

Talking About Knowledge^{1 2}

Johan van Benthem, University of Amsterdam & Stanford University

Parikh Volume, Springer Outstanding Logicians Series

Draft, February 2014

Abstract In current studies of knowledge at the interface of logic and epistemology, philosophical positions and logical systems lore meet in new ways. In this little piece, a programmatic sequel to van Benthem 2011 and a prequel to Baltag et al. 2014, I add some further perspectives and issues to this mix from dynamic-epistemic logics of information and inquiry. My aim is to show that we can have a yet richer agenda of epistemic themes, and a richer view of the interplay of logic and epistemology, when we make epistemic action a major focus on a par with knowledge or belief per se.

1 Background: epistemic logic and notions of knowledge

Modern encounters of logic and epistemology started with epistemic logic (Hintikka 1962).³ This system reads its key modality $K\varphi$ as truth of the proposition φ in the current range of epistemically accessible worlds for the relevant agent. The connection of this notion to knowledge in the philosopher's sense has long been a focus of discussion (cf. Stalnaker 1985, Williamson 2000). It may be viewed as an account of implicit ideal knowledge of ordinary agents, of explicit knowledge of ideal agents, or as a mere statement of the content of an agent's knowledge, without any claim to define knowledge in that way (Holliday 2012B). Be this as it may, among many epistemic logicians, $K\varphi$ has come to be viewed as standing for yet a different notion, equally important in philosophy and the sciences, namely the 'semantic information' that the agent has available. Especially in large open-ended spaces of epistemic options, this seems too demanding to allow for significant knowledge, and current philosophical accounts of knowledge have taken other roads.

¹ Rohit Parikh is an inspirational leader at the interface of logic, epistemology, mathematics, computer science and game theory whose influence (some might say, grace) has touched so many people. This piece does no justice to the depth of his thinking, but I hope it is in his spirit of free inquiry.

² I thank Shi Chenwei and, especially, Wesley Holliday for valuable comments on this paper.

³ Hintikka's views on the interface of logic and epistemology are still highly relevant, and they have kept evolving, witness Hintikka 1973 and the 'Socratic epistemology' of Hintikka 2007.

Indeed, there has been a veritable wave of creativity in the literature since the 1960s. Many philosophers analyze knowledge as belief, a notion that scans less than the full range of options for semantic information, and then upgrade belief with extra requirements of various kinds, such as tracking the truth (Nozick 1981), robustness under update with new information (Stalnaker 2006), temporal convergence to true belief (Kelly 1996), or defensibility in challenge games (Lehrer 1990).⁴ The resulting styles of reasoning about knowledge diverge considerably from classical epistemic logic, in line with philosophical intuitions that have been much discussed in the literature. In particular, there is no automatic closure of knowledge under simple logical inferences such as weakening the known proposition, or combining two known propositions to knowing their conjunction.⁵

2 Knowledge and logic of relevant alternatives

These richer philosophical accounts allow for significant senses in which one can know that φ at the present stage of inquiry even when not all possible $\neg\varphi$ -worlds (alternatives, states, scenarios) have been ruled out: knowledge ventures beyond the semantic information. In what follows I will be looking at the influential proposals of Dretske 1970, Dretske 1981, and Lewis 1996 that read knowledge of φ as a state of having ruled out all relevant potential counter-examples to the proposition φ . This is the view that I will be discussing as the running theme of this paper, since it seems the best suited for making my general points.⁶

Recently, logicians have looked in new ways at what makes such relevant alternatives ('RA', for short) theories tick, going beyond the basics of epistemic logic. A trail-blazing study is Holliday 2012B that presents a joint analysis of relevant-alternative and tracking theories of knowledge, and determines the complete valid closure principles governing reasoning about knowledge in these styles. Further analyses, using ideas from dynamic-epistemic logic (van Benthem 2011), are found in Cornelisse 2011, Xu 2010, Shi 2013, Holiday 2012A. I will not go into details of these systems here. The technical discussion of RA theory to follow is in the spirit of these papers, but the general issues I will raise may be all my own.⁷

⁴ Cf. Parikh et al. 2013 for another take on the striking entanglements of knowledge and games.

⁵ These are not the usual omniscience failures of bounded rationality that are often bandied against epistemic logic. As Dretske points out, they would even occur for "ideally astute logicians".

⁶ Personally, I am also attracted to accounts of knowledge that also involve belief and truth-tracking, but this second love on the side should not be a serious bias in the following presentation.

⁷ Analogies with the cited work may still be worth pursuing for their independent interest.

Relevant alternatives models For the purpose of this paper a minimal modeling suffices. We use worlds for alternative situations, with the current accessible worlds those that have not been ruled out yet as options for what the actual world is like, and thus potentially inviting further inspection.⁸ Equally crucially to the analysis, worlds come with a primitive order of relevance R that may have many plausible intuitive origins (cf. Hawke 2014 for a critical survey of current philosophical approaches). Putting together all these basic ingredients, we will work with simple pointed epistemic models

$$\mathbf{M} = (W, \sim, R, V, s), \quad \text{with a designated actual world } s.$$

In what follows, we will assume that \sim is an equivalence relation, while the relevance order is transitive and well-founded, allowing no loops or infinite upward sequences.⁹

Authors sometimes rush to the inferences validated by such models, but we will go slowly, more in model-theoretic mode. Prior to any specific definition of knowledge, and perhaps even more importantly than that, models such as these are not just semantic structures, they also embody an implicit proposal for a *conceptual framework* of notions, a core family that constitutes an epistemic practice of stating claims and rejoinders, and based on that, reasoning. It is a substantive view in itself that the epistemic core notions are the ones chosen here, and of course, it is not an entirely uncontroversial choice.

The present conceptual framework suggests a logical base language with two modalities:

- (a) $[\sim]\varphi$ for the semantic information of the agent, that we will write $I\varphi$,
- (b) $[\uparrow]\varphi$ for truth in all worlds that are more relevant than the current one.¹⁰

⁸ Another view of alternatives would use ‘possibilities models’ (Humberstone 1981, Holliday 2013).

⁹ This is highly simplified. A key choice point (cf. Holliday 2012B) is whether the relevance relation R is uniform, as in Lewis 1996, or depends on the proposition φ for which knowledge is claimed, as in Dretske 1970. Moreover, in a more sophisticated analysis, it becomes mandatory to make the epistemic and relevance ordering world-dependent instead of uniform across the models. An even more general framework of modal neighborhood models and ‘selection functions’ for relevant worlds to be inspected and alternatives already ruled out is developed in Holliday 2012B. The points to be made in this paper carry over to such richer settings, and even acquire more force there.

¹⁰ We will also use the dual existential modalities occasionally in the principles to follow.

Defining knowledge This is arguably the simplest setting in which a philosophically defensible notion of knowledge $K\varphi$ becomes an, if you wish, more subtle defined logical operator stating that all relevant counterexamples to φ have been ruled out. We define it as follows, in close analogy with Holliday 2012A:

$$K\varphi := I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$$

This definition requires some explanation. It says that all epistemic alternatives are either φ , or if not, then there exists some more relevant $\neg\varphi$ -alternative. Why does this capture the above account? What I have in mind here are finite, or at least upward well-founded models in the ordering R , where the following is going to happen.

Fact On pointed models with well-founded transitive relevance orders, $I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$ holds iff all the most relevant $\neg\varphi$ -worlds are epistemically inaccessible.

Proof This follows from analyzing $I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$. If I see a live $\neg\varphi$ -alternative x , then there is a more relevant $\neg\varphi$ -alternative y . Now either y is not epistemically accessible, and we are done (y has been ruled out), or it is. In the latter case, we repeat the argument. By well-foundedness, this process must stop short of creating an infinite upward chain, and so, by transitivity of the order R , some more relevant $\neg\varphi$ -alternative for x was ruled out. ¹¹ ■

Remark Veridicality.

This definition does not require the presence of φ -worlds in the epistemic range of the current world. Hence, in order to make sure that $K\varphi$ implies φ , authors in the area make assumptions such as the actual world being the most relevant one in the model. This raises some interesting issues about the justification for this. On externalist views of knowledge, where relevance order reflects real features of the actual situation, this may be just right. On internalist views of knowledge, the preceding assumption may be somewhat more ad-hoc. I will sidestep these issues here, and just enforce veridicality by adding a conjunct that φ is true, while the main body $I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$ says that the claim that φ is ‘defensible’. In what follows, I use veridicality, assuming that it has been introduced in some plausible manner.

There is more to this view of knowledge than meets the eye. For instance, it tolerates a lot of $\neg\varphi$ -worlds contradicting the claim. Depending on the relevance relation, one may know that φ even when the overwhelming majority of the current alternatives fails to satisfy φ .

¹¹ One can also find definitions that are more suitable to infinite or non-wellfounded models.

But there are other, less obvious points as well. For instance, knowledge now depends essentially on making relevance comparisons with worlds that have already been ruled out. If I recalibrate my epistemic horizon to just the accessible worlds as my current model, new most urgent counterexamples to φ would arise. In this model, it is easy to show that, with the same relevance ordering as above, the above reading of $K\varphi$ collapses into semantic information $I\varphi$. This contrast deserves some emphasis: the context where we make the relevance comparison is essential, and past effort still counts today.¹² In particular, then, even for something as simple as a single S5-agent with an epistemic equivalence relation, we cannot assume that models consist of one connected epistemic component: there may be different components, reflecting earlier stages of the process of ruling out.

Other epistemic notions Our framework does not exclude other legitimate intuitions about knowledge, and in fact, it facilitates their exploration. In particular, it may not be completely obvious in RA theory why we only ‘compensate’ counterexamples to the proposition φ by more relevant counterexamples to φ that have been ruled out. Given that our knowledge claim is in favor of the proposition φ , why not compensate a counter-example by some more relevant φ -world (epistemically accessible or not)? Taking the latter line would change the definition of knowledge from $I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$ to another notion studied in Holliday 2012A:

$$I(\varphi \vee \langle \uparrow \rangle T)$$

stating that all most relevant epistemically accessible worlds in the model are φ .¹³ The formula $I(\varphi \vee \langle \uparrow \rangle T)$ seems one more serious notion of knowledge with its own applications, though weaker than our first definition, from which it follows. It might fit another intuition about knowledge, viz. that one’s epistemic state is ‘in favor’ of the known proposition qua worlds supporting it, whence situations that are even more in favor do not disrupt this. Reasoning with this second notion differs from that with the first: for instance, it supports both weakening (upward monotonicity) and conjunction of known propositions. I think that a logical analysis should bring out such options and study them together.¹⁴

¹² Similar contrasts occur in logics of games when defining ‘non-dominated’ strategies (van Benthem 2007), or in the logic of knowledge and preference when defining ‘regret’ (van Benthem & Liu 2007).

¹³ This clause looks like a truth condition for absolute belief on plausibility orderings (see below).

¹⁴ While this may be the logician’s tendency towards tolerance, philosophers might object that this leads to easy ‘relativism’ concerning the notion of knowledge. However, one can also take an opposite

Summary We have identified two different components so far in RA theory. The first was a conceptual framework, consisting of basic notions of semantic information and relevance, that come with their own language. Next, within that framework, there was a proposal for a particular definition of knowledge. This distinction matters. One may endorse a conceptual framework for basic epistemics, even if one rejects the specific view of knowledge – but one might also reject the framework altogether. Next, prior to any account of valid and invalid inferences, definability in the resulting language turned out to be of interest in its own right. For instance, we saw how plausible alternative views of knowledge could be formulated.¹⁵

Once we have our models and formal language, we can investigate the logic that is validated, a style of epistemic reasoning. For our system, we start with a standard bimodal logic for the two static modalities I for semantic information and $[↑]$ for relevance.

Static base logic Depending on what constraints we put on the accessibility and relevance orders in our models, this will be a combination of known separate modal logics for these two modalities, plus perhaps bridge axioms reflecting connections between \sim and $<$.¹⁶ Inside this axiomatizable bimodal logic, there will then be an axiomatizable logic for any of its defined notions of knowledge, in particular, one for the RA notion $K\varphi = I(\varphi \vee <↑>\neg\varphi)$.

Closure principles Still, this does not imply that there is a perspicuous base logic of ‘closure principles’ for formulas with K and Boolean connectives only. That is why the completeness results in Holliday 2012B are such a combinatorial tour de force. But even without an explicit axiomatization in a closure principle format, we can understand some basic properties of inference with K merely in virtue of its syntactic form. The RA-formula for $K\varphi$ is neither upward nor downward monotonic with respect to φ , as it contains both a positive and a negative occurrence of φ . This fits the intuitions about closure principles stated in the

view on the plurality. Instead of forcing a choice between different views of knowledge, a logical framework allows for communication and cooperation between different philosophical positions.


¹⁵ There are also interesting more technical issues about the order of notions. For instance, we have defined RA knowledge in terms of semantic information. Conversely, can we also define semantic information in terms of RA knowledge? I think not, but a proof may be somewhat complicated.

¹⁶ For instance, we may want to insist that the relevance order is known to the agent in the sense expressed by the axiom $<↑>\varphi \rightarrow I<↑>\varphi$ (cf. van Benthem & Liu 2007).

RA literature.¹⁷ But there are further interesting inferential features to RA knowledge.¹⁸ Of these, we merely mention one curious observation that seems to have gone unnoticed in the literature. A knowledge claim may sometimes be strengthened in surprising ways.

Example Knowledge claims may ‘self-strengthen’.

Consider the following model with actual world 1, whose only epistemic alternatives are 1, 2 and 3, while world 4 is more relevant than both 2 and 3:



This supports the knowledge claim Kp in 1. But it also supports the stronger $K(p \wedge r)$ true only at 1: the $\neg(p \wedge r)$ -world 2 has a more relevant $\neg(p \wedge r)$ -world 4 that was ruled out. ■

Stated more generally, we have the following observation:

Fact The formula $K\varphi \rightarrow K(\varphi \wedge \neg \uparrow \neg \varphi)$ is valid.

Proof This follows since, even just in the minimal modal logic, $I(\varphi \vee \langle \uparrow \neg \varphi \rangle)$ implies $I((\varphi \wedge \neg \langle \uparrow \neg \varphi \rangle) \vee \langle \uparrow \neg \varphi \rangle)$, which in its turn implies $I((\varphi \wedge \neg \langle \uparrow \neg \varphi \rangle) \vee \langle \uparrow \neg(\varphi \wedge \neg \langle \uparrow \neg \varphi \rangle) \rangle)$.¹⁹ ■

The principle may look like something contorted that one would never say. But in a given model, $\varphi \wedge \neg \langle \uparrow \neg \varphi \rangle$ may well denote a set of worlds that also has a more concrete definition. We leave it to the reader to judge whether this closure principle, or others in the same vein, is a desirable or an undesirable effect of the relevant alternatives view of knowledge.

Remark Closure principles versus richer logic.

Closure principles are not necessarily the most revealing about the behavior of a knowledge operator. Consider an analogy with a more standard operator, Boolean conjunction. Taken by itself, this supports very few valid inference patterns, all of them rather trivial. But it would be a mistake to conclude from this fact that conjunction has no interesting logic. The point is just that the latter will show only in interactions with other operators, such as

¹⁷ By contrast, the alternative notion of knowledge $I(\varphi \vee \langle \uparrow \neg \varphi \rangle)$ given above has a syntax whose occurrence of φ predicts the validity of monotonicity and conjunction inferences automatically.

¹⁸ One sign of complexity is that standard *conditionalizations* such as $I(\varphi \rightarrow \psi)$ for $I\psi$ do not work. Conditional $K\psi = I(\varphi \rightarrow (\psi \vee \langle \uparrow \neg(\varphi \wedge \neg \psi) \rangle))$ has no reduction by substituting in the absolute notion.

¹⁹ This principle may fail in more general RA models than our simple setting with uniform relations.

disjunction and negation, witness the distribution and De Morgan laws. Likewise, the full inferential power of knowledge may only show in its inferential interactions with our other modalities. As an illustration, notice that, though Conjunction fails, we do have this

Fact The formula $(K\varphi \wedge I\psi) \rightarrow K(\varphi \wedge \psi)$ is valid.

It is easy to prove this, and related principles that govern the behavior of K -claims. ■

Much more can be said about the bimodal logic of RA models and its relation to philosophical intuitions about epistemic reasoning – but in this paper, our main concern is with an extension of the logical horizon to which we will turn now.

3 Dynamic logic of knowledge: semantic information

Epistemic actions The system of the preceding section may seem the logical core of RA theory. However, we are not done yet. Propositional attitudes are not the only driver in the functioning of knowledge as discussed here. There is a third epistemically important component, too, that needs to be made explicit. Our analysis so far leaves out what is perhaps the most striking aspect of RA theory. Intuitively, ruling out is an *activity* that we engage in, and a task that must be performed in order to deserve the appellation of knowledge for claims that we are making. The very flavor of RA theory is dynamic: current knowledge claims are based on things that we have done, or perhaps still need to do. Thus the framework comes with one more crucial ingredient: this time, a proposal for what are the basic epistemic actions that create or modify knowledge.²⁰

In what follows, I will explore this theme of epistemic action and show where it takes us. There will be two basic activities of ruling out and relevance raising, but the point made above is much more general. There is an amazingly rich repertoire of actions that make up epistemic practices, and ruling out and relevance raising are just a small slice.

Dynamic superstructure: ruling out as model update What dynamic superstructure fits the standard epistemic logic introduced in the preceding section? We need an explicit account of what it means to rule out a world or a set of worlds, where we assume these are defined by propositions φ . To do this we employ the framework of dynamic-epistemic logic

²⁰ This dynamics is not confined to RA theory, it is the main theme of my own work at the interface of logic and epistemology (van Benthem 1996, 2011, and Baltag, van Benthem & Smets, forthcoming).

(van Ditmarsch, van der Hoek & Kooi 2007, van Benthem 2011), where actions are treated by defining the corresponding model transformation, viewed as a change in epistemic state.

A first candidate for ruling out proposition φ would be the well-known transformation of ‘public announcement’ $(\neg\varphi)!$ that turns any current pointed model (\mathbf{M}, s) into its submodel $(\mathbf{M}/\neg\varphi, s)$ consisting of all worlds satisfying $\neg\varphi$.²¹ Public announcement is natural, and it may well be the common sense model of information growth as shrinking a current range. However, this update mechanism cannot quite work for us here, as RA theory makes essential reference to worlds that have been ruled out. They should be kept around, like the vanquished in a Roman triumph. Therefore, we will work with a slight variant:

Definition Test operation

The *test* $!\alpha$ cuts all links between α -worlds and $\neg\alpha$ -worlds in the current model. ■

This operation occurs in many places in the literature, for instance, as a way of describing what happens after performing an experiment, or when seeing a question answered.²² However, if the reader would prefer another account of the informal notion of ruling out, the techniques to follow will work for a broad variety of model-changing operations.

Dynamic logic of testing The test operation changes models, and hence assertions of knowledge that held before may not hold afterwards: knowledge may be acquired, or lost. It is straightforward to axiomatize this in a modal language that enriches our earlier setting with action expressions $\alpha!$ where α is any formula of the language. We add new action modalities $![\alpha]\varphi$ stating that φ holds at the distinguished world of the updated model:

$$(\mathbf{M}, s) \models [!\alpha]\varphi \quad \text{iff} \quad (!\alpha(\mathbf{M}), s) \models \varphi$$

In this way, we get a logic that can talk about semantic information, relevance, and RA-style knowledge, but also about actions of ruling out and what these do to the preceding notions.

To axiomatize the reasoning supported by this dynamic setting, we need to find the basic recursion laws telling us how the basic static operators change under testing. These are the

²¹ This is best understood more abstractly as a form of update with hard information that could be triggered by many events: public announcement, totally reliable observation, and so on.

²² There is an issue of whether the test operation respects the constraints that we have imposed on the orderings in our models. Such technical issues make sense, but we will ignore them here.

driving equations of the information dynamics at work. It is an application of standard dynamic-epistemic techniques to find these laws for a simple operation like testing:

Fact The following recursion laws are valid for semantic information and relevance:

- (a) $[!\alpha]I\varphi \leftrightarrow ((\alpha \wedge I(\alpha \rightarrow [!\alpha]\varphi)) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow [!\alpha]\varphi)))$
 (b) $[!\alpha][\uparrow]\varphi \leftrightarrow [\uparrow][!\alpha]\varphi$

Principle (a) reflects precisely how link cutting works, while the commutation in (b) shows that link cutting does not change the relevance relation – a topic that will return later.

Theorem The dynamic-epistemic logic of semantic information, relevance, and test is completely axiomatizable.

Proof By standard techniques from the literature, we can use the recursion laws to push dynamic modalities systematically inside until they hit atoms.²³ In the latter position, we use the evident validity of the equivalence $[!\alpha]p \leftrightarrow p$ to remove them altogether. The result is a provably equivalent dynamic-free formula in the base language for which we already had a complete logic. This reduction giving the dynamics an almost free ride on the statics works for many dynamic-epistemic logics. But we will discuss a generalization to ‘protocol models’ later on where no such reduction is possible – though completeness remains. ■

This dynamic logic still needs unpacking to see what it does with defined modalities. In particular, using the above recursion laws, we can calculate step by step how RA-style knowledge changes under testing, or equivalently, under ruling out:

$$\begin{aligned} [!\alpha]K\varphi &\leftrightarrow [!\alpha]I(\varphi \vee \langle \uparrow \rangle \neg\varphi) \leftrightarrow \\ &((\alpha \wedge I(\alpha \rightarrow [!\alpha](\varphi \vee \langle \uparrow \rangle \neg\varphi))) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow [!\alpha](\varphi \vee \langle \uparrow \rangle \neg\varphi)))) \leftrightarrow \\ &((\alpha \wedge I(\alpha \rightarrow ([!\alpha]\varphi \vee [!\alpha]\langle \uparrow \rangle \neg\varphi))) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow ([!\alpha]\varphi \vee [!\alpha]\langle \uparrow \rangle \neg\varphi)))) \leftrightarrow \\ &((\alpha \wedge I(\alpha \rightarrow ([!\alpha]\varphi \vee \langle \uparrow \rangle [!\alpha]\neg\varphi))) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow ([!\alpha]\varphi \vee \langle \uparrow \rangle [!\alpha]\neg\varphi)))) \leftrightarrow \\ &((\alpha \wedge I(\alpha \rightarrow ([!\alpha]\varphi \vee \langle \uparrow \rangle \neg[!\alpha]\varphi))) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow ([!\alpha]\varphi \vee \langle \uparrow \rangle \neg[!\alpha]\varphi)))) \end{aligned}$$

This is not a simple recursion law, and so we see that RA knowledge is not a trivial notion in a dynamic setting: something that also showed, of course, in its static inferential behavior.

²³ What we also need to proceed inside the formula are two simple valid recursion laws for the Boolean operations: $[!\alpha]\neg\varphi \leftrightarrow \neg[!\alpha]\varphi$, $[!\alpha](\varphi \wedge \psi) \leftrightarrow ([!\alpha]\varphi \wedge [!\alpha]\psi)$.

However, if we focus on base-level knowledge only, given by purely propositional ‘factual’ formulas φ , a simplification occurs. It is easy to show that in that case:

Fact $[!\alpha]\varphi$ is equivalent to φ itself for factual formulas φ .

Proof Ruling out does not change atomic facts in worlds, or their Boolean combinations. ■

The earlier complex recursion law for RA-knowledge then simplifies to something much closer to standard ways of conditionalizing knowledge:

$$[!\alpha]K\varphi \leftrightarrow ((\alpha \wedge I(\alpha \rightarrow (\varphi \vee \langle \uparrow \rangle \neg\varphi))) \vee (\neg\alpha \wedge I(\neg\alpha \rightarrow (\varphi \vee \langle \uparrow \rangle \neg\varphi))))$$

Dynamic perspectives on knowledge My interest here is not the technical dynamic machinery per se. I want to show what new themes are facilitated by this new analysis tool in a discussion about knowledge. I will consider a few topics that extend the usual focus on closure principles or other static properties that knowledge is supposed to satisfy. In one form or another, all of these themes can be found in the philosophical literature, but I will show that they can now be brought in to the open, and treated in an exact logical style.

What actions generate knowledge? Knowledge is good to have, but how does it arise? Let us first consider this for semantic information. An immediate intuition might be this. If a proposition φ is true, then a test $!\varphi$ gives us the semantic information that φ is true:

$$\varphi \rightarrow [!\varphi]I\varphi$$

One might view this as a basic intuition, on a par with the usual closure principles. However, the situation is delicate, and in fact, the principle is not valid in general. To see this, consider an evergreen: so-called ‘Moore sentences’ of the form $p \wedge \neg I p$. It may be true right now that p while I do not have this information. But after the test, I will have the information that p , and hence the conjunction has become false: a Moore sentence ‘refutes itself’. Of course, there is no inconsistency here, just a surprise. The recursion laws of our dynamic logic tell us precisely what Moore sentences achieve – and $\varphi \rightarrow [!\varphi]I\varphi$ is definitely not one of its valid principles for all formulas φ . Admittedly, there is a price to be paid for this sophistication: for, it is the reason why proofs usually have to be sprinkled with $[!\varphi]$ modalities penetrating inside formulas. As it happens, the principle $\varphi \rightarrow [!\varphi]I\varphi$ does hold for special classes of statement in our language, such as factual formulas containing no I -operators at all. ^{24 25}

²⁴ For a precise determination of all valid instances of our principle, see Holliday & Icard 2010.

The role of intuitions What should we make of this? We must be careful with our intuitions: their range of validity may have hidden restrictions. For instance, common intuitions about knowledge may not hold when we move beyond factual propositions to propositions that themselves refer to knowledge and ignorance. In that case, we could view our logic as an instrument for ‘bootstrapping’: starting from initial intuitions, we respect them as far as they go, and then refine them for larger realms. But we may also view the extension to epistemic statements in a language that is closed under epistemic and even all dynamic modalities as logical ‘system pressure’, that imports irrelevant concerns. Logicians might opt for the exciting dangerous road, but does it lead anywhere important? ²⁶

I conclude this theme with a simpler question. Since our notion of knowledge $K\varphi$ is weaker than semantic information $I\varphi$, what act would generate it, restricting ourselves to the safe realm of empirical propositions φ ? The answer is simple: we want to get $K = I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$ and by the preceding, a test for a weaker formula than φ will do it, viz. $!(\varphi \vee \langle \uparrow \rangle \neg\varphi)$. ²⁷

Knowledge and persistence under update Another basic aspect of knowledge, once in our possession, is what happens to it as we learn new things. RA theory gives a past-looking account: we have knowledge if we have made the relevant inquiries. But there is also a common intuition that knowledge gives us a guarantee for the future: it will persist. This, too, can be made sense of in our framework. We merely make one observation among many:

Fact The formula $K\varphi \rightarrow [!\alpha]K\varphi$ is valid for all factual φ and all α .

Proof Consider $I(\varphi \vee \langle \uparrow \rangle \neg\varphi)$ and apply the test $!\alpha$. This can only diminish the number of epistemic links from the current world: and the disjunction will still be true there, since factual formulas do not change their truth values under testing. For arbitrary formulas φ , the situation is more complex, though again the recursion laws are up to their analysis. ■

The preceding observation is reminiscent of the view of knowledge as stable belief that survives arbitrary new information (Stalnaker 2006, Rott 2005, Baltag & Smets 2008). Of

²⁵ It is also an interesting question if $!\varphi$ is the *minimal* action inducing $I\varphi$: I will not pursue this here.

²⁶ Here are two interesting technical problems in this setting. First, for every valid principle for factual formulas in this paper which possibly larger classes of modal formulas preserve its validity? But also, can we axiomatize the richer set of schematic valid principles that hold for factual formulas?

²⁷ Caveat: Testing for a weaker formula is not necessarily easier than testing for a stronger one.

course, as we saw with inference, many other questions make sense, of a more mixed type. For instance, we could ask what ‘upgrades’ knowledge to semantic information: and the above principle in fact gives a stronger mixed principle, using veridicality of $K\varphi$:

Fact The formula $K\varphi \rightarrow [!\varphi]I\varphi$ is valid for all factual φ .

My point is not that this is a deep difficult principle. It is not. But again it reflects a basic intuition, and one could investigate (as with the earlier Moore-type problems) how far it can be pushed for more complex propositions than merely factual ones.

From closure to epistemic work I conclude by returning to the closure principles that often serve to frame epistemological discussions. RA theory leads to a variety of closure failures, which tell us something important about the proposed account of knowledge. But in a dynamic setting, this is not the end of the story. A natural further question emerges. Closure principles say that, if we know certain things, then we also know other things *automatically*. While this sort of gain may suffice for the idle rich, I submit that a natural broader accompanying question would be what an ordinary working class agent would have to *do* in order to get new knowledge from available known propositions.

Van Benthem 2011 analyzes the resulting *dynamic closure principles* in terms of a ‘gap’ for actions: the correct epistemic principles should not be, say, $K\varphi \rightarrow K(\varphi \vee \psi)$, but

$$K\varphi \rightarrow [act] K(\varphi \vee \psi), \quad \text{where } act \text{ is some appropriate epistemic action.}$$

This approach has merits for analyzing omniscience, when *act* might stand for an act of drawing an inference, or engaging in introspection. However, such scenarios of bounded rationality are ruled out here, since RA theory assumes Dretske’s ‘ideally astute logicians’. But even enviable beings like that may still profit from other, non-inferential actions, such as receiving new information. Here is one relevant observation:

Fact The formula $K\varphi \rightarrow [!\psi]K(\varphi \wedge \psi)$ is valid.²⁸

Given that many such facts hold in a dynamic setting, a lack of standard closure principles may lose its sting, in RA theory or even less closure-tolerant accounts of knowledge.

²⁸ I have not been able to find an interesting action *act* that validates $K\varphi \rightarrow [act] K(\varphi \vee \psi)$.

Digression I even feel the pull of an opposite intuition, an ‘Anti-Closure Principle’ that makes knowledge a supremely delicate epistemic state, whose high status is as fragile as that of Andersen’s heroine in *The Princess and the Pea*. More concretely, on this intuition, an assertion $K\varphi$ should only imply knowledge of propositions that are equivalent to φ in the current model. In particular, this would make knowledge so delicate that it admits no weakening or strengthening at all.²⁹ On such a view, *all* genuine extensions of available knowledge would have to come by epistemic actions.

Summary I have shown how adding a dynamic component to epistemic logics is easy to do, and gets us a much richer view of intuitions that might hold for knowledge. Let us now see whether this proposal also works in extended settings beyond the bare minimum.

4 Dynamic logic of knowledge: relevance

The dynamic logic defined so far is not in balance. We had two static base modalities: I for semantic information and $[R]$ for relevance, but only one of them was ‘dynamified’ in the sense of adding explicit test actions $!\alpha$ that affect semantic information to the logic. But relevance, too, is a dynamic entity that can change as people engage in conversation about knowledge claims. In particular, the relation R can be affected by *raising an issue* – the way a critic might draw attention to some overlooked class of worlds in a conversation about a knowledge claim. This seems crucial to many epistemic scenarios, such as encounters with skeptics who try to change our view of what should count as relevant possibilities. In fact, it is this dynamics of making, defending, or modifying knowledge claims that seems the most interesting aspect of many epistemic scenarios in the literature (cf. van Benthem 2006).

Again this new action can be made precise as a model transformation in several ways, and there are different ways of raising an issue, with different force. We will explore only one.

Raising an issue Let us define an operation $+\alpha$ that transforms a current model so as to make α relevant. This operation will change the given relevance order, and it might do so in many ways. We borrow a proposal made for preference change in van Benthem & Liu 2007:

²⁹ Intuitions like this place a well-defined technical constraint on the proof power of epistemic logics, be it of a somewhat unusual sort. It would be interesting to see what modal logics satisfy this.

Definition Suggestion

A *relevance upgrade* $+α$ changes the relevance ordering R in the current epistemic model by deleting all upward links running from $α$ -worlds to $¬α$ -worlds.³⁰ ■

Dynamic logic of relevance This way of making $α$ an issue is again a simple relation transformer whose complete dynamic logic can be determined using our earlier methods:

Theorem The dynamic logic of semantic information, relevance, test, and issue raising is completely axiomatizable using recursion laws.

Here are the basic recursion laws, showing how relevance is affected:

$$(c) \quad [+α][↑]φ \leftrightarrow ((α \wedge [↑](α \rightarrow [+α]φ)) \vee (¬α \wedge [↑][+α]φ))$$

$$(d) \quad [+α]Iφ \leftrightarrow I[+α]φ$$

This full system seems much more congenial to analyzing epistemic scenarios. Information can change, but so can relevance, and the laws of the logic describe how the basic static modalities are modified. Again there are also many interesting technical issues here, some of them new. For instance, there can now be an interplay of the two actions, and the reader may want to ponder the potentially different effects of orders $!φ; +ψ$ versus $+ψ; !φ$.

Epistemic dynamics once more Recall the epistemic dynamics that we discussed before. It is easy to see that all our earlier issues still make sense here. We merely consider two examples. One is persistence of knowledge under epistemic events. We first note that knowledge does not change if we raise the issue of whether its proposition is true.

Fact The formula $Kφ \rightarrow [+φ]Kφ$ is valid for all factual $φ$.

Proof The reason is that the ‘compensation’ clause in our definition of RA knowledge runs inside the set of $¬φ$ -worlds where relevance order is unaffected by the model transformation of issue raising as defined above. Moreover, factual formulas have their truth value unaffected by issue raising, and the same is even true in the presence of I modalities. ■

But unlike the case of tests, knowledge is not preserved under arbitrary issue changes – something well realized by those skeptics prowling our streets looking for victims.

³⁰ Another option, out of many, might be adding the set of worlds satisfying $α$ to the current range of the relation R , creating the new relation $R \cup (\{s\} \times α^M)$.

Example Knowledge may be lost after issue change.

Consider the following model with actual world 1, whose only epistemic alternatives are 1, 2, while world 3 is more relevant than 2:



The knowledge claim Kp is true in world 1. However, after an issue change $+q$, there will be no upward relevance link any more from 2 to 3, and hence Kp is no longer true at 1. ■

This power of issue changes makes reasoning much more delicate. One way in which this shows are relations between knowledge and semantic information. For instance, in a given model (M, s) , let I_s be the proposition that defines the set of all worlds that are epistemically accessible from s . This encodes the semantic information that the agent possesses. Merely making this relevant leads to the following connection between our two notions. Under the right issue change, achieving RA-style knowledge amounts to having semantic information:

Fact $I\varphi$ is equivalent with $[+I_s]K\varphi$ for factual propositions φ .

Proof By earlier observations, $I\varphi$ implies $[+I_s]I\varphi$ and hence $[+I_s]K\varphi$. Conversely, once I_s has been made relevant in our sense, no compensation links can run from it into its complement, and hence $K\varphi$ can only be true because all epistemic alternatives satisfy φ . ■

This is only one of several dynamic connections between epistemic notions, and it shows the entanglements all around. Thus, even Dretske's ideally astute logician can still learn a lot under epistemic actions that arise naturally in the RA framework.³¹ And more generally, we see how the earlier picture of epistemic statics and dynamics gets even more appealing when we consider the full range of actions supported by our conceptual framework.

Summary We have defined a dynamic version of RA theory, and shown how this enriches the sort of logical-epistemological discussions one can have about epistemic practices.

³¹ For instance, what I have not investigated so far is the question whether and how relevance upgrades $+\alpha$, in addition to tests, could help validate, or invalidate closure principles. What is easy to see, however, is that further relevance transformers in our framework have this effect. For instance, an operation that adds, rather than removes, relevance links can link a current relevant alternative to a more relevant counterexample that has already been ruled out, thereby introducing knowledge.

5 Dynamic logic of knowledge: issues

We said at the start of this paper that our style of analysis is not restricted to RA theory. This section is a digression adding one more pilot for the above style of logical analysis, now in a somewhat deviant setting going beyond the basics of what was defined above.³²

Inquiry and issues Just postulating a primitive relevance relation R seems ad-hoc: can we say more? Major RA theorists have in fact said a lot more (cf. Dretske 1981, Lewis 1996), and Hawke 2014 is an investigation of possible approaches. Here we explore just one very simple technical idea, viz. that relevance for knowledge claims is created by a *current issue*, that can be created and modified through the typical act of raising issues, i.e., by asking *questions*.³³ While the idea of relating knowledge to questions is well known in itself (Hintikka 1973, Stalnaker 2006), the implementation to follow may be new.

Issue models It is common to model an issue as a partition of the set of possible worlds (Groenendijk & Stokhof 1984, Kelly 1996). The purpose of inquiry is to find out in which cell the actual world is located. But this suggests the following sense of relevance. The issue tells us which alternatives are relevant to the claim that we know a proposition φ . We do not need to inspect all $\neg\varphi$ -worlds to rule them out, we only need to rule out those alternative hypotheses (entire partition cells in the current issue relation) that contradict φ .³⁴

Questions and issue change Epistemic models $\mathbf{M} = (W, \sim, \approx, V)$ with an added issue relation \approx support test dynamics as before, but they also have crucial actions of issue management. The current issue may change in the course of inquiry, or a conversation. The most basic action of this sort may well be asking a Yes/No question. Its obvious dynamic transformation resembles the earlier $!\alpha$ -style link cutting for the epistemic relation:

³² What follows here is a brief précis of van Benthem 2012, building on van Benthem & Minica 2012.

³³ Admittedly, this is an agent- and inquiry-related approach that may not sit well with, for instance, Dretske's more realist construal of epistemic relevance relations.

³⁴ Issues can be much more structured than partitions, often with *priorities*. 'Priority graphs' in the style of Liu 2011 can model such issues and their more fine-grained dynamics. One can also structure the *evidence* underlying what is ruled out in similar ways, keeping a record of earlier tests, as is done in van Benthem & Pacuit 2011 in terms of modal neighborhood models.

Definition Issue change by a question

The model transformation $?\alpha$ takes the current issue relation \approx to a new relation \approx/α where all links between α -worlds and $\neg\alpha$ -worlds have been cut. ³⁵ ■

Again one can introduce a modal base language here, ³⁶ and add dynamic superstructure.

Theorem With explicit modalities added for the actions $[!\alpha]$ and $[?\alpha]$, the dynamic logic of information and issues can be axiomatized completely.

Issue-based knowledge A notion of knowledge close to RA accounts is what can be claimed on the basis of our semantic information, without being contradicted by a live alternative hypothesis for the current issue. However, to ground this in reality, we add a proviso: if we were to resolve the current issue with our current information, the claim holds. This intuitive description is expressed compactly in our language by the formula

$$K\varphi := R\varphi \wedge I\langle R \rangle\varphi$$

This may be closer to what an agent can reasonably claim than what she knows. Its inferential properties include upward monotonicity, but not conjunction of known propositions. Moreover, the above dynamic logic has recursion laws that track its dynamics.

Dynamic features Of the new dynamic phenomena that arise in this setting, we list a few. Recall our earlier idea of deserving closure through epistemic action. Consider conjunction of knowledge that failed here: $K\varphi$ and $K\psi$ together did not imply $K(\varphi \wedge \psi)$. In issue cells where φ and ψ are both possible, obviating a need for ruling them out for $K\varphi$ and $K\psi$ separately, their conjunction need not be true for any world. Thus, to know the conjunction, we have to rule out such trouble zones, and here is one action achieving this:

Fact The formula $(K\varphi \wedge K\psi) \rightarrow [![(\langle R \rangle\varphi \wedge \langle R \rangle\psi) \rightarrow \langle R \rangle(\varphi \wedge \psi)]] K(\varphi \wedge \psi)$ is valid.

Further interesting phenomena arise with issue change. Merely asking a question can thoroughly reshuffle the knowledge picture as the partition cells change. This need not always achieve closure: even raising what looks like a crucial question may not help. For instance,

³⁵ Note that this just like the earlier $!\alpha$ -style link cutting for the epistemic relation.

³⁶ A useful set of base modalities for issue models has $I\varphi$ for semantic information, $Q\varphi$ for what is true in the current cell of the issue partition, plus a modality $R\varphi$ for the intersection of \sim and \approx , stating what holds when the issue is *resolved* in combination with what we know.

it is easy to see that the implication $K\varphi \rightarrow [?\psi]K(\varphi \vee \psi)$ is not valid. Even more startlingly, our sense of knowledge is not preserved under raising the issue whether it holds:

Fact The implication $K\varphi \rightarrow [?\varphi]K\varphi$ is not valid.³⁷

But a slight variation on this observation does something that is perhaps more interesting: it captures a basic notion in dynamic terms. If knowledge in our sense persists under all ways of raising issues, then it must be the above semantic information I .

Fact The principle $[?\varphi]K\varphi \leftrightarrow I\varphi$ is valid in the dynamic logic of issues.

Summary These are just rather sketchy illustrations, but we hope that the reader has seen how the dynamics of questions and answers, too, systematically affects epistemic notions.

6 Adding further structure

Despite the new level of epistemic actions, our analysis is still quite austere in terms of basic epistemic notions and models. It is entirely possible to add further structure that matches the intuitive presentation of relevant alternatives theory, or other epistemic frameworks. In this section, we list a few directions that can be taken with the machinery we have set up.

Richer core families One can wonder whether knowledge should be studied in isolation, or whether we should start at once by identifying a richer core family of epistemic attitudes that govern our behavior.³⁸ One obvious candidate is then *belief*, itself an ingredient of many proposed definitions of knowledge, but also an important epistemic notion in its own right. The notions and techniques of the dynamic logics in this paper extend immediately to belief, for instance, in the dynamics of plausibility models found in Baltag & Smets 2008, van Benthem 2011. Basic actions then include softer forms of information change where no worlds are eliminated, but the plausibility order gets modified.³⁹ Soft dynamics might be a good way of construing natural weaker ways of ruling out relevant alternatives.

³⁷ It is easy to see how this can fail. Consider any issue zone inside which the epistemically accessible worlds include both φ -worlds and $\neg\varphi$ -worlds. After raising the issue whether φ , this zone splits into two, one of which now needs to be ruled out in order to have knowledge of φ by our definition.

³⁸ Interestingly, Dretske 1970 itself contains many more notions than those that have made it into RA theory. In particular, he emphasizes that propositions come with ‘normal criteria’ determining what are their relevant alternatives. This suggests further model structure representing such criteria, but also normality as a sort of default assumption summarizing past experience with knowledge claims.

³⁹ Plausibility is even one way of reading RA relevance orderings, but of course, not the only one.

Similar techniques work for other notions tied to knowledge, such as the counterfactuals used in tracking accounts of knowledge (Baltag, van Benthem & Smets, forthcoming).

Richer epistemic actions But still more important to the special focus of this paper is the choice of basic dynamics. There are many more epistemically relevant actions than testing, modifying relevance, or asking questions. Just to go one step further than we have done: possibilities can be ruled out, but later also ‘ruled in’– and issues may be raised in inquiry, but it is highly beneficial that they can also be ‘dropped’. If one steps back and looks at what really drives epistemic scenarios in the philosophical literature, one finds actions of detecting mistakes, repairing beliefs, raising doubts and objections, and many others. I do not have anything like a taxonomy of our basic repertoire, but it is certainly quite rich.

One particular type of action that falls outside the purely semantic scope of this paper is *inference* and reasoning in general. One might think that the valid principles of the earlier logical systems take care of this, but they are the meta-laws that govern test or issue raising, not actual acts of inferring. To deal with the subtler inferential aspects of epistemic activity, we need to move to another notion of information, more fine-grained than the possible worlds modeling we have provided. Dynamic logics of inference and other actions involving syntax and code are studied in van Benthem and Velazquez-Quesada 2010, but there are other paradigms for fine-grained epistemic dynamics, some closer to argumentation theory. In terms of closure principles, this is the realm where one can study what actions close knowledge under inference, or what forms of reflection validate introspection principles.⁴⁰

Knowledge, decision, action One can even wonder to which extent the purely informational focus of this paper, and much of the epistemological literature, is appropriate. Many authors point out that knowledge functions in a context of action, but then drop this theme. Relevant alternatives theory itself provides an obvious angle here. Relevance as used in daily discourse seems a relational notion, since our way of construing relevant options usually depends on something else. In particular, that something else is often reflection on possible outcomes before taking a *decision* or performing an action.⁴¹ It would be easy to

⁴⁰ Closure or closure failure hardly does justice to the subtleties of legal reasoning. One of the judges in the novel *Summerscale 2012*, when told that he cannot let the man in an affair go free, while convicting the woman involved of having had that affair, answers that he can convict the women of having had an affair *with someone* while absolving the man of having had an affair with *her*.

⁴¹ The view that knowledge is intimately tied to potential for action is widespread in the literature.

add this in the above models by giving the agent a set of possible actions a that can be performed in worlds s to produce new worlds sa that can then be compared as to the agent's preferences. One way in which an alternative might be relevant then is if taking it into account would change the optimal action that the agent would perform. In this way, justifiable knowledge claims would be tied intimately to possible action: they 'matter'.

Social aspects A striking feature of many epistemic scenarios in the philosophical literature is their multi-agent character. Encounters with skeptics are about people getting into a discussion about knowledge claims and relevant alternatives. And in Socratic dialogues, knowledge is what one can defend against others having perhaps different information. This multi-agent character usually gets lost in logical treatments, but it should be there. Now dynamic-epistemic logic was originally designed to deal with multi-agent scenarios, so there is no obstacle in principle. We can make all orderings agent-dependent and study levels of interaction, and forms of group knowledge. I think it would be interesting to give such types of analysis for skeptical scenarios, where there are two relevance orders, and negotiation has to take pace to change either. Or perhaps differently, there is social choice involved: the two discussants may merge their individual relevance orders into an ordering for the group consisting of the persons involved (effectively, a third epistemic agent), which can then serve as a basis for their further discussion. The multi-agent aspect may also help make something visible that tends to be ignored when focusing on success features of knowledge for agents: the fact that disagreement, not assent, is often the engine of epistemic progress.⁴²

Time, protocols, and games Another important epistemic aspect ignored in this paper is the nature of the process where knowledge functions. Epistemic actions are single steps in something bigger. There may be a procedure or *protocol* allowing us to perform certain actions, and not others – as is true in any process of inquiry, experimentation, or conversation (Hoshi 2009). A major aspect of such protocols is the regimentation of available tests: we can only rule out alternatives by some legitimate means of inquiry that is at our disposal. But a process limitation can apply equally well to questions that can be asked, or to admissible acts of changing relevance. If we encode this into our models, the laws of the dynamic logics that we had before will change in interesting ways (in particular, a static base reduction will no longer be possible), but the general style of analysis remains the same.

⁴² It should be emphasized though that this social agent-dependent perspective is not that of the founding fathers of relevant alternatives theory, who see relevance as something objective in nature.

Another major feature of processes is their global *temporal horizon*: what is now also important epistemically is what happens in the long run. Perhaps the best known example of this is formal learning theory (Kelly 1996) and its notions of inductive knowledge in the limit. Another relevant procedural framework, combining temporal horizon with social structure, is that of game theory where strategic interactive aspects of knowledge come to the fore. The knowledge games of Lehrer 1990 and Hintikka 2007 show these at work in epistemology (cf. Fiutek 2013, Hamami 2014). Learning theory and game theory fit well with the dynamic logics in this paper (cf. van Benthem 2014) and add further concerns of their own.

7 Conclusion

This paper is a sample of a particular logical style of looking at epistemology. I started with a simple model of relevant alternatives theory and then gradually added forms of epistemic dynamics. In doing so, I viewed epistemic theories as having three different components that can be evaluated separately: a conceptual framework of the proposed basic epistemic notions, a definition of knowledge, and often implicit, but crucial here: a repertoire of basic epistemic actions that constitute an epistemic practice. By providing a family of definitions, systems, and formal results, I hope to have shown how dynamic-epistemic logics can help think about all these aspects of knowledge in a coherent and suggestive manner. Finally, I have pointed at broader social and temporal dimensions where all this fits.

I see this as proposing an extended agenda for investigations at the interface of logic and epistemology, and for eliciting a wider range of philosophical intuitions about knowledge. To me, this involves cooperation between partners bringing different assets. Logic and epistemology are not the same mindset, and this can be an engine of progress. For instance, on my own agenda of intelligent interactive agency (van Benthem 2011, 2014), the core target are our rich information-driven epistemic practices, involving an array of different epistemic core notions, and different perspectives on knowledge,⁴³ and it is a major task of logical analysis to bring out their fruitful co-existence. But so far, most of the really innovative guiding ideas about knowledge that are around these days have come from the philosophers, and their intuitions may well come to affect how we develop epistemic logic. This paper just follows a few threads in this encounter. The results of teaming up epistemology and logic are not entirely predictable, and hence their collaboration becomes even more alluring.

⁴³ For instance, my own recent work approaches knowledge in terms of invariants and equilibria in longer-term epistemic activity (cf. van Benthem 2011, Baltag, van Benthem & Smets, forthcoming).

7 References

- A. Baltag, J. van Benthem & S. Smets, forthcoming, *The Music of Knowledge: Dynamic Epistemic Logic Meets Epistemology*, book manuscript, Institute for Logic, Language and Computation, University of Amsterdam.
- A. Baltag & S. Smets, 2008, 'A Qualitative Theory of Dynamic Interactive Belief Revision', in G. Bonanno, W. van der Hoek, M. Wooldridge, eds., *Texts in Logic and Games Vol. 3*, Amsterdam University Press, Amsterdam, 9–58.
- J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- J. van Benthem, 2006, 'Epistemic Logic and Epistemology: the state of their affairs', *Philosophical Studies* 128, 49–76.
- J. van Benthem, 2007, 'Rational Dynamics', *International Game Theory Review* 9:1, 13–45.
- J. van Benthem, 2011, *Logical Dynamics of Information and Interaction*, Cambridge University Press, Cambridge UK.
- J. van Benthem, 2014, *Logic in Games*, The MIT Press, Cambridge MA.
- J. van Benthem, P. Girard & O. Roy, 2009, 'Everything Else Being Equal: A Modal Logic for Ceteris Paribus Preferences', *Journal of Philosophical Logic* 38, 83–125.
- J. van Benthem & F. Liu, 2007, 'Dynamic Logics of Preference Upgrade', *Journal of Applied Non-Classical Logics* 17, 157–182.
- J. van Benthem & S. Minica, 2012, 'Toward a Dynamic Epistemic Logic of Questions', *Journal of Philosophical Logic* 41, 633–669.
- J. van Benthem & E. Pacuit, 2011, 'Dynamic Logics of Evidence-Based Beliefs', *Studia Logica* 99, 61–92.
- J. van Benthem & F. Velazquez-Quesada, 2010, 'The Dynamics of Awareness', *Synthese (Knowledge, Rationality & Action)* 177:1, 5–27.
- I. Cornelisse, 2011, *Context Dependence of Epistemic Operators in Dynamic Evidence Logic*, Master of Logic Thesis ML-2011-09, Institute for Logic, Language and Computation, University of Amsterdam.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, 2007, *Dynamic Epistemic Logic*, Springer Publishers, Dordrecht.
- F. Dretske, 1970, 'Epistemic Operators', *The Journal of Philosophy* 67:24, 1007–1023.
- F. Dretske, 1981, 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363–378.

- K. Fine, 2013, 'Truthmaker Semantics for Intuitionistic Logic', *Journal of Philosophical Logic* on-line, May 2013, DOI 10.1007/s10992-013-9281-7.
- V. Fiutek, 2013, *Playing with Knowledge and Belief*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.
- J. Groenendijk & M. Stokhof, 1984, *Studies on the Semantics of Questions and the Pragmatics of Answers*, Philosophical Institute, University of Amsterdam.
- Y. Hamami, 2014, 'The Interrogative Model of Inquiry Meets Dynamic Epistemic Logic', Centre for Logic and Science-Philosophy, Brussels. To appear in *Synthese*.
- P. Hawke, 2014, *The Problem of Epistemic Relevance*, Dissertation, Department of Philosophy, Stanford University.
- J. Hintikka, 1962, *Knowledge and Belief*, Cornell University Press, Ithaca NY.
- J. Hintikka, 1973, *Logic, Language Games and Information*, Clarendon Press, Oxford.
- J. Hintikka, 2007, *Socratic Epistemology: Explorations of Knowledge-Seeking by Questioning*, Cambridge University Press, Cambridge UK.
- W. Holliday, 2012A, 'Epistemic Logic, Relevant Alternatives, and the Dynamics of Context', in *New Directions in Logic, Language, and Computation*, Springer Lecture Notes in Computer Science 7415, 109–129.
- W. Holliday, 2012B, *Knowing What Follows. Epistemic Closure and Epistemic Logic*, Dissertation, Department of Philosophy, Stanford University & Dissertation DS-2012-09, ILLC, University of Amsterdam. A substantial part relevant to this paper is to appear in the *Journal of Philosophical Logic* under the title 'Epistemic Logic and Epistemic Closure I: Relevant Alternatives and Subjunctivism'.
- W. Holliday, 2013, 'From Worlds to Possibilities for Knowledge and Belief', Department of Philosophy, University of California at Berkeley.
- W. Holliday & Th. Icard, 2010, 'Moorean Phenomena in Epistemic Logic', in L. Beklemishev, V. Goranko & V. Shehtman, eds., *Advances in Modal Logic*, Vol. 8, College Publications, London, 178–199.
- Th. Hoshi, 2009, *Epistemic Dynamics and Protocol Information*, Ph.D. thesis, Department of Philosophy, Stanford University (Dissertation ILLC-DS-2009-08).
- L. Humberstone, 1981, 'From Worlds to Possibilities', *Journal of Philosophical Logic* 10, 313–339.
- K. Kelly, 1996, *The Logic of Reliable Inquiry*, The Clarendon Press, Oxford.

- K. Kelly, 2013, 'A Computational Learning Semantics for Inductive Empirical Knowledge', in A. Baltag & S. Smets, eds., *Van Benthem Volume*, Springer, Dordrecht.
- K. Lehrer, 1990, *Theory of Knowledge*, Westview Press, Boulder.
- D. Lewis, 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549–567.
- F. Liu, 2011, *Dynamic Logics of Preference Change*, Springer Science Publishers, Dordrecht.
- R. Nozick, 1981, *Philosophical Explanations*, Princeton University Press, Princeton NJ.
- P. van Ormondt, 2012, *Finite Narrative Modeling, Contextual Dynamic Semantics and Elusive Knowledge*, Master of Logic Thesis, Institute for Logic, Language and Computation, University of Amsterdam.
- R. Parikh, Ç. Taşdemir, and A. Witzel, 2013, 'The Power of Knowledge in Games', *International Game Theory Review* 15, 1340030 (2013) [28 pages].
- H. Rott, 2005, 'Stability, Strength and Sensitivity: Converting Belief into Knowledge', *Erkenntnis* 61, 327–351.
- C. Shi, 2013, 'Logic of Evidence-Based Knowledge', *Lecture Notes in Computer Science* 8196, Springer Verlag, 347–351.
- R. Stalnaker, 1985, *Inquiry*, The MIT Press, Cambridge MA.
- R. Stalnaker, 2006, 'On Logics of Knowledge and Belief', *Philosophical Studies* 128:1, 169–199.
- K. Summerscale, 2012, *Mrs Robinson's Disgrace*, Bloomsbury, London.
- T. Williamson, 2000, *Knowledge and its Limits*, Oxford University Press, Oxford.
- S. Yablo, 2014, *Aboutness*, Princeton University Press, Princeton NJ.
- Z. Xu, 2010, *Knowledge, Evidence and Relevant Alternatives: A Logical Dynamical Approach*, Dissertation, Department of Philosophy, Peking University.