue epistemologists are offering theories of *knowledge itself*, a relation between cognizers and facts or truths, whose essential features might be absent or misrepresented in the ordinary knowledge concept. If virtue epistemologists were doing this, then presumably they would motivate and evaluate theories in light of relevant findings from cognitive science on how organisms acquire, store, and use knowledge. This methodology is sometimes used in contemporary Anglophone epistemology (e.g. Machery 2009 on concepts; Carruthers 2011 on self-knowledge), but it is almost entirely absent from the mainstream virtue epistemology literature.

I am aware of one potential exception to this generalization (Kornblith 2002: chapter 2; see p. 56 n. 39 for an explicit connection to virtue-epistemological theories of knowledge). On this exceptional approach, knowledge is defined as reliably produced true belief. This definition is motivated by the role knowledge attributions play in the systematic scientific study of animal behavior. Animals have biological needs, including access to food, shelter, and mates. The satisfaction of these needs is promoted by the evolution of capacities for receiving, storing, and processing information about the environment. Cognitive psychologists classify such informational states as “knowledge” that generates behavior contributing to fitness. A key passage summarizes the basic argument for this approach to knowledge, illustrated with reference to piping plovers’ remarkable nest-protection behavior:

Notice that these explanations require more than just the category of true belief. If we are to explain why it is that plovers are able to protect their nests, we must appeal to a capacity to recognize features of the environment, and thus the true beliefs that particular plovers acquire will be the product of a stable capacity for the production of true beliefs. The resulting true beliefs are not merely accidentally true; they are produced by a cognitive capacity that is attuned to its environment. In a word, the beliefs are reliably produced. The concept of knowledge which is of interest here thus requires reliably produced true belief. (Kornblith 2002: 58).

I applaud the careful, valuable, and insightful engagement with relevant scientific literature embodied by this approach. However, its conclusion does not follow, for two reasons. First, if “belief” is being used in its ordinary sense, an informational state is not automatically a belief. A representational state that was not a belief could still provide plovers with the relevant information, so belief is not essential. But arguably this is a tolerable extension of “belief.” Second, and more importantly, *unreliable* cognitive capacities can play the relevant role. Reliability, in the sense of producing more true representations than false ones, is not needed. In particular, unreliable cognitive capacities can serve informational needs, guide behavior, and thereby promote fitness.

This can be illustrated with a different avian example: predator alarm calls among willow tits. In one study, over 80% of calls were false alarms. Despite viewing this “high proportion” as “remarkable,” the ornithologist concluded that evolution “selected for such a low threshold” because it increases survival in the long run (Haftorn 2000: 445). Despite its unreliability, the ca-

capacity for receiving predator alarm calls serves an important informational need. This fits a general pattern in evolutionary analysis: in order to be selected for, a trait need only promote a beneficial outcome often enough to enhance fitness. If the benefit of a single success is great enough, even a highly unreliable trait could be selected for.

There is a better account of the knowledge concept operative in the science of animal behavior: knowledge is an accurate representation produced by cognitive ability. The relevant ability could be reliable or unreliable. The representation might be a belief or it might take some other form. Interestingly, the experimental results discussed above suggest that this same account of knowledge also fits the ordinary knowledge concept very well. In particular, it would explain why people overwhelmingly attribute knowledge to agents who get the right answer through unreliable ability. It would also explain why information about reliability does not affect knowledge attributions whereas information about ability does. Thus I propose that knowledge ordinarily understood and scientifically understood amounts to the same basic thing: accurate representation produced by cognitive ability. I call this theory of knowledge *abilism* (for defenses, see Turri 2015a, Turri 2016a, Turri 2016b).

Relatedly, there are also positive theoretical arguments for the possibility of knowledge produced by unreliable ability (Turri 2015b). One such argument begins with the observation that achievements can be unreliably produced. For example, human toddlers around twelve months old are highly unreliable walkers, but many of their early steps are genuine achievements reflecting their blossoming bipedalism. More generally, achievements populate the road to proficiency in many spheres, despite our unreliability. This is true in art, athletics, politics, oratory, music,

science, and elsewhere. If achievement in all these spheres tolerates unreliability, then knowledge probably does too.

So far I have argued that virtue epistemology fails as a theory of knowledge ordinarily understood and scientifically understood. But sometimes virtue epistemologists seem to be aiming for something different from either of those things. In particular, sometimes they seem to be prescribing a technical concept to serve some other purpose. For example, some theorists “analyze” the knowledge concept, and other epistemological concepts, in order to improve education and intellectual culture (Roberts & Wood 2007: 27-9). Unfortunately they provide no evidence that their analyses can accomplish such laudable goals, or that philosophical analysis fares better than, say, research in the learning sciences to promote beneficial change (Turri 2011). Sensible prescriptions are accompanied by relevant cost-benefit analysis, which in the present case would presumably include results from social scientific and psychological studies.

Relatedly, some theorists aim for a “metaphysical analysis” improving on the ordinary knowledge concept (Sosa 2015a; Sosa 2015b). “Conceptual improvement” involves carving a “domain more closely at the joints” than is accomplished in “ordinary speech and thought” (Sosa 2015b: 33). Empirical science provides many examples of conceptual improvement. For instance, biologists improved our ordinary fish concept by removing cetaceans from its extension. Similarly virtue epistemologists might “find in the phenomena themselves differences” important enough to warrant revising the ordinary knowledge concept. The central phenomenon here is “human knowledge,” which is “a state that people host” (Sosa 2015b: 9). But in order to find differences in a phenomenon itself, a field must study the phenomenon itself or be closely informed

by research that does. Biologists improved our fish concept by studying marine organisms' anatomy and physiology, not by entertaining thought experiments about marine organisms. By contrast, virtue epistemologists do not study humans and their states; nor is virtue epistemology closely informed by relevant research in the cognitive, social, or life sciences. As currently practiced, virtue epistemology, including the variant aiming for metaphysical analysis, proceeds by collating judgments about thought experiments. This will not produce conceptual improvement of the sort advertised.

To take another example, some theorists propose an analysis of the knowledge concept that addresses “contemporary concerns” while remaining “faithful to the history” of epistemology, is “practically useful,” “permits a natural extension” to analyses of understanding and wisdom, and “links” epistemology with “the general study of value” (Zagzebski 1996: xvi, 262-4, 275, 336). But it is unclear whether this project is well motivated or even viable. Consider the challenge involved in defining a coherent concept that remains “faithful” to the history of epistemology. Some accounts of knowledge — or, rather, something translated as “knowledge” — require certainty; others do not. Some accounts require belief; others do not. Some accounts rule out perceptual knowledge of our immediate environment; others treat it as a paradigm case of knowledge. Focusing just on this last contrast, one reaction is to “favor a compromise position” that treats perceptual knowledge as degenerate or “low-grade” (Zagzebski 1996: 280). Another reaction is to favor a vague account that takes no stand on the matter, implying neither that perceptual knowledge occurs, nor that it does not occur (Zagzebski 1996: 282). A third reaction is to abandon the idea that there is a single concept at issue: some philosophers theorized about an al-

leged status that cannot be based on perception, whereas others theorized about a different status that can be. Little if anything important is gained by crafting a concept that enables us to treat all these philosophers, who lived in different epochs and spoke different languages, as engaged in a single conversation or resolvable dispute. And even if we could craft such a concept, there is no reason to expect that it would serve any practical use.

There are many possible purposes for prescribing a technical concept. A virtue-epistemological knowledge concept could turn out to be well suited for some purposes. At present, I am unaware of any such purpose that has been explicitly identified.

In sum, virtue epistemologists define knowledge as true belief produced by intellectual virtue. I have argued that this definition fails in three important ways. First, it fails as an account of the ordinary knowledge concept, because neither belief nor reliability is essential to knowledge ordinarily understood. Second, it fails as an account of the knowledge relation itself, at least insofar as that relation is operationalized in scientific study of animal behavior. Third, it serves no prescriptive purpose identified to this point. Thus, at this point, the evidence suggests that the virtue-epistemological definition is false and serves no other useful purpose. An alternative theory, abilism, provides a superior account of knowledge as it is ordinarily and scientifically understood. According to abilism, knowledge is an accurate representation produced by cognitive ability.

Acknowledgments — For feedback I thank Heather Battaly, Ashely Keefner, David Rose, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research

Council of Canada, the Ontario Ministry of Economic Development and Innovation, and the Canada Research Chairs program.

References

- Buckwalter, W., & Turri, J. (2016). In the thick of moral motivation. *Review of Philosophy and Psychology*. <http://doi.org/10.1007/s13164-016-0306-3>
- Buckwalter, W., & Turri, J. (in press). Modest scientism in philosophy. In J. R. Ridder, R. Peels, & R. van Woudenberg (Eds.), *Scientism: prospects and problems*. Oxford University Press.
- Buckwalter, W., Rose, D., & Turri, J. (2015). Belief through thick and thin. *Nous*, 49(4), 748–775. <http://doi.org/10.1111/nous.12048>
- Carruthers, P. (2011). *The opacity of mind: an integrative theory of self-knowledge*. Oxford: Oxford University Press.
- Code, L. (1987). *Epistemic responsibility*. Hanover, NH: University Press of New England.
- Craig, E. (1990). *Knowledge and the state of nature: an essay in conceptual synthesis*. Oxford: Oxford University Press.
- Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, Mass.: MIT Press.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel.
- Goldman, A. I. (1993). Epistemic folkways and scientific epistemology. *Philosophical Issues*, 3,

271–285.

- Goldman, A. I. (2012). *Reliabilism and contemporary epistemology*. Oxford: Oxford University Press.
- Greco, J. (1993). Virtues and vices of virtue epistemology. *Canadian Journal of Philosophy*, 23(3), 413–432.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Greco, J., & Turri, J. (2011). Virtue epistemology. *Stanford Encyclopedia of Philosophy*.
- Haftorn, S. (2000). Contexts and possible functions of alarm calling in the willow tit, *parus montanus*; the principle of “better safe than sorry.” *Behaviour*, 137, 437–449.
- Hume, D. (1748/1993). *An enquiry concerning human understanding*. (E. Steinberg, Ed.) (2nd ed.). Indianapolis: Hackett.
- Jackson, F. (1998). *From metaphysics to ethics*. Oxford: Oxford University Press.
- James, W. (1889). The psychology of belief. *Mind*, 14(55), 321–352.
- Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Machery, E. (2009). *Doing without concepts*. Oxford: Oxford University Press.
- Montmarquet, J. A. (1993). *Epistemic virtue and doxastic responsibility*. Lanham, MD: Rowman & Littlefield.
- Murray, D., Sytsma, J., & Livengood, J. (2013). God knows (but does God believe?). *Philosophical Studies*, 166(1), 83–107. <http://doi.org/10.1007/s11098-012-0022-5>

- Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Nous*, 47(2), 371–384.
- Phillips, J., Knobe, J., & Cushman, F. (2015). Knowledge before belief: response-times indicate evaluations of knowledge prior to belief. Presented at the 41st annual meeting of the Society for Philosophy and Psychology, Durham, NC.
- Plato. (380 BCE/1997). *The Republic*. In J. M. Cooper (Ed.), G. M. A. Grube & C. D. C. Reeve (Trans.), *Plato: Complete Works*. Indianapolis: Hackett.
- Pritchard, D., & Turri, J. (2014). The value of knowledge. Retrieved January 4, 2014, from <http://plato.stanford.edu/entries/knowledge-value/>
- Radford, C. (1966). Knowledge—by examples. *Analysis*, 27(1), 1–11.
- Roberts, R. C., & Wood, W. J. (2007). *Intellectual virtues: An essay in regulative epistemology*. Oxford: Oxford University Press.
- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge, volume I*. Oxford University Press.
- Sosa, E. (2015a). On metaphysical analysis. *Journal of Philosophical Research*, 40, 309–314.
- Sosa, E. (2015b). *Judgment and agency*. Oxford University Press.
- Turri, J. (2011). Review of *Intellectual virtues: an essay in regulative epistemology*. *Philosophy and Phenomenological Research*, 82(3), 793–797.
- Turri, J. (2015a). From virtue epistemology to abilism: theoretical and empirical developments. In C. B. Miller, M. R. Furr, A. Knobel, & W. Fleeson (Eds.), *Character: new directions*

- from philosophy, psychology, and theology (pp. 315–330). Oxford University Press.
- Turri, J. (2015b). Unreliable knowledge. *Philosophy and Phenomenological Research*, 90(3), 529–545.
- Turri, J. (2016a). A new paradigm for epistemology: from reliabilism to abilism. *Ergo*, 3(8), 189–231.
- Turri, J. (2016b). Epistemic situationism and cognitive ability. In M. Alfano & A. Fairweather (Eds.), *Epistemic situationism*. Oxford: Oxford University Press.
- Turri, J. (2016c). How to do better: toward normalizing experimentation in epistemology. In J. Nado (Ed.), *Advances in experimental philosophy and philosophical methodology* (pp. 35–51). London: Bloomsbury Academic.
- Turri, J. (2016d). Knowledge judgments in “Gettier” cases. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy* (pp. 337–348). Wiley-Blackwell.
- Turri, J. (2016e). Vision, knowledge, and assertion. *Consciousness and Cognition*, 41(C), 41–49.
- Turri, J. (in press). Knowledge attributions and behavioral predictions in social cognition. *Cognitive Science*.
- Turri, J., & Buckwalter, W. (in press). Descartes’s schism, Locke’s reunion: completing the pragmatic turn in epistemology. *American Philosophical Quarterly*.
- Turri, J., Buckwalter, W., & Rose, D. (2016). Actionability judgments cause knowledge judgments. *Thought: A Journal of Philosophy*, 5(3), 212–222.
- Zagzebski, L. T. (1996). *Virtues of the mind: an inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.

Zagzebski, L. T. (2009). *On epistemology*. Belmont, CA: Wadsworth.