Finite narrative modelling, contextual dynamic semantics and Elusive Knowledge

MSc Thesis *(Afstudeerscriptie)*

written by

Peter van Ormondt
(born 30 December 1976 in Delft, the Netherlands)

under the supervision of **Prof. dr. B. Löwe**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
30 November 2012

Members of the Thesis Committee:
Prof. dr. J.F.A.K. van Benthem
Prof. dr. D.H.J. de Jongh
Prof. dr. B. Löwe
Dr. S.J.L. Smets
Prof. dr. M.J.B. Stokhof

Institute for Logic, Language and Computation

Abstract

The central point of this thesis is the semi-formal approach to contexualist semantics of knowledge by David Lewis. Lewis introduces a set of rules that allow us to ignore certain parts of the space of all possibilities when we evaluate a knowledge claim. These rules are dependent on the context of utterance of the knowledge claim and therefore give rise to a contextualist notion of knowledge. We focus in particular on the Rule of Attention. The open nature of this rule poses problems for a formalisation in the usual framework of *dynamic epistemic logic*.

We make a distinction between two modelling approaches: *open modelling* and *closed modellling*. The latter approach, called *finite narrative modelling*, is a methodology that relieves us from the duty of using a language that can represent arbitrary facts, since we as modellers know which facts will become relevant when we start with the modelling.

We introduce two extensions of *dynamic epistemic logic* and conclude by considering Lewis' rules as restrictions of this model.

Contents

1 Introduction

This thesis consists of three parts: a philosophical part (sections 2 and 3) in which we discuss the position of contextualism in epistemology, a methodological part (sections 4 and 5) in which we discuss a modelling approach that we will use in the formalisation, and a formal part (sections 6 and 7) in which we apply the techniques of *dynamic epistemic logic* to the problems outlined in the first part.

The central point of this thesis is the semi-formal approach to contextualist semantics of knowledge by David Lewis [Lewis, 1996]. In this paper Lewis introduces a set of rules that allow us to ignore certain parts of the space of all possibilities when we evaluate a knowledge claim. These rules are dependent on the context of utterance of the knowledge claim and therefore give rise to a contextualist notion of knowledge. Although Lewis discusses these rules only informally, he claims that his approach could be adequately formalised.

> I could have said my say fair and square, bending no rules. It would have been tiresome, but it could have been done. The secret would have been to resort to 'semantic ascent'. I could have taken great care to distinguish between (1) the language I use when I talk about knowledge, or whatever, and (2) the second language that I use to talk about the semantic and pragmatic workings of the first language. If you want to hear my story told that way, you probably know enough to do the job for yourself. If you can, then my informal presentation has been good enough. [Lewis, 1996, p. 566–567]

In this thesis we accept this challenge.

There is one of Lewis' rules that receives particular focus and that is the *Rule of Attention*. This rule says that if we are not ignoring a certain possibility then it is part of the contextually determined domain. Moreover:

> If it is an uneliminated possibility in which not-P, then it will do as a counter-example to the claim that P holds in every possibility left uneliminated by S's evidence. That is, it will do as a counter-example to the claim that S knows that P. [Lewis, 1996, p. 559]

This means that our conversation partner can potentially destroy knowledge by raising an issue that was previously ignored.

It seems therefore that this conversationally controlled mechanism has *dynamic* properties in the sense of logics of agency and action studied by the *Logic, Language and Information* community. This makes it natural to consider our attempt to formalise Lewis in a system of *dynamic epistemic logic* (*cf.* section 5 for definitions).

The open nature of Lewis' *Rule of Attention* does pose problems for the usual framework of *dynamic epistemic logic*. This framework is a propositional modal logic whose main applications (and explanatory power) come from easily surveyable finite models: complex facts about the real world are encoded in single propositional variables to simplify the formalism. However, the *Rule of Attention* seems to require a language that allows us to represent virtually arbitrary connections and situations as they might be brought up in conversation and therefore must be attended to.

In order to deal with this discrepancy, in the second part of this thesis, we will introduce a distinction between two fundamentally different modelling approaches: an open modelling approach that leaves all formal possibilities undecided in order to deal with changes of situation flexibly, and a more restrictive approach that makes decisions at the beginning of the modelling what the scope of the model should be and fixes the vocabulary accordingly.

The latter approach is called *finite narrative modelling*. We think of this as a method where the modeller gets the entire (finite) narrative of the developing situation up front, then makes the modelling decisions, and finally reestablishes the dynamics via a sequence of models that represent the narrative. In fact, the method of *finite narrative modelling* is the main methodology used in most theoretical applications of logics of actions (whereas the open modelling approach is used by people who implement AI and build robots).

The methodology of *finite narrative modelling* relieves us from the duty of using a language that can represent arbitrary facts, since we as modellers know which facts will become relevant when we start with the modelling. This will allow us to use the language of *dynamic epistemic logic* to formalise Lewis' rules. This we shall undertake in the final part of this thesis. We explain the basics of *dynamic epistemic logic* (section 5), we then introduce two extensions of *dynamic epistemic logic* (section 6) and finally we will use these frameworks to give a formalisation of several of Lewis' rules.

By restricting our methodology to *finite narrative modelling*, some of Lewis' rules seem to lose their force. We will discuss what this means for the interplay between formal ontology of the modelling framework and the philosophical theory.

Since works like [Stalnaker, 1999] and [Lewis, 1979c] context has been a prominent theme in the research area of logic, language and information. Ranging from natural language semantics, where in dynamic semantics interpretation is taken to both depend on context as well as influence it (*cf.* [van Eijck and Visser, 2010]), to AI, introduced by John McCarthy in his [McCarthy, 1993] and found further development in context logics such as [Buvač et al., 1995]. Dynamic context logics have been studied in [Aucher et al., 2009]. The idea of specifically formalising contextual notions of knowledge has recently attracted interest witnessed by publications [Holliday, 2012a, Holliday, 2012b, Holliday, 2010]. Similar to attending and unattending to possibilities is awareness and unawareness of possibilities. [Fagin and Halpern, 1988] make a distinction between implicit and explicit knowledge in order to deal with the problem of *logical omniscience*. A dynamic variant of the latter system can be found in [van Benthem and Velázques-Quesada, 2009]. We finally point to [de Jager, 2009] that concentrates on unawareness as an epistemic attitude.

2 Epistemic Contextualism

2.1 Introduction

Epistemic contextualism is the position that the truth of knowledge statements, or the propositions that are expressed by them, are in some way dependent on the context in which they are uttered. It is a position that is rooted in the analytical tradition and has been developed in the second half of the 20th century.[1] Arguably, epistemic contextualism is a position also to be found in other traditions dating back further but those will not concern us here.[2]

On the face of it, the claim that the truth of the proposition expressed by the utterance of a sentence like 'a knows that p' depends on context, seems hardly controversial. Surely, the truth of 'I know that *Beatrix is the Queen of the Netherlands*' depends on *when* this sentence is uttered and by *whom*. But it is not in this respect that knowledge claims are said to be context-sensitive. What is meant is that knowledge claims are context-sensitive with respect to *changing epistemic settings*. Different contexts have, or determine, different epistemic settings and therefore impose different truth conditions on propositions expressed by claims of knowledge. Alternatively, the content of the word "know" is a function from epistemic settings in contexts to relations between individuals and propositions.[3] What these

[1] *Cf.* [DeRose, 1998, section 5].

[2] See [Norman, 1999] who traces contextualist claims in other philosophical traditions. Also see [Bianchi and Vassallo, 2005, p. 42] for a claim that contextualism is already to be found with Descartes, Hume and Locke.

[3] These claims assume a lot of theory. For instance that sentences express propositions

epistemic settings are, or what it is that changes with context is a matter of debate and different frameworks have been developed in this respect. For instance, some people claim that the truth conditions of a knowledge claim will be more easily met in some contexts and more difficult in others. A claim which is denied by so-called "invariantists" who will not agree that the truth conditions of the knowledge claims vary according to these changing settings.

Although most people working in the field of epistemic contextualism would agree with the general claim of the contextual dependence of knowledge statements, positions start to curdle when you ask whose contexts we are talking about? Well, the ascriber's, say some.[4] Not so, it is the subject's context, say others.[5] Positions change more when you ask what these epistemic settings are, and when they are met. And when we talk about knowledge, do we talk about propositions that are expressed when knowledge claims are uttered or are we talking about knowledge *per se*? Following [Rysiew, 2011], there are two general distinctions we can make in the field of epistemic contextualism. One distinction is between a semantic and a substantive view of contexts. The other between subjective and attributive contexts, or whose contexts we should consider. We will look at these distinctions in turn.

Semantic and Substantive contextualism

A central concern for epistemology has always been what the structure of knowledge and justification is. Traditional theories include foundationalism and coherentism and are about knowledge *per se*. By this we mean that these theories try to give an account of what knowledge is, rather than when we can adequately utter knowledge statements.[6] By the term *substantive contextualism* we refer to the position that tries to give an account of knowledge and justification in this respect but argues for the contextual depedency of knowledge and justification, respectively. For instance, David Annis formulates an alternative to the view of foundationalism by pointing out the

---

and that propositions are the semantic contents or meanings of sentences. This is not the place to go into this century-old philosophical history. Instead we refer to [Stalnaker, 1984], [Stalnaker, 1999] for a contemporary view and to [Nuchelmans, 1973] for an overview of the historical development.

[4] Most notably, Stewart Cohen [Cohen, 1999], Keith DeRose [DeRose, 1998], David Lewis [Lewis, 1996].

[5] For instance, Fred Dretske in [Dretske, 1981].

[6] We could have also phrased the opposition as a distinction between understanding knowledge in dialectical or structural terms. *Cf.* [Norman, 1999, p. 386].

contextual parameters that have been overlooked in these theories [Annis, 1978]. In particular, he stresses the *social* nature of justification by arguing that the question whether an agent S is justified in believing h is always relative to a so-called *issue-context*:

> The issue-context is what specific issue involving h is being raised. It determines the level of understanding and knowledge that S must exhibit, and it determines an appropriate objector-group. [Annis, 1978, p. 215]

The point being that given a set of basic statements in a foundationalist framework, the foundationalist should recognize that it is only in specific contexts that these basic statements do not stand in need of further justification, but they might require justification in others. But it is not the particular structure of knowledge the contextualist is objecting to, for instance, coherentism is also frequently attacked, it is the idea that there is a perspective that can reveal the correct structure of a justified true belief. It is this assumption of a structural unity of the object of investigation in epistemology that is also argued against by Michael Williams in [Williams, 1991]. This "epistemological realism", as he calls it, assumes that knowledge has a context-independent structure.

In contrast to these investigations that concentrate on the structural properties of knowledge, the semantic position is looking at the proposositions that are expressed when knowledge claims are made and tries to provide adequate truth conditions for them. Such investigations are carried out within particular (semi-)formal frameworks, of which we will see a few examples in section 2.3, and it is within these frameworks that the semantics of the propositions are analysed. Semantic contextualism can therefore be said to only be epistemelogical because it concerns itself with epistemic *terms*.[7] This is a clear difference from the substantive position we described above, where knowledge, and the structure and justification of knowledge, are concerned.

In recent years people have investigated the linguistic implications of the claims made by semantic contextualists, for instance by looking how the verb "to know" as a linguistic item *actually* functions in natural language.[8] To give an example, one of the claims made by contextualists is that the verb "to know" like the adjective "flat" is *gradable*, in the sense that what counts as flat in one context (the claim that the surface of my desk behind

[7] *Cf.* [Rysiew, 2011, section 2].

[8] See [Ludlow, 2005] for an overview.

which I am sitting is flat) does not so in another (when one is building an advanced space telescope). In [Stanley, 2004], the author investigates if there is linguistic evidence for such a claim. The linguistic underpinnings of epistemic contectualism is a different and large field of research altogether and will not be explored in this thesis.

Rather, in what follows, we will be concerned with the semantic position of contextualism and not with substantive contextualism[9] nor the linguistic underpinnings of these claims. In this thesis we will introduce a methodology for modelling knowledge claims in narratives and build on the concepts introduced in semantic contextualism and has therefore no bearing on questions about knowledge of a substantive nature nor does it explore a linguistic analysis of narratives.

Attributive and Subjective contextualism

A second distinction we can make in the field of epistemic contextualism is a distinction between subjective and attributive contextualism. When we talk about contexts we can ask, whose context? When we say that the truth conditions of the proposition that an agent a knows that p are dependent on context we might mean the context of the subject of knowledge, *i.e.*, the *knower*. We might also mean the person attributing the knowledge that p to a, *i.e.*, the knowledge *ascriber*. As Keith DeRose puts it:

> The basic issue here is whether the varying standards a subject must live up to to count as knowing are relative to the context of that subject or rather to the context of the attributor – the person describing the subject as a knower or a non-knower. [DeRose, 1998, section 4]

Between the theories of epistemic contextualism that are around there exist many subtle differences. In order to interpret these correctly, the above distinctions provide a good guiding principle. In this text we will focus on the distinction between what we have called subjective contextualism and attributive contextualism. Both are a semantic thesis in the way pointed out above, but as the naming suggests some find the contextually relevant factors in facts about the subject of knowledge, others also in that of the attributor of knowledge.

[9] See [Pritchard, 2002] for a comparison between substantive and semantic contextualism.

2.2 The scepticist claim

Much of the motivation to develop contextualist theories comes from trying to neutralize sceptical arguments and because it is such an important motivation for contextualism we consider the argument and its responses.[10] In some of the literature of epistemic contextualism the scepticist claim is presented as a paradox. The argument can be rendered as follows:

(1) I know that I have hands.

(2) I don't know that I am not a brain-in-a-vat.

(3) I don't know that I have hands if I don't know that I am not a brain-in-a-vat.

Individually, all propositions seem plausible up to some extent yet they cannot all be true. (1) seems true because if we are to know anything then surely these are the mundane things in life we usually take for granted we know. (2) is true, we could say, by design. The brain-in-a-vat argument is manufactured in such a way that we cannot differentiate between being a brain-in-a-vat and not. Therefore, we cannot know that we are not a brain-in-a-vat. (3) In order to know that I have hands I must of course be able to know that I am not a brain-in-a-vat.

The sceptic will reason on the basis of (2) and (3) that

(4) I don't know that I have hands.

which would be contradicting (1). The disagreement is the following: because this reasoning works the same way if we substitute "I have hands" with some other proposition which we normally assume we know and substitute "being a brain-in-a-vat" with some other exotic mind construct, like continuously deceiving demons, nefarious scientists on Alpha Centauri, *etc.*, the sceptic seems to have his victory: we really don't know the things we ordinarily suppose we know.

Those people who do not wish to accept this result should point out where the reasoning above goes wrong. Because (2) is true by design, many have considered (3) to be the faulty premiss in the sceptic's reasoning. But behind it is a principle that few would care to give up and is known as the *closure principle of knowledge*. This principle can be stated as follows:

[10] This seems to be true especially for theories formulated after the 1970s. *Cf.* [Pritchard, 2002, p. 19].

(K) If agent a knows p and a knows that $p \to q$, then agent a knows that q.

Let p be the proposition that I am a brain-in-a-vat and let q be the proposition that I have hands. As agreed to above, I don't know that I am not a brain-in-a-vat, $\neg K \neg p$. Also, it's reasonable to assume that I know that $\neg p$ is a consequence of q, *i.e.*, if I have hands, then I am not a brain-in-a-vat, $K(q \to \neg p)$. We must then conclude that I do not know that I have hands, $\neg Kq$, or (4) above. Because, were we to conclude that I *do* know that I have hands, Kq, then by virtue of (K) above we have Kq & $K(q \to \neg p)$, hence $K\neg p$, which contradicts our first assumption, or (2) above. Here we see clearly where the sceptic gets his strength from. By the mere mentioning of "You don't know that not —", where we can insert any sceptical hypothesis for —, the sceptic can cast a doubt over all that seems certain.

It seems we have three choices. We could of course just be content with believing that we do not know the things we normally think we know, *i.e.*, throw out (1) above, a position sometimes called 'fallibilism'. We can deny (3), *i.e.*, deny the closure principle for knowledge, an option which is argued for by Fred Dretske, which we will discuss in the next section. We could also deny (2) and insist that we do know that we are not brains-in-a-vat. Something which is known as the *Moorean response* after the philospher G.E. Moore.

Contextualists, however, particularly contextualists of the attributive kind, believe that the truth conditions depend on context and they will try to show that in the argument we have presented above, there is a shift in context taking place which is the reason why the argument leads to apparent contradiction. As we shall see in the next section, they will use this strategy to dodge these three choices altogether.

A lot of the epistemological contextualist theories are developed with the precise intent to neutralize the sceptical argument. This plays an important role in the background and for this reason we have brought it to our attention. In the next section we will have a look at two contextualist solutions to this scepticist paradox.

2.3 Frameworks for contextualism

There seems to be general consensus among contextualists that contextualism is right when arguing for the solution of the skeptical paradox but the treatments among contextualists differ. We will examine two such treatments and their responses to the paradox presented above. We will see in section 3.2 how David Lewis further develops these ideas.

1. Relevant alternatives theory
2. Theory of changing epistemic standards

We will look at these theories in turn.

Relevant alternatives theory

We will first consider *relevant alternatives theory* as a response to the sceptical paradox mentioned in the previous section. On the one hand *relevant alternatives theory* is a contextual theory in it's own right which has been primarily developed by Fred Dretske[11] but it can also be seen as a (technical) framework and as such has been adopted by others, as we shall see, to express their views which might not be in agreement with the theory. The reason for this popularity might be that *relevant alternatives theory* is easily expressed in modal logical terms because the central idea of eliminating possibilities is the same in epistemic logic, with the difference that basic epistemic languages do not normally distinguish between relevant and irrelevant possibilities.[12] However, the framework and the theory are two issues that need to be kept separate. The strategy of Dretske to counter the sceptical argument is to argue that "knowledge" is an absolute concept which can be analysed in terms of *relevant alternatives* and to deny the closure principle of knowledge, as seen in the previous section, a claim backed-up with a device called *truth-tracking*.

Factual knowledge is an absolute concept and does not admit of degree. It makes no sense to say that one person knows that Amsterdam is the capital of The Netherlands *better than* another person. "Knowledge is an all or nothing affair, like being pregnant", Fred Dretske claims in [Dretske, 1981]. Dretske agrees with Peter Unger[13] on this point, but unlike him does not draw sceptical conclusions from this idea. Dretske, with his approach of *relevant alternatives theory*, promises a way to keep sceptical claims at bay. Indeed, as we have said before, many contextualist theories were developed with the aim to do just that, and *relevant alternatives theory* is no different.

[11] See [Dretske, 1981, Dretske, 1970]. Other people who developed *relevant alternatives theory* include Alvin Goldman [Goldman, 1976] and Gail Stine [Stine, 1976].

[12] This last remark will be made more clear when we present formal epistemic languages in section 5.2.

[13] *Cf.*, [Unger, 1975].

We therefore first recall the discussion of absolute concepts put forward by Peter Unger.

Unger makes his point by comparing the concept of *flat* to that of *knowledge*. *Flat* like *knowledge* is an absolute term. Unger contends that when we say that a surface is flat we mean that there are no bumps or dents on it whatsoever. In case there would be irregularities on the surface, however minute these might be, we have to conclude that the surface is not flat. Perhaps we would say that it is almost flat but not flat *tout court*. Were we to put the surface in question underneath a (strong enough) microscope we would find all sorts of irregularities and the sad conclusion is that almost nothing is really flat. So too, Peter Unger concludes, it is the case with knowledge (*cf.*, [Unger, 1975, chapter II]).

Dretske does not share the sceptical conclusions Unger draws from the charateristics of absolute concept but agrees that *knowledge*, like *flat*, is an absolute concept which does not admit of degree. He does not agree, however, that nothing is really flat, or that we cannot be said to ever really know anything. The argument is as follows. Given an absolute concept like flatness, for an object to be flat, it must not have any bumps. This is an absolute matter. What is not an absolute matter is what *counts* as a bump. Similarly for knowledge, if I say that I know that p, it means that I know this on the basis of my evidence or justification being sufficient or adequate. Once the threshold of justification has been passed, only then knowledge, the one kind there is, is established. My evidence eliminates all possibilities that conflict with the knowledge statement, therefore I can truthfully be said to know and this does not admit of degree. What actually counts as a possibility that conflicts with the knowledge statement and has to be eliminated by my evidence is another matter. Dretske says:

> 'Such concepts, we might say, are *relationally absolute*; absolute, yes, but only relative to a certain standard. We might put the point this way: to be empty is to be *devoid of all relevant things*, thereby exhibiting, simultaneously, the absolute (in the word 'all') and relative (in the word 'relevant') character of this concept.' [Dretske, 1981, p.366–367]

One can of course wonder if one should still call such a concept *absolute*, if it is allowed to change its standard for different applications. Talking about an 'absolute concept' implies there is one standard for all applications. The absoluteness Dretske seems to have in mind has more to do with the evaluation of a rule that sets a standard rather than the concept itself. Calling such concepts absolute therefore seems a bit off. Having said this, we don't see it affecting the content of Dretske's claims. We will see below that Lewis, who also chooses Unger's infallibility of knowledge as a starting point, wisely avoids this terminology altogether.

According to *relevant alternatives theory* for a knowledge claim to be true one has to be able to rule out the *relevant* alternatives that are incompatible with it. Given the knowledge claim that p there are those possibilities that are necessarily excluded by p. Those possibilities form what Dretske calls the *contrasting set*. Then there is the *relevancy set* which consists of those possibilities that the person who knows must be able to exclude. Or as Dretske says:

> In saying that [a person] must be in a position to exclude these possibilities I mean that his evidence or justification for thinking these alternatives are *not* the case must be good enough to say he *knows* they are not the case.[14] [Dretske, 1981]

The intuitions of *relevant alternatives theory* are captured formally as follows. Define a *contrasting set*, S, of epistemic possibilities and a subset of S, the *relevancy set*, R. The elements of S represent those epistemic possibilities that, given a knowledge claim, the knower should be able to exclude, by justification or evidence, in order for it to be true to say he knows. The elements which are in S but not in R, $S \setminus R$, contain those possibilities that are irrelevant and need no justification for exclusion. Amongst others $S \setminus R$ contains the malicious demons, and evil scientists manipulating our brain and perception. In this formalism a sceptic would be represented by identifying R and C, in which case one would need to exclude all possibilities which is a tremendous or even impossible burden on evidence.

This does not mean that R is always the same given a knowledge claim that p. In different situations, or contexts, what is a relevant alternative to p can differ. But this does not change what we know, although it might change the evidence that is required.

How is this theory intended to prevent the sceptical paradoxes outlined in section 2.2? According to Dretske, knowledge is only transmitted if the entailment is relevant.[15] If we again consider principle (K), knowing that the animal in front of me is a dog may transmit by entailment to knowing that it is not a sheep, but not to knowing that it is not a hologram generated by a martian in order to fool epistemologists, because this is not in the

[14] *Cf.*, Lewis in [Lewis, 1996, p. 553] defines the excluding of possibilities as a conflict a possbility has with the propositional information that exists in the subject's entire perceptual experience and memory, denying that justification plays any role. We will return to this issue in section 3.2.

[15] In [Dretske, 1970], Fred Dretske calls epistemic operators 'semi-penetrating'.

relevancy set for the knowledge claim. (K) then fails in the sceptic paradox and knowledge is saved.

What remains to be explained, and what of course is essential for this theory to go through, is an answer to the question *what makes a possibility relevant*. What are the (relevant) possibilities that need to be excluded given a knowledge claim that p? For instance, why is the possibility that I am standing next to my desk a relevant alternative to the claim that I am sitting behind it and why is the possibility that I am a brain-in-a-vat not? In order to determine this, Dretkse introduces a counterfactual account of knowledge: an agent knows some proposition p iff had p been false then the agent would not have believed it.[16] This works as follows. Let us consider that I know I have two hands, then according to the definition I would not believe this if I, as a result of a tragic accident, would only have one, whereas I would still believe it if I was a brain-in-a-vat.

Usually, this account is given a modal logical interpretation in such a way that knowledge requires in addition to a true belief in the actual world a belief which is *sensitive* to the truth of the proposition in question. This means that in the nearest possible world where this proposition is not true, I no longer believe it.[17] This requirement for knowledge to be sensitive to the truth of the proposition in question is sometimes called *truth-tracking*.

To make this more clear, suppose John is at the zoo and is looking at a striped animal in a cage in front of him. In normal circumstances, Dretske will say, John can claim he knows that the animal in front of him is a zebra. Let's call this fact p. The relevant alternatives to p, the possibilities that John wil need to rule out, are for instance that the animal is not some kind

[16] Initially this idea was presented in [Dretske, 1971] and later revived by Nozick. See [Nozick, 1981], especially part 3 on epistemology.

[17] Truth-tracking depends on the analysis of counterfactuals. The general form of a counterfactual statement is:

(1) If it had been the case that φ, then it would have been the case that ψ.

which we write as

(2) $\varphi \rightsquigarrow \psi$

Analysis of conterfactual statements as a material implication fails because utterence of (2) is done usually when the antecedent φ is false, so if we would analyse counterfactual statements as a material implication all counterfactual statements would be true.

Robert Stalnaker in [Stalnaker, 1968] and David Lewis in [Lewis, 2008] define (2) to be true in the actual world w iff the consequent ψ is true in all accessible worlds in which the antecedent φ is true and which differ minimally from w. "Differing minimally" is then interpreted by means of a comparative relation on the domain, which makes it possible to compare how much worlds differ from actuality, or in the case of truth-tracking how near they are to the actual world.

of horse, gazelle or cow. The possibility that the animal is a cleverly painted mule to fool the zoo-going public is a possibility which is not relevant because we would still not believe it even if it were true.

Relevant alternatives theory advanced by Dretske has been criticised. Especially, the rejection of the closure principle of knowledge has been the focus of the attack.[18] First noticed by Gail Stine [Stine, 1976], the account by Dretske and Nozick depends on sensitivity of belief *in the nearest worlds*. This means that knowledge of p depends solely on being able to rule out *relevant* alternatives to p. As a result, this means that one cannot know the denial of a sceptical hypotheses, because one cannot rule out that it actually obtains, but as long as it is not a relevant alternative, Dretske would say, this does not harm us. So, where we can be said to know in a normal situation that the animal in the zoo is *not* a gazelle, horse *etc.* One does not know, however, that it is not a cleverly painted mule, even though it being a zebra implies this. It is here where the K-axiom fails. The problem with this approach is that, as we can see, it allows for the conjunction "I know that that is a zebra but I do not know it is not a cleverly painted mule", and this is hard to stomach. The following section will focus on a position that tries to circumvent these difficulties.

Theory of changing epistemic strength

The most notable proponents of attributive contextualism are Keith DeRose, David Lewis and Stewart Cohen, *cf.*, footnote 4. A central observation of attributive contextualism is that it seems perfectly legitimate to ascribe knowledge in ordinary (non-sceptical) contexts whereas it does not seem legitimate to do so in contexts where sceptical possibilities have been raised. Furthermore, it is claimed that the change from an ordinary context to

[18] A concern in epistemic logic is what is sometimes called *logical omniscience*. Together with the rule of necessitation, the K axiom allows agents to know all the logical consequences of what they know. The authors of [Fagin and Halpern, 1988] acknowledge that human reasoning is limited to some extent making a distinction between *implicit* and *explicit* knowledge. However Dretske would deny K to hold in general even if humans would not have limits to their reasoning capabilities:

> Were we all ideally astute logicians, were we all fully appraised of all the necessary consequences [...] of every proposition [...] That is, assuming that if P entails Q, we know that P entails Q, then every epistemic operator is a penetrating operator: the epistemic operators penetrate to all the known consequences of a proposition. It is this latter, slightly modified, claim that I mean to reject. [Dretske, 1970, p. 1010]

a context where sceptical possibilities have been raised is determined by conversational factors. This is why in ordinary contexts it is alright to say somebody knows this or that mundane fact *and* knows the denial of some sceptical hypothesis, whereas in a sceptical context this knowledge is lost.

A subjective contextualist will argue that if an ascription sentence is true, then it is true in every conversational context, because somebody's knowledge depends entirely on the subject's information status. For this reason a subjective contextualist will deny at least one of the above observations. An attributive contextualist will argue against this and say that an ascriber responds to changing "epistemic standards" due to a changing conversational context and that whether or not a knowledge ascription is true depends on the conversational context of the ascriber. In what follows we will mostly focus on the explanation given by DeRose in [DeRose, 1995].

With regard to explaining why it does not seem right to ascribe knowledge to a subject in a conversational context where sceptical possibilities have been raised, DeRose relies on the *counterfactual insensitivity* of beliefs we have explained in the previous section:

> [W]e have a very strong general, though not exceptionless, inclination to think that we don't know that P when we think that our belief that P is a belief we would hold even if P were false. Let's say that S's belief that P is *insensitive* if S would believe P even if P were false. [...] We tend to judge that S doesn't know that P when we think S's belief that P is insensitive. [DeRose, 1995, p. 18]

DeRose cannot leave things like this because this would involve a denial of closure, in the way pointed out in the previous section, and this is exactly what he is trying to prevent. DeRose too expresses the oddness of maintaining that in normal contexts I can know some ordinary fact expressed by p and not know the denial of a sceptical hypothesis q even if I might know that p implies $\neg q$. So, the notion of *sensitivity* needs to somehow be limited in the explanation. In order to do this DeRose points out that the reason why this situation is odd is because my "epistemic postion" with regard to p is no stronger than with regard to $\neg q$, where *epistemic strengh* is understood as follows:

> [...] one's belief should not only be true, but should be non-accidentally true, where this requires one's belief as to whether P is true match the fact of the matter at nearby worlds. The further away one can get from the actual world, while still having it be the case that one's belief matches the fact at worlds that far [sic] away and closer, the stronger a position one is in with respect to P. [DeRose, 1995, p. 34]

Note that for Dretske and Nozick counterfactual sensitivity required the agent to *lack* the belief that p in nearby worlds where $\neg p$ holds, or tracking the truth in *not-p* worlds. For DeRose, epistemic position simply means that if p is the case the belief that p should match it. As a result I can be in a strong epistemic position with regard to the denial of a sceptical hypothesis q even though my belief that $\neg q$ is not sensitive in the sense outlined by Dretske and Nozick.

With the notion of epistemic strength in hand we are in a position to look at what mechanism DeRose proposes to explain how epistemic standards vary across conversational contexts. It comes down to determining how strong the epistemic position of an agent a must be with respect to a proposition p in order for an ascriber A's claim that 'a knows that p' is true. This will involve both sensitivity and epistemic strength:

> When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards of how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge. [DeRose, 1995, p. 36]

According to DeRose, in order for a knowledge claim that p pertaining to a subject S to be true, S has to believe p, p has to be the case and S's epistemic position must be strong enough in order for S's believe to be sensitive. What the scepticist does, according to DeRose, is to set high standards according to which the requirement for knowledge is that my beliefs must be sensitive if they are to count as knowledge. In these sceptical contexts, I fail to know that I am *not* a brain-in-a-vat because this is not a sensitive belief. But it is only in these contexts that the statement that I do not know this enjoys any plausibility. In other, more relaxed contexts, the requirement that our beliefs must be sensitive is not required and therefor we can reject that I do not know that I am a brain-in-a-vat. And it is only with respect to an agent's belief which are part of the ascriber's conversational context that need to be sensitive in order for the ascription sentence to be true.

A common way to picture this model is to think of logical space ordered by a similarity relation inducing a set of nested spheres. If a world v is in sphere i we denote that world with a corresponding index: v_i. A world v_i is closer, *i.e.*, more similar, to the actual world w than a world v_j if $i < j$.

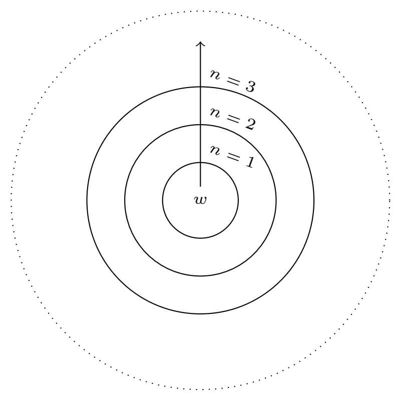

We can now think of the notions of sensitivity and epistemic strength as follows. We say an agent is sensitive to p if in all v_i, where $i < k$, if p does not hold in v_i, then neither does $B_a p$. Similarly, we can put a measure on epistemic strength by saying that the epistemic strength of an agent with repsect to a proposition p is i if for all $j < i$, if in v_j p holds then so does $B_a p$.

2.4 Conclusions

In this section we have discussed two frameworks in which semantics for contextual knowledge statements are developed. The important difference between the two positions are in whose context to look for the contextual relevant parameters. We have seen that subjective contextualism as advanced by Dretske has a flaw called "abominal conjunctions" by DeRose who has adopted an attributive contextual view to overcome these problems. This does not mean that the latter view, which we shall further explore in the next section, is without critics. *Cf.* [Rysiew, 2011, §4] for an overview of the arguments against the attributive thesis. This article also makes clear that the debate is far from over.

3 Some concrete issues in contextualism

3.1 Introduction

In the previous chapter we have set the stage to look more closely at some of the aspects of contextualism. In his 1996 paper [Lewis, 1996] David Lewis discusses a context-dependent notion of knowledge. A traditional definition of knowledge says that an agent a knows p iff a has eliminated all possibilities where not-p. Either we interpret this definition in its strictest sense, *i.e.*, we consider *all* possibilities imaginable, which has the unfavorable consequence that we are rid of all *everyday* knowledge that we normally suppose we have.[19] Or, we relax the interpretation and admit that we are mere fallibilists, allowing that we cannot possibly eliminate all possibilites where not-p, ridding the word 'knowledge' of its content.

As we will show in the next session, in order to sidestep these problems David Lewis adds a "proviso" to the above definition, making knowledge a context-dependent notion. By introducing a set of rules he aims to determine exactly which possibilities are relevant and which ones we can properly ignore. In the next section we will have a a look at Lewis' paper in more detail. We will then take up some issues which were left in section 2 and see how his views relate to other positions in epistemic contextualism. Do Lewis' suggestions indeed solve the problems he addresses, *i.e.*, does he succeed in defining a notion of knowledge that saves the big bulk of everyday knowledge, which we normally assume we have, from the arguments of the sceptic, without admitting to the fallibilistic position?

3.2 David Lewis' Elusive knowledge

"We know a lot", Lewis starts out in [Lewis, 1996]. Though he also admits that all the things that we think we know are endangered by the infallibilistic character of the definition of knowledge, mentioned in the introduction above. Because as soon as we start doing epistemology we find that there are a lot of possibilities not eliminated ("let your paranoid fantasies rip"[20]) and we know next to nothing. A solution to this could be that we only have *fallibilistic* knowledge. But, Lewis implores us, doesn't knowledge despite uneliminated possibilities just sound wrong? Fallibilism is not an option and he seeks to "dodge the choice" altogether.[21]

Assuming that knowledge ascriptions are context-dependent, Lewis claims that it is only in the context of epistemology, or other demanding contexts, that we are robbed of our everyday knowledge we normally assume we have. If we are in an informal setting we normally assume that we know a lot, however as soon as we engage in the systematic philosophical examination of knowledge, thereby raising the standards of evaluation, things go bad. What this makes clear is that Lewis is considering the context of

[19]Recall the discussion on page 11 where we showed that the strategy of the sceptic consists in letting us agree that we do not know that the sceptical hypothesis is false. It's this unrestricted quantification over all possibilities imaginable that causes trouble.

[20] [Lewis, 1996, p. 549]

[21] [Lewis, 1996, p. 550]

the knowledge ascribers and the propositions uttered by them.

Lewis further disagrees that justification is necessary. We rely on perception, memory and testimony and often don't know how we know. So he insists that justification should not be part of the definition of knowledge. Dretske in fact says that the way we come to know is always the locus of irrelevant possibilities. If you claim to know the car manufacturers went on strike because you read this in the newspaper you are not claiming to know that this is a reliable source even though this is assumed.[22]

Like Fred Dretske, Lewis follows Peter Unger and makes the infallibility of knowledge his starting point:[23]

> Subject S *knows* that P iff P holds in every possibility left uneliminated by S's evidence; equivalently, iff S's evidence eliminates every possibility in which not-P. [Lewis, 1996, p. 551]

What does it mean to say that every possibility in which not-P is eliminated, Lewis asks. Normally, a quantifier is restricted to a specific domain. When I say in a bar that I will buy everybody a drink, I do not mean everybody on the face of the planet. It is restricted to the people I am having a drink with, the people relevant to that specific situation. So too, it is the case with the restriction of the quantifier in the definition above. 'Every possibility' does not mean every possibility in logical space, but only those that are relevant for the truth of the knowledge statement.[24] But we cannot just ignore any possibility we choose. We can properly ignore some uneliminated possibilities, we may not properly ignore others. The definition Lewis started out with needs revision. He adds a '*sotto voce* proviso':

> 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. [Lewis, 1996, p. 554]

Lewis restricts the universal quantifier to domains determined by context, *i.e.*, in specific situations we may ignore some possibilities. This gives the verb "to know" a contextualist reading. Of course, we cannot just ignore any possibility we want, otherwise true ascriptions of knowledge would be very easy to come by. In order to distinguish between possibilities that cannot be ignored, the 'relevant alternatives' as Dretske would call them, and the irrelevant ones, Lewis formulates a set of rules that determine just that. In the next session we will explore what these rules are.

---

[22] [Dretske, 1981, p. 374]
[23] See section 2.3 for a discussion on Unger's claims.
[24] [Lewis, 1996, p. 553]

3.3 Elusive knowledge: the rules

The rules that Lewis introduces are a modal treatment of the *sotto voce proviso* we mentioned above. They can be seen as rules for the modeller or acriber of knowledge given a certain epistemic setting. In the following we will briefly introduce and discuss the rules. There are seven rules in total of which there are three rules, or prohibitions, which tell us which possibilities cannot be properly ignored and four which can be properly ignored. However, Lewis mentions that some rules could possibly be subsumed under others.

We begin with the rules that include possibilities.

Rule of Actuality

> The possibility that actually obtains is never properly ignored; actuality is always a relevant alternative; nothing false may properly be presupposed. [Lewis, 1996, p. 554]

But Lewis asks, *whose* actuality cannot be ignored? This might seem like a silly thing to ask, he says, there is but one actual world and the ascribers and the subjects of knowledge share this world, hence actuality is the same for all. However, in other situations it is not so easy.

Contrary to the standard view, Lewis insists that the objects of belief are not propositions, *i.e.*, sets of possible worlds, but properties, *i.e.*, sets of possible objects.[25] Lewis claims that every *de dicto* (or propositional) belief can be explained in terms of the self-ascription of a property. Believing that p is the self-ascription of the property that one inhabits a world where p. Some of our beliefs are about specific individuals and their location and these *de se* beliefs determine *where* one is in the world rather than in which world one thinks one is. Lewis claims that de se beliefs subsume de dicto ones.

Going back to The Rule of Actuality. When Lewis asks *whose* actuality, he means that there are some situations where information about who and where one is can be relevant.

Rule of Belief

> A possibility that the subject believes to obtain is not properly ignored, whether or not he is right to so believe. Neither is one that

[25] He argues this point at length in [Lewis, 1979a].

> he ought to believe to obtain – one that evidence and arguments justify him in believing – whether or not he does so believe. [Lewis, 1996, p. 555]

Lewis mentions here that the rule as it stands needs some refinement in terms of degrees of belief, where the degree depends on the situation. Situations where the stakes are very high and error would be especially disastrous would result in properly ignoring fewer possibilities. But even when the stakes are very high, let's say a court of law, some possibilities may still be properly ignored. For instance "that it was the dog, marvellously well-trained, that fired the fatal shot." But if the world's greatest dog-trainer was the victim's mortal enemy, the dog hypothesis would get a higer degree of belief and be relevant after all.[26]

Rule of Resemblance

> Suppose one possibility saliently resembles the another. Then if one of them may not be properly ignored neither may the other. [Lewis, 1996, p. 556]

A famous riddle in epistemology is that of the Land of Fake Barns.[27] This is a fictitious country side teeming with fake barn façades made out of paper mâche that cannot be distinguished from a real barn and we happen to gaze upon one of the few real barns. Do we know that that is a barn? According to Lewis, what is at play here is the Rule of Resemblance, and we do not know that it is a barn because I cannot ignore the possibility that I am seeing another fake one. "This possibility saliently resembles actuality in respect of the abundance of bogus barns, and the scarcity of real ones, hereabouts."[28]

However, this rule should be applied with care, Lewis warns us. The resembling should be in virtue of the rules other than this one. Otherwise, "enough little steps of resemblance can take us from anywhere to anywhere".[29] Lewis also warns that the actual world w which may not be properly ignored and any other world v not yet eliminated by the subject resembles actuality at least with respect to the subject's

[26] [Lewis, 1996, p. 556]

[27] Normally Alvin Goldman is cited in [Goldman, 1976, p. 22] who says it is originally due to Carl Ginet.

[28] [Lewis, 1996, p. 557]

[29] [Lewis, 1996, p. 556]

evidence. This can be the case even if v is very dissimilar to w, for instance in the case of a deceiving demon or evil scientists on Mars. The Rule of Actuality combined with the application of the current rule would have us consider those possibilities we tried to (properly) ignore. We therefore need to apply this rule with care.

The next rules are presumptive, they tell you which possibilities you may properly ignore.

Rule of Reliability

Processes that deliver information to us, perception, memory and testimony, may be expected to work without fault: "Defeasibly – *very* defeasibly! – a possibility in which they fail may properly be ignored". This of course does not mean that knowledge that we arrive at through one of these processes is always right, we could be hallucinating, or somebody might be intentionally lying to us:

> We do not, of course, presuppose that nowhere ever is there a failure of, say vision. The general presupposition that vision is reliable consists, rather, of a standing disposition to presuppose, concerning whatever particular case may be under consideration, that we have no failure in that case. [Lewis, 1996, p. 558]

Rules of Method

Similar to the Rule of Reliability, this rule tells us that

> We are entitled to presuppose –again, very defeasibly – that a sample is representative; and that the best explanation of our evidence is the true explanation. [Lewis, 1996, p. 558]

Rule of Conservatism

> Suppose that those around us normally do ignore certain possibilities, and it is common knowledge that they do. [...] We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us. [Lewis, 1996, p. 559]

Rule of Attention

> When we say that a possibility is properly ignored, we mean exactly that; we do not mean that it could have been properly ignored. Accordingly, a possibility not ignored at all is ipso facto not properly ignored. What is and what is not being ignored is a

> feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative. [Lewis, 1996, p. 559]

This rule can be read as the driving force behind the sceptical argument: the making salient of possibilities we may properly ignore by introducing them in the conversational context, with the result that they are no longer properly ignored. Because of this context-change potential, it makes sense to consider the perspective of dynamic semantics to describe this rule: what is the result of adding a previously ignored possibility to the context?

The generality of this rule is rather large. The rule of course applies to "the mere mentioning" of possibilities, but even goes beyond that. Consider for instance the following sentence:

(4) The mere mentioning of the name of her friend John made her wonder about his whereabouts.

or even

(5) All of a sudden she considered the possibility that he might be in America.

This underlines not only the open nature of the Rule of Attention, but also its fluidity. There need not be definite reasons for possibilities to enter the context. We will see in the formal part of this thesis that as a result of this, it is not always possible to encode in our language the result of announcements or actions.

After introducing these rules, Lewis invites us, to "[d]o some epistemology. Let your fantasies rip."[30] And you will find yourself in a context where there are uneliminated possibilities everywhere and by attending to them you are no longer ignoring them and very quickly almost no ascription of knowledge, either to yourself or to others, is true anymore.

3.4 Conclusions

In this section we have seen a semi-formal framework of attributive contextualism. What makes this framework special is the systematic modal

[30] [Lewis, 1996, p. 559]

treatment of a set of rules that govern given a knowledge claim which possibilities are relevant and need to be considered and which possibilities can be (properly) ignored. The driving force behind the rules is the Rule of Attention. This rule because of its context-change potential makes it a good candidate to find a representation by dynamic logics. However, because of the very open nature of the rule – anything mentioned gets put in the context and therefore no longer is ignored – this demands a decision on a methodological level in order to decrease the overhead a representation of the Rule of Attention would place on the formal system.

4 Finite narrative modelling

4.1 Introduction and motivation

In the previous section we have seen that the Rule of Attention allows us to bring anything to the foreground and make it part of the conversational domain. When attempting to formalise this rule one is faced with the problem of representing that which is ignored and that which we attend to. Because of the open nature of the Rule of Attention, this means that a language should be able to represent any possibility.

In this section we will discuss two approaches to this problem. We can distinguish two general modelling methodologies called *open modelling* and *closed modelling*. It is a distinction between models that continuously interact with that which it models, *i.e.*, the target, and models which targets are complete, in a sense to be determined, when the modelling starts. We will discuss the differences between the two approaches in section 4.2.

In this thesis we will follow the approach of *closed modelling* by use of a specific meta-technique we will call *finite narrative modelling*. We will give a description of this semi-formal framework in section 4.3.

After considering an example (section 4.4) we will discuss some structural properties of *finite narrative modelling* in section 4.5.

4.2 Open and closed modelling

We can distinguish between two general modelling methodologies we denote by *open modelling* and *closed modelling*.[31] These modelling methodologies

[31] Sometimes we will we use the terms *online* and *offline*, respectively, to denote the two. We have in mind here models used in formal sciences ranging roughly from philosophical logic to AI and robotics and we wish to avoid confusion with the terms "open system" and "closed system" found for instance in thermodynamics or other areas of science.

are providing the meta-level of the modelling.

Let's call that which needs to be modelled the *target* or *target system* and the result a *model* or *formal representation*. The target can be any part of reality, but, without defining this explicitly, we will roughly assume this to be a scenario, situation or story consisting of events and actions. The key difference between the two approaches is not so much the target but rather *when* the modelling is done with respect to the target. *Open modelling* is typically done whilst the target unfolds or progresses, whereas *closed modelling* is typically done when the target is complete in some relevant sense. For instance, when our target is a narrative *open modelling* will model a narrative *as the narrative progresses*, *closed modelling* will have the possibility to go *offline*, and construct a respresentation of the narrative when the narrative has reached its conclusion.

The choice of methodology has important consequences for what we might call *overhead* of the (formal) system doing the modelling. Suppose we want to model a narrative and we are confronted with an initial fact, for instance, that the story is set "in fair Verona, where we lay our scene". Whether or not the location of the story will become relevant cannot be known by the system in advance. This fact might become negligible, but it also might not. As a result this fact needs to be represented in the formal language employed by the system, together with representations of other locations, where the story might have been set. In contrast, closed modelling will consciously leave some of the modelling to the modelling practice, in the sense that the modeller will decide in advance, if the location is relevant or not. We stress again that the closed modeller only has this opportunity if the target is given to him or her in a complete form up front. If the modeller would for instance decide it is not an important fact (the story works just as well in the back alleys of Manhattan) he or she could leave out location entirely in the representation. Languages employed by open modellers tend to be much richer than those of closed modellers since they need to be able to deal with new and unforeseen situations.

Admittedly, this is a distinction between idealized approaches: hardly any system in practice is using a completely open modelling methodology. In the area of AI that studies formal representation of narratives an example that comes close can perhaps be found in [Dyer, 1983], but the model presented there can only deal with narratives that are similar to each other.[32]

[32] In the 1960s researchers in AI became intrigued by notions from narratology and started to build formal representations for narratives. [Dyer, 1983] is part of this research area of "story understanding", as it is also known, which is situated on the area where AI and computational linguistics overlap. For a more extensive overview see [Mueller, 2000].

It is also clear that even Curiosity, the mars rover, does not have a module installed that allows it to interact with traffic lights, and this is a decision made on the level of the modelling practice. Typically, one will find the open modelling approach used by people who implement AI software or build robots. The method of closed modelling, in fact, is the main methodology used in most theoretical applications of logics of action, which we shall argue for at more lenght below.

Because of the open nature of the Rule of Attention, by which anything can be introduced in the conversational discourse by mere mentioning, a decision with respect to the methodological approaches described above needs to be taken. In what follows we imagine that the dialogues, stories *etc.*, to which we can apply Lewis' rules, are given to us in a finite narrated form that the modeller can read from beginning to end before having to decide on the modelling tools.

4.3 Finite narrative modelling as a modelling technique

The general methodology of *finite narrative modelling* can be described as follows. A human modeller is given a (finite) narrative in some form (natural language text, movie, libretto, a play *etc.*) and the aim is to develop a framework in which a formal respresentation can be given of the narrative.

Narratives are in a general sense a description of the world, where 'world' is not restricted to the actual, every-day world. Narratives can be as diverse as a movie about a ficticious heroine in a fictitious part of the universe and the minutes of last Thursday's board meeting.

A common definition of *narrative* is *a description of a sequence of events.* Also in linguistics the term *narrative* seems to refer to texts in which events occur in a temporal order. We shall not define precisely what we mean by *narrative* here[33] but we will roughly assume that it means a coherent sequence of events featuring one or more agents in some narrated context, where *event* is taken to mean a change in the state of affairs.

Even before a narrative is selected to be formalised, we may assume the modeller wants to focus on something particular. For instance, he or she could be interested in the belief structure or the deontic structure or other structure of the narrative. We call this *contextual focus* and simply is the particular concepts a modeller tries to formally capture and determines, up to a point, which components of the narrative become relevant for the

See also [McCarthy, 1976] for an early programmatic memo.

[33] *Cf.* [Fisseni and Löwe, 2012] for an argument of defining *narrative* as a sequence of events.

modeller. We believe that all research is embedded in the customs and traditions of a particular research community, and that this will influence the contextual focus, although of course, not in a strict, necessary way. We note that even though this is perhaps somewhat obvious it is an important feature of *finite narrative modelling*, because it will limit the range of frameworks.

We imagine that next the modeller transforms the narrative into list form, creating a finite list of descriptions of the relevant stages of the narrative. Let $\mathcal{N}$ be a narrative, then we will denote the narrative list by $L_{\mathcal{N}}$. It could be that not all the sentences are relevant for the representation of the narrative and the modeller could decide to leave out certain sentences. Sentences could also be added to make explicit certain implicatures that could be hidden in the text. This process of transformation should probably be more likened to a soft massage rather than the execution of a computer script.

Once the modeller is satisfied with the list $L_{\mathcal{N}}$, the modeller will go through the list to look for the building blocks of his language. For instance, if the decision was made the language is based on propositional logic, then a propositional variable has to be added for each of the facts listed in $L_{\mathcal{N}}$. The gathering of the building blocks of the language, for instance, selecting propositional variables to represent facts, or constants to represent objects, is an iterative process, which is characteristic of the structure of *finite narrative modelling* as a whole, something we will point out in more detail below, and is done with one eye focussed on the framework to be developed.

At this point we have a *formal framework* $\mathscr{F}$ to our disposal, which is given by a syntax and corresponding semantics. The syntax determines a formal language, consisting of elements we have gathered in the previous step, and will also determine a class of mathematical models for the language and a satisfaction relation, in the usual sense of mathematical logic.

The penultimate step of the process is to construct models $\mathcal{M}$ of the framework $\mathscr{F}$ for each stage i of a narrative $\mathcal{N}$ given in $L_{\mathcal{N}}$, thereby giving a sequence of models $\mathcal{M}_i$, which is intended to represent the narrative $\mathcal{N}$. What exactly the procedure is a modeller follows when he or she constructs a model $\mathcal{M}$ at a particular stage i is a difficult question and is more of an art than of a science, according to some. In [van Benthem, 2009], the author, considering the construction of a model for an epistemic situation, remarks the following:

> First there is no algorithm for producing it – but most people would agree that it fits the situation, and most students are quite capable of finding models like these with just a little training. This *art of modeling*

> is a cognitive abstraction skill that many people have, a serious fact in itself. [van Benthem, 2009, p. 3]

This art or skill depends on the experience and creativity of the modeller and the more complex the situations become the more experience and creativity is required to construct an adequate representation.[34]

Finally, the modeller needs to confront the representation of the narrative with the narrative itself and check if the representation is faithful.[35] If the modeller is satisfied he or she stops, if not he or she can go back to the beginning and rework his steps and make changes anywhere in the process.

4.4 An example

Example 4.1 (A small narrative)**.** *John's mother complained over the phone that she hasn't been feeling well lately, so John decides to visit her. He takes the train and walks over to his mother's house which is quite near to the train station. He notices along the way that a new bakery has opened shop and suddenly is reminded that his mother asked him to bring some cookies to go with the coffee. He quickly buys some and arrives at his mother's house to find here in much better shape than she made seem over the telephone.*

Although we use the above as an example of a narrative we note that narratives normally tend to be much longer and are not always in a ready form to serve as an object for modelling. For instance, in a literary story, there could be flashbacks or flashforwards, stories within stories, and so on.[36] In the case of this simple example, we can easily create a list of single statements in chronological order, but this might not always be so easy and in those cases the modeller needs to do some ordering, perhaps leave out sentences and do some interpretation.

Example 4.2 (An example of the narrative in list form)**.**

[34] Commenting on this point, the authors of [van Eijck et al., 2011] try to show that some of the art is actually a science.

[35] What exactly a *faithful* representation is, is a difficult matter. It is something that cannot be defined context-independently and depends on what you aim for. We will discuss this in some more detail in the next section.

[36] This seems to be an important difference between 'humanities-centered' narratology and formal or computational narratology. Where the former will look at the structure of large bodies of text, the latter is bounded by the computational capabilities of the systems employed. [Mani, 2012, §4] names three obstacles for formal systems. As a result, in the literature on computational narratology we will normally only find narratives of a couple of sentences.

1. *John's mother complains over the phone to him that she has not been feeling well.*

2. *John decides to visit his mother.*

3. *Luckily, on the way over to her house, John notices a new bakery and is reminded by this to bring cookies, which he would have otherwise forgotten.*

4. *John brings the cookies to the visit of his mother who is in fact quite well.*

Considering the example above, there are several formal frameworks we could think of, depending on what it is exactly we are trying to represent, *i.e.*, the contextual focus. We could be interested in the beliefs of John and his mother. For example, how John's mother believes John will bring cookies but that it is only after he sees the bakery he is reminded. Similarly, we could be interested in a simple temporal representation of the events that take place, or the various relations between the two characters that are in the story, and so on.

Suppose we would proceed naively and we wish to formally represent the basic properties and relations of the narrative $\mathcal{N}$ in example 4.1. We would then go through the list $L_{\mathcal{N}}$ to gather the elements for our language.

Example 4.3 ("Formal representation" of example 4.1 just to make a point)**.** *Let $\mathcal{L}$ be a language consisting of the following components. We have a set of variables $A = \{a_1, a_2, \ldots, a_m\}$ denoting agents, a set of variables $O = \{o_1, o_2, \ldots, o_n\}$ for objects and the following sets of predicate letters $P = \{\mathsf{Son\_of}, \mathsf{Not\_feeling\_well}, \mathsf{Small\_distance}\}$ and $Q = \{\mathsf{Remembers}, \mathsf{Visits}, \mathsf{Brings}\}$. If a_1, a_2 are agents and o_1, o_2 objects, then*

$$\mathsf{Son\_of}(a_1, a_2)$$
$$\mathsf{Not\_feeling\_well}(a_1)$$
$$\mathsf{Small\_distance}(o_1, o_2)$$

are states. If a_1, a_2 are agents and o_1 is an object, then

$$\mathsf{Remembers}(a_1, o_1)$$
$$\mathsf{Visits}(a_1, a_2)$$
$$\mathsf{Brings}(a_1, o_1)$$

are events. A basic semantics for this language, based on elementary set theory, now seems straightforward. Given this toy language we could represent the story of example 4.1 as follows:

$$\boxed{\begin{array}{l}\mathsf{Son\_of}(j,m)\\ \mathsf{Not\_feeling\_well}(m)\\ \mathsf{Visits}(j,m)\end{array}} \tag{1}$$

Of course, this is at best an abbreviated summary of example 4.1, however for the moment this will be enough to discuss some important issues.

Here are some things we can remark about our examples:

1. There are elements of the narrative that cannot be represented by the language, but the framework still captures the right structure of the narrative, so enriching the language is a minor revision. For instance, if we look at (1), we see that the fact that John almost forgot to bring cookies but remembered because he saw the new bakery, is not represented. One could argue that this is an essential part of the narrative and any representation of it should represent this fact (see below).

2. There are elements of the narrative that the language is fundamentally not able to express, *i.e.*, a revision would require a change of the whole framework. For instance, if we look at our "framework", we see that we cannot express the motivation of why John is visiting his mother. The story strongly suggest that John is visiting his mother because he is concerned about her wellbeing. Adding "motivations" to the framework would require an extension of the ontology.

3. There are elements of the narrative that are intentionally left out by the modeller because they are deemed irrelevant for the representation of the narrative. In (1), the fact that the house of John's mother is near to the train station is not represented. On the face of it this does not really seem to matter for the story and although it could be represented this fact is deemed irrelevant.

With these remarks we touch on the difficult problem of determining when a certain formalisation is a faithful representation of the narrative. *Faithfulness of the representation* applies to the structure of the narrative and to the information contained in the narrative. With respect to the former we can ask what level of detail should be represented and do we have access to an objective structural core on the basis of which we could decide these matters? This notion of *granularity*, *i.e.*, how finely-grained our representation is, and hence how much of the detail of the narrative is included in the representation, hardly ever seems to be a real issue when we are considering examples like the one we have, but things change rapidly when the narratives become longer. Determining the relevant aspects for longer narratives is not a given and needs empirical grounding (*cf.* [Löwe, 2011, §4.3, §6], [Block et al., 2012]).

When we say that *faithfulness of the representation* also applies to the information that is contained in the narrative we mean the following. Our example narrative seems to imply that John's mother gave a somewhat negative report of her wellbeing in order to persuade her son to come and visit. This information is not really explicit so one can argue if it should be represented (we don't in our example). What the intuitions are and if they are shared by everyone who reads the narrative is also something which would need to be investigated.

If we compare (1) and (3), we see that in both cases a particular component of the narrative *could* be represented but is not. The difference, however, is that in case (1) this renders an incomplete representation whereas in case (3) not representing it probably does not matter.

If we compare (1) and (2), we see that in both cases it could be argued that the particular facts should be represented, but looking at our language we could probably only (easily) accommodate (1) and not (2).

In order to incorporate point (1) we would need to expand the language of example 4.3, for instance by adding a one place predicate Forget. After this a new representation would look as follows:

$$\boxed{\begin{array}{l} \mathsf{Son\_of}(j,m) \\ \mathsf{Not\_feeling\_well}(m) \\ \mathsf{Forget}(j,c) \\ \mathsf{Remembers}(j,c) \\ \mathsf{Visits}(j,m) \\ \mathsf{Brings}(j,c) \end{array}} \tag{2}$$

At this point the modeller could complain (or somebody else might) that we did not represent the reason why John decided to visit his mother (*i.e.*, the emotional blackmail from the part of the mother, (2) above) and that this is essential for the representation. If the modeller would agree that this should be represented he or she would again look for adjusments of the framework although in this case this would require a more fundamental revision: we need a way of representing that it is not the case that the mother is not feeling well, but she makes it seem that that is the case.

This iterative procedure, where we check the representation against the narrative, or target, is an example of *conceptual modelling*, and we will have a closer look at it in the following section when we discuss in more detail the structural aspects of the modelling process.

In a recent paper *finite narrative modelling* has been applied when finding formalizations for episodes of a TV crime series and can be seen as a more "real-life" example of narratives and how the technique of *finite narrative modelling* can be applied. In [Löwe et al., 2009, §3] there is a discussion on how to make modelling decisions and there examples are given showing that these decisions can be hard to make.

4.5 Structure of the formalisation process

Having looked at a description and an example of *finite narrative modelling* we now take a step back and look at the structural properties of the process. The process of *finite narrative modelling* can roughly be divided in the following three stages:

Stage 1

- Transform the narrative into list form, creating a finite list of sentences, representing the stages of the narrative.
- From these sentences "read off" building blocks for the object language: for instance propositional variables, predicate letters, *etc.*
- Define a formal framework given by a syntax and corresponding semantics.

Stage 2

- Devise an initial model and models for each stage of the narrative, thereby creating a sequence of models representing the narrative.

Stage 3

- Check the formalisation against the narrative.

Central to the methodology of finite narrative modelling is the possibility of reading the narrative in entirety up front, and the iterative nature of the process. If the formalisation (stage 3) does not give a faithful representation, the modeller can rework the steps of the modelling procedure by starting again with first stage. Figure 1 is a picure of this structure.

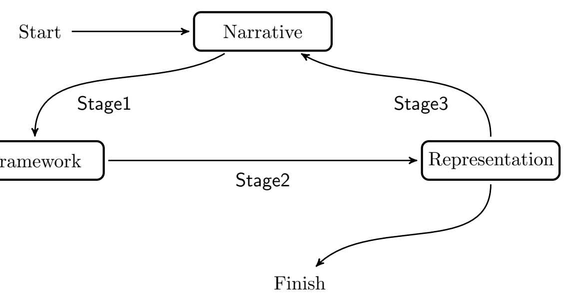

Figure 1: Structure of *finite narrative modelling*

In the papers [Löwe and Müller, 2011, Löwe et al., 2010] the authors discuss the widespread methodology of *conceptual modelling*. This methodology can be seen as a generalisation of *mathematical modelling*, but is also applied to non-quantitative research areas, for instance in philosophy.

> *Conceptual modelling* is an iterative process through which a stable equilibrium is reached between a concept or a collection of concepts as explanandum and a (somewhat) formal representation of it. [Löwe and Müller, 2011, p. 133]

This idea of using iteration towards an equilibrium involves three natural steps:[37]

1. Formal representation
2. Phenomenology
3. Assessment

The last step is where the model is confronted with the phenomena. In our setting of finding "the right" representation for a narrative this is the step where we test the model against our intuitions. For instance, suppose we model the epistemic setting of tossing a coin, where H is heads and T is

[37] [Löwe and Müller, 2011, p. 133]

tails, we would want (unless the story indicates otherwise), that in our model $\mathsf{H} \leftrightarrow \neg\mathsf{T}$ is true everywhere in the model. This procedure of translating our intuitions into the formal language and testing the model is an important part to see if the model will hold up under scrutiny. If errors occur we can analyse them and return to the first stage for revision. We point again to [Löwe et al., 2009, §3] for a more extensive discussion on this point.

4.6 Finite narrative modelling and dynamic epistemic logic

Finite narrative modelling was introduced as a technique for rendering a formal model of a narrative. The idea is that at each stage of the narrative the modeller constructs a model thereby creating a sequence of models. This sequence would then serve as a model for the entire narrative. In the following section we will discuss *dynamic epistemic logic* (DEL), which is a family of logics specifically designed to reason about information change and the flow of information. The semantics for these languages differ from other, so-called, 'static' languages: the meaning of a formula of DEL is determined by the change it has on the model in which the formula is evaluated. Without going into detail yet, given a particular epistemic action which is represented by a formula (or an action model, as we shall see) an operation on the model takes place, potentially changing it.

Give a language with dynamic semantics gives us an additional tool we can use in *finite narrative modelling.* In addition to constructing a model for the narrative at each particular stage, we can let the semantics of the dynamic language do part or all of the work. In the latter case we could build an initial model at the first stage and then construct action models, representing the epistemic action at all subsequent stages and use the semantics of the language to calculate the model for the next stage. A model for the narrative then is an initial model together with an ordered tuple of action models.

We shall make this idea precise in section 6 but we first turn to DEL and introduce the language.

4.7 A case study

In section 4.2 we have explored the differences between open and closed modelling. Open modelling is used by systems modelling targets that incomplete in a particular sense. Typically, we will find applications of this approach in AI, robotics, *etc.* In this section we will take a closer look at how closed modelling is the methodology most commonly used in applications of logics of action. Most examples in the literature of epistemic logic come in the form of small stories. These can be simple descriptions of certain (epistemic) situations to more complex situations in which a series of events occur. Logics of action and agency (see section 5 for an introduction) have been developed specifically to model the information flow resulting from action.

We will explore the paper [Baltag et al., 2008] which is representative for the DEL paradigm. Many examples are used in this survey paper to introduce and explain the different frameworks. Simple examples of a coin being tossed (§2.1) to public announcements in card games and more complex epistemic actions (§5.4–§5.5).

Although other scenario's are discussed as well, we will focus on the one centering the card games. In example 2 of §4.4 (p. 27–28) we have a setting where three agents, Amina, Bao, Chandra, denoted by a, b and c, respectively, are being dealt one card each from a deck that consists of exactly three cards: clubs, spades and hearts. In order to represent this situation a basic language is constructed consisting of atoms representing the fact that a has hearts ($Hearts_a$), b has hearts ($Hearts_b$) and so on, resulting in a language with nine atoms. After this a model (*Hexa*) is constructed, consisting of a domain representing the possible deals

$$\begin{aligned} W := \{ & \heartsuit\clubsuit\spadesuit, \\ & \heartsuit\spadesuit\clubsuit, \\ & \clubsuit\spadesuit\heartsuit, \\ & \clubsuit\heartsuit\spadesuit, \\ & \spadesuit\heartsuit\clubsuit, \\ & \spadesuit\clubsuit\heartsuit\} \end{aligned}$$

and epistemic relations representing what agents know and don't know immediately after the first deal:

$$\begin{aligned} R_a := & \{\langle\heartsuit\clubsuit\spadesuit, \heartsuit\spadesuit\clubsuit\rangle, \langle\clubsuit\spadesuit\heartsuit, \clubsuit\heartsuit\spadesuit\rangle, \langle\spadesuit\heartsuit\clubsuit, \spadesuit\clubsuit\heartsuit\rangle\} \\ & \ldots \end{aligned}$$

Semantics should for instance then reflect that in the possible world where a got hearts, $Hearts_a$ is true:

$$Hexa, \heartsuit\clubsuit\spadesuit \models Hearts_a$$

i.e., there is a valuation function V from the propositions to the powerset of the domain such that $V(Hearts_a) = \{\heartsuit\clubsuit\spadesuit, \heartsuit\spadesuit\clubsuit\}$.

It is shown that this situation is fairly straightforward to represent and results in a nice compact model (p. 28). But this is in part due because many modelling decisions have already been taken when the model is constructed. The fact that it was decided that there are three agents with a deal of one card each out of a three-card deck with three suites allows us to consider a language with nine propositional letters. The definition of the general language (§4.3, p. 24) says that "our set of atomic propositions is taken to be arbitrary". But when we actually model this situation we consider a finite fragment because the modeller knows this will suffice.

This scenario is being used as a basis for more complicated settings: public announcements, whispering, showing cards, and so on. And in the process new languages and models are introduced and explained. This way of introducing a subject should of course not be confused with the modelling technique, and if we take a look §5.5 we can make our point more clearly.

In §5.5 (p. 52) of the paper, the full general set up is defined (pp. 54–55) following [Baltag and Moss, 2004] and it is a generalization of the examples previously considered in the paper. Although this general framework was built up using concrete examples, once it is defined the way back to applying this framework to concrete settings is not a given. Suppose we take these definitions, blank our memories and are then confronted with Amina, Bao and Chandra playing cards, in order to give precise representations of this situation we would need to let them play for a while to see what happens. What game are they playing? Is anyone cheating? and so on. We stress that this is not a matter of *expressibility* where we say that the framework defined cannot model certain things (that would be cheap indeed), we say that ending up with a nice, uncluttered model needs knowledge about the situation up front.

The authors write:

> Our tasks as modelers are (1) to provide an adequate representation of this scenario; (2) to use the representation as part of a formal account of "knowledge" and related terms; (3) to see where the representation and formal account run into problems; (4) to then "scale up" all of the previous points by considering more complicated scenario's models, and accounts, with the same goals in (1)-(3). [Baltag et al., 2008, p. 4]

and it is reminiscent of the stages we have described in sections 4.3–4.5 with the important distinction that where we set out to find a framework for modelling narratives, in this paper the examples that are considered are used as part of an account of a particular concept (*i.e.*, "knowledge"). But it is clear that in order to find the right models for particular situations "one must test models and semantic definitions against intuitions, that the proof of the pudding is in the eating." (p. 76) and this is mostly done using closed modelling.

5 Dynamic Epistemic Logic

With *dynamic epistemic logic* we denote a collection of (logical) languages that are used to describe *information change*. The term 'dynamic epistemic logic' can be split up in 'dynamic' and 'epistemic logic'. Epistemic logic is the systematic investigation of the concept of knowledge of individuals. Since the 1930s, starting with the work of Pierce, a lot of research was done on modalities and possible world semantics, and it is with the work of von Wright [Von Wright, 1951] and subsequently the work of Hintikka [Hintikka, 2005] that the formal investigation of reasoning about knowledge was initiated.[38] Hintikka's work can hardly be overestimated. The publication of [Hintikka, 2005] not only initiated research in philosophy, but the impact stretches to computer science [Fagin et al., 1995], artificial intelligence [Meyer and van der Hoek, 1995] and game theory [Aumann and Brandenburger, 1995], all of which are research areas that are still very much alive today.

The system found in [Hintikka, 2005] is based on modal logic, where the epistemic and doxastic operators are analyzed in Kripke frames. What makes this system different from *dynamic epistemic logic*, is that 'plain' epistemic logic is 'static'. This means that given a situation it deals with the epistemic content of a set of subjects at a specific *moment* in time, whereas *dynamic epistemic logic* sets out to model the change of epistemic information available to a set of subjects over a specific *period* in time. This dynamic turn was inspired by developments in computer science, logical semantics and belief revision and a first concrete step towards a *dynamic* epistemic knowledge was made by Jan Plaza who introduced a logic of *public announcement* in [Plaza, 1989].

In what follows we will first introduce basic definitions and *common knowledge* (section 5.1). In section 5.2 we will look at epistemic actions and action models.

[38] *Cf.* [Hendricks and Symons, 2009]. Also see [Copeland, 2002] in which a historic reconstruction of the emergence of ('modern') possible world semantics can be found.

5.1 Epistemic logic with common knowledge

As mentioned in the introduction, *dynamic epistemic logic* is concerned with reasoning about change of knowledge of multiple agents and is essentially dynamic in nature. However, we will first introduce the basic language which is a *static* language. In what follows we fix a set $\{p_0, p_1, \ldots, \}$ of propositional letters we denote with PROP and a set $\mathcal{G}$ of agents $\{a_1, \ldots, a_n\}$. The language of multi-agent epistemic knowledge, the most basic system, which we will denote with $\mathcal{L}_0$, is defined by the following rule:

Definition 5.1 (The basic epistemic language $\mathcal{L}_0$)**.**

$$\varphi := p \mid \bot \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid K_a\varphi$$

where a ranges over $\mathcal{G}$ and p a propositional variable ranging over PROP.

Basically, $\mathcal{L}_0$ is a propositional language to which an epistemic operator K_a is added for each $a \in \mathcal{G}$. Throughout the work we will use common abbreviations like $\psi_1 \to \psi_2 := \neg\psi_1 \vee \psi_2$, $\top := \neg\bot$ *etc.* Intuitively, we will take $K_a p$ to mean 'agent a knows that p'. Other constructions we can make is $\neg K_a \neg p$ which would be a translation of 'agent a considers it possible that p and $K_a p \vee K_a \neg p$, which would be a translation of 'agent a knows *if* p'.

Formulas of $\mathcal{L}_0$ are evaluated on *Kripke frames* to which a valuation function is added.

Definition 5.2 (Kripke frame)**.** A *Kripke frame* $\mathcal{F} = \langle W, \{R_a \mid a \in \mathcal{G}\}\rangle$ is a pair consisting of the following elements:

- A non-empty domain W,
- A set of binary relations $\{R_a \mid a \in \mathcal{G}\}$ on W, for each $a \in \mathcal{G}$, sometimes called *indistinguishability relations* or *epistemic relations.*

Definition 5.3 (Epistemic model)**.** An epistemic model $\mathcal{M} = \langle \mathcal{F}, V\rangle$ for $\mathcal{L}_0$ is a pair consisting of a Kripke frame $\mathcal{F}$ and a valuation function $V : \mathsf{PROP} \to \wp(W)$, where $V(p) = X$, some $X \subseteq W$, is taken to mean 'p is true in all $w \in X$'. If w is an element of the domain and $\mathcal{M}$ an epistemic model we call $\langle \mathcal{M}, w\rangle$ a *pointed epistemic model.* The epistemic relations are commonly thought of as equivalence relations rendering an S5 model.[39]

Truth of a formula is then defined as follows:

[39] An equivalence relation is a relation which is symmetric, reflexive and transitive and corresponds to accepting the following axioms (for all $p \in \mathsf{PROP}$) describing properties

Definition 5.4. Let $\mathcal{M}$ be a model for $\mathcal{L}_0$, w an element of the domain, and $\varphi \in \mathcal{L}_0$. We recursively define the notion of a formula φ to be true at w in $\mathcal{M}$, which we shall write as $\mathcal{M}, w \models \varphi$, as follows:

$$\begin{array}{llll}
\mathcal{M}, w & \models p & \text{iff} & w \in V(p) \\
\mathcal{M}, w & \models \bot & \text{iff} & \text{never} \\
\mathcal{M}, w & \models \neg\varphi & \text{iff} & \text{not } \mathcal{M}, w \models \varphi \\
\mathcal{M}, w & \models \varphi \vee \psi & \text{iff} & \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi \\
\mathcal{M}, w & \models K_a\varphi & \text{iff} & \text{for all } v \text{ such that } R_awv : \mathcal{M}, v \models \varphi
\end{array}$$

With these definitions we are well equipped to describe first-order and higher-order knowledge of agents given an epistemic situation. The epistemic operators that we have introduced so far do not change the information states of the agents involved but only describe them. This language is therefore a static language. With this language we are able to express what agents consider possible. This includes facts and knowledge of other agents.

Common knowledge

Apart from making statements about individual agents' knowledge we often make claims about the knowledge of agents *in a particular group*. Some of the properties that hold for individual agents do not necessarily hold for the knowledge of a group of agents. For instance, if everybody in a group $E \subseteq \mathcal{G}$ knows that p, it is not without saying that all members of E know that all members of E know that p. John and Mary might both know the election outcome but they might not know of eachother that they know. They might not even know eachother. This example can be seen as an instance of a *general knowledge*. We say that E has *general knowledge* of φ iff for all $a \in E : K_a\varphi$. All members of the group know that φ, and we will write $N_E\varphi$ for this notion.

Another important group notion of knowledge is common knowledge. A formula φ is common knowledge if everybody knows that φ, everybody knows that everybody knows that φ, and so on. Mathematically, this could

---

knowledge is thought to have and defining the class of equivalence frames:

$$\begin{array}{rr}
K_ap \to K_aK_ap & \text{Positive introspection} \\
K_ap \to p & \text{Veridicality} \\
\neg K_ap \to K_ap\neg K_ap & \text{Negative introspection}
\end{array}$$

However, in section 6, we will not require the epistemic relations to be transitive on the level of frames, something which is needed for the specific set-up.

be described by an infinite conjunction

$$C_E\varphi = \bigwedge_{n=0}^{\infty} N_G^n\varphi$$

Extending $\mathcal{L}_0$ by adding C_E as an operator would give a new rule for generating a language:

Definition 5.5 ($\mathcal{L}_C$: the language $\mathcal{L}_0$ with common knowledge)**.**

$$\varphi := p \mid \bot \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid K_a\varphi \mid C_E\varphi$$

where, a ranges over $\mathcal{G}$, $p \in \mathsf{PROP}$ and $E \subseteq \mathcal{G}$.

In order to interpret this new formula we need to consider that common knowledge can be expressed in terms of knowledge which in turn is defined by the epistemic relations. Considered this way, general knowledge of a group E, N_E, can easily be obtained by taking the union of the relations R_a, for all $a \in E$:

$$R_{N_E} := \bigcup_{a\in E} R_a$$

In order to interpret common knowledge we first define the notion of *transitive closure* of a relation R:

Definition 5.6. Let $R \subseteq W \times W$ be a binary relation. We recursively define the *transitive closure* of a relation R, denoted by $(R)^{\text{tcl}}$, as follows:

$$\begin{aligned}
(R)^0 &:= R \\
(R)^{n+1} &:= (R)^n \cup \{\langle w,u\rangle \mid \exists v \in W(\langle w,v\rangle \in (R)^n \text{ and } \langle v,u\rangle \in (R)^n)\} \\
(R)^{\text{tcl}} &:= \bigcup_{n\in\mathbb{N}} (R)^n
\end{aligned}$$

If no confusion can arise we will sometimes omit the brackets and write the transitive closure of R as R^{tcl}.

We can now extend definition 5.4 with the following clause for common knowledge:

$$\mathcal{M}, w \quad \models C_E\varphi \quad \text{iff} \quad \text{for all } v \text{ such that } R_{N_E}^{\text{tcl}}wv : \mathcal{M}, v \models \varphi$$

Although perhaps well equipped to talk about what agents and groups consider possible, the languages $\mathcal{L}_0$ and $\mathcal{L}_C$ are *not* dynamic because they are not able to model the effects of an information-bearing event. In the next section we will have a look at such epistemic update languages and the epistemic events and actions they aim to model.

5.2 Epistemic events and action models

We mentioned in the introduction (section 5) that Dynamic Epistemic Logic is the study of how to reason about epistemic information over a given period of time. The basic language is a modal language and is extended in various ways to capture specific epistemic actions or events. The prime example here being *communication*.

Public announcement

Let's consider the following situation.

Let a be an agent who does not know that p. In symbols we can write this as follows

$$\neg K_a p$$

We know from definition 5.4 that this means that agent a cannot distinguish between a world where p is false and a world p is true. Figure 2 is a model of this situation.

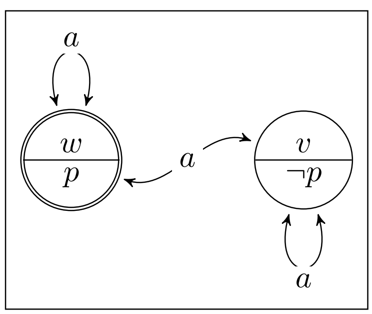

Figure 2: Agent a does not know that p

We see that the epistemic relation is represented by arrows between the worlds w and v. The arrows are labeled with the name of the agent, in this case agent a. The epistemic relation is an equivalence relation, here only the reflexive and symmetric relation is drawn. In future, in order to prevent our pictures from clodding too much with arrows we sometimes omit the reflexive and transitive arrows, unless this is a possible source of misunderstanding. We assume that world w represents the actual world, *i.e.*, the way things in fact are. In the picture, this is represented by the double circle.

Next, suppose that the (truthful) public announcement is made that p. This means that after the announcement agent a knows that p (because we assumed p to be the case). What happens is that the information conveyed by the public announcements causes the agent to change from a state where he or she does not know p to a state where he or she does. Figure 3 is a depiction of this event.

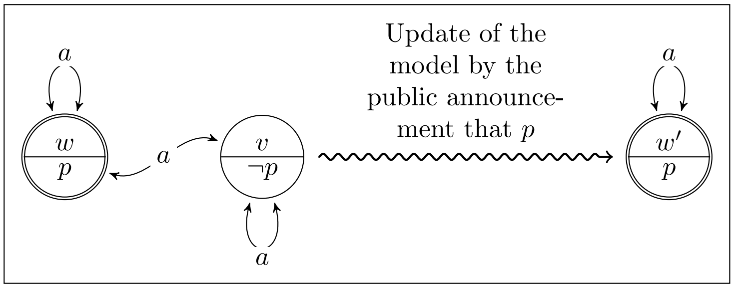

Figure 3: Update by public announcement

We see here a change from a model where in world w, $K_a p$ is false to a model where in world w', $K_a p$ is true. As we can see this is achieved by cutting away the world where p is false.

Public announcements have first been studied in [Plaza, 1989], but it is not the only epistemic action we can think of. For instance, we have assumed the public announcement to be truthful. What would be the result if we drop this assumption and consider (in)deliberate, false public announcements? Or announcements only to specific groups of agents?

In what follows we will introduce a specific strand of *dynamic epistemic logic*, namely *action model logic*. *Action model logic* allows us to model more complex epistemic actions than public announcement. Although we have used public announcement as an example of an epistemic action we will not give a full treatment of public announcement logic. Instead we refer the reader, in addition to already mentioned literature, to [van Ditmarsch et al., 2007, chapter 4] for a full development. The reason for this decision is that in section 6 and thereafter, we will use *action model logic* (and not public announcement logic). Also, we will be able to express public announcements with *action model logic*.

Action model logic

This logical system was first introduced in [Baltag et al., 1998], and subsequently developed in papers such as [Baltag, 1999] and [Baltag and Moss, 2004]. The idea behind action models is that actions are Kripke models representing the agents' uncertainty about the current action and it's effect, much in the same way as epistemic models represent the uncertainty of an agent about the facts of the world at a particular moment.

A simple example would be the following. Suppose the men's final of *Roland Garros* was between Roger Federer and Rafael Nadal and agent a

does not know if Federer has won. (*Cf.* figure 2 for such a situation.) A friend has (truthfully) written the outcome of the test on a piece of paper and put it in an envelope and a knows that the friend has done this. Opening the envelope and reading the piece of paper results in either a knowing that Federer has won, if Federer did in fact win, or a knowing that Federer has not won, if he lost. In other words, a can *distinguish* the possible outcomes of the action of reading the piece of paper. Figure 4 is an example of an action model for the action of reading the piece of paper.

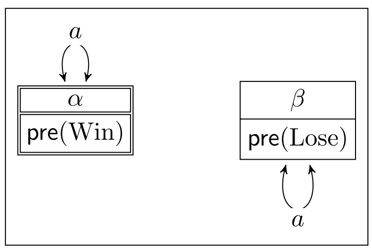

Figure 4: Action model of reading the result of the tennis match.

We see that the model is a Kripke model. The worlds are so-called *action points*. In this particular example α represents the action of reading that Federer won, wheras β that of Federer losing. Because in our example a can distinguish between the outcome of the actions there is not relation between the two action points. This would be different if, for instance, a would see another agent open the envelope and read the piece of paper without a being able to see. The preconditions of those actions should match the facts, *i.e.*, the precondition for reading that Federer won is that he in fact did so. Again, the double lines here represent the assumption of the actual state of affairs.

The main operation is the product update between an epistemic model and an action model resulting in a model in which the effects of the epistemic action have been executed. It should be the case that after a's reading that Federer won, a has gained knowledge of this fact.

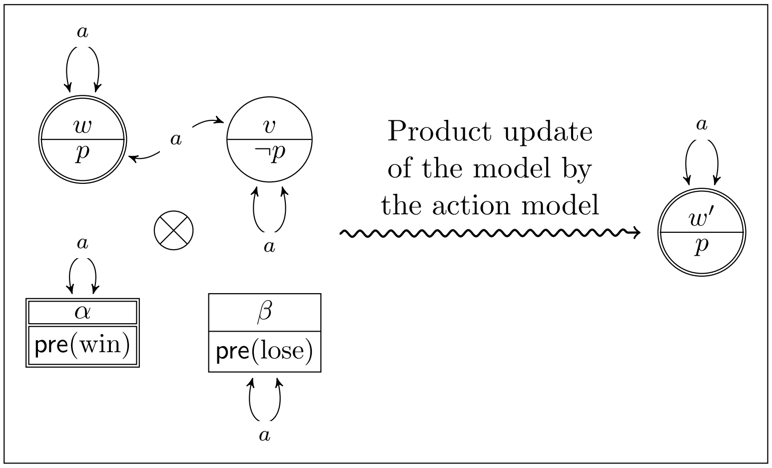

Figure 5: Product update

Definition 5.7 (Action model)**.** Let $\mathcal{L}_{C\otimes}$ be the language of action models we will define in definition 5.8. An action model $\mathcal{A} = \langle A, S_a, \mathsf{pre}\rangle$ is a triple where

- A is a set of *action points*, denoted by lower case greek letters,
- S_a is an equivalence relation on A, for each $a \in \mathcal{G}$,
- $\mathsf{pre} : A \to \mathcal{L}_{C\otimes}$ is a *precondition funtion* that maps actions $\alpha \in A$ to preconditions $\mathsf{pre}(\alpha) \in \mathcal{L}$.

We call a *pointed action model* a structure $\langle \mathcal{A}, \alpha\rangle$, where $\alpha \in A$.

Next we define the language of action models.

Definition 5.8. [The language $\mathcal{L}_{C\otimes}$] Let $\mathcal{L}_C$ be the logical language defined in definition 5.1 on page 40 for a fixed sets of agents $\mathcal{G}$ and propositional letters PROP. Then the language is generated by the following rule.

$$\begin{array}{ll} \varphi & := p \mid \bot \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid K_a\psi \mid C_B\psi \mid [\mathfrak{A}]\psi \\ \mathfrak{A} & := \langle \mathcal{A}, \alpha\rangle \end{array}$$

where $p \in \mathsf{PROP}$, $a \in \mathcal{G}$, $B \subseteq \mathcal{G}$.

The semantics of formulas and actions are then the same as in definition 5.4 with the following clause added for actions:

Definition 5.9.

$$\mathcal{M}, w \models [\mathfrak{A}]\varphi \quad \text{iff} \quad \mathcal{M}, w \models \mathsf{pre}(\alpha) \Longrightarrow \mathcal{M} \otimes \mathcal{A}, \langle w, \alpha \rangle \models \varphi,$$

where $\mathcal{M} \otimes \mathcal{A} = \langle W', \{R'_a\}_{a \in \mathcal{G}}, V' \rangle$ is a modal product defined as follows:

- $W' = \{\langle w, \alpha \rangle \mid w \in W \ \& \ \alpha \in A \ \& \ \mathcal{M}, w \models \mathsf{pre}(\alpha)\}$
- $\langle w, \alpha \rangle R'_a \langle v, \beta \rangle :\Leftrightarrow wR_a v \ \& \ \alpha S_a \beta$
- $\langle w, \alpha \rangle \in V'(p) :\Leftrightarrow w \in V(p)$

We refer to [Baltag and Moss, 2004] for very specific examples.

What is important to remark at this point is that although action model logic is designed to model actions and the change actions can have on a model, this might no always be enough when we model context change. We remember the discussion of the Rule of Attention in section 3.3.

5.3 Conclusions

We have seen a brief introduction to a *dynamic epistemic logic*, which is one of the main frameworks for reasoning about knowledge and information change. In particular, we focussed on *action model logic*. We have seen that the semantics of *dynamic epistemic logic* uses Kripke models, where we have a set of worlds, an accesibility relation and a valuation function. Similar to the the *relevant alternatives theory* we discussed in section 2.3, *dynamic epistemic logic* analyzes knowledge by considering the possibilities consistent with a statement. If an agent a knows p, this means that p is true throughout the worlds he or she can access. An alternative to p has to be eliminated in order to gain knowledge that p. Knowledge of an agent a is seen as what is true throughout the current range of uncertainty.[40]

We have focussed on this particular framework because it is designed to reason about information change, making this very apt for the context change potential of the Rule of Attention.

6 Contextual models

6.1 Introduction and notation

In section 3.3 we have seen that according to Lewis that a mention of a possibility can make that possibility part of the conversational context and thus

[40] *Cf.* [van Benthem, 2008, section 2].

potentially overthrow knowledge. With DEL in mind, the Rule of Attention, as it was called, can be seen as the contextual effect of announcements: if a possibility p is raised, it can no longer be ignored, hence all actions that explicitly involve p, will have a contextual change as a result.

Looking at real life examples or narratives in general we notice that this is not the only way contexts might change. For instance, in narratives it is quite common to encounter phrases like "all of a sudden the heroine remembered that ...", or "the smell of freshly baked bread made him aware that ...". Or that a in story line things are all of a sudden the case and have been the case the whole story without the reader being able to be aware of this: "By looking into his eyes she realized he was in fact her long lost son."

When these examples are considered in the context of narratives they are not unusual and not even that problematic. Apparently, when we follow a narrative we without problem are able to accept new facts or perspectives. But from a modelling perspective this is different. Because anything can happen the model would need to somehow be able to represent even the most unexpected. This is difficult for any approach of *open modelling* and this is why we have decided to follow the methodology of *finite narrative modelling*.

In what follows we will introduce two different settings. In section 6.2 we will explore *context models* that are able to deal with the full freedom of the narrative where "anything can happen". In section 6.3 we will consider a restriction to this setting by investigating *assumption function models* that only allow change on the level of propositional variables.

6.2 Context function models

The intuition we have is that we have a model and some notion of context that changes as the narrative progresses. So, proceeding naively we will define a context function that selects subsets of the domain for each stage of a narrative. The idea is that the models give 'contextual snapshots' at each stage of the narrative by selecting a subset of the domain. These states will be used to model the epistemic states in the narrative in chronological order.

Throughout we will fix a finite propositional language with N propositional variables $\{p_1, \ldots, p_N\}$ and a set $\mathcal{G}$ agents. We recall *action models* from definition 5.7. We further recall *pointed epistemic models* from definition 5.3, section 5.1, with the exception that we do not require the epistemic relations to be transitive. Formally, we are considering the class of models whose underlying frames are reflexive and symmetric. We will explain below why we drop transitivity and it will be made clear that we do this only ‘temporarily’.

We call a sequence $\vec{\mathcal{A}}, \vec{\alpha} := \langle\langle \mathcal{A}_1, \alpha_1\rangle, \ldots, \langle \mathcal{A}_n, \alpha_n\rangle\rangle$ of pointed action models a *DEL narrative.*

Definition 6.1 (Narrative model)**.** Let $\langle \mathcal{M}, w\rangle$ be a pointed epistemic model and $\langle \vec{\mathcal{A}}, \vec{\alpha}\rangle$ a DEL narrative. We call a tuple $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}\rangle$ a *narrative model.* If $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}\rangle$ is a narrative model, we can read off n epistemic states recursively by:

$$\begin{aligned} \mathcal{M}_0 &:= \mathcal{M}, & w_0 &:= w, \\ \mathcal{M}_{i+1} &:= \mathcal{M}_i \otimes \mathcal{A}_{i+1}, & w_{i+1} &:= \langle w_i, \alpha_{i+1}\rangle. \end{aligned}$$

where $\mathcal{M}_i \otimes \mathcal{A}_{i+1}$ just is the product update we have defined in definition 5, section 5.2. We write W_i for the underlying set, or domain, of $\mathcal{M}_i$ and R_i^a for the epistemic relation of agent $a \in \mathcal{G}$ in $\mathcal{M}_i$.[41]

Definition 6.2 (Context function model)**.** Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}\rangle$ be a narrative model such that the length of the DEL narrative is n. We use the notation of the last paragraph for the sequence of epistemic models generated by $\langle \mathcal{M}, w, \vec{A}, \vec{\alpha}\rangle$. Then any sequence $\langle C_0, \ldots, C_n\rangle$ is called a context function if for all i, $C_i \subseteq W_i$ and $C_i \neq \emptyset$. We sometimes call C_i the *context set* C of $\mathcal{M}$ at i.

We interpret C_i as that subset of the model at i that contains the worlds that cannot be ignored, whereas the worlds v outside C_i, *i.e.*, $v \in W_i \setminus C_i$, are the worlds that can be ignored. This distinction on the domain will be our basis for determining whether an agent knows a particular fact or not.

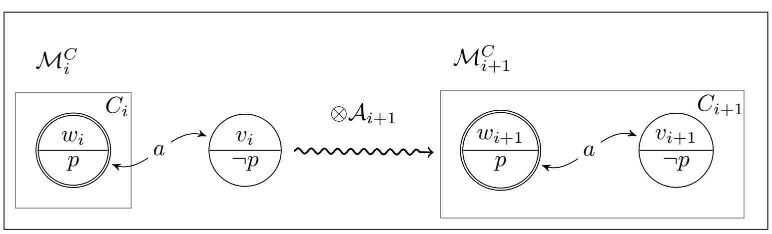

Figure 6: Contextual change in context function models

As mentioned in the beginning, the context function selects subsets of the domain at the different stages of the narrative. Figure 6 is a small example

[41]In order to have consistency with the indices for epistemic states we have changed the agent index of relation symbols from subscript to superscript.

where the context set expands from a situation where only w is included in the context set to a situation where also v is included. We see that the product update did not change the model, only the context set changed.

Given a *context function model* $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C\rangle$, determining truth of a formula φ in w comes down to considering the model restricted by the context function.

Definition 6.3 (Semantics of context function models)**.** Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C\rangle$ be a *context function model*. For each i we can now determine a new epistemic state as follows

$$\mathcal{M}_i^C := \langle C_i, \{(R_i^a \upharpoonright C_i)^{\mathrm{tcl}}\}_{a \in G}\rangle.$$

where

- $R_i^a \upharpoonright C_i := \{\langle w, v\rangle \in R_i^a \mid w \in C_i \text{ and } v \in C_i\}$,
- $(R_i^a \upharpoonright C_i)^{\mathrm{tcl}}$ is the transitive closure according to definition 5.6.

We now define truth of a formula φ in a *context function model* as follows

$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C, i \models \varphi :\Leftrightarrow \mathcal{M}_i^C, w_i \models \varphi$$

where the latter are the semantics defined in definition 5.4.

Remember that our epistemic models do not have transitive accessibility relations. In the following example, we shall show why we need this slightly unusual set-up.

Example 6.1 (Don't forget the coffee!)**.**

1. *John is walking in the supermarket. That morning, he was given strict instructions by his wife not to forget to buy coffee.*
2. *Walking down the aisle he receives a text message from her: "DON'T FORGET THE SUGAR!!!", it reads. Because John's wife was rather explicit that morning that they needed coffee, poor John was now quite confused and no longer was certain what they were out off. Maybe they needed both? Or neither? But what then?*
3. *After gazing at the many kinds of candy for a moment he considered that John's wife was quite overworked finishing a project and he had noticed that sometimes she would make small mistakes like that. He then decided that what she meant was "DON'T FORGET THE COFFEE!!!" and so, for a treat and because she clearly needs it, he bought some of the expensive Italian coffee to support her.*

This small narrative, already broken up in three stages, shows two changes in context. The first is caused by the confusing text message by John's wife, seemingly contradicting her earlier instructions to buy coffee. The second contextual change is caused after John reflects on the situation and decides that his wife simply sent an erroneous message. In figure 6 we see the changing of the context set. In the first model only the worlds w and v are included in C_1. Taking the transitive closure of this set gives us a set where $\mathcal{M}_1^C, w \models K_J c$.

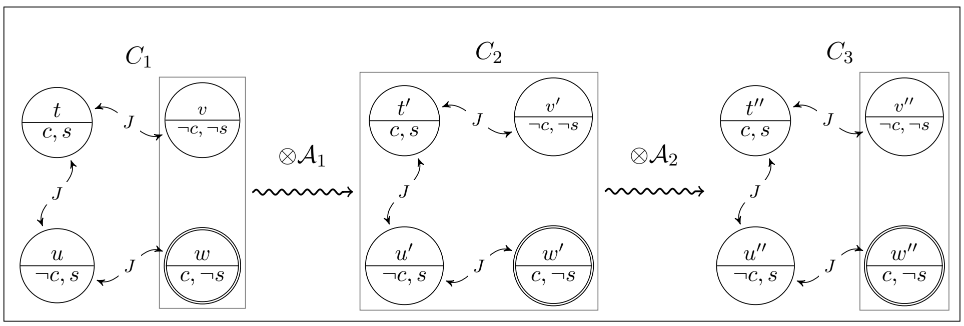

Figure 7: Contextual change in a small narrative

The first contextual change widens the context set to include all worlds of the domain. This represents the situation where John lost all knowledge. Taking the transitive closure of C_2 gives us a model where $\mathcal{M}_2^C, w \models \neg K_J c \wedge \neg K_J s$. The second contextual change restores the initial situation when John considers that exhaustion must have caused his wife to send a wrong message.

There are two points to make here. The first point is that the change of John's knowledge is caused solely by the contextual change, *i.e.*, the underlying model stays the same in all i. This means that we did not use the capabilities of the action model language to update the model and we could just have well used a simpler language like $\mathcal{L}_C$. We did not specify the action model but the "empty action" or "identity action" *where nothing happens* can for instance be given by the following action model $\mathfrak{A} := \langle \{\alpha\}, S_J, \top \rangle$. Of course, "nothing happening" here means that the product update does not have any effect on the epistemic model. However, in principle it is possible to have change on two levels: on the level of the epistemic models by means of the product updates and on the level of the context set by means of the context function. For the moment we note that the two function completely independently.

The second point we need to make is about the non-transitivity of the frames. Although our underlying model does not have transitive relations, the evaluation of formulas occurs on the transitive closure of these models, this is the case in particular for the K-operators, and therefore standard semantics apply.

Observe that if our models would be **S5** we would have $\mathcal{M}_i^C, w \not\models K_J c$, for all $i \in \{1, 2, 3\}$. In $\mathcal{M}_2$, John is completely confused by the text message and no longer knows whether c or s is true. In particular, John cannot distinguish between the worlds w' and v'. Because there is an R_J-path from w' to v' via u' and t', the transitive closure of the context function model ensures the indistinguishability of w' and v'. If there would be an R_J-path from w' to v', then w' and v' would also be indistinguishable in $\mathcal{M}_1$ and $\mathcal{M}_3$, but in those contexts we actually need w', v' to be distinguishable for John in order to represent that he knows what to buy from the store.

Because the first contextual change represents that more possibilities are being considered, these possibilities need to already be present in the model and would already be tangled up in indistinguishability and representing knowledge would not be possible. This means that the initial models would need to be constructed with the entire narrative and change of context in mind. This will be made clear in this and subsequent sections.

We noticed in the previous example that because we have put no constraints on the context function, the context *change* at a particular stage is not accounted for, *i.e.*, there is no connexion between the epistemic action and the change of context. This means that the context set at a particular stage can be any subset of the domain. We will now consider another example and show that even though we have complete freedom with our context function this is nevertheless an unwanted situation.

Example 6.2.

1. *Bob is sitting on a chair. He knows his own name.*
2. *Absolutely nothing happens.*

Admittedly this example is a bit contrived, but still shows that in our set up we could have the following sequence, where the action is again the empty action:

$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C, 1 \models [\alpha_1] K_b p$$
$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C, 2 \models \neg K_b p$$

Figure 6 on page 49 could be used as an illustration of this situation. According to our intuition and, more importantly, the definition of the product update, Bob should know his name at $i = 2$. This is not the case because apparently the context function does not listen to any epistemic events that might occur but has a mind of itself.

This freedom is nice, as we have seen, if we want to model the effect of phrases like "all of a sudden ..." and the like, but can also work against us, as the above example clearly shows. But we can just define the context function, so why should it be a problem? This is of course true, but this also means that we would need to always define the context function entirely dependently on the specific situation we are trying to represent. Instead what we would like is that we are able to define the context function independently of the specific application at hand.

In the next section we will introduce *assumption function* that can "work together" with action models.

6.3 Assumption context models

In order to bring the context functions under control we will consider the context sets to be specific sets of propositions, namely the propositions whose truth value is assumed. Widening or restricting context will correspond to lifting or taking up assumptions respectively. The extreme cases being the case where the agent has an assumption about all values of the propositional letters, in which case the context set would consist of a single world, and the case where the agent would have no assumptions whatsoever about the values of propositional letters, in which case the context set would equal the entire domain of the model.

Definition 6.4 (Assumption function)**.** An *assumption function* is a function

$$D : \{p_1, \ldots, p_N\} \to \{p_i \mid 1 \leq i \leq N\} \cup \{\neg p_i \mid 1 \leq i \leq N\} \cup \{\top\}$$

such that $D(p_i)$ is either p_i, $\neg p_i$ or $\top$. We interpret $D(p_i) = p_i$ as "p_i is assumed to be true", $D(p_i) = \neg p_i$ as "p_i is assumed to be false" and $D(p_i) = \top$ as "there is no assumption about p_i".

Using the assumption functions we next add them to our narrative models.

Definition 6.5 (Assumption context model)**.** Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha} \rangle$ be a narrative model and $\vec{D} := \langle D_1, \ldots, D_n \rangle$ a sequence of assumption functions. We define a context function $C_i^{\vec{D}}$ as follows:

$$C_i^{\vec{D}} := \{w \in W_i \mid \text{for all } j, \mathcal{M}_i, w \models D_i(p_j)\}.$$

We then call a tuple $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, C^{\vec{D}} \rangle$ an *assumption context model.*

Observe that given an assumption context model, for every i, $C_i^D \subseteq W_i$, and therefore the notion of an assumption context model is in fact a special case of the context function models we defined in definition 6.2, section 6.2.

Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha} \rangle$ be a narrative model and suppose we have N propositional variables p. Because C simply maps to the subsets, all possible susbsets are in its range. This is different for the assumption context function that selects possible worlds that satisfy the current assumptions at i.

For instance, let $N = 2$ and let $X \subseteq W_i$ consist of the worlds w, v such that

$$
\begin{aligned}
\mathcal{M}, w &\models p \\
\mathcal{M}, w &\models q \\
\mathcal{M}, v &\models \neg p \\
\mathcal{M}, v &\models \neg q
\end{aligned}
$$

Then there exists no D_i such that $C_i^D = \{w, v\}$. The only possibility including to include worlds w and v is to let $D_i(p) = D_i(q) = \top$, but then in fact $C_i^D = W_i$, the entire domain. Setting the assumption function to make an assumption about p or q would result in either w or v.

In general then, exactly those subsets X are not captured by the assumption functions, which have worlds $w, v \in X \subsetneq W_i$ s.t. $\mathcal{M}, w \models p_j$ and $\mathcal{M}, v \models \neg p_j$, for all p_j.

Claim. Let $C^{\vec{D}}$ be a sequence of assumption functions. Let $X \subsetneq W_i$ and $X \neq \emptyset$. Note that $D_i = W_i$, by defining $D_i(p_j) = \top$, for all $p_j \in \{p_1, \ldots, p_N\}$. Let w be the world where all propositions are true, i.e., w such that $\mathcal{M}_i, w \models p_j$ for all p_j, $1 \leq j \leq N$, and similarly let v be the world where all propositions are false. If $X \subset W_i$ such that $w, v \in X$, then there is no D and no i such that $X = C_i^{\vec{D}}$.

Proof of claim. Let w be the world where all propositions are true, and v the world where all propositions are false. Let $w, v \in X \subsetneq W_i$. Suppose there is a D and an i such that $C_i^D = X$. If it is not the case that $D_i(p_j) = \top$ for all j, then there is a p_k such that $D_i(p_k) = p_k$ or $D_i(p_k) = \neg p_k$. Suppose the former, then $\mathcal{M}_i, v \not\models D_i(p_k)$ hence $v \notin X$. Contradiction. Suppose that $D_i(p_k) = \neg p_k$, then, similarly, $\mathcal{M}_i, w \not\models D_i(p_k)$ hence $w \notin X$. Contradiction. So, it must be the case that $D_i(p_j) = \top$ for all j. But then, $C_i^D = X = W_i$, contradicting our assumption that $X \subsetneq W_i$. $\square$

Note that while the class of assumption function models is more restrictive than that of the context function models, we still retain the generality mentioned in section 6.2: there is no direct link between the action and the context change it produces. For instance, if we consider again example 6.2 of section 6.2 above but only this time apply the assumption context function by defining $D_i(p) = p$ and $D_{i+1}(p) = \top$, we would have a similar situation, where the following would still be true:

$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, D, i \models [\alpha_{i+1}]K_a p$$
$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, D, i+1 \models \neg K_a$$

In this particular situation we see that the epistemic action of "absolutely nothing happening" should be reflected by the assumption context function. What is needed therefore is to make the assumption context function dependent on the action model. In the next section we will explore this possibility.

Finally, we would like to mention that we now have defined assumptions as *one* set pertaining to *all* agents. In future work, it would of course make sense to index the assumption function D, *viz.* specify an assumption function D^a for every $a \in \mathcal{G}$. We will return to this issue when we discuss the Rule of Conservatism in section 7.2.

6.4 Contextual models

As we have seen in the previous sections, we need to extend the definition of the DEL narrative so it includes the change of the context set. Differently, the epistemic action that takes place in a narrative should be linked to the context change in the model. A 1-1 correspondence between pointed action models and context change does not seem likely this is why we make the contexts explicit in the new definition.

Definition 6.6 (Contextual action model)**.** We call $\langle \mathcal{A}, \alpha, D \rangle$ a *contextual action model*, if $\langle \mathcal{A}, \alpha \rangle$ is an action model and D an assumption function.

With this notion we can now define *contextual narrative models.*

Definition 6.7 (Contextual narrative model)**.** We call a sequence

$$\vec{\mathcal{A}}, \vec{\alpha}, \vec{D} := \langle \langle \mathcal{A}_1, \alpha_1, D_1 \rangle, \ldots, \langle \mathcal{A}_n, \alpha_n, D_n \rangle \rangle$$

a *contextual narrative.* A tuple $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, \vec{D} \rangle$ is a contextual narrative model if $\langle \mathcal{M}, w \rangle$ is an epistemic model and $\langle \vec{\mathcal{A}}, \vec{\alpha}, \vec{D} \rangle$ a contextual narrative.

If $\mathcal{M}, w$ is a pointed epistemic model and $\vec{D}$ an assumption function, we define

$$W_i^{\vec{D}} = \{w \in W_i \mid \text{for} \quad \text{all } j, \mathcal{M}, w \models D_i(p_j)\}.$$

We define the semantics as follows:

$$\mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, \vec{D}, i \models \varphi :\Leftrightarrow \mathcal{M}_i^{\vec{D}}, w \models \varphi,$$

where $\mathcal{M}_i^{\vec{D}} = \langle W_i^{\vec{D}}, \{(R_i^a \upharpoonright W_i^{\vec{D}})^{\text{tcl}}\}_{a \in \mathcal{G}} \rangle$. The latter is using semantics we will define below in definition 6.10.

Reading these definitions, one might wonder what the difference is between this model and the assumption-context model defined in section 6.3? With respect to the class of models they define there is no difference and the two models are equivalent. The assumption-context model, defined in section 6.3, was defined using a context function *restricted by* an assumption function. The model defined in this section just uses the assumption function and restricts the domain of the model itself. Because the context function was defined to be a function assigning subsets of the domain, this means that $C_i^{\vec{D}}$ and $W_i^{\vec{D}}$ are the same set.

So, model theoretically we might not have made that much of a difference, however, by including the assumption function in the action model, we have made it possible to include the change of context in the language.[42] The general idea being that $\mathcal{M} \models [\mathcal{A}, D]\varphi$ is understood to mean that the model $\mathcal{M} \otimes \mathcal{A}$ in the context given by D satisfies φ.

Before we redefine the product update for these models, we first extend definition 5.8:

Definition 6.8 (The language $\mathcal{L}_{\otimes A}$)**.** Let PROP and $\mathcal{G}$ be given. The language *contextual action logic*, $\mathcal{L}_{\otimes A}$, is defined by the following rule:

$$\begin{aligned} \varphi \ &:= p \mid \bot \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid K_a\psi \mid [\mathfrak{A}]\psi \\ \mathfrak{A} \ &:= \langle \mathcal{A}, \alpha, D \rangle \end{aligned}$$

Definition 6.9. If $\langle \mathcal{M}_i, w_i \rangle$ is a pointed epistemic model and $\langle \mathcal{A}_{i+1}, \alpha_{i+1}, \vec{D} \rangle$ a pointed contextual action model, the product model $\mathcal{M}_{i+1} = \langle W_{i+1}, R_{i+1}^a, V_{i+1} \rangle$ obtained by the product update $\langle \mathcal{M}_i, w_i \rangle \otimes \langle \mathcal{A}_{i+1}, \alpha_{i+1}, \vec{D} \rangle$ is given by:

[42] Note that action points in the action models are *possible* actions, and are thus dynamic objects. What is more is that the preconditions map the actions into formulas, hence making the action models partition formulas, or syntactic objects. But we could also view action models as semantic objects, i.e., as functions operating on the semantic structures. [van Ditmarsch et al., 2007, sections 6.1.1–6.1.4] has a more extensive discussion on this subtle point.

- $W_{i+1} = \{\langle w_i, \alpha_{i+1}\rangle \mid w_i \in W_i \ \&\ \alpha_{i+1} \in A_{i+1} \ \&\ \mathcal{M}_i, w_i \models \mathsf{pre}(\alpha_{i+1})\}$
- $\langle w_i, \alpha_{i+1}\rangle R^a_{i+1} \langle v_i, \beta_{i+1}\rangle :\Leftrightarrow w_i R^a_i v_i \ \&\ \alpha_{i+1} S^a_{i+1} \beta_{i+1}$
- $\langle w_i, \alpha_{i+1}\rangle \in V_{i+1}(p) :\Leftrightarrow w_i \in V_i(p)$

Definition 6.10 (Semantics for the language $\mathcal{L}_{\otimes A}$)**.** Let w be a state in the contextual narrative model $\mathcal{M}^{\vec{D}}$ and $\mathcal{G}$ a set of agents. We define recursively what it means for a formula φ to be satisfied by the model $\mathcal{M}^{\vec{D}}_i$ at state w_i, as follows:

$$
\begin{array}{lll}
\mathcal{M}^{\vec{D}}_i, w_i & \models p & \text{iff} \quad w_i \in V(p),\\
\mathcal{M}^{\vec{D}}_i, w_i & \models \bot & \text{iff} \quad \text{never},\\
\mathcal{M}^{\vec{D}}_i, w_i & \models \neg\varphi & \text{iff} \quad \text{not } \mathcal{M}^{\vec{D}}_i, w_i \models \varphi,\\
\mathcal{M}^{\vec{D}}_i, w_i & \models \varphi \vee \psi & \text{iff} \quad \mathcal{M}^{\vec{D}}_i, w_i \models \varphi \text{ or } \mathcal{M}^{\vec{D}}_i, w_i \models \psi,\\
\mathcal{M}^{\vec{D}}_i, w_i & \models K_a\varphi & \text{iff} \quad \text{for all } v \text{ such that } (R^a_i{\restriction}W^{\vec{D}})^{\mathrm{tcl}} w_i v :\\
 & & \qquad \mathcal{M}^{\vec{D}}_i, v \models \varphi,\\
\mathcal{M}^{\vec{D}}_i, w_i & \models [\mathfrak{A}_{i+1}, D_{i+1}]\varphi & \text{iff} \quad \mathcal{M}^{\vec{D}}_i, w_i \models \mathsf{pre}(\alpha_{i+1}) \Longrightarrow \mathcal{M}^{\vec{D}}_{i+1}, w_{i+1} \models \varphi,
\end{array}
$$

where $\mathcal{M}^{\vec{D}}_{i+1} := (\mathcal{M}_i \otimes \mathcal{A}_{i+1})^{D_{i+1}}$ and $w_{i+1} := \langle w_i, \alpha_{i+1}\rangle$.

Given these semantics, we can see that contrary to the examples in the previous sections, it cannot happen that the following will both be true:

$$
\begin{array}{rl}
\mathcal{M}^{\vec{D}}_i, w_i & \models [\mathfrak{A}_{i+1}, D_{i+1}]\varphi\\
\mathcal{M}^{\vec{D}}_{i+1}, w_{i+1} & \not\models \varphi,
\end{array}
$$

where $\mathcal{A}_{i+1}$ is the identity event.

There are two differences between this and standard DEL. The syntax of $\mathcal{L}_{C\otimes A}$ is the same as a standard language for DEL, except that the contextual models take the place of the standard action models. The other difference is that normally a language for DEL is interpreted on equivalence frames. In our set-up we do not require the relation to be transitive. Having said this, we have seen that at a particular stage formulas are evaluated on the transitive closure of the submodel determined by the assumption function. And in the case of the action update, we have encoded in the truth condition that a formula is true after the execution of an action model if it is true in the resulting product update *and the context in that resulting model.*

7 The destruction of knowledge – the case of Lewis

In section 3.2 we have seen that, according to Lewis, knowledge can be overturned by the mere mentioning of a possibility. The way it has been set up is that the possible doubt a contender of knowledge can raise is already included in the model, because the modeller has already flipped ahead all the way to the last page, and knows what will happen in the narrative. This knowledge the modeller has is exactly what he or she is basing his model on. For instance he or she will assign propositional letters in order to represent all the facts which will be relevant in the narrative.

Now that we have our methodology and language in place it is time to revisit Lewis and his contextual notion of knowledge. We recall from section 3.3 that Lewis introduces a set of rules to determine which possibilities are relevant and which ones are not, given a particular knowledge claim.

Although we would have liked to also consider the Rule of Belief in our framework by extending our language with a belief operator, unfortunately, we will have to leave this for future endeavours. With respect to the Rules of Method and Reliability we will say this. A failure of, for instance "vision", would show up in a narrative as a relevant possibility and be dealt with on the level of the meta-level.

In what follows we will in particular focus on the Rule of Conservatism, Resemblance and Attention. We will consider Lewis' rules as restrictions to the model we defined in the previous section (section 6.4).

7.1 The Rule of Actuality

We recall the Rule of Actuality from section 3.3:

> The possibility that actually obtains is never properly ignored; actuality is always a relevant alternative; nothing false may properly be presupposed. [Lewis, 1996, p. 554]

The rule of actuality says that we can never assume something that is false. In order to ensure this we define the assumption function such that for each stage i the actual world is included in the context set.

Definition 7.1 (The Rule of Actuality). Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, \vec{D} \rangle$ be a contextual narrative model. We say a contextual narrative model is **actual** iff for every $i, w \in W_i^{\vec{D}}$.

This is pretty straightforward because in epistemic logic we usually consider pointed epistemic models. Assuming a world to be the actual world is already entrenched in the practice of modelling in epistemic logic and we have been doing this in all the previous sections.

7.2 The Rule of Conservatism

We recall the Rule of Conservatism from section 3.3:

> Suppose that those around us normally do ignore certain possibilities, and it is common knowledge that they do. [...] We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us. [Lewis, 1996, p. 559]

We briefly noted at the end of section 6.3 that a natural extension of our set-up would be to consider assumption functions D^a for every $a \in \mathcal{G}$, in order to express the assumptions of each agent individually. The set-up we have now is one funtion for all agents expressing the collective assumptions of a group of agents. This is also the reason why our language, definition 6.8, does not have common-knowledge operator, contrary to for example the language we considered in section 5.1.

7.3 The Rule of Resemblance

We recall the Rule of Resemblance from section 3.3:

> Suppose one possibility saliently resembles the another. Then if one of them may not be properly ignored neither may the other. [Lewis, 1996, p. 556]

Although Lewis does not mention this himself when he discusses the Rule of Resemblance in [Lewis, 1996], it is hard not to think of Lewis' *(centered) comparative similarity system* from [Lewis, 2008, §2.3]. Of course in [Lewis, 2008] Lewis uses this system to provide counterfactual statements with truth conditions,[43] but we can of course consider other applications.[44]

The idea is that for possible world w all the other possible worlds can be ordered according to a similarity relation inducing a set of spheres around the world w (see figure 8). Worlds $v, u \in S_i$ are both equally similar to w. And for $j < k$, world in S_j are more similar to w than worlds in S_k are.

[43] We already ran into this point earlier, *cf.* footnote 17 on page 15.

[44] *Cf.* [Grove, 1988], [van Ditmarsch et al., 2007, §3.3] where this *system of spheres* is used to model theory change. We have also seen this model when we discussed different frameworks of contextualism in section 2.3.

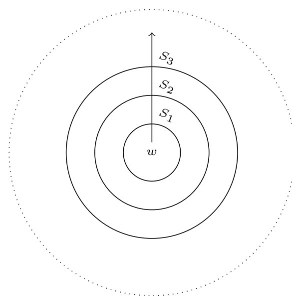

Figure 8: System of Spheres

Intuitively, there seems to be some objectivity to the notion of similarity. We have no problem with agreeing to a statement like “a cat is more similar to a dog than it is to a beetle (car or insect)”. It’s this intuition that underlies the similarity ordering on worlds, but that doesn’t mean that this ordering is a given, but needs to be re-assessed in each context.[45]

Returning to the Rule of Resemblance, what Lewis seems to have in mind is that whenever a world v is not properly ignored and v is in some sphere S_i, then no world in S_i can be properly ignored. So, suppose we have the actual world w and worlds v and u that only differ from w with respect to the particular location of a particular mosquito flying in the Brasilian rainforest, then if v cannot be ignored, then u cannot either. Lewis famously defended a particular form of realism about possible worlds, or *modal realism* as he called it. According to Lewis the actual world, the world we live in, is but one among infinitely many possible worlds of the same kind . And all worlds combined make up reality.[46] It seems that it is only because Lewis considers all the possible ways the world could have been, or might be, because we do not always know, as given, the notion of a Rule that includes those worlds that connot be properly ignored seems to become necessary. In the case of *finite narrative modelling* those worlds in one sphere would not show up in the model. The Rule of Resemblace seems to be necessary when *all* possibilities are given and need to be considered in order to lump together those that differ from the actual world in the same degree.

[45] *Cf.* [Lewis, 1979b, p. 466–467] for some ideas about how we may find this similarity relation.

[46] *Cf.* [Lewis, 2008, §4.1]. The classic critique of this view is [Stalnaker, 1984, Chapter 3].

7.4 The Rule of Attention

We recall the rule of attention from section 3.3:

> When we say that a possibility is properly ignored, we mean exactly that; we do not mean that it could have been properly ignored. Accordingly, a possibility not ignored at all is ipso facto not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative. [Lewis, 1996, p. 559]

We have set out to formalise this rule in its full generality, where the change of context might not necessarily be the direct result of an epistemic action. However, the context model of section 6.2, where the context function and the product update function independently of eachother proved too wild and needed to be restricted to a more behaved model we introduced in section 6.3.

With the machinery we have developed so far we can model the archaic case where an agent knows p and after a widening of the context no longer does so (*cf.* example 6.2). By defining contextual models, where an assumption function was made dependent on the action model, we were able to find adequate semantics to represent these cases.

Definition 7.2 (The Rule of Attention)**.** Let $\langle \mathcal{M}, w, \vec{\mathcal{A}}, \vec{\alpha}, \vec{D} \rangle$ be a contextual narrative model. We say a contextual narrative model is **attentive** iff at stage i an agent $a \in \mathcal{G}$ considers a possible world v, then $v \in W_i^{\vec{D}}$.

8 Future work

This thesis suggests the following future research.

- A natural next step would be to find axioms for the system we defined in section 6.4 in order to prove completeness for the semantics.
- We have noted in the course of this thesis that several extension of our model would be interesting to investigate. First of all, assumption functions for all agents instead of one assumption function for the set of all agents. This way we would have a better chance of finding a

formalisation of the Rule of Conservatism. We have not formalised the Rule of Belief and it would be interesting to consider an extension of our models with a belief operator.

- We have found little to no literature on the historic overview of the development of (contextual) epistemology and epistemic logic. How are the developments of these research areas related exactly? Also, a more thorough embedding in the existing literature on contextual dynamic logic is needed.

References

[Annis, 1978] Annis, D. (1978). A contextualist theory of epistemic justification. *American Philosophical Quarterly*, 15(3):213–219.

[Aucher et al., 2009] Aucher, G., Grossi, D., Herzig, A., and Lorini, E. (2009). Dynamic context logic. In He, X., Horty, J. F., and Pacuit, E., editors, *Logic, Rationality and Interaction: Second International Workshop (Proceedings LORI 2009)*, volume 5834 of *Lecture notes in artificial intelligence*, pages 15–26. Springer.

[Aumann and Brandenburger, 1995] Aumann, R. and Brandenburger, A. (1995). Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):1161–1180.

[Baltag, 1999] Baltag, A. (1999). A logic of epistemic actions. In van der Hoek, W., Meyer, J., and Witteveen, C., editors, *(Electronic) proceedings of the ESSLLI 1999 workshop on Foundations and Applications of Collective Agent-Based Systems*. Utrecht University.

[Baltag and Moss, 2004] Baltag, A. and Moss, L. S. (2004). Logics for epistemic programs. *Synthese*, 139:165–224.

[Baltag et al., 1998] Baltag, A., Moss, L. S., and Solecki, S. (1998). The logic of common knowledge, public announcements, and private suspicions. In Gilboa, I., editor, *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)*, pages 43–56.

[Baltag et al., 2008] Baltag, A., van Ditmarsch, H., and Moss, L. (2008). Epistemic logic and information update. In Adriaans, P. and van Benthem, J., editors, *Handbook on the Philosophy of Information*, pages 361–456. Elsevier Science Publishers.

[Bianchi and Vassallo, 2005] Bianchi, C. and Vassallo, N. (2005). Epistemological contextualism: A semantic perspective. In *Modeling and Using Context. Proceedings of the 5th International and Interdisciplinary Conference CONTEXT 2005*, volume 3554 of *Lecture Notes in Computer Science*, pages 41–54. Springer.

[Block et al., 2012] Block, A., Fisseni, B., León, C., Löwe, B., and Sarikaya, D. (2012). Narrative summarization and its correspondence to proppian functions. *ILLC Publication*, (PP-2012-10). The Institute for Logic, Language and Computation.

[Buvač et al., 1995] Buvač, S., Buvač, V., and Mason, I. A. (1995). Metamathematics of contexts. *Fundamenta Informaticae*, 23:412–419.

[Cohen, 1999] Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13:57–89.

[Copeland, 2002] Copeland, B. J. (2002). The genesis of possible worlds semantics. *Journal of Philosophical Logic*, 31(2):99–137.

[de Jager, 2009] de Jager, T. (2009). *"Now that you mention it, I wonder ..." Awareness, Attention, Assumption*. PhD thesis, Institute for Logic, Language and Computation. Universiteit van Amsterdam. DS-2009-10.

[DeRose, 1995] DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1):1–52.

[DeRose, 1998] DeRose, K. (1998). Contextualism: An explanation and defense. In Greco, J. and Sosa, E., editors, *The Blackwell Guide to Epistemology*, pages 187–205, Oxford. Blackwell.

[Dretske, 1970] Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67(24):1007–1023.

[Dretske, 1971] Dretske, F. (1971). Conclusive reasons. *Australiasian Journal of Philosophy*, 49:1–22.

[Dretske, 1981] Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40(3):363–378.

[Dyer, 1983] Dyer, M. G. (1983). *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press, Cambridge, MA, USA.

[Fagin and Halpern, 1988] Fagin, F. and Halpern, J. (1988). Belief, awareness and limited reasoning. *Artificial Intelligence*, 34:39–76.

[Fagin et al., 1995] Fagin, R., Halpern, J., Moses, Y., and Vardi, M. (1995). *Reasoning About Knowledge*. MIT Press, Cambridge, Massachusetts.

[Fisseni and Löwe, 2012] Fisseni, B. and Löwe, B. (2012). Event mappings for comparing formal frameworks for narratives. Manuscript in preparation.

[Goldman, 1976] Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73:771–791.

[Grove, 1988] Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170.

[Hendricks and Symons, 2009] Hendricks, V. and Symons, J. (2009). Epistemic logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Spring 2009 edition. `http://plato.stanford.edu/archives/spr2009/entries/logic-epistemic/`.

[Hintikka, 2005] Hintikka, J. (2005). *Knowledge and Belief: An introduction to the logic of two notions*, volume 1 of *Texts in Philosophy*. College Publications.

[Holliday, 2010] Holliday, W. H. (2010). Epistemic logic and relevant alternatives. In Slavkovik, M., editor, *Proceedings of the 15th Student Session of the European Summer Schools of Logic, Language and Information*, pages 4–16.

[Holliday, 2012a] Holliday, W. H. (2012a). Epistemic logic, relevant alternatives, and the dynamics of context. In Lassiter, D. and Slavkovik, M., editors, *New Directions in Logic, Language and Computation*, volume 7415 of *Lecture Notes in Computer Science*, pages 109–129. Springer.

[Holliday, 2012b] Holliday, W. H. (2012b). *Knowing What Follows: Epistemic Closure and Epistemic Logic*. PhD thesis, Stanford University. DS-2012-09.

[Lewis, 1979a] Lewis, D. (1979a). Attitudes de dicto and de se. *The Philosophical Review*, 88:513–543.

[Lewis, 1979b] Lewis, D. (1979b). Counterfactual dependence and time's arrow. *Noûs*, 13:455–476.

[Lewis, 1979c] Lewis, D. (1979c). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359.

[Lewis, 1996] Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567.

[Lewis, 2008] Lewis, D. (2008). *Counterfactuals*. Blackwell Publishing.

[Löwe, 2011] Löwe, B. (2011). Methodological remarks about comparing formal frameworks for narratives. In Allo, P. and Primiero, G., editors, *Third Workshop in the Philosophy of Information*, pages 10–28, Brussel. Contactforum van de Koninklijke Vlaamse Academie van België voor Wetenschappen en Kunsten.

[Löwe and Müller, 2011] Löwe, B. and Müller, T. (2011). Data and phenomena in conceptual modelling. *Synthese*, 182:131–148.

[Löwe et al., 2009] Löwe, B., Pacuit, E., and Saraf, S. (2009). Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from csi: Crime scene investigation$^{\mathrm{TM}}$. In Duvigneau, M. and Moldt, D., editors, *Proceedings of the Fifth International Workshop on Modelling of Objects, Components and Agents, MOCA'09, Hamburg 2009 [FBI-HH-B-290/09]*, pages 45–63.

[Löwe et al., 2010] Löwe, B., Thomas, M., and Müller-Hill, E. (2010). Mathematical knowledge: a case study in empirical philosophy of mathematics. In Van Kerkhove, B., De Vuyst, J., and Van Bendegem, J. P., editors, *Philosophical Perspectives on Mathematical Practice*, number 12 in Texts in Philosophy. College Publications, London.

[Ludlow, 2005] Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning and Truth*, pages 11–50. Clarendon Press.

[Mani, 2012] Mani, I. (2012). Computational narratology. In Hühn, P. et al., editors, *The Living Handbook of Narratology*. Hamburg University Press. `http://hup.sub.uni-hamburg.de/lhn/index.php/Computational_Narratology#Trends_in_the_Field`.

[McCarthy, 1976] McCarthy, J. (1976). An example for natural language understanding and the AI problems it raises. Memo, unpublished.

[McCarthy, 1993] McCarthy, J. (1993). Notes on formalizing context. In Bajcsy, R., editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, pages 555–562. `http://www-formal.stanford.edu/jmc/context3/context3.html`.

[Meyer and van der Hoek, 1995] Meyer, J.-J. C. and van der Hoek, W. (1995). *Epistemic Logic for AI and Computer Science*. Cambridge University Press.

[Mueller, 2000] Mueller, E. T. (2000). Prospects for in-depth story understanding by computer. *CoRR*, cs.AI/0003003. `http://arxiv.org/abs/cs.AI/0003003`.

[Norman, 1999] Norman, A. P. (1999). Epistemological contextualism: its past present and prospects. *Philosophia*, 27(3–4):383–418.

[Nozick, 1981] Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.

[Nuchelmans, 1973] Nuchelmans, G. (1973). *Theories of the proposition: ancient and medieval conceptions of the bearers of truth and falsity*. North-Holland Publishing Company.

[Plaza, 1989] Plaza, J. (1989). Logics of public communications. In Emrich, M., Pfeifer, M., Hadzikadic, M., and Ras, Z., editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216.

[Pritchard, 2002] Pritchard, D. (2002). Two forms of epistemological contextualism. *Grazer Philosophische Studien*, 64:19–55.

[Rysiew, 2011] Rysiew, P. (2011). Epistemic contextualism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2011 edition. `http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/`.

[Stalnaker, 1968] Stalnaker, R. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in Logical Theory*, pages 98–112. Oxford.

[Stalnaker, 1984] Stalnaker, R. (1984). *Inquiry*. MIT Press.

[Stalnaker, 1999] Stalnaker, R. (1999). *Context and Content: Essays on intentionality in speech and thought*. Oxford University Press.

[Stanley, 2004] Stanley, J. (2004). On the linguistic basis for contextualism. *Philosophical Studies*, 119:119–246.

[Stine, 1976] Stine, G. C. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29:249–261.

[Unger, 1975] Unger, P. (1975). *Ignorance: a case for scepticism*. Oxford University Press.

[van Benthem, 2008] van Benthem, J. (2008). ‘Tell It Like It Is’: information flow in logic. Technical Report PP-2008-08, The Institute for Logic, Language and Computation.

[van Benthem, 2009] van Benthem, J. (2009). The great art of modeling. *ILLC Publication*, (PP-2009-06). The Institute for Logic, Language and Computation.

[van Benthem and Velázques-Quesada, 2009] van Benthem, J. and Velázques-Quesada, F. (2009). Inference, promotion, and the dynamics of awareness. *ILLC Publications*, (PP-2009-43). The Institute for Logic, Language and Computation.

[van Ditmarsch et al., 2007] van Ditmarsch, H., van der Hoek, W., and Kooi, B. (2007). *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer.

[van Eijck et al., 2011] van Eijck, J., Sietsma, F., and Wang, Y. (2011). Composing models. *Journal of Applied Non-Classical Logics*, 21(3–4):397–425.

[van Eijck and Visser, 2010] van Eijck, J. and Visser, A. (2010). Dynamic semantics. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Fall 2010 edition. `http://plato.stanford.edu/archives/fall2010/entries/dynamic-semantics/`.

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

[Williams, 1991] Williams, M. (1991). *Unnatural Doubts; Epistemological realism and the basis of scepticism*. Blackwell, Cambridge, MA.