Assessing the Measure of Knowledge

Abstract

A fundamental aim of education is that the student come to 'know more'. Similarly, we assess students to make claims regarding their knowledge states over time, and provide comparisons between students. No doubt some such comparisons are of dubious educational and ethical nature, nonetheless there seem to be many contexts in which it would be entirely legitimate to wish to claim that one agent knows more than another. Consider, for example, medical expertise, a lookouts position in identifying approaching threats, or making decisions regarding which jobs to give to which candidates. Indeed, in education claims regarding the relative and absolute knowledge states of students are pervasive; but by what measure are they made? This paper discusses recent work by Treanor to address this epistemological issue, applying it to the educational context and providing an expanded perspective on Treanor's own proposal, motivated by a virtue-epistemic perspective on similarity. Such an approach respects 'truth' as an epistemic aim, while emphasising the importance of 'salience' in the measure of knowledge. Finally, and in agreement with Davis' work on the subject, I note the importance of philosophers of education considering such epistemological issues with regard to assessment.

Introduction

"Standards are slipping, it is easier now to get a pass than it used to be"; "Our students lack the knowledge and skills that students educated in other parts of the world routinely gain"; "exams do not adequately discriminate between students". By what measure are such claims made? What is the measure of knowledge here, which allows us to probe the differences in knowledge states across time, between countries, and between individual students?

"Assessment is one area where notions of truth, accuracy and fairness have a very practical purchase in everyday life." (Williams, 1998, p. 221). Perhaps one of the most widespread, and accepted, ways in which this manifests is with respect to the delimiting of 'knowledge' claims, and "The Measure of Knowledge" (Treanor, 2012) – the means through which we answer questions such as "who knows more?" Indeed, it is difficult to see how questions such as those above may be resolved, except with reference to some such 'measure of knowledge'. Moreover, such measure seems inseparable from various aims in formal education. This judgement or quantification of knowledge appears to play a role in assessment systems via at least three means:

1. Criterion based assessments stake the claim that a student has reached a certain level of knowledge, or that they have knowledge of particular know-that or know-how deemed important within a particular scheme of work. They prioritise assessing a particular level of knowledge, over discriminatory potential.
2. Rank based (including norm referenced) assessments stake the claim that a student Sarah knows more than a student Katherine in cases in which more creditable

knowledge claims are made by Sarah than Katherine, or Sarah scores higher on some measure (perhaps psychometric) than Katherine. They prioritise the ability to discriminate between the knowledge of individual students, and are sometimes criticised for 'shifting the goalposts' for what it means to obtain a particular grade (for example, between years or classes across which the norm may vary).

3. League tables implicate that teachers (or/and schools) have more (professional) 'knowledge' just in case the students within those schools achieve better grades than their peers at other (perhaps, particularly within 'value added' measures other similar) schools. Such systems are intended to allow for comparability across 'similar' schools.

Andrew Davis' monograph (Davis, 1998) and subsequent work (Davis, 2005a, 2005b; see for example, Davis & Williams, 2002) has given an excellent philosophical account of the difficulties with the high stakes testing regimes which are especially associated with the third issue above. Of course, it is also the case that one need not consider a 'measure of knowledge' to address the first and second of these issues. In both cases, a simple comparison between the agent and a pre-defined standard is satisfactory. That standard need not account for the breadth, depth, or quantity of knowledge. With regard to concerns around the suitability of this approach I refer readers to Davis' earlier work on the subject.

Nonetheless, it seems spurious to reject completely the claim that one might make comparison between individual agent's knowledge states either with regard to each other, or a pre-defined criterion schema. What then, do we engage in when we make claims such as "I know more than you", "she knows more about history than anyone", "that guy's got more train-knowledge in his little finger than you've got knowledge, period". Of course, hyperbole, exaggeration, and positioning play a role in such claims. But it is should also be clear that in many cases there is some legitimate – quantity related, if not quantifiable – statement regarding knowledge states being made in such claims. Indeed, while the examples given at the start of the paper may be ethically and epistemologically questionable, there do seem to be cases in which knowing whether one individual or group knows more than another is a legitimate epistemic aim; developing an account of the social-epistemic function of such comparisons.

We take it as given that *"it is uncontroversial, pre-philosophically, that education aims at the imparting of knowledge: students are educated in part so that they may come to know things"* (Siegel, 1998, p. 20). However, the nature of the assessment of knowledge is of interest, and presumably one means through which we understand that learners have developed their understanding – they have come to know – is through assessments of that knowledge. In this I draw on Nick Treanor's 2012 paper in which he discusses just this issue – 'The Measure of Knowledge' – dismissing two possibilities (cardinality and counterfactual distance) before proposing a third (similarity) which he advocates. This paper will first provide a brief overview of Treanor's claim, highlighting the importance of this philosophically rich analysis of an account of "knowing more" to the educational endeavour. I will then go on to suggest a modification to Treanor's approach, based on a comparison of knowledge states with an archetypal virtuous agent. I explain the advantages of this approach – specifically, that it avoids the concerns raised by Treanor, while providing a suitable account in both cases of similarly virtuous agents with differing knowledge states, and differing virtuous agents with similar knowledge states. Finally, I reiterate Davis' call for philosophers, and educators (and especially philosophers of education) to explore epistemological underpinnings of assessment and education systems.

Treanor and the Measure of Knowledge

The Measure of Knowledge: Cardinality

One intuitive 'measure of knowledge' might involve cardinality of knowledge states – that is, the counting of truths held by any given agent. In such a model, knowing whether Sarah knows more than Katherine simply involves an assessment of the number of truths each holds; assessing their knowledge states, then, is simply a procedural matter of counting tokens of knowledge held by each. However, as Treanor outlines, cardinality cannot be the way we compare agent's knowledge states. The reasons for the inadequacy of cardinality as a measure of knowledge can be summarised briefly by the following three concerns:

1. If we take it that any agent 'knows' infinitely many truths then the cardinality of any agent's knowledge (i.e. the numerable truth states) will be equal to any others; infinite. Similarly, if we take it that there are an infinite number of possible truths (i.e. things I do not know, but are true) then I am just as ignorant now as I always have been. Clearly I know more now than I did when I was 10, yet simple appeals to the counting of my knowledge states fails.
2. Cardinality fails to address the holistic nature of knowledge. Any atomic claim is associated with (perhaps entails) many other truths and decomposing this relation may be impossible (the holistic argument); it may not make sense to claim that we know 'some number' of truths. (See also, Davis' work on this subject).
3. Cardinality fails to address the 'natural language' representation of knowledge. Specifically, decomposing natural language claims into their atomic truth veritic (i.e., verifiable) parts may be impossible. Moreover, clearly there are issues of type in comparison of knowledge states: my knowledge of 1000 trivial truths (for example, hair 1 is in position x in relation to hair 689, etc.) is not the same as your knowledge of a smaller number of significant ones.

Treanor thus, rightly in my view, rejects cardinality as an adequate 'measure of knowledge'. As I note above though, this is a measure often used in educational contexts, particularly with regard to standardised assessments in which tokens of knowledge are added up to provide a score to any particular student.

The Measure of Knowledge: Counterfactual distance

Another possibility relates to the notion of counterfactual distance (see, Garson, 2014). This area of philosophy views states of the actual world as representing one set of states among the complete set of possible states. Conditions which are true across all possible worlds are 'necessary', while those which are true in only some possible worlds are 'possible'. Thus for any given state of the world, a number of possible belief states might occur: An agent's belief might reflect the state of the actual world (either because the belief is necessary, i.e. obtains across all worlds, or because the belief is possible and actually true in this world); an agent's belief might reflect some other possible world; an agent's belief might reflect some state which – necessarily – cannot obtain (i.e., an impossible state). Of particular interest here are the mid-kind, possible world beliefs which are counterfactual (i.e., not true in this world). In these cases, a finer grained analysis indicates that some counterfactuals are closer to the actual world than others. For example, if when asked "What is the weather?" you respond (reflecting your belief) "drizzly", when in fact it is raining rather hard, this belief can be said

to be counterfactually 'closer' to the actual world than had you responded (with a literal intent) "raining fire and brimstone", and indeed "sunny" would sit somewhere between.

Within the analysis of knowledge states, an approach might be to take one's knowledge to relate to the epistemic space in which one works – that is, as one comes to know more, one's space of possibilities around counterfactual worlds compatible with one's beliefs (one's epistemic space) becomes smaller. For example, as one grows up, some fundamental laws of physics are learnt, or (ignoring epistemic issues around normativity) that two wrongs don't make a right, and so on. One way of dealing with such a claim is to use the language of counterfactuals or possible worlds. In this case, we imagine two agents – Sarah and Katherine. The claim, then, is that if it is the case that the possible-world that is furthest away from the actual world while still being compatible with Sarah's epistemic-space is closer to the Actual world than Katherine's, then Sarah knows more than Katherine. That is, if Katherine's beliefs are compatible with views that are further away from actuality than Sarah's beliefs, then Sarah knows more than Katherine.

Such an approach is attractive for various reasons, including that it allows that some beliefs are more important than others because some beliefs narrow the set of possible worlds, bringing them closer to the actual world, to a greater extent than others. However two issues can be summarised regarding this proposal as follows:

1. Counterfactual distance as a measure cannot deal with knowledge about necessary truths, because such truths are true in all possible worlds and as such having or failing to have such beliefs would not make a difference to one's furthest away possible world[1]. This is a feature of possible worlds semantics rather than human psychology; that is to say, because possible worlds cannot represent impossible beliefs, it is not possible to represent an 'epistemic space' containing them, thus a) such beliefs are not represented by this model, and b) all such necessary truths are entailed in *any* 'epistemic space' because they are necessary (and thus obtain across all possible worlds).
2. Counterfactual distance as a measure leads to the consequence that once one knows that 'p', coming to know what 'p' entails has no bearing on the measure of one's knowledge (because once one knows 'p' one's furthest possible world has already been altered). Again, this is a feature of possible world semantics rather than psychology. The issue here is as in '1' above, with the obvious concern that while for some coming to know that 'p' will entail knowing 'q', for others this is not the case.

Moreover, we don't really know what such 'distance' would look like, and it is entirely possible that such distance would require a quantification of the kind discussed above. Clearly in educational contexts we make claims regarding the importance of some tokens of knowledge over others. To some extent such claims could be couched in the language of counterfactuals (we want students to have beliefs such that the distance between the actual world, and the possible world furthest from that world which is consistent with the student's beliefs is minimised), however, alternative approaches provide a better analysis of our epistemic-aims.

The Measure of Knowledge: Similarity

[1] That is, where impossible or inconsistent worlds are concerned we note that possible world semantics cannot distinguish between them because they cannot represent such belief states – see Berto (2013).

Another approach is to look at similarity. Here analysis turns on the similarity between the content of a representation, and what it is representing. Thus, a painting of a thing (Katherine) is in many ways different to the thing (Katherine herself), yet despite this there are similarities in form. This approach is stronger than that of possible worlds, because similarity is sub-maximal – representations do not need to be taken to represent complete possible worlds. Thus, some similarities are more important than others – a global difference in the skin tone of Katherine's painting may reduce similarity less than the replacement of Katherine's nose for a carrot in the image.

Now, at this point Treanor's claim is not that the problem has magically gone away but rather that 'similarity' is a term that is already widely accepted in everyday language and philosophy, and that this shifts the 'bubble' under the rug in an acceptable way. His claim is that "…it is a reductive move: the measure involved in knowing more is the measure involved in similarity." However, on this turn there is an obvious concern, 'similarity' measures are opaque, and at least to my mind some of them rely on counterfactual distance and/or cardinality. For example, counting the number of similarities, or assessing the distance of a representation from its target in counterfactual terms.

Moreover, one can imagine counterexamples of the following kind. Imagine again our two agents – Sarah, and Katherine. Both are theorising around some scientific notion, and using a set of false propositional beliefs to do so. To give a stark example, we might imagine that Sarah has conducted a number of experiments and concluded that "Water boils at 100 centigrade", while Katherine has read the same thing in a book. That is, in neither case are their knowledge states directly similar (or, mirroring) the actual world (in which the boiling point varies), and neither is any closer than the other – they are both some distance from the actual world (cf counterfactual similarity above). In such a case, in addition to concerns around counterfactual distance, we might well still be willing to say that one 'knows' more than the other. Yet, such a claim could not be placed at the door of 'similarity' of representation to represented, rather it might be to do with the kinds of activities one engages in, the ways they are framing the question, or some other set of peripheral knowledge. Indeed, the argument might be made here that the 'target-truth' is incorrect; that one *does* know more than the other but that the target knowledge relates to methods of investigation (or some such) – a call to cardinality. It seems, then, that similarity too is fatally flawed.

Similarity and the virtuous agent

As I note above, similarity between 'thing' and 'representation' is problematic, and one can easily imagine cases in which few truths are known, but one agent can be said to know more. Such claims appear to be an appeal to either reliablism (one's methods for acquiring truth reliably lead to believing truth as opposed to falsehood) or responsibilism (one's knowledge is contingent on one's epistemic-capabilities in acquiring that knowledge) (see, Greco & Turri, 2013). In such a perspective, one might imagine a measure of knowledge in which analysis focuses on how well one finds truths. That is, an analysis of the capabilities of the agent's behaviours for the detection of truths, and an assessment thereof.

However, we can imagine cases in which both Sarah and Katherine are equally virtuous with respect to their epistemic behaviour, yet Katherine is more epistemically unlucky than Sarah. We imagine, for example, cases in which both agents diligently go away and check their facts, but while Katherine uses a copy of a book with a random typographical error Sarah

does not. In such a case Sarah's greater knowledge is not due to some method for the acquisition of knowledge, rather, it is due simply to the luck of situation – and indeed many such disparities will be of this kind (see, Jr & Mylan, 2011; and Pritchard, 2007 for discussions of epistemic luck). Therefore, a simple analysis of the *processes* of knowledge acquisition is not enough; we need a way to distinguish between Sarah and Katherine that bears *some* relation to their propositional knowledge states.

In this section, I wish to propose that an approach which takes as the comparator "similarity to the knowledge state of a virtuous agent" may prove fruitful in offering a 'measure of knowledge'. Under this proposition, *knowledge states,* are the focus. However, analysis cannot be reduced to the methods described above around counting propositional-knowledge statements. This is because, in order for a proposition to be 'knowledge' (rather than belief, or statement) conditions must be met around the agent's creditworthiness for the claim. I do not here need to outline a 'manual' or step-by-step guide for such comparisons, rather an account should be given indicating that this is in fact what we do when we properly engage in 'measuring knowledge'. Briefly, though, we imagine that such a comparison might consist in:

1. Understanding the knowledge states relevant to any epistemic-situation, including the conditions for credit (the processes one must engage in to be creditworthy for any given knowledge-state) in that situation
2. A comparison of the knowledge-states of the target agent to those of an epistemically-virtuous agent, herein lies the similarity relation. This similarity-relation comparison consists in understanding the particular knowledge states – as inseparable from their credit-conditions – and the similarities and differences between agents. That is, such a comparison does not rely on a similarity to the 'real world' (as Treanor implies) nor some sort of quantification of epistemic-methods alone (in which the more knowledgeable agent would be the one who engaged in 'better' epistemic practices regardless of their propositional belief-state).

Such an approach allows, for example, geo-temporally situated (or contextualised) comparisons, making use of the most *salient* knowledge comparator (rather than just propositional claims). It is in this context that similarity between knowledge-practices, and (propositional) knowledge-displayed of different agents should be analysed. To give an example, if we imagine again our two agents – Sarah and Katherine – who exist in almost identical worlds (W1/W2). Sarah and Katherine have exactly the same belief states – they are, in that respect, epistemic twins across W1 and W2. As a result, the methods discussed above for describing their knowledge states (cardinality, etc.) will have little to say.

We imagine, then, one difference between W1 and W2: In W1 there exists some epistemic norm, a credit-condition – 'to phi' – where 'phi-ing' just describes carrying out whatever this particular practice is; in W2, no such norm exists. The practice in question might be something more likely to assure the truth of some particular claim 'x' but not actually related to it, perhaps an experimental control. Moreover, neither Sarah nor Katherine in fact engaging in 'phi-ing' around some particular knowledge-claim 'x'.

Applying the similarity notion described above, we thus consider the 'measure of knowledge' across our agents. At a simple level, in W1 neither Sarah or Katherine knows as much as an agent who phis (and is thus creditworthy for 'x'), in W1 (where no such norm exists), they know the same amount. A more interesting comparison is between Sarah and Katherine in different worlds – and indeed such comparisons are not entirely abstract, we make them when we compare knowledge states over historic agents, or between geographic locations. In such a comparison – insofar as comparison across W1/W2 makes sense – despite the identical knowledge states of S1 and S2 we should be prepared to say that a comparison of the agent to their respective comparators reveals a different similarity relation, despite the fact that S1 and S2 have the same propositional belief (both holding that 'x'). Thus, Sarah *knows* less than Katherine despite having identical belief states. This has the consequence that, for example, an agent who holds a belief 'x', without phi-ing in recent years (in which to phi is a norm) knows less than agents who did not phi in bygone years (where it was not the norm). Note that it need not be the case that Sarah does not *know how* to phi (although this would of course be an absence of knowledge), but rather that some environmental feature (an epistemic-norm) has an impact to our similarity relation in understanding the measure of knowledge.

We imagine, then, our agents Sarah and Katherine, are writing an essay on a scientific topic. In this case, our comparison is not the number of truths in each, or their counterfactual distance from the actual world. Nor is it 'similarity', where in this case the similarity would be how close to the thing being discussed (climate change theories, say) the essays achieve. Rather, the comparator would be the similarity to the knowledge state of a virtuous agent – the similarity of the true, creditworthy, beliefs stated to those of the virtuous agent, comprising understanding the salient truths and epistemic-practices engaged in for creditworthy knowledge claims. Again, this shifts the bubble, but this time in a more transparent way. We might ask questions, such as: what epistemic-environment the agent is in; how old they are; how much background knowledge do we expect of them; and what is the impact of all this on their behaviour and subsequent knowledge states? It is important to note that, the comparison is shifted from a similarity-relation between the represented and representing to a similarity-relation between representings. Alongside a shift from propositional-claims while avoiding a full shift to knowledge-practices, this approach avoids the concerns raised above.

Pre-empting complaints

I envisage three possible concerns with this approach. First, that it shifts the bubble to measurement of maximal virtue, second that it shifts the bubble to measurement of a given agent's virtue, and third that it eliminates, or at least reduces the importance of, the belief-world similarity-relation.

Measuring maximal virtue

The first concern, broadly stated, is that under the proposal given, the understanding of 'how much one knows' comes down to quantifying how much some virtuous agent knows and

comparing one's knowledge states to that quantification. That is, the concern is that by making the comparator a 'virtuous agent', somehow one must quantify the nature of virtue (how else, the complainant asks, can I decide whether agent 'a' or agent 'b' is *more* virtuous, and thus which knowledge state I ought to compare to). The weak version of this complaint is addressed above – that we are not interested in quantifying how virtuous an agent is in comparison to another, but rather, understanding their salient knowledge states. However, the strong version of the complaint – that in order to select the correct comparator, we ought to have some idea of what conditions might maximise the virtue of an agent, requires further consideration. It might, for example, be the case that multiple agents, with multiple comparator knowledge states, could all equally be imagined. However, while certainly this is a concern, the issue has been shifted to one of general epistemology and in particular notions of 'credit' – as Treanor notes in advocating similarity, at least this is a concept we have the tools to deal with.

Measuring an agent's virtue

A second concern is that such an approach requires us to have a measure of an agent's virtue, in order to assess either their similarity to the maximally virtuous agent, or to some other agent. However, again this is not the case. The focus of 'similarity' in this proposal is the salient knowledge states of the agent. Indeed, there is also no need to compare (or count) some set of virtuous or epistemic-methods between agents. Although it is the case that sometimes a 'less virtuous' agent (one with less reliable, or responsible methods for knowledge acquisition, for example) will 'know more' by virtue of their epistemic-environment, the ability to account for such epistemic luck through the consideration of the epistemic-environment in selecting appropriate comparators is a benefit of the proposal.

Removing the world-similarity relation

The final concern is broadly that by shifting 'the measure of knowledge' a step away from representations of the world, a crucial component of 'knowledge' (as reflection of actual states of affairs) is lost. That is, that the proposed approach does not represent 'knowledge' in the way Treanor intends in suggestion a comparison between representation and the thing being represented. Indeed, a specific benefit of Treanor's similarity approach is that it allows some sort of comparison between an agent's knowledge states, and some maximal set of knowledge – the similarity between what is known, and what there is. This approach has appeal, but it's hard to see how it avoids falling into various familiar traps, two of which I note. Firstly, it appears to be rather Platonistic in nature, relying on a notion of similarity of beliefs to the 'ideal form'. Secondly, it appears unlikely to account for normative knowledge or pragmatic features of knowledge, that is, that social features of the environment play a role in truth states and thus similarity – representing does not occur in a social vacuum. Moreover, the proposed approach is agnostic on how belief-states reflect states of the world – again, the bubble is shifting to one of general epistemology.

Conclusions: Implications for Education and Inquiry

Treanor ends his paper noting that his metaphysical discussion has implications for epistemology – what it is to know more, or be less ignorant, and that these are surely standard aims of epistemology. In particular he notes that while it is 'now commonplace' to hold that

the aim of inquiry is not 'truth' (because if it were truth alone, we would not discriminate with regard to the types of truths held counting threads would be equal in value to quantum mechanics), just because 'inquiry aims at truth' it does not follow that inquiry need treat all truths equally. Treanor thus notes that the conditional statement: 'if gold mining aimed at gold, it would aim at all gold (flakes, dust, nuggets, veins) equally', is patently false. However, its falsity is not an indicator that gold mining does not aim at gold; it is an indicator that it does!

The approach presented in this paper respects the value of 'truth', as key to knowledge, but places it in the context of social and virtue-epistemology, providing a comparator for our 'measure of knowledge'. In so doing, it emphasises that we are not simply interested in uncovering any truly held beliefs (our metaphorical flakes and dust), but in substantive, contextually-salient, and creditworthy knowledge (our metaphorical nuggets and veins). This stance has implications for educational systems. Perhaps most obviously, the rejection of cardinality and counterfactual distance as measures implies a falsity in perspectives focussed on the counting of atomic truths. Clearly we make decisions regarding the importance of knowledge taught. Similarity can obviously be assimilated into educational language, and indeed various forms of performance assessment may already attempt this through, for example, assessment criteria.

Thus, we should care about inquiry (and virtue, cognitive-ability, or epistemic-norms depending on the flavour of social-epistemology you favour), but the approach described in this paper gives us the means to describe a similarity relation which accords knowledge a status over epistemic-methods. Such an approach might motivate a shift in favour of assessments focussed on challenging epistemic-contexts (over recall of single knowledge tokens), while not advocating for 'contentless' 'skills based' assessments (a rather strawman proposal in any case). The approach recognises Treanor's analogy between 'inquiry for truth' and 'gold mining', while also leaving space for educators to aim at *inquirers* and equipping them to undertake inquiry. This is not so surprising – the aim of gold mining is the discovery of gold; a successful path to this end is through teams of skilled gold miners, more effective teams will hold a certain set of skills including (as discussed above) requisite knowledge that some gold (or knowledge) is worth more than others, and a disposition to gold seeking. The paper began with some common claims around assessment. The ways assessments instantiate epistemic assumptions, and the ways in which epistemology might inform our assessment policy are both areas for further work. This paper, though, has provided a similarity comparator, allowing analysis of knowledge states across time and place to be contextualised appropriately.

References

Berto, F. (2013). Impossible Worlds. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2013.). Retrieved from http://plato.stanford.edu/archives/win2013/entries/impossible-worlds/

Davis, A. (1998). Monograph on Educational Assessment. *Journal of Philosophy of Education, 32*(1), 1–152.

Davis, A. (2005a). Learning and the Social Nature of Mental Powers. *Educational Philosophy and Theory*, *37*(5), 635–647. doi:10.1111/j.1469-5812.2005.00148.x

Davis, A. (2005b). The Measurement of Learning. In R. Curren (Ed.), *A Companion to the Philosophy of Education* (pp. 272–284). Oxford, UK: Blackwell Publishing Ltd. Retrieved from http://www.blackwellreference.com.libezproxy.open.ac.uk/subscriber/uid=59/tocnode?id=g9781405140515_chunk_g978140514051522

Davis, A., & Williams, K. (2002). Epistemology and curriculum. In N. Blake, P. Smeyers, & R. Smith (Eds.), *The Blackwell guide to the philosophy of education*. Blackwell Reference Online. Retrieved from http://books.google.co.uk/books?hl=en&lr=&id=J_9WaYKxEygC&oi=fnd&pg=PA253&ots=DTsa4HxNxP&sig=DFNFobZy4E4227Up2K_CgqvE3DM

Garson, J. (2014). Modal Logic. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2014.). Retrieved from http://plato.stanford.edu/archives/sum2014/entries/logic-modal/

Greco, J., & Turri, J. (2013). Virtue Epistemology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2013.). Retrieved from http://plato.stanford.edu/archives/win2013/entries/epistemology-virtue/

Jr, E., & Mylan. (2011). Epistemic Luck. *Internet Encyclopedia of Philosophy*. Retrieved from http://www.iep.utm.edu/epi-luck/

Pritchard, D. (2007). *Epistemic Luck*. Clarendon Press.

Treanor, N. (2012). The Measure of Knowledge. *Noûs*, 1–28. doi:10.1111/j.1468-0068.2011.00854.x