

Epistemology Axiomatized

UNESCO / Paris / November 2003

Vincent F. Hendricks

Department of Philosophy and Science Studies

Roskilde University, Denmark

vincent@ruc.dk <http://www.ruc.dk/~vincent>

Forcing epistemology is a trendy way of defeating the skeptics who since the days of old have cited *prima facie* error-possibilities as some of the most devastating arguments against claims to knowledge. The idea of *forcing* is to delimit the set of possibilities over which the inquiring agent has to succeed: If the agent can succeed over the relevant possibility set then the agent may still be said to have knowledge even if he commits grave and many errors in other but irrelevant possibilities. Many contemporary mainstream and formal epistemologies pay homage to the forcing strategy. The aim of this article is to *systematically* demonstrate that the mainstream and formal epistemologies have much in common both epistemologically and methodologically. The two approaches may profit from one another paving the way for a new unifying program in 'interactive' axiomatized epistemology. The paper is partly excerpts from chapters 4, 5 and 6 of *Forcing Epistemology*, Vincent F. Hendricks, New York: Cambridge University Press, forthcoming 2004 (see <http://www.ruc.dk/~vincent/forcing.htm>).

1 Forcing Mainstream and Formal Epistemology

Agents inquire to replace ignorance with knowledge. Knowledge is a kind of epistemic commitment or attitude held towards propositions or hypotheses describing some aspect of the world under consideration.¹ Agents may in general hold a host of different propositional attitudes like belief, hope, wish, desire etc. But there is a special property that knowledge enjoys over and above the other commitments— Plato already pointed out that a distinct property of knowledge is truth: Whatever is known must be true otherwise it is not knowledge even though it very well may qualify as belief or some other propositional attitude.

The contemporary notions of knowledge are often *modal* in nature. Knowledge is defined with respect to other possible states of affairs besides the actual state of affairs (Figure 1). The possibility of knowledge seems ruled out when it is possible that we err. Introducing other possible state of affairs is an attempt to preclude exactly these error-possibilities. Knowledge must be *infallible* by definition as Lewis puts it:

To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* like a contradiction [78], p. 367.

A fallible notion of knowledge is not much different from a concept of belief potentially allowing the agent to ‘know’ a falsehood severing the connection between knowledge and truth.

Plato also observed that knowledge, as opposed to mere true belief, is stable in nature. Knowledge has steadfastness and indefeasibility attached to it. True belief is quite useful as far as it goes, but in the light of true evidence it may vanish. In the light of true evidence knowledge will not evaporate. No evaporation makes for the more robust usefulness of knowledge in contrast to simply beliefs which are true. True belief in the actual world is not necessarily preserved if circumstances were to change but slightly. On the other hand knowledge in the actual world is assumed stable over even quite radically varying circumstances. Thus, among both informally and formally minded epistemologists there is an agreement that knowledge is defined with respect to other ‘possible worlds’ as Hintikka notes:

In order to speak of what a certain person *a* knows and does not know, we have to assume a class (‘space’) of possibilities. These possibilities will be called scenarios. Philosophers typically call them possible worlds. This usage is a symptom of intellectual megalomania. [52], p. 19.

There is an immediate difference between a philosophical logician and a philosopher. The logician often remains rather agnostic about the ontological significance of the possible worlds and may just refer to them as scenarios, situations, states, contexts or conceptual constructions. The philosopher is usually quite concerned with the metaphysical baggage that comes along with the notion.²

¹The terms ‘hypothesis’ and ‘proposition’ will be used interchangeably unless otherwise stated.

²To stay in the currently adopted jargon, other scenarios, situations, states or contexts




Figure 1: Modal knowledge is defined with respect to other possible worlds

Be that as it may, the stability and robustness of knowledge over other state of affairs still leaves open the question of which space of worlds should be considered relevant for epistemic success. The classical conception of infallibilism is taken to require, that for an agent to have knowledge of some hypothesis, he must be able to eliminate *all* the possibilities of error associated with the hypothesis in question. The set of *all* worlds is considered. This set of possible worlds is too big for knowledge to have scope over. The set includes some rather bizarre worlds inhabited by odd beasts from demons to mad and malicious scientists who have decided to stick your brain in a tank of nutritious fluids to systematically fool you. Or worlds in which contradictions are true. If these worlds were to be considered relevant all the time skepticism would have the upper hand all the time. Epistemology may just end up with a fallibilistic notion of knowledge after all: There may not be a way for an agent to determine that he is not in the world of the beast or the brain. But then again, a fallibilistic notion of knowledge hardly qualifies as knowledge at all. At most it amounts to a description of knowledge-seeking practices. Consequently, *if infallibilism is to be a viable reply to the skeptic, then infallibilism cannot be defined with respect to all possible worlds.* This is where the epistemological forcing feature comes in.

The bizarre and extravagant possibilities of error may under the right circumstances be ignored even though they are logically possible for all the agent knows. Knowledge may accordingly remain infallible but with world restrictions imposed. Forcing is more of an heuristic principle than a thesis:

Whenever knowledge claims are challenged by alleged possibilities of error, the strategy is to show that the possibilities of error fail to be genuine in the relevant sense

Contemporary epistemologists choose to speak of the *relevant* possible worlds

will be referred to as 'possible worlds' but nothing metaphysical is necessarily implied by the usage. Possible worlds are not to be understood as ontological or semantical totalities complete in their spatio-temporal history. It is possible to endow possible worlds with enough formal structure to actually facilitate important epistemological results [39], [40], [41].

as a subset of the set of all possible worlds.³ The philosophical logicians and other formal epistemologists consider an *accessibility* relation between worlds in a designated class out of the entire universe of possible worlds. It will become apparent that there is no principled difference between relevance and accessibility. Informal epistemologies differ by the way in which relevance is forced given, say, perceptual equivalence conditions, counterfactual proximities or conversational contexts circumscribing the possible worlds. Formal epistemologies differ by the way in which the accessibility relation is defined over possible worlds.

A philosophical logician obtains different epistemic modal systems valid for a knowledge operator by varying (adding, dropping or relativizing) the properties of the accessibility relation from, say, reflexive and transitive to a reflexive, symmetric and transitive accessibility relation. Algebraic constraints on the accessibility relation are the forcing foundation for a formal approach to the theory of knowledge like logical epistemology, or better known as epistemic logic. Forcing clauses in logical epistemology will subsequently be compared to forcing restrictions from more mainstream epistemic frameworks like Nozickian counterfactual epistemology, and Lewisian modal epistemology. In the end, constraints on accessibility relations between possible worlds is a crucial *unifying* parameter for both mainstream and formal epistemological axiomatizations.

2 Logical Epistemology

Logical epistemology dates back to von Wright [110] and especially the work of Hintikka [46] in the early 1960's.⁴ Epistemic logics have since then grown into powerful enterprises enjoying many important applications.⁵ The general epistemological significance of the logics of knowledge has to some extent been neglected by mainstreamers and formalists alike. The field is in a rather awkward position today: On the one hand it is a discipline of immanent importance to for example theoretical computer scientists, linguists and game theorists but they do not necessarily have an epistemological ambition with their use of epistemic logic. On the other hand it is a discipline devoted to the logic of knowledge and belief but alien to epistemologists and philosophers interested in the theory of knowledge.

Recent results and approaches have fortunately brought the logics of knowledge quite close to the theories of knowledge. One may even identify a set of *logical epistemologies* having essential features in common with the mainstream epistemological proposals surveyed so far.

Epistemic logic was early on greatly influenced by the advances stemming from modal logic. Standard systems of modal logic were given epistemic in-

³Explicit forcing proposals in the epistemological literature are sometimes referred to as '*relevant alternatives proposals*'. Cf. Bernecker and Dretske [7].

⁴It has been pointed out by Boh [14] and Knuuttila [60] that the study of epistemic logic may be tracked back to at least scholastic philosophy.

⁵Refer to Gochet and Gribomont's forthcoming paper for an excellent survey of epistemic logic and its key issues [35]. For another broader survey of the significance of logic for epistemology refer to van Benthem [11].

terpretations, and some main technical results of epistemic logic could then be extracted.

Syntactically, the language of propositional epistemic logic is obtained by augmenting the language of propositional logic with a unary epistemic operator K_{Ξ} such that

$$K_{\Xi}A \text{ reads 'Agent } \Xi \text{ knows } A'$$

for some arbitrary *proposition* A .⁶ This formalization of knowledge is an interpretation of $\Box A$ in alethic logic reading ‘It is necessary that A ’. Interpreting modal logic epistemically is crudely a reading of modal formulae as epistemic statements expressing attitudes of certain agents towards certain propositions.

The semantics of modal logic is likewise given a novel interpretation. Hintikka came up with the following semantic interpretation of the epistemic operator [46], [47]:

$K_{\Xi}A \approx$ *in all possible worlds compatible with what Ξ knows it is the case that A*

The basic assumption is that any ascription of propositional attitudes like knowledge and belief, requires *partitioning* of the set of possible worlds into two compartments: The compartment consisting of possible worlds compatible with the attitude in question and the compartment of worlds incompatible with it.

Based on the partition the agent is capable of constructing different ‘world-models’ using the epistemic modal language. He is not necessarily required to know which one of the world-models constructed is the real world-model. All the same, the agent does not consider all these world-models equally possible or accessible from his current point of view. Some world-models may be incommensurable with his current information state or other background assumptions. These incompatible world-models are excluded from the compatibility partition. This is a variation of the forcing strategy. In logical epistemology, as in many mainstream epistemologies, it is typically stipulated that the smaller the set of worlds an agent considers possible, the smaller his uncertainty, at the cost of stronger forcing assumptions.

Consider an epistemic logic for a single agent—*mono-agent systems*. The set of worlds considered accessible by an agent depends on the actual world, or the agent’s actual state of information. It is possible to capture the forcing dependency by introducing a relation of accessibility, R , on the set of compatible possible worlds. To express the idea that for agent Ξ , the world w' is compatible with his information state, or accessible from the possible world w which Ξ is currently in, it is required that R holds between w and w' . This relation is written Rww' and read ‘world w' is accessible from w ’. The world w' is said to be an *epistemic alternative* to world w for agent Ξ . Given the above semantical interpretation, if a proposition A is true in all worlds which agent Ξ considers

⁶The discussion of epistemic logic is restricted to the propositional case. Epistemic logic has also been developed for first-order languages.

possible then Ξ knows A .⁷

Formally, a *frame* \mathcal{F} for an epistemic system is a pair (W, R) for which W is a non-empty set of possible worlds and R is a binary accessibility relation over W . A *model* \mathbb{M} for an epistemic system consists of a frame and a denotation function φ assigning sets of worlds to atomic propositional formulae. Propositions are taken to be sets of possible worlds; namely the set of possible worlds in which they are true. Let *atom* be the set of atomic propositional formulae, then $\varphi : \text{atom} \rightarrow P(W)$ where P denotes the powerset operation. The model $\mathbb{M} = \langle W, R, \varphi \rangle$ is called a Kripke-model and the resulting semantics Kripke-semantics [61]: An atomic propositional formulae, \mathbf{a} , is said to be true in a world w (in \mathbb{M}), written $\mathbb{M}, w \models \mathbf{a}$, iff w is in the set of possible worlds assigned to \mathbf{a} , i. e. $\mathbb{M}, w \models \mathbf{a}$ iff $w \in \varphi(\mathbf{a})$ for all $\mathbf{a} \in \text{atom}$. The formula $K_{\Xi}A$ is true in a world w , i.e. $\mathbb{M}, w \models K_{\Xi}A$, iff $\forall w' \in W : \text{if } Rww', \text{ then } \mathbb{M}, w' \models A$. The semantics for the Boolean connectives are given in the usual recursive way. A modal formula is said to be *valid* in a frame iff the formula is true for all possible assignments in all worlds admitted by the frame.

There is a formula or axiom which is valid in all possible frames

$$K_{\Xi}(A \rightarrow A') \rightarrow (K_{\Xi}A \rightarrow K_{\Xi}A'). \quad (1)$$

The axiom amounts to the contentious closure condition for knowledge, to be discussed further below. Logical epistemology unproblematically accepts (1) but for a formal reason.

One rule of inference which is valid in all possible frames is the rule of *necessitation* (N)

$$\frac{A}{K_{\Xi}A} \quad (2)$$

which says that if A is true in all worlds of the frame, then so is $K_{\Xi}A$. The rule was found valid on both the counterfactual and modal accounts of knowledge.

Neither (1) nor (2) require any assumptions to be made pertaining to the accessibility relation between the possible worlds considered compatible with the knowledge attitude. It actually turns out that (1) together with (2) comprise the characterizing axiom and rule for possible world semantics with binary accessibility relations. All modal logics in which (1) and (2) are valid are called *normal* modal logics.

Besides (2) other rules of inference are valid in normal modal systems like

$$\frac{A \leftrightarrow A'}{K_{\Xi}A \leftrightarrow K_{\Xi}A'} \quad (3)$$

called *congruence*.⁸

⁷ The quantifier restriction on knowledge in logical epistemology is comparable to the quantificational restriction found in Lewis' modal epistemology below. Knowledge claims are circumscribed by the compartment of possible worlds in accordance with the epistemic attitude, not the incompatible compartment and not the set of all possible worlds.

⁸ There are other valid rules of inference but for the current purposes the ones listed suffice.

There exist a whole range of formulae which are not characterized by all possible frames. Put differently, there are many epistemic axioms which are not valid in all possible frames. These modal axioms may however be valid in all frames of a certain sub-class of frames.

A nice feature of possible world semantics is that many common epistemic axioms correspond to certain algebraic properties of the frame in the following sense: A modal axiom is valid in a frame if and only if the accessibility relation satisfies some algebraic condition. For an example, the axiom

$$K_{\Xi}A \rightarrow A \quad (4)$$

is valid in all frames in which the accessibility relation is *reflexive* in the sense that

$$\forall w \in W : Rww.$$

Every possible world is accessible from itself. Similarly if the accessibility relation satisfies the condition that

$$\forall w, w', w'' \in W : Rww' \wedge Rw'w'' \rightarrow Rww''$$

then the axiom

$$K_{\Xi}A \rightarrow K_{\Xi}K_{\Xi}A \quad (5)$$

is valid in all *transitive* frames. Other axioms require yet other relational properties to be met in order to be valid in all frames.

A nomenclature due to Lemmon [68] and later refined by Bull and Segerberg [17] is helpful while cataloguing the axioms typically considered interesting for epistemic logic (Table 1).

K	$\overline{K_{\Xi}(A \rightarrow A') \rightarrow (K_{\Xi}A \rightarrow K_{\Xi}A')}$
D	$K_{\Xi}A \rightarrow \neg K_{\Xi}\neg A$
T	$K_{\Xi}A \rightarrow A$
4	$K_{\Xi}A \rightarrow K_{\Xi}K_{\Xi}A$
5	$\neg K_{\Xi}A \rightarrow K_{\Xi}\neg K_{\Xi}A$
.2	$\neg K_{\Xi}\neg K_{\Xi}A \rightarrow K_{\Xi}\neg K_{\Xi}\neg A.$
.3	$K_{\Xi}(K_{\Xi}A \rightarrow K_{\Xi}A') \vee K_{\Xi}(K_{\Xi}A' \rightarrow K_{\Xi}A)$
.4	$A \rightarrow (\neg K_{\Xi}\neg K_{\Xi}A \rightarrow K_{\Xi}A)$

Table 1: Common epistemic axioms

- Axiom K, also called the *axiom of deductive cogency*: If the agent Ξ knows $A \rightarrow A'$, then if Ξ knows A , Ξ also knows A' .
- Axiom D, also referred to as the *axiom of consistency* requires Ξ to have consistency in his knowledge: If an agent knows A , he does not simultaneously know its negation.⁹

⁹This axiom also has a deontic motivation in the sense that if an agent is obligated to do whatever A prescribes, he is not at the same time obligated to do $\neg A$.

- Axiom **T**, also called the *axiom of truth* or *axiom of veridicality*, says that if A is known by Ξ , then A is true.
- Axiom **4** is also known as the *axiom of self-awareness, positive introspection* or *KK-thesis*. They all refer to the idea that an agent has knowledge of his knowledge of A if he has knowledge of A .
- Axiom **5** is also known as the *axiom of wisdom*. It is the much stronger thesis that an agent has knowledge of his own ignorance: If Ξ does not know A , he knows that he doesn't know A . The axiom is sometimes referred to as the *axiom of negative introspection*.
- Axiom **.2** reveals that if Ξ does not know that he does not know A , then Ξ knows that he does not know not A .
- Axiom **.3** maintains that either Ξ knows that his knowledge of A implies his knowledge of A' or he knows that his knowledge of A' implies his knowledge of A .
- Axiom **.4** amounts to the claim that any true proposition *per se* constitutes knowledge and is sometimes referred to as *axiom of true (strong) belief*.

These axioms in proper combinations make up epistemic modal systems of varying strength depending on the modal formulae valid in the respective systems given the algebraic properties assumed for the accessibility relation.

The weakest system of epistemic interest is usually considered to being system **T**. The system includes **T** and **K** as valid axioms. Additional modal strength may be obtained by extending **T** with other axioms drawn from the above pool altering the frame semantics to validate the additional axioms. By way of example, while $K_{\Xi}A \rightarrow A$ is valid in **T**, $K_{\Xi}A \rightarrow A$, $K_{\Xi}A \rightarrow K_{\Xi}K_{\Xi}A$ and $\neg K_{\Xi}A \rightarrow K_{\Xi}\neg K_{\Xi}A$ are all valid in **S5** but not in **T**. System **T** has a reflexive accessibility relation, **S5** has an equivalence relation of accessibility. The arrows in table 2 symbolize that the system to which the arrow is pointing is included in the system from which the arrow originates and hence reflect relative strength. Then **S5** is the strongest and **S4** the weakest of the ones listed.

Epistemic Systems		
KT4	=	S4
KT4 + .2	=	S4.2 ↑
KT4 + .3	=	S4.3 ↑
KT4 + .4	=	S4.4 ↑
KT5	=	S5 ↑

Table 2: Relative strength of epistemic systems between **S4** and **S5**

What Hintikka recently dubbed ‘first generation epistemic logic’ [52] is characterized by the ambition that cataloguing the possible complete systems of such logics would allow for choosing the most ‘appropriate’ or ‘intuitive’ ones(s).

These ‘appropriate’ logics often range from **S4** over the intermediate systems **S4.2-S4.4** to **S5**. By way of example, Hintikka settled for **S4** [46], Kutschera argued for **S4.4** [62] while Lenzen advocated a system of knowledge at least as strong as **S4.2** and at most as strong **S4.4** [69], [70]. Computer scientists like van der Hoek has proposed to strengthen knowledge according to system **S4.3** [56] and later **S5** [57] just like Fagin, Halpern, Moses and Vardi [27] assume knowledge to be **S5** valid. One may at the same time quote Halpern for the following instructive insight:

My own feeling is that there is no unique right notion of knowledge; the appropriate notion is application dependent. [37], p. 483.

While the sentiment is certainly to appreciate it has recently been argued, as will become apparent below, that some systems are more tenable than others both philosophically but likewise for applications.

Reflexivity is the characteristic frame property of system **T**, transitivity is the characteristic frame property of system **S4**, equivalence the characteristic frame property of **S5**, etc. From an epistemological point of view, the algebraic properties of the accessibility relation are really forcing conditions (Figure 2).




Figure 2: Logical epistemology requires success in all worlds depending on the algebraic properties of the accessibility relation

The cognitive rationale of logical epistemology must be something like this: The more properties the accessibility relation is endowed with, the more access the agent has to his epistemic universe, and in consequence the more epistemic strength he will obtain. The stronger knowledge, the stronger forcing clauses.

Modal epistemic axioms and systems may be viewed as measures of infallibility and replies to skepticism. For instance, knowing your own knowledge is a way of blocking the skeptic, but knowledge of your own ignorance in terms of axiom 5 is better still. One motivation for the plausibility of axiom 5 is in data-base applications: An agent examining his own knowledge base will be let to conclude that whatever is not in the knowledge base he does not know and hence he will know that he does not.

The axiom of wisdom or negative introspection is a sort of closed world assumption. A closed world assumption is a forcing assumption if anything is,

‘shutting the world down’ with the agent, leaving the skeptic nowhere to go. To know the truth, to know of your knowledge, and to know of your own ignorance as in **S5** requires ‘full’ of epistemic access which is exactly why the accessibility relation must be an equivalence relation. A theorem of **S5** is the following

$$\neg A \rightarrow K_{\Xi} \neg K_{\Xi} A \quad (6)$$

which states that if A is not the case, then Ξ knows that he does not know A —the ‘truly Socratic person’ as Girle explains ([32], p. 157) knowing exactly how ignorant he is.

A bit more ignorance, a bit more skepticism and accordingly a bit more fallibility is allowed in **S4**. Since axiom 5 is dropped and (6) is no longer a theorem,

$$\{\neg A, \neg K_{\Xi} \neg K_{\Xi} A\} \text{ and } \{\neg K_{\Xi} A, \neg K_{\Xi} \neg K_{\Xi} A\}$$

are not inconsistent in **S4**. It is possible for an agent to be ignorant of the fact that he does not know when actually he does know. Put differently, the agent is allowed false beliefs about what is known.

Yet more ignorance and skepticism are allowed in system **T** because while

$$\{K_{\Xi} A, \neg K_{\Xi} K_{\Xi} A\}$$

is inconsistent in **S4**, this set of epistemic statements is not inconsistent in **T**. The agent may thus know something without knowing that he does. All the same, a restricted kind of positive introspection is still prevalent in system **T**. Given the rule of necessitation (2), Ξ knows all the theorems of the epistemic logic. By iteration, $K_{\Xi} K_{\Xi} A$ is also known. Thus if A is a theorem, Ξ knows that he knows A .

Philosophers have raised the question of whether the logic of knowledge makes any epistemological sense. There are at least two ways of answering this question. One is to deny the presupposition that epistemic logic should uphold broader epistemological pertinence. The discipline is not obligated to hook up with more general epistemological concerns ranging from closure conditions to justification, methodology, reliability and rationality, as Lenzen argues:

The search for the correct analysis of knowledge, while certainly of extreme importance and interest to epistemology, seems not significantly to affect the object of epistemic logic, the question of validity of certain epistemic-logical principles. [69], p. 34.

The other way to answer this charge is to maintain that epistemic logic does carry epistemological significance but in a rather idealized sort of way: One restricts attention to a class of rational agents where rationality is defined by certain postulates. Thus, agents have to satisfy at least some minimal conditions to simply qualify as rational. This is by and large what Lemmon originally suggests [67]. One such condition would be that assuming an agent as rational entails that he should know the logical laws. For instance, if the agent knows

A and $A \rightarrow A'$, he should be able to use modus ponens to infer A' . Now these ‘rationality postulates’ for knowledge exhibit a striking similarity with the laws of modal and epistemic logic. One may in turn legitimately attempt to interpret the necessity operator in alethic axioms as a knowledge operator and then justify the modal axioms as axioms of knowledge.¹⁰ While Lemmon constructs the rational epistemic agent directly from the axiomatization of the logic, yet another way of justifying the epistemic axioms is by ways of meaning: Find a plausible epistemological story to tell about the semantics of epistemic logic.

This line of thought Hintikka pursued from the outset [46]. Hintikka stipulated that the axioms or principles of epistemic logic are conditions descriptive of a special kind of general (strong) *rationality*. The statements which may be proved false by application of the epistemic axioms are not inconsistent meaning that their truth is logically impossible. They are rather rationally ‘indefensible’. Indefensibility is fleshed out as the agent’s epistemic laziness, sloppiness or perhaps cognitive incapacity whenever to realize the implications of what he in fact knows. Defensibility then means not falling victim of ‘epistemic negligence’ as Chisholm calls it [18], [19]. The notion of indefensibility gives away the status of the epistemic axioms and logics. Some epistemic statement for which its negation is indefensible is called ‘self-sustaining’. The notion of self-sustenance actually corresponds to the concept of validity. Corresponding to a self-sustaining statement is a logically valid statement. But this will again be a statement which is rationally indefensible to deny. So in conclusion, epistemic axioms are descriptions of rationality.

There is an argument to the effect that Hintikka early on was influenced by the autoepistemology of G.E. Moore [82] and especially Malcolm [81] and took, at least in part, their autoepistemology to provide a philosophical motivation for epistemic logic. There is an interesting twist to this motivation which is not readily read out of Moore’s autoepistemology. Epistemic principles may be interpreted as principles describing a certain strong rationality. The agent does not have to be aware of this rationality, let alone able to immediately compute it from the first person perspective as Hintikka argues:

In order to see this, suppose that a man says to you, ‘I know that p but I don’t know whether q ’ and suppose that p can be shown to entail logically q by means of some argument which he would be willing to accept. Then you can point out to him that what he says he does not know is already implicit in what he claims he knows. If your argument is valid, it is irrational for our man to persist in saying that he does not know whether q is the case. [46], p. 31.

In Hintikka’s logical system knowledge is closed in the sense of (1). The closure is needed for driving the argument through even if the local agent is not immediately computing it. ‘I get by with a little help from my friends’ applies here.

¹⁰ For a more detailed discussion of this approach, refer to Girle [32].

The man in the local street may turn out to be rational after all when properly advised of directions *actually* available to him vindicating the first person interpretation.

The autoepistemological inspiration is vindicated while Hintikka argues for the plausibility of the *KK*-thesis as a governing axiom of his logic of knowledge as he refers to Malcolm:

This is especially interesting in view of the fact that Malcolm himself uses his strong sense of knowing to explain in what sense it might be true that whenever one knows, one knows that one knows. In this respect, too, Malcolm's strong sense behaves like mine. [48], p. 154.

Besides the requirement of closure and the validity of the *KK*-thesis, axiom T is also valid to which neither Moore nor Malcolm would object. A logic of autoepistemology is philosophically congruent with Hintikka's suggestion for a **S4** epistemic logic describing strong rationality.

It may be helpful to briefly review the fundamentals of autoepistemology. Moore's common-sense considerations on which autoepistemology is founded deflates the skeptical possibilities of error from various dialectic angles of which one is particularly pertinent to the current discussion. It is called the argument from *incoherence*. The idea is to demonstrate that skepticism has severe difficulties in formulating its own position coherently. As with any argument, a skeptical conclusion presupposes knowledge of a set of premisses. Moore then points to the fact that merely *asserting* these premisses imply at least a doxastic commitment, but most likely an epistemic commitment. The skeptics cannot be retreating to a statement like

*'There are 9 planets in our solar system
but it is not the case that I believe it.'* (7)

The statement in (7) is an instance of what later has become known as the *Moore-paradox*. Let it be granted that (7) only involves an error of omission. All the same it still sounds self-contradictory simply given mere assertion. No formulation of skepticism without incoherence.

The argument from incoherence is a first person point argument. Skepticism is thus rejected along these lines. A first person perspective is one of the very characteristics of autoepistemology. This is also suggested in the label 'autoepistemology' attaching the Moore-paradox to it: Whatever an agent may know or believe is partly fixed by the concern whether the epistemic or doxastic claim advocated by the inquiring agent fall victim of a Moore-paradox or not. As long as a thesis concerning epistemic commitments does not pan out in a Moore-paradox the inquiring agent is free to adopt it. As an autoepistemologist one may, by way of example, say

'If I believe that A, then I believe that I know that A' (8)

which has later been called the *Moore-principle* and sometimes the *principle of positive certainty*.¹¹ Formalized (8) amounts to:

$$B_{\Xi}A \rightarrow B_{\Xi}K_{\Xi}A. \quad (9)$$

According to Moore's theory, there is nothing self-contradictory or incoherent about asserting the principle. No more Moore paradox to the Moore principle than to the widely adopted principle that one knows that one knows if one does the plausibility of which Malcolm argues for above and elsewhere [81].

From Moore's first person autoepistemological perspective a statement like

$$'A \text{ is the case, but I don't believe whether } A' \quad (10)$$

is a paradoxical Moorean statement. There is however nothing paradoxical about

$$'A \text{ is the case, but } \Xi \text{ doesn't believe whether } A' \quad (11)$$

from a third person perspective. In consequence, what for sure may sound quite implausible from the first person perspective, may sound very plausible from the third person perspective on inquiry and vice versa.

The epistemic and doxastic commitments that an agent may hold in the course of inquiry is sensitive the epistemic environment and what the agent in these local circumstances is both willing to and capable of defending or maximizing. He does not necessarily have an over-all skepticism defeating method at his disposal: You are doing the best you can, so is the skeptic, but he is probably not doing as well as you are due to incoherence. Forcing in autoepistemology then means:

Whenever knowledge claims are challenged by alleged possibilities of error, the strategy is to show that on an individual basis one can do no better than what is being done in the current epistemic environment and attempt to show that the skeptic is doing at least as bad as you are but probably even worse

Moore's principle (8) is a kind of introspection axiom for rational belief or *subjective certainty*. In a combined epistemic and doxastic logical system in which knowledge and belief are approximately equally strong (save for a truth-condition) the agent will (while subjectively reflecting upon his own state of mind with respect to what he believes) be led to believe that he knows the proposition in question if he certainly believes it. Some contemporary logical epistemologists embrace Moore's principle (e.g. Halpern [37]). Early logical epistemologists are more reluctant insofar belief is taken to be fairly weak as Lenzen notes:

Remember that the operator B of 'weak' belief covers in particular such presumptions as are expressed by ' a believes (presumes) that Holland will

¹¹Lamarre and Shoham explain: 'To the agent, the facts of which he is certain appear to be knowledge', [65].

win the next world championship in football'. It ought to be evident that, unless a is a prophet, he does not believe that he knows Holland will win the next championship in football. [69], p. 80.

Hintikka likewise dismisses Moore's principle:

Hence ... and (C.BK) [Moore's principle] are acceptable only when an unrealistically high standard of defensibility is imposed on one's beliefs. The conditions would make it (logically) indefensible to suppose that anyone would have given up any of his present beliefs if he had more information than he now has. And this is clearly too stringent a requirement. [46], p. 52.

To both Lenzen and Hintikka belief is a significantly weaker commitment than knowledge. For good reason too: Consider a combined epistemic and doxastic logic in which belief is understood as subjective certainty such that (9) holds. Assume also that positive introspection

$$B_{\Xi}A \rightarrow K_{\Xi}B_{\Xi}A \quad (12)$$

holds for belief together with negative doxastic introspection

$$\neg B_{\Xi}A \rightarrow K_{\Xi}\neg B_{\Xi}A. \quad (13)$$

Even subjective certainty, as strong as it may seem in this system, implies a margin of error: The fact that Ξ is subjectively certain of A does not necessarily imply that A is true. Accordingly axiom **T** will be dropped for subjective certainty and replaced by the consistency axiom **D**

$$B_{\Xi}A \rightarrow \neg B_{\Xi}\neg A. \quad (14)$$

On the standard definition of knowledge, knowledge implies belief

$$K_{\Xi}A \rightarrow B_{\Xi}A \quad (15)$$

which is also an uncontroversially accepted assumption for knowledge and subjective certainty. The logic of subjective certainty is **KD45**. Knowledge will obviously have to be stronger than subjective certainty, so it must validate **S5**. On assumptions (9), (12)–(15) Lenzen was able to show that $B_{\Xi}A$ in the end is equivalent to $K_{\Xi}A$ [69]. So knowledge and belief collapse into each other!¹²

Many contemporary epistemic logics do nevertheless consider strong belief, rational belief or subjective certainty to be approximately as strong as knowledge. Assuming belief is taken to be approximately as strong as **S5** knowledge with the equivalence relation over worlds implies some attractive formal features like readily epistemic and doxastic partitions. This does not by itself make up for the result that the logic of knowledge and belief coincide.

¹²Stalnaker also discusses this issue in [106].

The plausibility of the epistemic axioms has occasioned quite stormy debates in the logico-epistemological literature. Axiom K has been under attack, many have found axiom 4 implausible like Nozick below and so forth. In the eye of the tornado has been axiom 5. A few objections to this axiom will be briefly reviewed.

Hintikka denies its plausibility as does Lenzen essentially because introspection alone should not license agents to ascertain whether some proposition in question is known [69]. Under special circumstances axiom 5 suggests that agents can even decide intractable problems as Binmore reveals in [13], and Shin in [103]. Williamson has launched two objections to models of knowledge and belief validating axiom 5. For **S5** knowledge Williamson disagrees with the ones interpreting knowledge in a data-base like fashion to justify the closed world assumption of axiom 5. Even under the closed world assumption it does not follow in general that an agent can ‘survey the totality of its knowledge’.¹³ Secondly, Williamson recently noted that the result to the effect that knowledge and belief collapse under the strong understanding of belief in a combined system points to the untenability of axiom 5, not to the unacceptable nature of subjective certainty *per se*. Moore’s principle is not too extravagant an assumption for rational belief, neither are axioms (12), D, (15) nor axioms T, 4 for knowledge. That leaves axiom 5 as the culprit responsible for collapsing the two notions and besides entails the infallibility of the agent’s beliefs: Whatever Ξ believes is true. On these grounds, Williamson abandons axiom 5 rather than any of the other principles used in the derivation [114]. Voorbraak makes the unusual move of sacrificing (15) accordingly challenging the intuitions of philosophers since antiquity [109].

If **S5** assumptions about knowledge and belief are dropped ideal rationality descriptions and autoepistemological considerations may supply a philosophical foundation and motivation for logical epistemology.¹⁴ The treatment of logical epistemology as a branch of modal logic is still quite costly also for less ambitious logics than **S5**. A further objection to the modal approach to the logic of knowledge is that, even without **S5** stipulations, implausible and unrealistic assumptions about the reasoning powers of agents still prevail. The reason is related to the, by now, notorious principles of closure. The principle of closure, axiom K (1), can under the certain circumstances be generalized to a stronger closure property of an agent’s knowledge considered still more unacceptable than (1) itself. *Logical omniscience*:

Whenever an agent Ξ knows all of the formulae in a set Γ and A follows logically from Γ , then Ξ also knows A .

In particular, Ξ knows all theorems (letting $\Gamma = \emptyset$), and he knows all logical

¹³See [113], p. 317.

¹⁴From the point of view of autoepistemology, one also suspects that Moore himself would be disinclined to advocate the axiom of negative introspection (axiom 5). Either because it could amount to a Moorean sentence or because it imposes too much rationality on the part of the singular agent—there is a difference between doing the best you can, and then outdoing yourself.

consequences of a formula which he knows (letting Γ consist of a single formula). Logical omniscience incorporates some generally weaker forms of omniscience like: Knowledge of valid formulae: Agent Ξ knows all logical truths (given rule 2); closure under logical equivalence: If agent Ξ knows A and if A and A' are logically equivalent (i.e. $A \leftrightarrow A'$ is valid), then agent Ξ knows A' (given the rule of congruence (3)) and so forth.

There are various ways of dealing with logical omniscience. One is to avoid omniscience technically, another is to attempt to make epistemological or cognitive sense of the property. Especially computer scientists have opted for the latter proposing that what is being modelled in epistemic logic is not knowledge *simpliciter* but a related, however different, concept: The basic epistemic operator $K_{\Xi}A$ should be read as ‘agent Ξ knows implicitly A ’ suggested by Levesque [71], [?]; ‘ A follows from Ξ ’s knowledge’ as discussed by Fagin and Halpern [27]; ‘agent Ξ carries the information A ’ considered by Barwise [6]; or ‘ A is agent Ξ ’s possible knowledge’ as Huang and Kwast propose [58]. Propositional attitudes like these should replace the usual ‘agent Ξ knows A ’. Albeit the variation, the locutions all suggest modelling *implicit knowledge* or what is implicitly represented in an agent’s information state. That is, what logically follows from his actual knowledge.

Furnished with this new interpretation, the epistemic operator really describes cognitive dispositions which potentially could be realized by reflecting upon one’s current information state. The concept of implicit knowledge is totally detached from particular agents’ actual cognitive capabilities. The agents neither have to compute knowledge nor can they be held responsible for answering queries based on their knowledge under the implicit understanding of knowledge. A distinction must accordingly be enforced between what an agent actually knows, the *explicit* knowledge, and the *implicit* knowledge in the sense just described.

Logical omniscience and its various derivatives no longer present problems under the implicit interpretation of the epistemic operator. Now logical omniscience is an epistemological condition for implicit knowledge but the actual agent may nevertheless explicitly fail to realize this condition. While the axioms and inference rules of epistemic logic may seem unreasonable for the explicit view, they may be acceptable for the implicit view of knowledge.

Technical solutions to logical omniscience are either facilitated on the syntactical or semantical level. On the syntactical level, Hintikka recently suggested [50] to place suitable syntactical constraints on deductive arguments which preserve knowledge.

Interesting philosophical solutions are to be found on the semantical level. The idea is here to introduce some semantical entities which account for why the agent could be accused of logical omniscience but by the end of the day is not guilty of logical omniscience. These entities are called ‘impossible, possible worlds’ by Hintikka [49]. Similar entities called ‘seemingly possible’ worlds represented by urn-models are introduced by Rantala [94]. Allowing impossible possible worlds in which the semantic valuation of the formulas in a certain sense is arbitrary provide the necessary means for dodging logical omniscience:

The logical laws do not pass muster in the impossible possible worlds. When knowledge is evaluated with respect to all possible worlds but the logical laws do not hold in some of them, logical omniscience is simply out. In an impossible possible world a tautology $A \rightarrow A$ may, as odd as it admittedly sounds, be false. Now the agent Ξ may all the same view that very world a possibility, so universally $K_{\Xi}(A \rightarrow A)$ fails. In consequence, the rule of necessitation (2) is invalid in impossible possible world models. Axiom K is the victim of failure as well. In the impossible possibility both A and $A \rightarrow A'$ may be true while simultaneously A' is false.¹⁵ From a strictly logical point of view the epistemic logics specified by impossible worlds models are not very exciting. No real epistemic statement is valid in a universal way. The validity of the various epistemic principles may however be obtained by imposing suitable constraints on the impossibly possible models.

From a forcing perspective the introduction of impossible possible worlds is a rather curious strategy. The idea is to first inflate the local circumstances of the agent in the sense that the agent may regard some models of the (real) world possible. Then afterwards deflate the local situation because of the limited reasoning capacities of the agent. The worlds in question are really logically impossible. For example, a logical contradiction cannot be true. An agent may nevertheless not have enough resources to determine the truth-value of that contradiction and simply assume it to be true. He will consider some worlds possible, although logically they are impossible. To avoid logical omniscience let more worlds in, worlds worse than the demon worlds since the latter are at least logically possible whereas the former impossible possible worlds are not.

The key debate of whether epistemic logical principles are plausibly describing agenthood or not seems much to depend on whether one subscribes to a first or a third person perspective on inquiry. Given an autoepistemological inspiration epistemic axioms describe a first person knowledge operator as Hintikka suggested. If epistemic axioms are describing implicit knowledge as Fagin et al. suggest, then what is being modelled is what follows from actual knowledge independently of agent computations. Agents can on this third person perspective not be held actually responsible for failing to exercise some reflective disposition. Logical omniscience is a problem from a first person perspective but not necessarily from a third person perspective in logical epistemology. Closure principles may be problematic from the point of view of the agent, not necessarily from point of view of the ones studying the agent. This allows for *very* strong third person logics in contrast to both counterfactual and modal epistemologies below.

The same goes for assuming **S5** for knowledge in mono-agent systems. It may seem unrealistically strong for a singular agent in his environment, unless his environment is defined solipsistically, e.g. the closed world assumption. Solipsism is not necessarily a human or a real agent condition but a philosophical thesis; a thesis making idealized sense standing outside looking at the agent in

¹⁵The failure of axiom K would satisfy Nozick although he probably would consider impossible possible worlds as weird as demon worlds if not weirder.

his, admittedly, rather limited epistemic environment. Being a stone-hearted solipsist on a first person basis is hard to maintain coherently as W.H. Thorpe once reported:

Bertrand Russell was giving a lesson on solipsism to a lay audience, and a woman got up and said she was delighted to hear Bertrand Russell say he was a solipsist; she was one too, and she wished there were more of us. [108], p. 224.

A reason for adopting the third person perspective and pay homage to **S5** for singular agents is that these assumptions provide some nice technical advantages/properties. There is now also a philosophical basis for doing things in this idealized way—epistemic solipsism and no false beliefs, e.g. infallibilism. Both of these philosophical theses have little to do with logic but plenty to do with the preconditions for studying knowledge from any point of view.

The upshot of this discussion is that the distinction between first and third person perspectives is just as important in a formal epistemology like epistemic logic as it is in mainstream epistemology. For an example, criticizing epistemic logic for implying logical omniscience may just turn on the apples and oranges situation again if the first-third dichotomy is not observed.

Besides the interpretational difference, another important difference between alethic logic and epistemic logic is the addition of the agent Ξ to the syntax. The interesting epistemological question is what roles are assigned the agents in the first generation epistemic logic. The agents are the ones who apparently have knowledge which is, say, **S4.3** valid. That agents hold the knowledge is also the natural understanding of the symbolic notation $K_{\Xi}A$:

Epistemic logic begins as a study of the logical behavior of the expression of the form ‘ b knows that.’ One of the main aims of this study is to be able to analyze other constructions in terms of ‘knows’ by means of ‘ b knows that.’ The basic notation will be expressed in the notation used here by ‘ K_b .’ This symbolization is slightly misleading in that a formula of the form K_bS the term b for the agent (knower) is intended to be outside the scope of K , not inside as our notation might suggest. [54], p. 2.

There is only one role left to agents in the ‘first generation epistemic logic’. They serve as indices on the accessibility relation between possible worlds. Epistemological principles or axioms building up modal systems are relative to an agent whom may or may not validate these principles. Indices on accessibility relations will not suffice for epistemological and cognitive pertinence simply because there is nothing particularly epistemic about being indices. The agents are *inactive* in the first generation epistemic logic.¹⁶

¹⁶ Active and inactive agenthood was first discussed in a paper ‘Active Agents’ and the current discussion follows along these lines [40]. Reference to the agent is sometimes dropped in the formalism of epistemic logic such that $K_{\Xi}A$ becomes KA and is read ‘It is known that A ’ exactly due to the inactive nature of first generation agents. See for instance Hintikka [52].

If epistemic logics are not to be pertinent to the knower who are they to be pertinent to? An agent may have knowledge which is **S4.3** valid. What bakes the epistemological noodle however is *how* the agent has to *behave* in order to gain the epistemic strength that he has. We need to activate the agents in order to make epistemic logic pertinent to epistemology, computer science, artificial intelligence and cognitive psychology. The original symbolic notation of a knowing agent also suggests this: An agent should be inside the scope of the knowledge operator—not outside. Inquiring agents are agents who read data, change their minds, interact or have common knowledge, act according to strategies and play games, have memory and act upon it, follow various methodological rules, expand, contract or revise their knowledge bases, etc. all in the pursuit of knowledge. Inquiring agents are *active agents*.

This is an interpretation of one of the characterizing features, and great virtues of, what Hintikka calls the ‘second generation epistemic logic’ [52]: The realization that the agents of epistemic logic should play an active role in the knowledge acquisition, validation and maintenance processes. In [51], [52] and elsewhere, Hintikka observes this obligation by emphasizing the strategies for his new application of epistemic logic as a logic of questions and answers and the search for the best questions to ask. In this new setting, logical epistemology augmented with an independence-friendly logic constitute the basis for an interrogative theory of inquiry.¹⁷ Answers to questions are in essence requests for knowledge, information or epistemic imperatives [52], [53]:

Another main requirement that can be addressed to the interrogative approach – and indeed to the theory of any goal-directed activity – is that it must do justice to the strategic aspects of inquiry. This requirement can be handled most naturally by doing what Plato already did to the Socratic *elenchus* and by construing knowledge-seeking by questioning as a game that pits the questioner against the answerer. Then the study of those strategies of knowledge acquisition becomes another application of the mathematical theory of games ... [53], pp. 10–11.

Game theory is about strategies for winning games—and it is an agent whom may or may not have a winning strategy among other agents. van Benthem, Fagin, Halpern, Moses and Vardi, Aumann, Stalnaker and others studying game theory have demonstrated how logical epistemology uncovers important features of *agent rationality*. They also show how game theory adds to the general understanding of notions like knowledge, belief and belief revision.¹⁸ Belief

¹⁷Independence-friendly logic (or IF-logic for short) is a first-order logic augmented with an independence operator ‘/’. The slash notation for a quantified statement of the form Q_2y/Q_1x expresses the independence of the two quantifiers. This independence may be captured by game-theoretical semantics as informational independence in the sense that the move performed or mandated by Q_2y is independent of the move performed Q_1x . Introducing the independence operator then allows for the unequivocal formulation of a fan of questions and answers without scope ambiguity, cross-world identity problems etc.

¹⁸van Benthem has also pointed out that there is an epistemic logic hidden in game-theory [8].

revision theorists like Rott model ‘informational economy’ or ‘conservatism’ and consider cognitive economics and the problem of rational choice for *agents* [98]. Baltag, Moss, Solecki combine epistemic logic with belief revision theory to study actions and belief updates in games [5]. Another way to add an active perspective to epistemic logic is pursued in non-monotonic logic starting notably with Reiter’s default logic [96] and R.C. Moore’s autoepistemic logic [83], [84], [85]. These logics model various reasoning patterns based on the agent’s *prima facie* assumptions about the world or model reasoning of the agent’s own current knowledge and belief.

The idea of combining the static first generation epistemic logic with the dynamics of belief revision theory dates back to the mid 1990’s. Alchourrón, Gärdenfors and Makinson’s seminal belief revision theory (AGM) from the 1980’s is a theory about the rational change of beliefs for expansions, contractions and revisions in light of new (possibly conflicting) evidence [1], [30]. In 1994 de Rijke showed that the AGM-axioms governing expansion and revision may be translated into the object language of dynamic modal logic [97]. Segerberg about the same time demonstrated how the entire theory of belief revision could be formulated in a modal logic.

A bit before but especially around the turn of the millennium Segerberg merged the static first generation doxastic logic with the dynamics of belief change into ‘dynamic doxastic logic’ [100]. Doxastic operators in the logic of belief like $B_{\Xi}A$ may be captured by AGM in the sense that ‘ A is in Ξ ’s belief-set T ’, or $\neg B_{\Xi}\neg A$ becomes ‘ $\neg A$ is not in Ξ ’s belief-set T ’. Similarly for other combinations of the belief operator with negation. An immediate difference between the two paradigms is that while AGM can express dynamic operations on belief-sets like expansions (‘ A is in Ξ ’s belief-set T expanded by D ’, i.e. $A \in T + D$), revisions (‘ A is in Ξ ’s belief-set T revised by D ’, i.e. $A \in T * D$), and contractions (‘ A is in Ξ ’s belief-set T contracted by D ’, i.e. $A \in T - D$), no such dynamics are immediately expressible in the standard language of doxastic logic. On the other hand, action languages include operators like $[\nu]$ and $\langle \nu \rangle$ which prefixed to a well-formed formula A , $[\nu]A$, respectively $\langle \nu \rangle A$ on Segerberg’s interpretation mean that ‘after [*every*] [*some*] way of performing action ν it is the case that A ’. By introducing three new operators $[+]$, $[*]$, and $[-]$ into the doxastic language the three dynamic operations on belief-sets may be rendered as $[+D]B_{\Xi}A$, $[*D]B_{\Xi}A$ and $[-D]B_{\Xi}A$.

After revising the original belief revision theory such that changes of beliefs happen in ‘hypertheories’ or concentric spheres enumerated according to entrenchment Segerberg has provided several axiomatizations of the dynamic doxastic logic together with soundness and completeness results [101], [102]. The dynamic doxastic logic paradigm may also be extended to iterated belief revision¹⁹ as studied by Lindström and Rabinowicz in [79] and accommodate various forms of agent introspection [80].

A related approach drawn up by notably van der Hoek, Linder and Meyer

¹⁹ A change in beliefs may either occur once in which case it is a one-shot revision or multiple changes may successively occur in which case it is an iterated revision.

at approximately the same time as the dynamic doxastic logic establishes a new way of distinguishing knowledge from belief [55]. Actions are responsible for bringing about changes of belief. The distinction between knowledge and belief is not just the static feature of paying homage to axiom T or not but likewise the dynamic features of defeasibility or indefeasibility. Knowledge is indefeasible under the belief-changing operations of expansion, revision, and contraction, belief is not. van Ditmarsch, van der Hoek and Kooi's new 'dynamic epistemic logic' is partly a continuation of this approach which studies how information changes and how actions with epistemic impact on agents may be modelled [57].

Attention has so far been restricted to mono-agent systems. The logics of knowledge have been developed with a single agent as the object. To get more feel for the new dynamics and possibilities of modern epistemic logics the mono-agent setup will be expanded to a multi-agent setup.

Syntactically the idea is to augment the language of propositional logic with n knowledge operators, one for each agent involved in the group of agents under consideration. The primary difference between the semantics given for a mono-agent and a multi-agent semantics is roughly that n accessibility relations are introduced. A modal system for n agents is obtained by joining together n modal logics where for simplicity it may be assumed that the agents are homogenous in the sense that they may all be described by the same logical system. An epistemic logic for n agents consists of n copies of a certain modal logic. In such an extended epistemic logic it is possible to express that some agent in the group knows a certain fact, that an agent knows that another agent knows a fact etc. It is possible to develop the logic even further: Not only may an agent know that another agent knows a fact, but they may all know this fact simultaneously. From here it is possible to express that everyone knows that everyone knows that everyone knows, that That is *common knowledge*.

A convention would hardly be looked upon as a convention if it was not for common knowledge among the agents to observe the convention as Lewis once noted. Other norms, social and linguistic practices, agent interactions and games presuppose a concept of common knowledge. A relatively simple way of defining common knowledge is not to partition the group of agents into subsets with different common 'knowledges' but only to define common knowledge for the entire group of agents. Once multiple agents have been added to the syntax, the language is augmented with an additional operator C . CA is then interpreted as 'It is common knowledge among the agents that A '. Well-formed formulas follow the standard recursive recipe with a few, but obvious, modifications taking into account the multiple agents. An auxiliary operator E is also introduced such that EA means 'Everyone knows that A '. EA is defined as the conjunction $K_1A \wedge K_2A \wedge \dots \wedge K_nA$.

To semantically interpret n knowledge operators, binary accessibility relations R_n are defined over the set of possible worlds W . A special accessibility relation, R° , is introduced to interpret the operator of common knowledge. The relation must be flexible enough to express the relationship between individual and common knowledge. The idea is to let the accessibility relation for C be the transitive closure of the union of the accessibility

relations corresponding to the singular knowledge operators. The model \mathbb{M} for an epistemic system with n agents and common knowledge is accordingly a structure $\mathbb{M} = \langle W, R_1, R_2, \dots, R_n, R^\circ, \varphi \rangle$ where W is a non-empty space of possible worlds, $R_1, R_2, \dots, R_n, R^\circ$ are accessibility relations over W for which $R^\circ = (R_1 \cup R_2 \cup \dots \cup R_n)^\circ$ and φ again is the denotation function assigning worlds to atomic propositional formula $\varphi : atom \rightarrow P(W)$. The semantics for the Boolean connectives remain intact. The formula $K_i A$ is true in a world w , i.e. $\mathbb{M}, w \models K_i A$ for agent i , iff $\forall w' \in W : \text{if } R_i w w', \text{ then } \mathbb{M}, w' \models A$. The formula CA is true in a world w , i.e. $\mathbb{M}, w \models CA$, iff $R^\circ w w'$ implies $\mathbb{M}, w' \models A$. Varying the properties of the accessibility relations R_1, R_2, \dots, R_n as described above results in different epistemic logics. For instance system **K** with common knowledge is determined by all frames while system **S4** with common knowledge is determined by all reflexive and transitive frames. Similar results are possible to obtain for the remaining epistemic logics [27].

A dynamic embodiment of multi-agent systems is found in [27]. In such a multi-agent system each individual agent is considered to being in some *local state*. This local state holds all the information available to the individual agent ‘now’. The whole system as the sum of the local agents is in some *global state*. A system like this is a dynamic entity given the global state of the system and local states of the involved agents for any particular time. The dynamics may be modelled by defining what is referred to as a *run* over the system which really is a function from time to global states. The run may in consequence be construed as an account of the behavior of the system for possible executions. This gives rise to *points* which are pairs of runs and times. For every time, the system is in some global state as a function of the particular time. The system may be thought of as series of runs rather than agents. What is being modelled here are the possible behaviors of the system over a collection of executions.

In multi-agent systems like the one just described it is possible to endow the agents with *epistemic capacities* facilitating special epistemic behaviors. Fagin, Halpern, Moses and Vardi consider ‘perfect recall’: Interacting agents’ knowledge in the dynamic system may increase as time goes by but the agents may still store old information. The agent’s current local state is an encoding of all events that have happened so far in the run. Perfect recall is in turn a methodological recommendation telling the agent to remember his earlier epistemic states.

There are other interesting structural properties of agents being studied in the literature of dynamic epistemic logics. In an epistemic logic suited for modelling various games of imperfect information van Benthem refers to such properties as ‘styles of playing’ [8]. Properties like ‘bounded memory’, various ‘mechanisms for information updates’ and ‘uniform strategies’ are analyzed in [10],²⁰ perfect recall and ‘no learning’ are studied by van Ditmarsch et al. as they relate to the change of knowledge given the execution of certain plans [57]. These and other properties of the agents are making them active agents.

²⁰ For an excellent survey of the logic in games refer to van Benthem’s recent lecture notes [9].

The modelling ‘record’ for the second generation logical epistemology is impressive: Multiple epistemic operators, multiple doxastic operators, common knowledge operators, temporal operators, mono-modal systems, multi-modal systems, dynamic modal systems, epistemic capacities of active agents, and this is not an exhaustive list. There is a vast fan of interesting applications and modelling using these advanced epistemic logics. Examples range from robots on assembly lines, social and coalitional interactions, card games, ‘live’ situations in economics, miscellaneous linguistic practices and so on.

The impressive expressiveness of the new epistemic logics re-invites some philosophically problematic issues—notably the **S5** stipulation. It is tricky to argue that **S5** is reasonable to assume in multi-agent setups. But when game theorists for instance model non-cooperative extensive games of perfect information an **S5** logic of knowledge is used to establish the backward induction equilibrium [12].

For game theory the untenability of **S5** in multi-agent systems is quite severe. The problem concerns the knowledge of action as Stalnaker has pointed out:²¹ It should be possible for a player Ξ to know what a player Θ is going to do. For instance it should be rendered possible in case Θ only has one rational choice, and Ξ knows Θ to be rational, that Ξ can predict what Θ is going to do. This should not imply however that it is impossible for Θ to act differently as he has the capacity to act irrationally. In order to make sense of this situation what is needed is a counterfactually possible world such that (i) Θ acts irrationally, but (ii) is incompatible with what Ξ knows. Now Ξ ’s prior beliefs in that counterfactual world must be the same as they are in the actual world for Θ could not influence Ξ ’s prior beliefs by making a contrary choice (by definition of the game, Ξ and Θ act independently). Then it has to be the case in the counterfactual world, that Ξ believes he knows something (e.g. that Θ is irrational) which he in fact does not know. This is incompatible with **S5**.

For single-agent systems the **S5** assumption is, at least to some extent, philosophically defensible as a stipulation about epistemic solipsism and infallibilism from a third person perspective. Even from a third person point of view, in multi-agent systems agents have to reason and act in the presence of, and based on, what yet other agents know and decide to do given again their information about these other agents. In this setting it becomes harder to maintain that the world shuts down with the individual agent exactly because he is not alone.

There is finally an obstacle tagged on to this problem which is more conceptual than technical. If what is being modelled by logical epistemology is the implicit knowledge of the respective agents how can this knowledge, and knowledge about the other agents, be common explicit knowledge to the group of agents: Whatever the singular agent knows about himself and others is not necessarily explicit to him. Common knowledge would have to be explicit otherwise it would seem that the agents base their actions on information they don’t really, or at least readily, have available.

²¹This problem was touched upon by Stalnaker in [106], and I’m indebted to him for explaining it to me in detail later during a conversation.

3 Counterfactual Epistemology

According to the counterfactual theories of knowledge, epistemology begins with facing the beastly skepticism. No knowledge possession, acquisition let alone maintenance before skepticism's claim about the impossibility of knowledge is defeated. To get epistemology off the ground it must be demonstrated that knowledge is in fact possible and this holds whether a formal or informal approach is adopted:

Our task here is to explain how knowledge is possible, given what the skeptic says that we do accept (for example, that it is logically possible that we are dreaming or are floating in a tank). [87], p. 355.

This is the starting point for the counterfactual epistemology developed by Dretske [24] and later refined by Nozick [87].²²

The often cited premiss in favor of the skeptical conclusion that agents do not know much of anything is this: If the agent cannot be guaranteed to be able to know the *denials* of skeptical hypotheses, then the agent cannot be ascribed knowledge on any other issues. The traditional understanding of infallibilism counting every possible world as relevant supports the pessimistic premiss presented. Some arbitrary skeptical hypothesis is a possibility of error the falsity of which must be known to the agent for him to acquire knowledge of some other common hypothesis in question. The inability to know the denials of skeptical hypotheses suffice for lacking knowledge of the ordinary hypotheses.

In order for skepticism to get the premiss off the ground skeptics mobilize the set of all possible worlds—the weird worlds included. In existence is however a thesis which is logically weaker than traditional infallibilism but all the same suffices for supporting the pessimistic premiss. It suffices that knowledge is *closed* in the following sense:

$$\text{If } \Xi \text{ knows } h, \text{ and knows that } (h \rightarrow h'), \text{ then } \Xi \text{ knows } h' \quad (16)$$

which is equivalent to (1) of logical epistemology. The epistemologically optimistic counterfactual idea is eventually that if the weaker closure condition can be denied, then infallibilism may be rejected as a matter of assumption.

The classical thesis of infallibilism supports the skeptical premiss by the demand that Ξ should be capable of knowing the denials of all the possibilities of error. The closure condition (16) demands that Ξ only is knowledgeable of the denials of those possibilities of error which in effect are known logical consequences of Ξ 's knowledge.²³ Suppose Ξ knows the hypothesis that he is currently sitting reading this article on forcing epistemology. Let it also be the case that Ξ knows that if he is sitting reading this paper, then he is not being

²² For a thorough investigation of the relationship between Dretske's and Nozick's approaches to counterfactual epistemology, refer to Gundersen [36].

²³ ... or perhaps rather known logical consequences of Ξ 's knowledge – including denials of all possibilities of error (the so-called contrast consequences, Dretske [24]).

fooled by the Cartesian demon. Then Ξ must also know that he is not being fooled by the demon.

If Ξ does not know that he is not being deceived by the demon then, given Ξ knows the implication, Ξ in turn lacks knowledge of the hypothesis that he is sitting reading forcing epistemology. Now this is exactly what the pessimistic premiss pushes for. But Ξ can know that he is sitting reading this article without knowing that there is no demon of deception seducing him into the false belief that he is sitting reading this paper. Being seated reading this paper implies that no Cartesian demon is leading Ξ to falsely believe that he is reading this very article. It follows that knowledge is not closed in the sense of (16) according to counterfactual epistemology.

Having denied the condition of closure the epistemological mission is still not completed. An explanation must still be provided describing how knowledge of common hypotheses is possible joined with an explanation of the failure to know the denials of skeptical hypotheses. This also goes for the situations in which it is known that the common hypothesis at issue implies relevantly rejecting the skeptical hypothesis.

Dretske's solution is to install a modal condition for knowledge imposing truth-conduciveness by *sensitivity*:

$$\text{If } h \text{ were not true, } \Xi \text{ would not believe } h. \quad (17)$$

A belief qualifying as knowledge is a belief which is sensitive to the truth: The hypothesis h is true in accordance with the standard definition of knowledge. Had h which is believed been false, the agent would not be led to the belief that h .

Condition (17) readily explains why closure fails. Proximity relations between possible worlds are introduced due to the semantics for the inserted subjunctive conditional. One may know both antecedents h and $(h \rightarrow h')$ relative to one set of relevant worlds accessible from the actual world, and yet fail to know the consequent h' relative to a different set of possible worlds. Now relative to a set of possible worlds with proximity 'close' to the actual world one knows h and simultaneously knows that h implies the denial of the skeptical hypothesis, say h' . But one may all the same fail to know the consequential denial of the skeptical hypothesis itself for knowledge of the skeptical hypothesis is relative to possible worlds with a 'way-off' proximity to the actual world. These possible worlds are radically different from the actual world by all means. 'Way-off' worlds are accordingly forced out, skepticism far away because closure fails, but the possibility of knowledge prevails.

In the monumental monograph on knowledge, skepticism, free will and other pertinent philosophical issues [87], Nozick formulates a definition of counterfactual knowledge along the Dretsian lines. The proposal furnishes virtually the same answer to the skeptic and also solves the original Gettier-problems.

Nozick's idea is to use (17) to efficiently solve the notorious Gettier-paradoxes ([31]): If it were not the case that Jones owns a Ford car, then Smith would not be believing that Jones indeed does. In contrast, suppose now that Jones in fact

has a Ford. The hypothesis h_1 regarding Jones' ownership of a Ford car would be true. It may be the case that Smith in fact believes h_1 . It does not follow, however, that Smith's believing h_1 in any way is connected or sensitive to the truth of h_1 . The proposal is augmented with yet another subjunctive conditional to the effect that if it were the case that Jones has a Ford, Smith would be led to believe that Jones does own one such automotive vehicle. Nozick's definition of counterfactual knowledge may be formulated as follows:²⁴

Ξ knows h iff

1. h is true,
2. Ξ believes that h ,
3. $\neg h \mapsto \neg(\Xi$ believes that $h)$,
4. $h \mapsto (\Xi$ believes that $h)$.

The truth and the belief conditions of the standard tripartite definition of knowledge are retained but the justification condition is dropped. Then again not exactly. Conditions (3) and (4) of the definition actually provide justification when justification is understood along reliable methodological lines as will become more apparent below.

To see how the definition works, the possible world semantics provides the following account of the truth-conditions for the subjunctive conditional: A subjunctive

$$A \mapsto B$$

for arbitrary statements A and B , is true, insofar, in all those worlds in which A is true that are in proximity 'closest' to the actual world, B is also true in these 'closest' worlds. More specifically of three worlds w, w', w'' if w' is closer to w than w'' , then $A \mapsto B$ will be true in w iff A is not true in any world or there exist a world w' in which A and B are true which is closer to w than any world w'' in which A is true but B is false.²⁵

The criterion of closeness or proximity in possible worlds semantics is in general troublesome. It hinges on a notion of similarity between possible worlds. The more similar one world is to another, the closer it is. Usually similarity is cashed out in terms of *ceteris paribus* clauses: The less tinkering with a possible world, everything else being equal, the more similar that world is and accordingly the closer it is to the actual world. From this perspective, floating around in a tank on a planet called Alpha Centauri, requires *ceteris paribus* more conditions to be changed than the world in which you were born in Copenhagen rather than New York City. In effect, the beliefs an agent may have on Alpha

²⁴' \mapsto ' denotes the subjunctive conditional.

²⁵This semantic account of the subjunctive follows rather closely Lewis in [75]. Nozick is however not committed to a particular understanding of the semantics and also discusses Stalnaker's subjunctive semantics from [104]. See furthermore [87], p. 680, footnote 8.

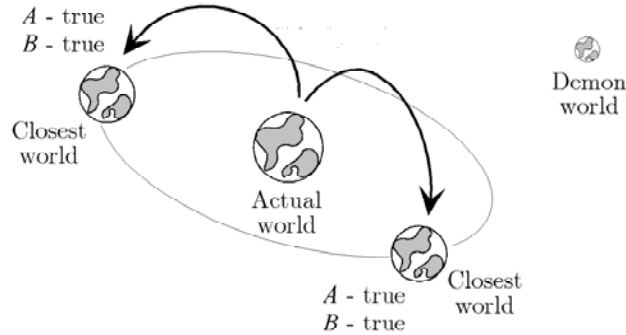


Figure 3: Counterfactual epistemology requires categorical success in *close* possible worlds

Centauri may sound like a far cry from the beliefs the agent would have had he just been born somewhere closer to home.²⁶

For knowledge possession, one does not have to consult all possible worlds as the skeptic would insist: Given the standard semantical analysis of the subjunctives it is enough that the consequent *B* holds in those possible worlds which are closest to the actual world such that the antecedent *A* holds (Figure 3). Speaking 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 third condition of the definition above is there to avoid error. The fourth is there to gain truth. The two conditions are collapsible into one condition: Ξ 's belief tracks the truth of *h*:

To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it. [87], p. 178.

The method resolving the question of whether *h* is a combined verification and refutation procedure and thus a *decision* procedure: Ξ must output *h* if *h* is true and must output $\neg h$ if *h* is false. The justification condition is absent in Nozick's definition of knowledge. When methodology exhausts justification as the study of truth-conducive methodological recommendations the justificational condition is re-introduced in terms of a recursive procedure in all close worlds.

Consider a common context in which brigadier-general Ξ is of the belief that his standard issued sidearm is in its holster. Skepticism would now have

²⁶ Some attempts have been made to specify what is meant by similarity and closeness. But it remains an open problem in philosophy and possible worlds semantics. Lewis has argued that our notion of similarity is notoriously vague and by its very nature resists specification [75].

it that Ξ dumps this belief because a demon may have set things up in a way responsible for Ξ believing that his gun is in its holster when in fact this is not the case. Reasoning in accordance with counterfactual epistemology yields the following result: Suppose that the closest possible world to the actual world in which Ξ 's gun is not in its holster is a world where no diabolical creature is around. In this case general Ξ would not perceive the gun to be in its holster as Ξ 's senses are still working normally in this world (*ceteris paribus*). Ξ would not entertain the belief that his gun is in its holster.

The counterfactual epistemological account of knowledge also guarantees that if Ξ has taken up the challenge of reading this article, then he knows that he is reading this article when he really is: All the conditions 1–4 of the definition are met with satisfaction: It is true that Ξ is sitting reading this article and he believes it to be so. If he were not reading this paper he would not be believing that he is. In the end, if Ξ is sitting reading this paper, he will believe that he is reading this article on *forcing epistemology*:

- Condition 3 ensures that in the closest possible worlds in which Ξ is not reading this article, he does not believe that he is.
- Condition 4 guarantees that in the closest possible worlds in which Ξ is reading this article, he believes that this is so.

Everyday knowledge is secured in these contexts. The idea of introducing the proximity relation is that the agent's local epistemic environment normally suffices for the truth witnessing Nozick's first person stance.

It is on the other hand impossible for Ξ to know that he is not this brain in a vat. Assuming the brain receives the same sensory patterns as it would was it not dumped in the vat, there would not be anything in the input revealing to Ξ that he was not a brain in a vat. In this devious scenario Ξ is also barred from knowing that he is sitting reading this paper on forcing. If Ξ claims to know that he is sitting reading this article, it must follow that he as a prerequisite tacitly approves of the hypothesis that he is not a brain in a vat. Given this prerequisite and *modus tollens* as Ξ does not know that he is not sunk into the vat he does not know that he is sitting reading this paper either.

Now the possible world in which Ξ is a brain in a vat is *ceteris paribus very* distant from the actual world. Failure of knowledge in these cases is not devastating to counterfactual epistemology. It hinges on the relevant possibilities of error. True beliefs are only required in possibilities closer to actuality than any $\neg h$ -possibilities: Picture a physicist measuring the voltage drop over some LRC-circuit. A student from epistemology class comes to him and asks whether a relevant possibility of error could be that the voltmeter is calibrated incorrectly. The physicist would probably answer 'yes' as calibration problems could lead to a measurement error. Then asking the scientist whether being a brain in a vat is a relevant possibility of error would likely result in the physicist asking the student to go back to his course and stop bothering him with silliness.

Gettier-cases are situations where the evidence available might lead to mistake. Knowledge is intimately related to an interest in gaining truth and avoid-

ing error. The method envisioned by Nozick for knowledge acquisition is a reliable method. The method is reliable because of its recursive workings in close other worlds besides the actual one. Close worlds in which the belief is true are important and close worlds in which the belief is false matter likewise. Relevant possibilities or error are belief-dependent. Reliability is categorical on Nozick's account as a result of a decision procedure operating in all close worlds. In this sense the proposal is normative in that avoiding Gettierization and skepticism require strong epistemic prescriptions to be met.²⁷ Reliability is a world-dependent property insofar as it is sensitive to which possible world is the actual world. Reliability is a belief-dependent property to the extent that the relevant possibilities of error depends on the belief. A method of inquiry cannot be determined to being *a priori* reliable on the counterfactual account. Reliability is sensitive to which world is actual and one may not know which world this is.

Counterfactual epistemology has close shaves with a formal approach like logical epistemology. Recall that epistemic logic devised a sliding scale for assessing the validity and epistemic strength of knowledge operators. To recapitulate, the scale is built up by axioms which in proper combinations make up modal systems of different strength in the following sense: There are propositions valid in one system not valid in another system for which less (or different) algebraic properties are assumed for the accessibility relation. The differences in strength are determined by which relational properties the accessibility relation between possible worlds is endowed with relative to the knowledge operators' scope. The reasonable property that a possible world is accessible from itself, i.e. that the relation between possible worlds is reflexive allows one, given the standard setup of modal propositional logic, to validate the axiom **T** which with an epistemic interpretation amounts to

$$\text{If } \Xi \text{ knows } h, \text{ then } h \text{ is true.} \quad (18)$$

(18) is the default epistemic assumption of far the majority of epistemologies since Plato. It is the second condition of the standard tripartite definition of knowledge, and the first condition of Nozick's definition of counterfactual knowledge.

Validating (18) together with two other principles again results in the modal system **T**. Now, if it is further assumed that the accessibility relation is both reflexive and transitive, then it is possible to validate the following axiom 4 or the *KK*-thesis:

$$\text{If } \Xi \text{ knows } h, \text{ then } \Xi \text{ knows that he knows } h. \quad (19)$$

Nozick rejects (19) for counterfactual knowledge:

Some writers have put forth the view that whenever one knows, one knows that one knows. There is an immediate stumbling block to this, however.

²⁷ It is often argued that the counterfactual definition of knowledge requires so strong epistemic prescriptions that merely mortal agents cannot meet them, cf. [26].

One may know yet not believe one knows; with no existing belief that one knows to do the tracking of the fact that one knows, one certainly does not know that one knows. [87], p. 246.

An agent may be tracking the truth of h without tracking the fact that he is tracking the truth of h . The stronger epistemic principle

$$\text{If } \Xi \text{ does not know } h, \text{ then } \Xi \text{ knows that he does know } h \quad (20)$$

was referred to as axiom 5, the axiom of wisdom or the principle of negative introspection in logical epistemology. An agent validating (18), (19) and (20) is said to have epistemic strength on the order of the modal system **S5** requiring an equivalence relation over possible worlds. (20) is valid in **S5** but not in **S4**, so **S5** is a genuinely stronger epistemic system than **S4**. Even though the principle of negative introspection is not discussed by Nozick it seems very likely that he would dismiss it all the same. If an agent is not tracking the truth of h it does not follow that he will be tracking the fact that he is not tracking h .

Accepting (18) while rejecting (20) and (19) result in a rather weak epistemic logical system. Counterfactual epistemology furthermore rejects the closure principle (16), i.e. (1). But the simplest modal system, and also the weakest modal system is obtained by augmenting the standard propositional logic with the closure principle. Whatever logic holds for counterfactual knowledge is very weak.

The following rule of necessitation

$$\text{If } h \text{ is a theorem, then } \Xi \text{ knows } h \quad (21)$$

is also adopted by logical epistemology. The rule of necessitation holds on the counterfactual account of knowledge as well. If h is a logical truth in all possible worlds, the fourth verification condition of the counterfactual definition of knowledge is going to ensure that Ξ believes h , and since h never turns false the third condition never comes into play. So h is true, Ξ believes h , and since h is true in all possible worlds, h is also true in close worlds so Ξ knows h .²⁸

Counterfactual and logical epistemology view the importance of the epistemic axioms almost diametrically opposite with respect to the threat from skepticism. A logical epistemologist considers axioms (19) and (20) as ways of blocking the skeptical challenges. An agent may know of his knowledge and sometimes even his ignorance leaving the skeptic with less room to maneuver. The stronger the system the less room for error.

On the other hand it seems that the only logical axiom acceptable to the counterfactual epistemology is the axiom of veridicality (18) together with the rule of necessitation (21). Accepting the axiom and the rule hardly amounts to a logical system at all in any *normal* sense of modal logic. Closure under implication is denied but closure does not require any algebraic properties to be assumed for the accessibility relation between possible worlds. For this reason

²⁸I'm indebted to Robert Stalnaker for bringing this to my attention.

a logical epistemologist always assumes (16) to be valid under normal circumstances. Skepticism is then defeated subsequently by imposing more and more accessibility structure on the set of relevant possible worlds. Nozick pursues the opposite direction. Skepticism is defeated by very limited accessibility between possible worlds. The more access to worlds, the more freedom for the skeptic to move around in (distant) worlds citing apparently relevant possibilities of error. The cost is a very weak first person logic of knowledge but at least knowledge is possible however not closed.

The counterfactual epistemology in general accommodates elements of the contextualistic epistemology of the next section. Dretske's view of the closure lets knowledge transfer work across known implications insofar as the implications in question are close or relevant. Knowing that one is sitting down reading this article transfers immediately through the known implication to the 'close' hypothesis that one is not standing on a street corner doing the same. This knowledge will at the same time not run through the known implication to the 'way-off' hypothesis that one is not being fooled by a malicious demon. Dretske's point seems to be that knowledge acquisition of a hypothesis in some common *context* assumes by default the very falsity of particular 'way-off' and irrelevant possibilities of error [24]. These possibilities of error are skirted, or their falsity presupposed in many everyday knowledge acquisition contexts. Lewis strongly subscribes to this contextualistic forcing feature in his modal epistemology.

4 Modal Epistemology

Contextualistic epistemology, or *modal epistemology*,²⁹ as one may label Lewis' recent version of contextualism [78] starts much closer to home. Agents in their local epistemic environments have knowledge—and plenty of it in a variety of (conversational) contexts. Knowledge is not only possible as counterfactual epistemology demonstrates, it is real human condition.

The general contextualistic template for a theory of knowledge is crisply summarized in DeRose's description of the attribution of knowledge. The description also embodies many of the epistemological themes central to the contextualistic forcing strategy:

Suppose a speaker *A* says, '*S* knows that *P*', of a subject *S*'s true belief that *P*. According to contextualist theories of knowledge attributions, how strong an epistemic position *S* must be in with respect to *P* for *A*'s assertion to be true can vary according to features of *A*'s conversational context. [22], p. 4.

²⁹ Admittedly, Lewis does not refer directly to his theory as modal epistemology. The term figures in reference to Stalnaker's distinction between modal and hyperintensional epistemology [105] (according to which all 'forcing' epistemology is modal epistemology). To differentiate Lewis' build of contextualism from other contextualistic models, modal epistemology is reserved for Lewis' version.

For excellent surveys of contextualism refer to Pritchard [92], [93].

The incentive to take skeptical arguments to knowledge claims seriously is based on an exploitation of the way in which otherwise operational epistemic concepts, notably knowledge, can be gravely disturbed by sudden changes of the linguistic context in which they figure.

The standards for the possession of knowledge vary from context to context depending on what is at stake. In a course on epistemology the standards for knowledge possession fixed by the interlocutors (teacher and students) are usually very high. The conclusions that we know very little, if anything at all, may by the end of class be true. In a discussion after class a fellow student says ‘I know that *Matrix Reloaded* plays in the Park & 86th Street Theater on 125 E. 86th St.’. The circumstances have now changed and the standards for knowledge possession in this new, presumably, non-skeptical conversational context are lower. The relatively lower standards put us in the comfortable position of maintaining that we know most of what we think we know. It is admittedly to this low epistemic standard but it surely suffices for going to the movies.

Not only may knowledge attributions fluctuate with contexts, they may also be sensitive to who ascribes knowledge to whom. As indicated by DeRose there is a delicate issue to be addressed pertaining to the strength of the position an agent has to be in order for the epistemic commitment to truthfully pan out. This position is context-sensitive, not only to the agent in the environment, but also to possible ascribers of knowledge to the very agent in question. The first-third person dichotomy is immanent in contextualistic epistemologies.

Finally, the strength of the epistemic position is responsible for turning the contextualistic theory of knowledge into a modal account according to DeRose. For every local environmental ‘time-slice’ the epistemic position of the agent remains constant. The epistemic position the agent however *were* to be in to warrant possession of knowledge is a subjunctively defined spatio-temporal function of the context. A strong epistemic position with respect to some hypothesis h is to have belief as to whether h is the case and tracking this fact not only through the actual world but through close worlds as well. Maintaining that one’s belief still tracks the truth at long distances increases the strength of the epistemic position with respect to the hypothesis in question. For belief to become knowledge it should be ‘non-accidentally’ true in the actual world and in close ones as well.³⁰ This way of realizing the forcing relation resembles the construction advanced by the counterfactual epistemology of the previous section using sensitivity or tracking.

There are garden varieties of contextualism developed by Cohen [20], [21], DeRose [22], [23], Lewis [78], and recently Williams [112] advanced a new build which diverges somewhat from the standard foundation of contextualism. Cohen’s contextualism focus on justificational issues and the structure of supportive reasons for knowledge possession. Both DeRose and Lewis’ accounts serve to spell out the modal contextualistic program in general. Williams’ version diverges from this program by individuating contexts in terms of inferential

³⁰See further [22], p. 34.

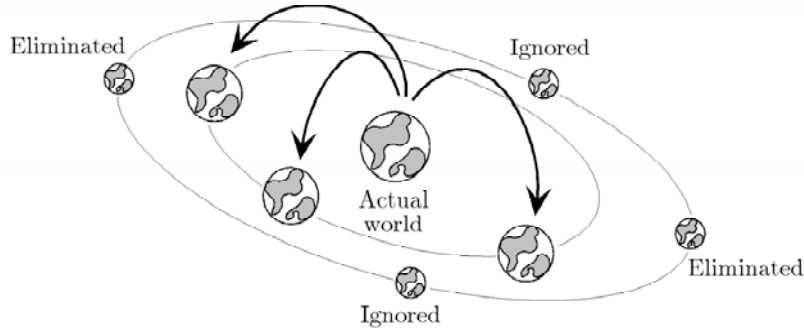


Figure 4: Modal epistemology requires success in all worlds not either eliminated or ignored

structure of a context rather than through conversational measures.

Lewis' modal epistemology is particularly engaging as it balances elegantly between mainstream and formal *modi operandi*. This is not too surprising since Lewis through his career was concerned with modal logic, in particular the logic of counterfactuals [75], modal ontology [77] and almost consequently modal epistemology [78]. Modal logics, epistemic logics in particular, are much about partitioning the set of all possible worlds into classes that are in close proximity, similar, relevant or accessible from the actual world and into classes which are not.

As humans we force for knowledge on a daily basis and obtain it. This means partitioning the set of all possible worlds into relevant, irrelevant and extravagant possibilities of error determined by the current context. To obtain knowledge eliminate the relevant possibilities of error, ignore the extravagant ones, and succeed over the remaining possible worlds where the hypothesis in question is true (Figure 4). Everything dictated by the current context. There are rules for elimination, ignoring and success. On a new definition of knowledge yet to be formulated, these rules are what Lewis' modal epistemology is about.

To kick off the modal epistemology, the standard tripartite definition of knowledge as true justified belief is abandoned. Justification is neither necessary nor sufficient for the possession of knowledge. Instead of citing clairvoyance as a relevant possibility undermining the sufficiency of justification (e.g. Bonjour [15], [16]), something less far-fetched suffices: Observing some consequences of the *lottery paradox*. Formulated by Kyburg [63] the lottery paradox goes to show how otherwise intuitively appealing justificational principles may be rendered logically inconsistent. The paradox rests on the following three assumptions about justification and rational acceptance: (i) One is justified in accepting a conclusion which is likely to be true; (ii) justification is closed under conjunction,³¹ and (iii) one is never justified in accepting logically inconsistent

³¹ If one is justified in accepting h , and justified in accepting h' , one is justified in accepting the conjunction $h \wedge h'$.

statements.

Consider next a fair lottery with a 1000 tickets and one winning ticket. In such a lottery the probability of each particular ticket losing is .999, thus very high. Enumerating the tickets from 1–1000, and examining the probability of ticket 1 should, given assumption (i), lead to the conclusion that the ticket will lose. This line of reasoning can be repeated for ticket 2, 3, ..., 1000. Given the assumption of closure under conjunction one is justified in believing that all 1000 tickets will lose which is the same as committing to the belief that no ticket will win. By the rules of the lottery one ticket will win which *a fortiori* also is a justified belief. One is now faced with being justified in believing the logical inconsistency that no ticket will win and that some ticket will win. This is in conflict with assumption (iii) and the lottery paradox is generated.

Some attempts to retain justification but at the same time avoid the lottery paradox focus on showing that at least one of the initial assumptions (i)–(iii) are incorrect.³² One way out is to reject assumption (i): Justification strong enough to keep the skeptic quiet requires deductive certainty and not likelihoods no matter how high these likelihoods may be. A justification for belief is not licensed if there is just an infinitesimal chance of mistake (cf. Dretske [25]).³³

The way to boost the degree of justification for believing in losing is by increasing the number of tickets in the lottery. The odds of losing may be set arbitrarily high short of 1. In a one winning ticket setup there will always remain the chance of winning admittedly growing smaller and smaller as the number of tickets increases. An arbitrary increase in the number of tickets implies an arbitrary increase in the degree of justification for believing in losing. But an arbitrary high degree of justification will not suffice for converting fallible belief about losing into knowledge of losing. Lewis concludes that justification is not sufficient for knowledge.

The necessity of justification is also dismissed. A glance at the actual epistemic practice is enough to reveal that justification is not always necessary for knowledge:

What (non-circular) argument supports our reliance on perception, on memory, and on testimony? And yet we do gain knowledge by these means. And sometimes, far from having supporting arguments, we don't even know how we know. We once had evidence, drew conclusions, and thereby gained knowledge; now we have forgotten our reasons, yet still we retain knowledge. [78], p. 368.

Knowledge is being stripped of its classically distinguishing features by modal epistemology: True justified belief does not amount to knowledge; knowledge is allowed without justification; thirdly, knowledge of something is not always accompanied by belief in something. Lewis cites Radford's list of cases in which

³² For instance, Kyburg suggests abandoning assumption (ii) in [64].

³³ A more lenient version maintains that high probability does confer justification given additional constraints of say, coherence of belief sets. Justification is not warranted with respect to singular beliefs in losing on individual tickets (cf. Lehrer [66], and Pollock [90]).

an agent may be said to have knowledge but insecurity, lack of confidence, or timidity prevents him from believing what is known [95].

A substitute for the standard tripartite definition of knowledge is in order. It should explain why what is allegedly known is known infallibly. In the ears of modal epistemology knowledge in light of possibilities of error sounds like a contradiction.

Taking infallibility as a basic epistemological condition, for an agent to know a hypothesis, all possibilities of error must be eliminated given the agent's available information. That is, all the possible worlds in which the negation of the hypothesis is the case must be eliminated. This forcing relation is given by different measures. One measure is simply to ignore possibilities *extravaganza*, another is to use the available evidence to force such that the uneliminated possible worlds are determined by perceptual equivalences over these alternatives with the actual world as the fix-point. The perceptual experience (and memory) the agent has in the actual world fixes the set of uneliminated possible worlds insofar the agent's cognitive apparatus functions the same in these worlds. Suppose that a perceptual experience has the propositional content A . The perceptual experience with content A (memory included) eliminates a certain world w' if and only if the content of the experience the agent has in w' differs from A .

Quantifiers are usually restricted to run over some domain of interest. This also goes for the universal quantifier over possible worlds that would lead to error. Every uneliminated world in which the hypothesis holds is restricted to a sub-domain of properly all uneliminated worlds. Saying that the surface is 'clean' in a certain conversational context is to properly ignore the microscopic dust particles laying on the surface. If somebody was to disagree it would have to be because the new interlocutor in the conversational context means clean in a more restrictive sense. The microscopic dust balls in this case suffice for making the assertion about the clean surface false. Words like 'flat' or 'round' behave in the same way, as does the word 'knowledge'. They are context-sensitive.³⁴

Alterations of the conversational context occur when a new hypothesis is introduced which for its part is more demanding than any of the other hypotheses currently explicit in the particular context. Such a *non-uniform* introduction implies an increase in the range of possible worlds to be considered for attribution of knowledge. The strength of the required epistemic position mentioned above is increased accordingly. In a context where the usage of 'knowledge' remains *uniform* throughout the conversation, the range of possible worlds to be considered remains stable. Given the context-sensitive nature of knowledge, in every context where knowledge attribution is at stake some uneliminated possible worlds are not rendered relevant by the *current* context. The universal quantifier is restricted accordingly.

These considerations essentially pave the way for the colloquially stated but forceful knowledge definition of modal epistemology:

³⁴The context-sensitivity of various words including 'knowledge' was noted by Lewis much earlier in [76].

S knows that P iff S 's evidence eliminates every possibility in which not- P —Psst!—except for those possibilities that *we* are properly ignoring. [78], p. 378.




Figure 5: Forcing filters in modal epistemology

Knowledge is on this definition forced through a two-step context-sensitive filtering process – eliminating and ignoring – leaving the hypothesis or proposition true in every remaining possible world determined by the agent's evidence and other relevant information (Figure 5). These are the worlds over which to succeed in order to attribute knowledge. To carry out the filtration, the agent follows prohibitive rules about which worlds may not be properly ignored, and presumptive rules about which worlds may be properly ignored while ascribing knowledge to himself and to others.

The envisioned order of the filtration process is not entirely clear. Reading 'Which of all the uneliminated alternative possibilities may not be properly ignored' [78], p. 371 suggests eliminating before ignoring. One argument for this order would be that elimination fixes a set of perceptually equivalent possibilities after which a subset of this set is ignored. Yet ignoring may also supersede eliminating if the question of what is ignored and what is not ignored is connected to the agent's view of the world. In a particular context it is not a part of agent Ξ 's view of the world that he is a brain in a vat so the elimination is not sensitive to this possibility. The two filtrations may also be so intimately connected making them virtually indistinguishable during the actual epistemic processing.³⁵ For expositional reasons they are separated here.

During the individual knowledge attribution process, the possible world which the agent takes to be the actual state of affairs is never ignored. Actuality is by reflexivity always a relevant possible world alternative although indexical. It follows that falsity may not properly be supposed. If falsity is

³⁵I'm indebted to Louise Westmark for pointing this out.

never to be presupposed whatever in the end will turn up knowledge must be true, so the classical condition of truth for knowledge is derived. Never ignoring the actual world is referred to as the *rule of actuality*.

Turn next to the ascription of knowledge to others. The way in which the modal knowledge definition is given on page 35 italicizes '*we*'. What *we* may properly ignore is going to be dependent on whose actuality is being referred to in the context in question. Assuming that there is only one actual world-index in play in non-modal contexts one should expect that the world considered actual by the agent coincides with the world indexed 'actual' by the ascribers.

In counterfactual situations referring for instance to what an agent would have known today had he read the paper yesterday, or whether an agent knew yesterday who he was then, fixing the index of actuality is more complicated. Had the agent read the paper yesterday he would presumably have known more than he in fact knows today. The agent is ascribing knowledge and ignorance to himself now as the one not having read the paper last night. The ascriber, say Ξ' , of knowledge to agent Ξ has an index of actuality demonstratively different from Ξ 's index. The index on actuality for Ξ' is what Ξ' would have been like knowledge-wise had he read the paper yesterday. Actuality indices differ for Ξ and Ξ' in this situation. Similarly for the attribution of knowledge to Ξ knowing yesterday who he was. For Ξ 's reality is defined for his spatio-temporal history up until yesterday; for Ξ' reality is defined for his spatio-temporal history up to today when the question is popped whether Ξ knew yesterday who he was. The two world stories are different. Ξ 's actuality yesterday is different from Ξ' 's actuality today. Similarly for a host of other situations involving say iterated modal constructions like knowledge of knowledge etc.

The rule of actuality applies both to the ascriber and the ascribed. What may not be properly ignored is the local agent's actuality. Epistemologists considering what Ξ knows from a third person perspective will attend to whatever possible worlds that Ξ himself attends to as possible and then some. The set of possible worlds ignored by a third person knowledge attributor for Ξ will properly be a superset of the possible worlds Ξ ignores. An agent may know more than what may be ascribed to him because his actuality in some cases differs from the ascribers and his range of viable worlds does as well. Applying the principle of 'epistemic' charity means that while attributing knowledge to an agent in his local epistemic environment, the third person ascriber may ignore fewer possibilities than Ξ .

From the actual world a certain set of alternative worlds are *believed* to be accessible by the agent. These may not be properly ignored either. This is the content of the *rule of belief*: A possible world which the agent believes to obtain may not be ignored independently of whether he is justified in believing that the state of affairs obtains or not. The rule of belief is an epistemic correlate to the \Diamond -rule in modal tableaux. A sort of existential quantifier rule introducing other worlds in the modal knowledge tree. Not only may such belief-dependent worlds not be properly ignored but possible worlds which the agent *should* acknowledge to obtain given evidence and other information may not be ignored either. This deontic clause need not be explicitly realized by the agent witnessing an

externalistic bias.

Belief comes in degrees and some possible worlds will be doxastically more accessible than others given the measure of opinion. Some possible worlds will be tossed out because the measure of opinion pertaining to them is too low, infinitesimal or even 0. Whatever worlds may not be properly ignored should have a ‘sufficiently high’ degree of belief attached to them. Sufficiently high is context-dependent and will vary with what is at stake. For contexts in which error would be especially ‘disastrous’ ([78], p. 373) a quite low measure of opinion may do for not tossing out the possibility. To avoid miscarriage of justice in a court of law, the accused is only convicted when guilt has been proved beyond reasonable doubt. In this context a very moderate degree of belief is enough to warrant ignoring fewer possible worlds, though some possible worlds are still ruled out. Other contexts are not as serious. Sufficiently high degree of belief will be higher in these contexts tossing out more worlds. The degree of belief for legitimately ignoring worlds requires a justificational story. Justification is in a certain sense re-invited on the way to knowledge.

Having a rule for introducing new worlds calls for a rule of accessibility between the introduced worlds. A world w' which ‘salient resembles’ another world w enforces a kind of symmetry. If w may not be properly ignored in virtue of the other rules neither may w' and vice versa. This accessibility clause is referred to as the *rule of resemblance*, the last of the prohibitive rules. The rule is dangerous and powerful at the same time.

The rule of resemblance is dangerous because not applied carefully invites skepticism and global underdetermination back in. The actual world is left uneliminated by the agent’s available evidence. It follows that any other uneliminated world resembles the agent’s actual world in one important respect, namely, with respect to the agent’s evidence. This will continue to hold even in worlds which otherwise are radically different from the agent’s actual world including the demon world. By application of the rule of actuality together with the rule of resemblance leads to the conclusion that these worlds are relevant alternative worlds as well!

There is no obvious remedy to this problem and it reappears with respect to knowledge closure. Agreeing with counterfactual epistemology that closure over arbitrary contexts amounts to a fallacy driving skeptical arguments, modal epistemology holds that closure is possible locally without skepticism. Knowledge is closed for a fixed context. Knowing that you are reading this paper implies that you are reading this paper and not being deceived (by a demon or a mad scientist) in this particular uniform context c_1 . If the context is non-uniformly changed right after the antecedent conditions obtain to a new context c_2 , ‘all bets are off’ [78], p. 382:

$$\underbrace{K_{\Xi}h \rightarrow K_{\Xi}(h \rightarrow h')}_{c_1} \xrightarrow{\text{change!}} \underbrace{\neg K_{\Xi}h'}_{c_2} \quad (22)$$

Closure fails because the strength of the epistemic position now required in c_2 to attribute knowledge has been increased way beyond c_1 by the increase in

possible worlds at issue dictated by c_2 . The range of possible worlds may now include the demon world which is a whole different context. Knowledge is closed under implication because implication preserves truth in a fixed context not over arbitrary contexts.

Similarly to Lewis, DeRose's contextualistic closure argument relies on the Dretsikian sensitivity or Nozickian tracking of the previous section. For closure the problem is that sensitivity increases the required epistemic position because of the switch in the possible world scope due to the contextual change. For an agent to know that he is not a brain in a vat, sensitivity or tracking demands that if being a brain in a vat is true, the agent would not believe that he was not so. This takes the issue to another level in terms of epistemic position. The required epistemic position is boosted because the agent has now to believe that he is a brain in a vat in a world where his brain in fact is rigged to an experience simulating computer. The closest such world is *ceteris paribus* very remote from the actual world. In the actual world the agent is barred from finding out that his brain is sunk into the tank of fluids and being stimulated. When sensitivity elevates the standards for such hypotheses it raises it for all other beliefs in the same context. In a world in which the agent is a brain in a vat, the agent will believe that he is reading this article even though this in fact is false. When the agent entertains a proposition and a skeptical hypothesis at the same time the standards for sensitivity are raised to such an exorbitant level that common hypotheses are unknowable.

There is not any obvious way to ensure that such a contextual change is not taking place. The rules of actuality and resemblance combined immediately permit for such a change to occur. The demon world resembles saliently the actual world with respect to agent's evidence and should accordingly not be ignored. Lewis readily admits to an *ad hoc* modification of the rule as to exclude this resemblance.

Although the rule of resemblance may lead to error in certain circumstances, its forcefulness is demonstrated by its ability to deal with both the lottery paradox and the Gettier-cases. The rule reveals why the agent is barred from knowing losing the lottery independently of the arbitrary high odds of losing in turn responsible for the agent concluding that he will lose. Every ticket describes a possibility of winning. Every one of these possible worlds resembles the other with respect to this feature. By the symmetry of resemblance all of them should be ignored or none of them should be tossed out. The trouble is then that the winning ticket describes a possible world which may not be properly ignored.

Gettier-cases are handled in a similar fashion. Smith has ignored the possibility that Jones drives a rented car as much as he has ignored the possibility that Brown is in Barcelona. By assumption the actual world is not ignored. The Gettier-scene world matches the actual world as far as Smith is concerned. The error is committed by ignoring a world which the accessibility relation of resemblance requires you to take into account. Smith has knowledge for the wrong reasons and consequently does not know having been too eager to ignore. The story is the same with the paper-marche barns. Seeing a real barn in the country of mixed barns does not license a claim to knowledge about seeing a

real barn. The possibility of seeing a paper-marche version must not be ignored because of the salient resemblance.

The failure to know of a real barn could also be due to the mode of acquisition, i.e. that visual perception is unreliable. This possibility is however emphatically denied in modal epistemology. In fact the possible worlds in which say perception, memory and testimony are unreliable may be properly ignored. This is content of the *rule of reliability* which, as opposed to the three discussed above, is a presumptive rule about which worlds may be tossed out or ignored.

It is not entirely clear what sort of reliability is envisioned besides the fact that it has a modal scope. The three cognitive processes (perception, memory and testimony) are described as ‘fairly reliable. Within limits we are entitled to take them for granted.’ [78], p. 375. This may suggest a variant of reliability with stochastic success over other possible worlds with a high truth-ratio. To determine ‘fairly reliable’ references are then made in a footnote to Armstrong’s causal theory of knowledge [3] and Goldman’s early entertained theory of a similar nature [33]. The nomic sufficiency account of knowledge utilizes a categorical concept of reliability in terms of success not a stochastic one. For Lewis it could be a notion of reliability ranging from stochastic success over other worlds to categorical success over other possible worlds. The confusion is not really resolved at any point because ‘we do not presuppose, of course, that nowhere ever is there a failure of, say, vision’ [78], p. 375. This could mean with respect to sense of success but it could also be with respect to the range of possible worlds. Visual experience and acquired belief based on vision are at the same time described in strict accordance with the thermometer model where the causal dependence covers other possible worlds.

Hallucination is described as an unreliable method of knowledge acquisition except when the hallucinations are so powerful that it is impossible to tell the difference between real perceptual experience and hallucinated experience. In many contexts the worlds in which hallucinations are so powerful are ignored. Since hallucination is a real human condition it cannot always be tossed out. The rule of reliability may from time to time be overridden by the rule of actuality possibly in conjunction with the rule of resemblance. Whatever reliability is supposed to mean it follows, as it did for Nozick’s counterfactual epistemological account, that reliability is a world-dependent property.

Turning to inferential beliefs vindicates the problem. Sampling and inference to the best explanation comprise the *rule of method*. It is a rule which licences tossing out the possible worlds in which these two methods are unreliable non-deductive inference engines. The general methodological problem with inference to the best explanation is exactly whether it is reliable (stochastically or categorically) for some enumeration of possible explanations reflecting audacity, simplicity, unification, predictive power etc. It could moreover be the case that for some enumerations based on for instance simplicity, consistency and boldness, inference to the best explanation could not possibly be reliable. Neither in some desired stochastic nor categorical sense.

Two presumptive rules remain to be accounted for. Both seem trivial yet furnish insight as to Lewis’ view of the situation in epistemology today. One

is a *rule of conservatism* or convention. If a certain set of possible worlds is ignored by the epistemic environment to which the agent belongs, and it is common knowledge among the interacting agents that precisely these possible worlds are tossed out, then they may be ignored. During routine knowledge attribution, the common and mutually expected presuppositions are adopted by the members of the epistemic community. Knowledge attribution is partly a socially determined process forced by conventional means to be taken seriously.

This seriousness is reflected in the final *rule of attention*. Which worlds are ignored is context-dependent. When ignored in a specific context these worlds are *really*, not only counterfactually so, tossed out and not to be considered. Attending to even far-fetched possible worlds in a different context make them relevant possibilities again. Relevant possibilities of error undercut infallible knowledge claims and knowledge flies away—becomes elusive.³⁶

The prohibitive and presumptive rules of knowledge attribution delicately balance between descriptive and normative epistemology. They are descriptive to the extent they describe how we have the knowledge we think we have, yet normative to the effect that were the rules not to be followed the alleged knowledge would vanish or never come about. They serve as the justifications for knowledge claims yet justification is explicitly dismissed. If justification is tantamount to guidelines for getting to the truth, i.e. methodology, then justification is normative and may as well be formulated in rules. The application manual for the methodological rules is not fixed once and for all but fluctuates with the context. The rules may also meet with conflict. Conflict resolution involves priorities with respect to which rule should override for knowledge attribution to take place. Choosing the application of one rule over another may be a matter of social convention or majority vote. Knowledge is a accordingly a normative attribution but also a factual human condition.

Buying into too many uneliminated possibilities of error often makes epistemologists end up with buyers regret. Potential counterexamples to knowledge ascriptions are waiting everywhere in the wings of rich domains making the required epistemic position impossible to reach for anybody. No first persons have knowledge in these particularly demanding contexts, no third persons either. Unfortunately, as a discipline epistemology is one such demanding context. The foe of epistemology is not really skepticism but epistemology itself:

That is how epistemology destroys knowledge. But it does so only temporarily. The pastime of epistemology does not plunge us forevermore into its special context. We can still do a lot of proper ignoring, a lot of knowing, and a lot of true ascribing of knowledge to ourselves and others the rest of the time. [78], p. 377.

Contextualism may all the same be accused of, if not buying into skepticism, then being unusually hospitable to skeptical challenges for an epistemology to

³⁶ Ignoring worlds may from this perspective be seen as a necessary last resort because the available evidence may always be insufficient to block global underdetermination. Ignoring is a precondition for knowledge—love it or leave it.

be. This accommodation leaves first person agents as well as third person ascribers no better off than they were before the time of contextualism and modal epistemology.

On the way to such objections one may initially note a relativization of crucial epistemological concepts. The notion of truth fluctuates with the context. A hypothesis held true by one group of interlocutors may be false for another such that a single hypothesis may be both true and false. Truth is substituted for 'true for' or 'true in the current context' leaving no general notion of truth behind. Contextualism does not have to take the relativistic bait: ' Ξ knows h ' has a definite truth-value fixed by Ξ 's current context; ' Ξ knows h in context c ' has a definite truth-value across contexts. A general trans-contextual notion of truth is preserved.

Preservation of trans-contextual truth still accommodates skeptics. Modal epistemology concedes to skepticism the high epistemic standards on which the skeptical position operates. These epistemic standards are exceedingly harder to meet than those required for everyday attributions of knowledge. Admitting this much to skepticism licences the concern that these elevated standards are in fact the correct standards to be met for genuine knowledge ascriptions and acquisitions. When push comes to shove, the everyday knowledge attributions do not stand up to these standards, so knowledge attributions on a daily basis are bogus as discussed by Pritchard [91]. Skepticism can never be dodged. Presumptive and prohibitive rules may conflict in such a way that skeptical possibilities like hallucinations become relevant. Applying the prohibitive rule of resemblance merely escapes skepticism by *ad hoc* qualifications. This leaves us again 'caught between the rock of fallibilism, and the whirlpool of skepticism' as Lewis puts it [78], p. 367. Modal epistemology was supposed to come to the rescue.

As bogus as these ascriptions may seem, they may also be as good as it gets. A similar response to skepticism following 'smooth' lines may be found in Levi's formal epistemology [72], [73]. To gain truth and avoid error beliefs should be chosen carrying the highest 'epistemic utility'. The epistemic utility embodies truth as well as content. Significant possibilities of error are forgivable just the agent settles for the belief with the highest epistemic utility in the particular context. This may not exactly add up to real knowledge but it is good enough for decision and action. The elevation of the skeptical standards for knowledge is immaterial for common epistemic practice. Infallibilism with respect to all worlds cannot be reached anyway and agents are doing the best they can quantifying over less reaching at least a workable impasse with skepticism. That is the epistemic balance; Agents can act on their 'discount' infallible knowledge, but skeptics can do very little with their high standards. Turning the tables, skeptics are the *real* epistemologists.

Denials of skeptical hypotheses cannot be known on the modal conception of knowledge trans-contextually. So an objection would be that knowledge is not even possible, much less real. A defense would be to simply admit that the logics of knowledge are rather weak at least for the third person knowledge operator and in case of contextual changes. As opposed to counterfactual epistemology's

denial of closure it holds for a first person operator in a uniform context in Lewis' modal epistemology. Closure may fail from the third person perspective because the set of worlds to be considered is strictly a superset of the set of worlds the first person operator has to consider leaving room for radical context change, and a failure. There is support to be found for such a defense.

Levi's epistemological program is a version of a first person perspective emphasizing a distinction between *the logic of truth* and *the logic of consistency* and not the first and third person perspectives [74]. Even though related the two distinctions are not exactly the same. Levi denies the validity of various epistemic axioms as axioms of an epistemic logic of truth. This crudely means to reject these axioms as axioms for a third person knowledge operative. An axiom like the *KK*-thesis ((19), p. 28) found to be invalid in counterfactual epistemology is here valid as an axiom serving regulative purposes of maintaining consistency for a rational epistemic agent. The logic of truth for an epistemic agent on the other hand is not necessarily regulated by a principle like the *KK*-thesis. Lewis seems to follow suit because knowledge of knowledge introduces a discrepancy of actualities for the first and the third person operator. Because of the subject-based contextualism enforced by the rule of actuality, the third person operator is to ignore fewer worlds leaving more room for error. The agent may perhaps know that he knows, the third person may not necessarily be able to determine that the *KK*-thesis holds for the agent, nor that it holds for himself pertaining to the agent in question. The agent in the local environment may have more knowledge than a third person is able to ascribe to him or to the third person himself. If there is a trans-contextual third person logic of knowledge, such a logic is probably rather weak seems to be the suggestion of Levi and Lewis.

While Lewis may consider a universal third person logic rather weak there is nothing in the way of arguing for a much stronger first person logic. This is in stark contrast to the counterfactual proposal of the previous section in which the first person logic was quite weak. On the modal epistemological account all of (18)–(21) may be valid in uniform contexts for a first person knowledge operator.

5 Epistemological Axiomatizations

There is a distinct formal feature to both the counterfactual and modal theories of knowledge. They are in a sense 'formal mainstream' theories as they both observe the significance of epistemic axioms drawn from logical epistemology and their intimate relations to the algebraic properties of the accessibility relation between possible worlds. Nozick considers the accessibility relation to be reflexive while Lewis takes it to be at least reflexive and a sort of symmetric given the rule of actuality and the rule of resemblance respectively. Now, closure holds in uniform contexts, the *KK*-thesis holds, and the rule of necessitation will also immediately hold for a first person modal epistemological logic. Using the sliding scale devised by logical epistemology to determine validity will make

the first person modal epistemological logic at least have epistemic strength on the order of **S4**, perhaps even **S5** is acceptable to Lewis under certain conditions although not discussed. The third person logic of Lewis' contextualism seems to be no stronger than Nozick's first person logic validating (18) and (21) which by being so weak is a non-normal modal logic.

Table 3 below summarizes the results pertaining to the validity of common epistemic axioms given the first and third person perspectives on inquiry for logical, counterfactual and modal epistemology

	CE	ME	LE
N: $\frac{A}{K_{\Xi}A}$	1	1/3	1/3
K: $\frac{K_{\Xi}(A \rightarrow A') \rightarrow (K_{\Xi}A \rightarrow K_{\Xi}A')}$		(1)/(3)	1/3
T: $K_{\Xi}A \rightarrow A$	1	1/3	1/3
4: $K_{\Xi}A \rightarrow K_{\Xi}K_{\Xi}A$		1	1/3
5: $\neg K_{\Xi}A \rightarrow K_{\Xi}\neg K_{\Xi}A$		(1)	3

Table 3: CE: Counterfactual Epistemology, ME: Modal Epistemology, LE: Logical Epistemology. 1: First person perspective, 3: Third person perspective. (,): context-sensitive validity

The axioms are in turn answers to skepticism as their validity is sensitive to the forcing restrictions entertained by the various paradigms of knowledge considered above. Nozick's strategy to combat the skeptic is to impose very little relational structure on the universe of possible worlds leaving the skeptic with very little room to manoeuvre, thus limiting the skeptic's movement. The strategy of modal and logical epistemology is the opposite: To impose much more relational structure on the universe of worlds (in uniform contexts) leaving the agent with much room to manoeuvre, thus enhancing the agent's movement. To combat skepticism, force the skeptic out, either by not giving him a chance to cite distant possibilities of error as relevant, or by making sure that whatever he cites you can reach truthfully at least from the first person perspective, and possibly from the third if you ask logical epistemology.

The common epistemic axioms now furnish a challenging meeting point for mainstream and formal epistemologies ... and there are many others. Some more are to be found in *Forcing Epistemology* [41] others yet uncovered. Let's join the forces to create an interactive epistemology of value to the interdisciplinary study of knowledge.

References

- Alchourrón, C.E., Gärdenfors, P. and Makinson, D. (1985). ‘On the Logic of Theory Change’, *Journal of Symbolic Logic* **50**: 510–530.
- Arló-Costa, H. (1998). ‘Qualitative and Probabilistic Models of Full Belief’, *Proceedings of Logic Colloquim’98, Lecture Notes on Logic* **13**, Buss, S., Hajék, P. and Pudlák, R. (eds.), ASL.
- Armstrong, D. (1973). *Belief, Truth and Knowledge*. New York: Cambridge University Press.
- Aumann, R.J. (1976). ‘Agreeing to Disagree’, *Annals of Statistics* **4**: 1236–1239.
- Baltag, A., Moss, L.S., Solecki, S. (1998). ‘The Logic of Public Announcements, Common Knowledge, and Private Suspicion’. Presented at TARK98. Accepted for publication in *Annals of Pure and Applied Logic*, 2002.
- Barwise, J. ‘On the model theory of common knowledge, the situation in logic’, *CSLI Lecture Notes*, Center for the Study of Language and Information: 201–220.
- Bernecker, S. and Dretske, F. (eds.) (2000). *Knowledge: Contemporary Readings in Epistemology*. Oxford: Oxford University Press.
- Benthem, J.F.A.K. v. (2000). ‘Logic and Game Theory—close encounters of the third kind’, in *Proceedings of LLC99*, Loon, I.v., Mints, G. and Muskens, R. (eds). CSLI Publications.
- Benthem, J.F.A.K. v. (2000). *Logic in Games*. Electronic Lecture Notes: <http://turing.wins.uva.nl/~johan/>, and manuscript, ILLC Amsterdam.
- Benthem, J.F.A.K. v. (2001). ‘Dynamic-Epistemic Logic of Games’, to appear in *Proceedings of LOFT4*, Bonanno, G. (ed.), *Bulletin of Economic Research*.
- Benthem, J.F.A.K. v. (2003). ‘Fifty Years: Changes and Constants in Logic’, in [42]: 35-56.
- Bicchieri, C. (1993). *Rationality and Coordination*. New York: Cambridge University Press.
- Binmore, K. and Shin, H. S. (1992). ‘Algorithmic Knowledge and Game Theory’, in *Knowledge, Belief, and Strategic Interaction*, Bicchieri, C. and Dalla-Chiara, M.-L. (eds). Cambridge: Cambridge University Press: 141–154.
- Boh, I. (1993). *Epistemic Logic in the Middle Ages*. Routledge.

- Bonjour, L. (1980). 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* **5**: 53–73.
- Bonjour, L. (1988). *The Empirical Structure of Knowledge*. Cambridge: M.I.T. Press.
- Bull, R. and Segerberg, K. (1984). 'Basic Modal Logic', in ([29]): 1–88.
- Chisholm, R. (1963). 'The Logic of Knowing', *Journal of Philosophy* **60**: 773–795.
- Chisholm, R. (1977). *The Theory of Knowledge*, 2nd edn. First edition 1966. Englewood Cliffs, NJ: Prentice Hall.
- Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* **13**: 57–90.
- Cohen, S. (2000). 'Contextualism and Skepticism', *Philosophical Issues* **10**: 94–107.
- DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review* **104**: 1–52.
- DeRose, K. (2000). 'How Can We Know that We're not Brains in Vats?', *The Southern Journal of Philosophy* **38**: 121–148.
- Dretske, F. (1970). 'Epistemic Operators', *The Journal of Philosophy* **67**(1970): 1007–1022.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Everitt, N. & Fisher, A. (1995). *Modern Epistemology – A New Introduction*. McGraw Hill.
- Fagin, R., Halpern, J. Y., Moses Y. and Vardi, M. Y. (1995). *Reasoning about Knowledge*. Cambridge: MIT Press.
- Fagin, R., Halpern, J.Y., Moses, Y. and Vardi, M.Y. (1996). 'Common Knowledge Revisited', *Annals of Pure and Applied Logic* **96**(1999): 89–105. Reprinted in [45].
- Gabbay, D. and Guenther F. (1984). *Handbook of Philosophical Logic*, vol. II: *Extensions of Classical Logic*. Dordrecht: D. Reidel Publishing Company.
- Gärdenfors, P. (1988). *Knowledge in Flux—Modelling the Dynamics of Epistemic States*. Cambridge: MIT Press.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* **23**, 6: 121–123.

- Girle, R. (2000). *Modal Logics and Philosophy*. Acumen Publishing.
- Goldman, A. (1967). 'A Causal Theory of Knowing', *The Journal of Philosophy*, vol. 64, **12**: 355–372. Reprinted in [7]: 18–30.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy*, **73**: 771–779. Reprinted in [7]: 87–102.
- Gochet, P. and Gribomont, P. (2003). 'Epistemic Logic,' forthcoming in *Handbook of the History and Philosophy of Logic*, edited by Gabbay, D.M. and Woods, J. Amsterdam: Elsevier Science.
- Gundersen, L.B. (2002). *Dispositional Theories of Knowledge*. Ashgate Publishing.
- Halpern, J. Y. (1995). 'Should Knowledge Entail Belief?', *Journal of Philosophical Logic* **25**: 483–494.
- Hendricks, V. F. (1999). 'Methodology in Epistemology', *Danish Yearbook of Philosophy* **34**: 43–64.
- Hendricks, V. F. (2001). *The Convergence of Scientific Knowledge—a view from the limit*. Trends in Logic: Studia Logica Library Series, Kluwer Academic Publishers.
- Hendricks, V. F. (2003). 'Active Agents', Φ NEWS (2002), vol. **2**: 5–40. A revised version of the paper is published in a special volume of the *Journal of Logic, Language, and Information*, van Benthem, J. and van Rooy, R. (eds.), (2003).
- Hendricks, V. F. (2004). *Forcing Epistemology*, forthcoming, New York: Cambridge University Press.
- Hendricks, V.F. and Malinowski, J. (eds.) (2003). *Trends in Logic: 50 Years of Studia Logica*. Trends in Logic: Studia Logica Library Series, Kluwer Academic Publishers. Contributions from J. van Benthem, W. Buszkowski, M.L. Dalla Chiara, M. Fitting, J.M. Font, R. Giuntini, R. Goldblatt, V. Marra, D. Mundici, R. Leporini, S.P. Odintsov, H. Ono, G. Priest and H. Wansing.
- Hendricks, V. F. and Pedersen, S. A. (2000). *The Companion to Modal Operator Theory—A Program in Philosophy Online*. Department of Philosophy and Science Studies, Roskilde University, Denmark: <http://www.mot.ruc.dk>.
- Hendricks, V. F. and Pedersen, S. A. (2002). *Moderne elementær logik*. Forlaget Høst & Søn.

- Hendricks, V.F., Pedersen, S.A. og Jørgensen, K.F. (eds.) (2003). *Knowledge Contributors*. Synthese Library, vol. 322. Kluwer Academic Publishers. Contributions from H.v. Ditmarsch, R. Fagin, J. Halpern, J. Hintikka, W. v. d. Hoek, B. Kooi, W. Lenzen, Y. Moses, H. Rott, J. Sowa, M. Vardi and R. Wójcicki.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell: Cornell University Press.
- Hintikka, J. (1969). ‘Semantics for Propositional Attitudes’, in *Models for Modalities*, Dordrecht: D. Reidel Publishing Company: 87–111.
- Hintikka, J. (1970). ‘Knowing that One Knows’ Revisited’, *Synthese* **21**: 141–162.
- Hintikka, J. (1975). ‘Impossible Possible Worlds Vindicated’, *Journal of Philosophical Logic* **4**: 475–484.
- Hintikka, J. (1989). ‘Reasoning about Knowledge in Philosophy: The Paradigm of Epistemic Logic’, in *The Logic of Epistemology, and the Epistemology of Logic*, J. and M. Hintikka (eds.). Kluwer Academic Publishers: 17–35.
- Hintikka, J. (1999). *Inquiry as Inquiry: A Logic of Scientific Discovery*. Kluwer Academic Publishers.
- Hintikka, J. (2003). ‘A Second Generation Epistemic Logic and its General Significance’, forthcoming in [45].
- Hintikka, J. (2003). ‘Epistemology without Knowledge and Belief,’ *draft*: 35 pp.
- Hintikka, J. and Halonen, I. (1998). ‘Epistemic Logic’, *Routledge Encyclopedia of Philosophy*, volume 1.
- Hoek, W. van der and Meyer, J.-J.Ch. (1995). *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41. Cambridge: Cambridge University Press.
- Hoek, W. van der, (1996). ‘Systems for Knowledge and Belief’, *Journal of Logic and Computation* **3**, 2: 173–195.
- Hoek, W.v.d., Ditmarch, H. and Kooi, B. (2003). ‘Concurrent Dynamic Epistemic Logic,’ forthcoming in [45].
- Huang, Z. and Kwast, K. (1991). ‘Awareness, Negation and Logical Omniscience’, *Logics in AI, Proceedings JELIA’90*, Eijck, J. van (ed.). New York: Springer-Verlag: 282–300.
- Kelly, K. (1996). *The Logic of Reliable Inquiry*. New York: Oxford University Press.

- Knuuttila, S. (1993). *Modal Logic in the Middle Ages*. Routledge.
- Kripke, S. (1963). 'Semantical Analysis of Modal Logic', *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **9**: 67–96.
- Kutschera, F. von. (1976). *Einführung in die intensional Semantik*. Berlin: W. de Gruyter.
- Kyburg Jr., H.E. (1961). *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Kyburg Jr., H.E. (1970). 'Conjunctivits', in *Induction, Acceptance, and Rational Belief*, Swain, M. (ed.). Dordrecht: D. Reidel Publishing Company.
- Lamarre, P. and Shoham, Y. O. (1994). 'Knowledge, Belief and Conditionalization', *Principles of Knowledge Representation and Reasoning. Proceedings of the 4. International Conference (KR '94)*, Doyle, J., Sandewall, E. and Torraso, P. (eds.). Morgan Kaufman: 415–424.
- Lehrer, K. (1980). 'Coherence and the Racehorse Paradox', *Midwest Studies in Philosophy* **5**: 183–191.
- Lemmon, E.J. (1959). 'Is There Only One Correct System of Modal Logic?', *Aristotelian Society Supplementary Volume XXXIII*: 23–40.
- Lemmon, E. J. (1977). *An Introduction to Modal Logic*, in collaboration with D. Scott. Oxford: Basil Blackwell Publishers.
- Lenzen, W. (1978). *Recent Work in Epistemic Logic*, in *Acta Philosophica Fennica* **30**: 1–219.
- Lenzen, W. (2003). 'Knowledge, Belief, and Subjective Probability', forthcoming in [45].
- Levesque, H. J. (1984). 'A logic of implicit and explicit belief,' *Proceedings AAAI-84* (Austin, TX): 198–202.
- Levi, I. (1983). *The Enterprise of Knowledge*. Cambridge: Cambridge University Press.
- Levi, I. (1991). *The Fixation of Belief and its Undoing*. Cambridge: Cambridge University Press.
- Levi, I. (1997). 'The Logic of Full Belief', *The Covenant of Reason: Rationality and the Commitments of Thought*, Peters, A. K. (ed.), Cambridge: Cambridge University Press: 40–69.
- Lewis, D. (1973). *Counterfactuals*. Cambridge: Harvard University Press.
- Lewis, D. (1976). 'Scorekeeping in Language Games', *Journal of Philosophical Logic* **8**: 339–359.

- Lewis, D. (1984). *On the Plurality of Worlds*. Oxford: Basil Blackwell Publishers.
- Lewis, D. (1996). 'Elusive Knowledge', *The Australian Journal of Philosophy* **74**: 549-567. Reprinted in [7]: 366-384.
- Lindström, S. and Rabinowicz, W. (1997). 'Extending Dynamic Logic: Accommodating Iterated Beliefs and Ramsey Conditionals within DLL', in *For Good Measure*, Lindahl, L., Needham, P. and Sliwinski, R. (eds.). *Uppsala Philosophical Studies* **46**: 123-153.
- Lindström, S. and Rabinowicz, W. (1999). 'DDL Unlimited: Dynamic Doxastic Logic for Introspective Agents', *Erkenntnis* **50**: 353-385.
- Malcolm, N. (1952). 'Knowledge and Belief', *Mind* **LXI**, 242. Reprinted in *Contemporary Readings in Epistemology*, Goodman, M. F. and Snyder, R. A. (eds.). Prentice Hall. New Jersey: Englewood Cliffs (1993): 272-279.
- Moore, G. E. (1959). 'Certainty'. *Philosophical Papers*. London: Allen and Unwin/Unwin Hyman. Reprinted in *Contemporary Readings in Epistemology*, Goodman, M. F. and Snyder, R. A. (eds.). Prentice Hall. New Jersey: Englewood Cliffs (1993): 257-271.
- Moore, R.C. (1982). 'A Formal Theory of Knowledge and Action', in *Formal Theories of a Commonsense World*, Hobbes, J. R. and Moore, R. C. (eds.). Norwood: Ablex.
- Moore, R.C. (1985). 'Semantical Considerations on Non-Monotonic Logic', *Artificial Intelligence* **25**(1): 75-94.
- Moore, R.C. (1995). *Logic and Representation*, CSLI Lecture Notes **39**.
- Moses, J. Y. and Shoham Y. O. (1993). 'Belief as Defeasible Knowledge', *Artificial Intelligence* **64**(2): 299-322.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Nozick, R. (1993). *The Nature of Rationality*. Cambridge: Harvard University Press.
- Osborne, M. and Rubinstein, A. (1994). *A Course in Game Theory*. Cambridge: MIT Press.
- Pollock, J. (1983). 'Epistemology and Probability', *Synthese* **55**: 231-252.
- Pritchard, D.H. (2001). 'Contextualism, Skepticism, and the Problem of Epistemic Descent', *Dialectica* **55**: 327-349.
- Pritchard, D.H. (2002a). 'Contemporary Skepticism', *Internet Encyclopedia of Philosophy*:
<http://www.utm.edu/research/iep/s/skepcont.htm>.

- Pritchard, D.H. (2002b). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* **39**(3): 215–257.
- Rantala, V. (1975). 'Urn Models: A New Kind of Non-Standard Model for First-Order Logic', *Journal of Symbolic Logic* **4**: 455–474.
- Radford, C. (1966). 'Knowledge—by examples', *Analysis* **27**: 1–11.
- Reiter, R. (1980). 'A Logic for Default Reasoning', *Artificial Intelligence* **13**: 81–132.
- de Rijke, M. (1994). 'Meeting Some Neighbours: A Dynamic Modal Logic Meets Theories of Change and Knowledge Representation', in *Logic and Information Flow*, Eijck, J.v and Visser, A. (eds.). Cambridge: MIT Press.
- Rott, H. (2003). 'Economics and Economy in the Theory of Belief Revision', forthcoming in [45].
- Russell, B. (1948). *Logic and Knowledge: Essays 1901–1950*, Marsch, R.C. (ed.). London: Allen and Unwin (1956).
- Segerberg, K. (1995). 'Belief Revision from the Point of View of Doxastic Logic', *Bulletin of the IGPL* **3**: 535–553.
- Segerberg, K. (1999a). 'The Basic Dynamic Doxastic Logic of AGM', *Uppsala Prints and Preprints in Philosophy* **1**.
- Segerberg, K. (1999b). 'A Completeness Proof in Full DDL', in *Philosophical Crumbs: Essays Dedicated to Ann-Mari Henschen-Dahlqvist on the Occasion of her Seventy-fifth Birthday*, Sliwinski, R. (ed.). *Uppsala Philosophical Studies* **49**: 195–207.
- Shin, H.-S. and Williamson, T. (1994). 'Representing the Knowledge of Turing Machines', *Theory and Decision* **37**(1): 125–146.
- Stalnaker, R. (1968). 'A Theory of Conditionals,' in *Studies in Logical Theory*, Rescher, N. (ed.). Oxford: Basil Blackwell: 98–112.
- Stalnaker, R. (1983). *Inquiry*. Cambridge: MIT Press.
- Stalnaker, R. (1996). 'Knowledge, Belief and Counterfactual Reasoning in Games', *Economics and Philosophy* **12**: 133–163.
- Stalnaker, R. (1999). 'Extensive and Strategic Form: Games and Models for Games,' *Research in Economics* **53**: 291–293.
- Thorpe, W.H. (1983). *Beyond Reductionism*. Irvington Publications.
- Voorbraak F. (1991). 'The Theory of Objective Knowledge and Rational Belief, *Logics in AI, Proceedings of the European Workshop JELIA '90*. New York: Springer Verlag: 499–515.

Von Wright, G., H. (1951). *An Essay on Modal Logic*. Amsterdam: North-Holland Publishing Company.

Von Wright, G., H. (1984). *Truth, Knowledge and Modality: Philosophical Papers Vol. II*. Oxford: Basil Blackwell Publishers.

Williams, M. (2001). 'Contextualism, Externalism and Epistemic Standards', *Philosophical Studies* **103**: 1-23.

Williamson, T. (2000). 'Tennant on Knowable Truth', *Ratio* 13, **2**: 99-114.

Williamson, T. (2001). 'Some Philosophical Aspects of Reasoning about Knowledge', *Proceedings of the 8th Conference on TARK*, van Benthem, J. (ed.). San Francisco: Morgan Kaufman Publishers: 115-121.

Williamson, T. (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.