Limiting Skepticism

Vincent F. Hendricks
Department of Philosophy and Science Studies
Roskilde University, Denmark
vincent@ruc.dk

John Symons
Department of Philosophy
University of Texas, El Paso, USA
jsymons@utep.edu

March 17, 2006

For an agent to have knowledge of some proposition he or she must be able to eliminate all relevant possibilities of error. Furthermore, since Plato it has been assumed that knowledge is robust insofar as it does not vanish in the light of new evidence or information. So, if one ascribes knowledge to an agent, one is proposing, as Jaakko Hintikka puts it,

> ...to disregard the possibility that further information would lead him to deny that p although he could perhaps imagine (logically possible) experiences which could do just that. [Hintikka 05]: 18

Notice that in the ascription of knowledge to an agent one is still recognizing that the agent could imagine the possibility of being wrong. Error is always conceivable, but conceivable error is not always relevant to knowledge. Knowing p involves the right to disregard irrelevant possible worlds in which it is not the case that p. Are the seeds of skepticism smuggled in via the putative knower's need to overlook allegedly irrelevant possibilities? As David Lewis notes:

> If you claim that S knows that P, and yet you grant that S cannot eliminate a certain possibility in which not-P, it certainly seems as if you have granted that S does not after all know that P. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory. [Lewis 96]: 419

Skeptics argue that the acquisition of knowledge is impossible given the standing possibility of error. Who would want to set such strict standards for knowledge acquisition? According to Lewis, either the skeptic or the epistemologist. In ordinary life, by contrast, Lewis claims that we know many things

with Moorean certainty. The fact that we know a lot, he writes, "is one of those things that we know better than we know the premises of any philosophical argument to the contrary" [Lewis 96]: 418. However, the Moorean strategy is not, nor was it intended to be, a direct response to the skeptical challenge. Directly confronting the problem would involve meeting the standards that the skeptic sets. Specifically the requirements of infallibility and certainty. Whereas Lewis challenged the reasonableness of these standards, and was content with emphasizing the "fact that we know a lot" (*ibid*). This investigation of the skeptical challenge begins by granting the skeptic the benefit of the doubt.

1 The Benefit of the Doubt

Perhaps nature is secretive, refusing to reveal itself to our senses or our scientific scrutiny. Even if nature does reveal itself, perhaps we are unable to grasp the meaning of the message. If the truth of an agent's knowledge claim depends on the 'underlying reality' or some other aspect of nature that transcends immediate experience then the truth of his or her knowledge claim is always *per definitionem* going to outstrip her power to certify the truth of her claim to know. A familiar cast of characters and scenarios relies on some version of this problem: Descartes' *malign genie*, Hume's hidden springs of nature, the Duhem-Quine thesis, Kuhn's incommensurability, Putnam's brains-in-vats, and Rorty's advocacy of edification over inquiry. Such skeptical scenarios depend on the assumption that any proposition is systematically underdetermined by any evidence. This is what is known as *global underdetermination*; two worlds ascribe contrary truth values to proposition p such that no evidence will favor the choice of one world over the other. This would leave the decision as to which world is actual underdetermined. If this is the case then global underdetermination would render rational inquiry futile.

Responding to the threat of global underdetermination, epistemologists have noted that agents will almost inevitably engage in some partitioning of worlds such that only relevant possibilities of error are dealt with in the process of knowledge acquisition. This nearly ubiquitous epistemic practice is known as *forcing* [Hendricks 06].

Skepticism plays on more than one string. In addition to global underdetermination the skeptic may also point to local underdetermination. Ever since the Pyrrhonian skepticism of Sextus Empiricus, the problem of induction and its various derivatives have presented a series of challenges to knowledge acquisition. Skepticism about induction is the result of the possibility of *local underdetermination* obtaining between evidence and the proposition. Kevin Kelly defines local underdetermination in the following way:

> A hypothesis is locally underdetermined by the evidence in a possible world if there is an infinite sequence of evidence possible for all the agent knows, such that each initial segment of this evidence sequence could arise independently of whether the hypothesis is true or false. [Kelly 96]: 24

The definition implies the lack of a determinate point in time after which the agent can reach a decision concerning the truth or falsity of the proposition in question. Skeptical arguments are designed to show that inquiry is in vain either from the outset – as in the case of global underdetermination or that a counterexample is to be found with certainty at some later stage, rendering further inquiry unnecessary. Global underdetermination suggests dropping inquiry apriori, while inductive skeptical worries based on local underdetermination purport to show with certainty that a counterexample will eventually appear.

2 Convergence

Knowledge, as characterized by infallibility and robustness exhibits convergence. The idea that scientific knowledge is convergent may be found in the works of American pragmatists like Charles Sanders Peirce and William James. Peirce held that scientific inquiry asymptocially converges to truth in the limit, and whatever the theories of science say in the limit exhausts truth. For the present purposes, it suffices to say, that convergence simply means that there is a time such that for each later time, the agent is not going to change his mind pertaining to the truth value of the proposition under consideration. Depending on what time that is, different notions of convergence will arise.
First, convergence with certainty means that there is a finite time after which the agent is not going to change his mind about the truth value of the proposition and that he clearly signals his success by going into his designated state of *halting*:

> *Agent S converges to proposition p with certainty iff there is a time n such that*

1. *S signals at n that he is ready to conjecture,*
2. *S conjectures p at $n+1$,*
3. *S does not signal earlier than n that it is ready to conjecture.*[1]

Convergence with certainty is generally viewed as the hallmark of convergence in epistemology. For instance, in response to Hume's problem of induction, hypothetico-deductivism is committed to formulating universal propositions and waiting around for incoming evidence to refute them. When a counterexample is encountered the proposition in question could not possibly be true – output 0, and halt! An existential proposition has a similar property but instead of being

[1] Note that immediately prior to the certainty conjecture the agent is required to produce a *signal* (say, Eureka!) of certainty. This is due to the fact that the agent (or the method he applies) may produce the sign of certainty more than once. Therefore, the certainty conjecture is taken to be the one following immediately after the first occurrence of Eureka! Subsequent signals of certainty will be 'ignored, as though the method has finished its job and is merely producing irrelevant noise thereafter'. [Kelly 96]: 48.

refutable it is verifiable with certainty – conjecture the existential hypothesis and wait for the first corroborating instance in the observed evidence. Eureka! The hypothesis is verified with certainty, so stop inquiry and output the truth.

As attractive as certainty convergence may be it is not always possible to obtain this kind of security. Real epistemological problems are not always amenable to convergence with certainty. In these cases one may choose to drop the halting condition but not the requirement of convergence. Limiting convergence allows the agent to oscillate pertaining to his conjecture some finite number of times. This number need not be specifiable in advance. At some point nevertheless the agent must reach a convergence modulus and stabilize his conjecture even if he does not know when stabilization has occurred. Thus, limiting convergence does not require the agent to report convergence. Peirce considered a similar idea insofar as he took it to be impossible to say anything about the direction of science in the short run while arguing that science may all the same asymptotically approach the truth in the long run. Similarly for James who recognized that knowledge of universal laws may become impossible to acquire if one is obliged to say when science has gotten it right [James 60]. Limiting convergence may be defined in the following way:

> *Agent S converges to proposition p in the limit iff*
> *there is a time n such that for each later time $n\prime$: S conjectures p at $n\prime$.*

Now, why entertain a notion of convergence but not of certainty of when convergence has occurred? As Phillip Kitcher asks:

> To be sure, there are [Bayesian] convergence theorems about the long run—but as writers from Keynes on have pointedly remarked, we want to achieve correct beliefs in the span of human lifetimes. [Kitcher 93]: 293

Kitcher's objection misses the mark. One can allow that we humans may not achieve many true beliefs in the 'span of human lifetimes' without thereby falling prey to skepticism. If, for instance our species goes extinct next Wednesday without achieving many true beliefs, the skeptic cannot claim victory. Skepticism, after all, is a judgment concerning the possibility of knowledge. Its success as an epistemological thesis should be independent of the date and time of our demise. The contingent fact of when our species ends its run has no bearing on the claim that knowledge will always be undermined by the possibility of error.

By contrast, treating inquiry as a matter of convergence in the limit is fully consonant with the mission of epistemology and science. For if we have to go to the limit to get the truth, then why not wait around for it, even if only in our philosophical imaginations? Therefore, *pace* Kitcher, reasoning about epistemic conditions in the limit is not vitiated by the possibility that our planet could be struck by an asteroid next Wednesday. Consider the birds in the trees. Note that 'There exists a black raven' is verifiable with certainty while 'All ravens are

black' is refutable with certainty. 'All ravens are black' is also verifiable in limit, because if the agent has not encountered the crucial example of a non-black raven leading him to change his mind, and he is not going to change his mind anymore, then the entire problem of whether all ravens are black becomes a trivial decision problem in the limit. Reasoning about the condition of knowers in the limit might sound like a cheap shortcut to the solution of epistemological problems. However, limiting convergence is a characteristic of any scientific practice in which our claims are subject to revision. As computational epistemologists like Martin and Osherson explain:

> The general point is that Ψ is not required to recognize or signal in any way that its conjectures have begun to converge. In this respect our paradigm is faithful to the situation of real scientists, whose theories remain open to revision by new, unexpected data. It is, of course, possible to define paradigms that require scientists to signal convergence. The prospects for success, however, are then diminished. [Martin & Osherson 98]: 12

For instance computational epistemologist Oliver Schulte proves how the identification of conservation principles for particle reactions is a limiting tractable problem and not one tractable with certainty [Schulte 00].

Based on certainty and limiting convergence one may formulate the following two notions of convergent knowledge:

Agent S may know proposition p with certainty iff

1. *p is true*
2. *S converges to p with certainty*
3. *in all possible worlds in accordance with one's choice of forcing clause.*

Agent S may know proposition p in the limit iff

1. *p is true*
2. *S converges to p in the limit*
3. *in all possible worlds in accordance with one's choice of forcing clause.*

Nozick's celebrated definition of counterfactual knowledge is a case in point for certainty convergence [Nozick 83]. The inherent decision procedure given by avoiding error and gaining truth together with the counterfactual semantics require the agent to converge in all possible worlds sufficiently close to the actual world in order to acquire knowledge of the proposition of interest. Thus, in terms of forcing, a subjunctive conditional is true just in case the consequent is forced among the closest worlds to the actual world in which the antecedent

holds. The idea of introducing the proximity relation is that when the agent's local epistemic circumstances suffice for the truth of the consequent, inquiry may as well just halt. Hence, Nozick's proposal uses decision with certainty as the convergence criterion. Consider now the following subjunctive:

1. *If the proposition "All ravens are black" were false, agent S would not believe "All ravens are black" 'now'.*

It seems that (1) would not be true unless

2. *If the proposition "All ravens are black" were false, S would have observed something different than he has up until 'now'.*

The problem of induction teaches that up until now, the evidence may be all the same, hence no answer with certainty seems to be forthcoming pertaining to this epistemic problem on the counterfactual account. Though Nozick has a forcing strategy to dismiss brains-in-vats and Cartesian demons as genuine possibilities of error given the proximity relation he has no immediate strategy for the problem of induction if the answer is to be had with certainty settled by local circumstances.

If an epistemic problem is solvable with certainty it is also solvable in the limit, but if a problem is solvable in the limit it is not necessarily also solvable with certainty. Consider two methods Q and R. Let Q be Popperian in nature in the sense that if the first observed raven is black, Q will conjecture that all ravens are black and will continue to project that all ravens are black unless a non-black raven is encountered. Method R by contrast is infallible in nature insofar as it does not conjecture anything which is not entailed by the evidence. Thus if a non-black raven is encountered R concludes that not all ravens are black but refuses to produce a conjecture otherwise.

Suppose the actual world is such that not all ravens are black. Then both methods will refute, with certainty, the proposition that all ravens are black. Suppose, on the other hand, that all ravens in the actual world are black, then Q will conjecture that all ravens are black after the first raven has been observed and will never alter its conjecture. Q will converge to the truth of the proposition in the limit. R, by contrast, will fail to generate the conjecture that all ravens are black, due to the requirement of infallibility. R will therefore fail to converge to the truth. The Popperian method has the virtue of converging to the truth in the limit, no matter what the truth might be, whereas the infallible method does not reliably converge to the truth in all cases.

Local underdetermination crippled Nozick's proposal because a particularly demanding criterion of success for inductive inferences was imposed, namely *decision with certainty*. Thus weakening the convergence criterion to a limiting one allows for more problems to come within the scope of rational inquiry.

3 Certain Doubts

Just as epistemologists have favored certainty convergence, so have skeptics. Both global and local underdetermination bring inquiry to a halt; a decisive possibility of error is either in place apriori or is forthcoming soon enough. Gettier-cases also terminate inquiry with certainty: If Jones does not own a Ford car but Brown all the same is in Barcelona then this suffices for getting it right wrongly, and knowledge as true justified belief is undercut with the same kind of certainty.[2] In sum, demonstrating doubt has generally been a short-run strategy to terminate with certainty:

> *Agent S may be in doubt concerning proposition p with certainty iff S produces a counterexample to p.*

This short-run strategy is insufficient to settle the case in favor of the skeptic since the knower may claim knowledge in the limit. Can the skeptic follow him there? How does the challenge of skepticism fare in the limit? First, consider what skepticism would look like in the limit:

> *Agent S may be in doubt concerning proposition p in the limit iff S produces a counterexample to p.*

This limiting version of skepticism seems to miss the point of the challenge. It says that there is a time such that for each later time, S produces a counterexample to p although S may not know when it is safe to produce the counterexample. In fact, it will never be safe, because if knowledge is defined as limiting convergence, such a time will never arise. If the knowing agent has knowledge of p in the limit, then p is true and nothing will ever again provoke him to change his mind, even though he may not know when the modulus of convergence has arisen. Once the agent has limiting knowledge and is thereby locked on to the truth forever after in all relevant worlds it seems that the skeptic is left with nowhere to go but to succumb to knowledge! Skepticism is then defeated by the very nature of limiting convergent knowledge.

Again, this easy victory for the knower does not pay proper credit to the skeptical challenge. If knowers move to the limit and the skeptic cannot provide a counterexample, what the skeptic is in need of is not a strategy for showing that the knowers are wrong if they are right *pace* limiting convergent knowledge, but rather an assurance that he, the skeptic himself, can limiting converge to doubt. In this case, the putative knower would be no better off than the skeptic in the limit, and the skeptical challenge stands. The question then is whether the skeptic can converge to a proposition witnessing the impossibility of knowledge in the limit.

This proposition is the Socratic dictum of epistemic modesty embraced by Academic skeptics like Carneades and Archilaus:

[2] Running through the remaining Gettier-derived counterexamples – including Russell's anticipation thereof with the stopped clock and Balfour being the prime minister when in fact it was Campbell Bannerman [Russell 56] – will reveal more doubt with certainty.

All I know is that I don't know.

Another skeptic, Sextus Empiricus, took the dictum of epistemic modesty to mean a universal generalization:

> The adherents of the New Academy, although they affirm that all things are non-apprehensible, yet differ from the Skeptics even, as seems probable, in respect of this very statement that all things are non-apprehensible (for they affirm this positively whereas the Skeptic regards it as possible that some things are apprehended). [Sextus 33]: 139.

In response to the adherents of the New Academy, Sextus Empiricus launched the classical pre-Cantorian diagonal argument against inductive inference to disprove the coherence of the Academic position. By this argument he attacked the Academic skeptics by concluding that their position was just as dogmatic as Sextus' reading of Plato's conception of the actuality of knowledge:

> For, when (the dogmatists) they propose to establish the universal from the particulars by means of induction, they will effect this by a review either of all or of some of the particular instances. But if they review some, the induction will be insecure, since some of the particulars omitted in the induction may contravene the universal; while if they are to review all, they will be toiling at the impossible, since the particulars are infinite and indefinite. Thus, on both grounds, as I think, the consequence is that induction is invalidated. [Sextus 33]: 283.

Sextus' argument is based on the assumption that the Academics are supposed to converge to their doubt with certainty. In other words doubt in inquiry is equivalent to stabilizing to the correct answer—no knowledge possible, halt! Sextus takes his argument to undermine this possibility because of local underdetermination.

4 Long-Run Doubt

In order to directly confront the skeptical challenge the skeptic must be permitted everything he needs. Hence, let the skeptic

- entertain a limiting convergence criterion

since in the short run he can do no better of proving himself right and the knowers wrong, and

- have use of the infallible method

which Sextus endorses, namely the method that never makes mistakes and only conjectures what is entailed by the evidence. In the limit, armed with the infallible method, the skeptic must prove his case by converging to doubt.

However, the skeptic may resist this way of articulating the conceptual situation. He may, for instance, suggest the possibility that inquiry may simply fail to converge. Failure to converge will take one of the following three forms: Quietism, oscillation or randomness. In queitism, the agent or method simply does not produce an output. No inquiry takes place. One example of the queitist strategy, as discussed above, is the method R. Academic skeptics advocated a queitist approach to inquiry. The relationship to inquiry here is problematic, for the familiar reason that, as soon as the skeptic asserts anything, he is abandoning his queitism and engaging in inquiry.

In the case of a method or agent which does not converge due to the production of an output that osscillates infinitely often between truth and falsity, the skeptic will be faced with accepting a scenario in which knowledge is not always subject to error, but which regularly produces true outputs.

The third, non-converging scenario is one in which the method or agent produces a random output. In this case, by definition, the method will produce true outputs. Therefore, the skeptic will not be able to appeal to the possibility of error at all points in the process of inquiry.

Given these three alternatives, the skeptic must accept that a non-converging line of inquiry does not license doubt. Doubt in the limit is not failure to converge. Therefore, the skeptic must accept the convergence condition for inquiry. At this stage, the skeptic is committed to the Socratic dictum of epistemic modesty. This is standardly rephrased as axiom 5 of epistemic logic:

$$\neg K_S p \rightarrow K_S \neg K_S p$$

This axiom says that if an agent does not know p then he knows that he does not know p. Since being in doubt about p implies not knowing p, by transitivity, being in doubt about p implies knowing that he does not know p for any arbitrary proposition p.

Here is the situation: The skeptic is entertaining the weakest convergence criterion together with the strongest method of infallibility, and the following theorem sets in:

Theorem 1

> *If knowledge is defined as limiting convergence, and S is infallible, then $\neg K_S p \rightarrow K_S \neg K_S p$ is impossible to validate.* [Hendricks 06]: 212

The theorem demonstrates that if knowledge is defined as limiting convergence, then it contrapositively follows that if agent S has not converged, S accordingly does not now even in the limit, and the use of the infallible method (or any other method for that matter) makes no difference to this result. So

the skeptic cannot, in the limit, converge to doubt, or more precisely cannot converge to axiom 5 which would witness the impossibility of knowledge.[3]

An immediate objection might be that the skeptic is required here to know that he does not know. Ignorance per se is not sufficient to make the case for skepticism, since in the limit mere doubt without knowledge of one's ignorance will not defeat knowledge. This is because in the limit the knower can do better than the skeptic, since the knower can come to *know that he knows* in the limit.

Theorem 2

> *If knowledge is defined as limiting convergence, then* $K_S p \rightarrow K_S K_S p$ *is possible to validate.* [Hendricks 06]: 205

In the limit the knower can know the he knows—so the KK-thesis is valid. This may seem quite surprising. William James, for instance, dismisses the KK-thesis in the limit, claiming that one may not infallibly know when one has converged to the fact that one has converged to the correct answer. Contemporary computational epistemologists are of the same opinion:

> This does not entail that Ψ knows he knows the answer, since (as observed above) Ψ may lack any reason to believe that his hypotheses have begun to converge. [Martin & Osherson 98]: 13.

So far, knowledge has been treated in light of the idea of limiting convergence, and yet limiting convergence is often cited as one of the primary reasons for not validating the KK-thesis. How is it possible to have the cake and eat it too?

First distinguish between two interpretations of the implicational epistemic axioms:

- **Synchronic Interpretation**: *An epistemic axiom is* **synchronic** *if the consequent obtains at the very same time the antecedent obtains.*
- **Diachronic Interpretation:** *An epistemic axiom is* **diachronic** *if the consequent either obtains later or would have obtained later than the antecedent.*

Most discussions of the KK-thesis (for and against) assume a synchronic interpretation. To date, there has not been a defense of the diachronic interpretation of the KK-thesis. However, as will be shown, a diachronic interpretation of KK is precisely what is needed in order to validate KK in the limit.

To have knowledge of a proposition is to have reached a modulus of convergence after which the method continues to project the conjecture over all later times and relevant possible worlds. For the purposes of this argument let a proposition p be the set of worlds in which p is true. The set of worlds making knowing-p true is a subset of the set of worlds that make p true. To

[3] For an elaboration of the technical details, refer to [Hendricks 06]: 212–215.

have knowledge of knowledge of a proposition p is to reach a modulus of convergence *after* convergence to knowledge of p. This is because the set of worlds making knowledge of knowledge of p true is a subset of the set of worlds making knowledge of p true. Therefore knowledge of knowledge can only happen once knowledge of the proposition has obtained. Hence the inclusion order

$$[K_S K_S p] \subseteq [K_S p] \subseteq [p].$$

One has to opt for a diachronic interpretation of the KK thesis in order to validate it. This falls out naturally given a method that respects the inclusion order defined above. First the method converges to knowledge of p. Then, the method must determine whether there are worlds in which it is true that $K_S K_S p$ which are not included in the set of worlds associated with $K_S p$. In short, this dependence simultaneously ensures that the necessary ordering is not violated and motivates acceptance of the diachronic interpretation of KK.[4]

In the long run, skeptics cannot know of their doubt, but knowers can diachronically come to know of their knowledge, so in the limit knowers are much better off than skeptics, or rather, ignorance is always only a short-term assurance, if any assurance at all.

In the End

Any objection that the skeptic might consider launching against the limiting convergence strategy will prove unsuccessful by virtue of the character of assertion and inquiry described above. Take for instance the claim that since there exists an epistemically cursed world, a possible world in which agents are always wrong and that therefore our beliefs are always subject to the possibility of error. The possible existence of such a state of affairs may be undeniable. The problem for the skeptic involves defending the claim that the actual world *is* the world of his conception. If he is to take the additional step of asserting that this identity holds, then he needs to play the game of inquiry and knowledge seeking. Once he enters into the epistemological project, he is subject to the kind of limiting convergence arguments presented above.

How might knowledge in the limit be challenged by the conceivability of possible worlds in which agents are always wrong? Consider an agent whose knowledge has converged in the limit not knowing that the point at which he will no longer change his mind has already passed. In fact one could also imagine him conceiving the possibility that he is wrong or that he is an inhabitant of the epistemically cursed world. The agent has every right to ignore the skeptical possibility entirely, in spite of not knowing his entitlement. This agent is (by stipulation) locked on to an unwavering path. His philosophical intuitions might lead him to conceive that this is not the case, but his intuitions are irrelevant, they have no bearing on the fact that his beliefs have all the properties demanded of knowledge.

[4]For an elaboration of the technical details, refer to [Hendricks 01]: 253–260.

What then of the skeptic who simply denies that we have knowledge? Assertions of this kind about the current state of our knowledge are similarly irrelevant. In addition to the trivially self-defeating aspect of such assertions, it has been shown by the foregoing argument that this self-ascription simply cannot be correct even given the benefit of the doubt in the most generous of ways. The skeptic will not converge towards knowledge of the futility of inquiry, not even in the limit.

The issue for the skeptic involves demonstrating that knowers are operating with a bad method or that we are in an epistemically cursed world. Such demonstration involves a significant step beyond conceivability or intuition and immediately draws the skeptic into the kind of convergence situations described above. Having knowledge then, is an objective property of agents that have converged in the limit. One could imagine an agent that has already converged doubting himself, or imagining ways that he could be wrong. That's just fine. His job at that point would be to continue inquiring by whatever reliable means are available to him. Of course, from a third-person perspective, he will simply continue down the same path that he was on before and will continue having the robust features of a knower, in spite of any skeptical worries that might afflict him.

References

[Hendricks 01] Hendricks, V. F. (2001). *The Convergence of Scientific Knowledge – A View From the Limit.* Dordrecht: Springer.

[Hendricks 06] [Hendricks 06] Hendricks, V. F. (2006). *Mainstream and Formal Epistemology.* New York: Cambridge University Press.

[Hintikka 05] Hintikka, J. (2005). *Knowledge and Belief: An Introduction to the Logic of the Two Notions.* Reprint edition prepared by Hendricks, V.F. and Symons, J. London: King's College Publications.

[James 60] James, W. (1960). "The Will to Believe". In *Essays in Pragmatism.* Hafner Publishing Company.

[Kelly 96] Kelly, K.T. (1996). *The Logic of Reliable Inquiry.* New York: Oxford University Press.

[Kitcher 93] Kitcher, P. (1993). *The Advancement of Science.* New York: Oxford University Press.

[Lewis 96] Lewis, D. (1996). "Elusive Knowledge", reprinted in *Papers in Metpahysics and Epistemology.* New York: Cambridge University Press.

[Martin & Osherson 98] Martin, E. and Osherson, D. (1998). *Elements of Scientific Inquiry*. Cambridge, MA: MIT Press.

[Nozick 83] Nozic, R. (1983). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

[Russell 56] [Russell 56] Russell, B. (1956). *Logic and Knowledge: Essays 1901-1950*. March, R.C. (ed.). London: Allen and Unwin.

[Schulte 00] Schulte, O. (2000). "Inferring Conservation Principles in Particle Physics: A Case Study in the Problem of Induction", *British Journal for the Philosophy of Science*, **51**: 771-806.

[Sextus 33] Sextus Empiricus (1933). *Outlines of Pyrrhonism*, Vol. 1. Bury, R.G. (transl.). Cambridge, MA: Harvard University Press.