THEORIA, 2012

doi:10.1111/j.1755-2567.2011.01119.x

Everything is Knowable – How to Get to Know *Whether* a Proposition is True

b

HANS VAN DITMARSCH

University of Sevilla, Spain

WIEBE VAN DER HOEK

University of Liverpool, UK

PETAR ILIEV

University of Liverpool, UK

Abstract: Fitch showed that not every true proposition can be known in due time; in other words, that not every proposition is *knowable*. Moore showed that certain propositions cannot be consistently believed. A more recent dynamic phrasing of Moore-sentences is that not all propositions are known after their announcement, i.e., not every proposition is *successful*. Fitch's and Moore's results are related, as they equally apply to standard notions of knowledge and belief (S 5 and KD45, respectively). If we interpret 'successful' as 'known after its announcement' and 'knowable' as 'known after some announcement', successful implies knowable. Knowable does not imply successful: there is a proposition φ that is not known after its announcement but there is another announcement after which φ is known. We show that all propositions are knowable in the more general sense that for each proposition, it can become known or its negation can become known. We can get to know *whether* it is true: $\lozenge(K\varphi \lor K\neg \varphi)$. This result comes at a price. We cannot get to know whether the proposition *was* true. This restricts the philosophical relevance of interpreting 'knowable' as 'known after an announcement'.

Keywords: modal logic, knowability, Fitch's paradox, dynamic epistemics, public announcements

1. Successful - the Historical Record

TO OUR KNOWLEDGE, the first wording of the Moore-sentence is from the chapter *A reply to my critics*, from Moore's own hand, in the 1942 *Library of Living Philosophers* volume *The Philosophy of G.E. Moore*.

'I went to the pictures last Tuesday, but I don't believe that I did' is a perfectly absurd thing to say, although *what* is asserted is something which is perfectly possible logically. (Moore, 1942, p. 543)

Moore's (1912) *Ethics* provides a clue to the meaning of assert: asserting a proposition implies that I believe ('think to be') or know it:

there is an important distinction, which is not always observed, between what a man *means* by a given assertion and what he *expresses* by it. Whenever we make any assertion whatever... we are always expressing... either that we *think* the thing in question to be so or that we *know* it to be so (Moore, 1912, p. 77).

© 2011 Stiftelsen Theoria. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA.

Although by asserting we express a belief, the meaning of an asserted proposition cannot be equated with belief in that proposition, as that would lead to infinite regress:

But thus to believe that somebody believes, that somebody believes, that somebody believes ... quite indefinitely, without *ever* coming to anything which is what is believed, is to believe nothing at all. (Moore, 1912, p. 77)

All this is in the context of a discussion on whether moral judgements are judgements about our feelings, or about our beliefs. We emphasize that Moore (1912) does not formulate a Moore-sentence; for that we had to wait another thirty years. The absurdity in the cited passage of Moore (1942) then follows from the implicature 'asserting p implies belief of p', pointed out in Moore (1912). Similar examples can be found in Moore (1944, p. 204).

Let us write p for a proposition and Kp for both knowing and believing that p in the continuation we will show that the puzzling phenomena of knowability and success apply to both notions equally, although they are of course different in many other respects. The cited passage of Moore (1912) demonstrates that 'p' cannot be said to *mean* 'knowing p', as this would cause, by substitution, an infinite sequence p, Kp, KKp, KKKp, ad infinitum. And neither in Moore (1942) nor in Moore (1944) does he claim that $K(p \land \neg Kp)$ is inconsistent (in his own words: 'self-contradictory'). He states that asserting $p \land \neg Kp$ implies Kp, which contradicts $\neg Kp$ (Moore, 1944, pp. 204–205). His reluctance to formalize or describe in English the expression $K(p \land \neg Kp)$ may be methodological (avoiding infinite regress?) or esthetic. Either way, by the time of Hintikka (1962), the issue associated with the Moore-sentence means the inconsistency of $K(p \land \neg Kp)$.

There are two ways to derive the inconsistency of $K(p \land \neg Kp)$, and this reveals why the schema is relevant for knowledge and also for belief, i.e., for the common S 5 notion of knowledge and the common KD45 notion of belief. The first proof uses two properties of belief, D (consistency of belief, corresponding to seriality) and 4 (positive introspection, corresponding to transitivity), and it therefore also holds for knowledge.

¹ Hintikka's *Knowledge and Belief* contains an excellent list of references to the Moore-sentence (1962, p. 64). An entire chapter is devoted to its analysis. Although Hintikka (1962) cites Moore (1912) as a source, in fact the oldest Moore-sentence we found is in Moore (1942).

² The modal operator K models S 5 knowledge if it satisfies the axiom schemata $Kp \to p$ (T), $Kp \to KKp$ (4), and $\neg Kp \to K\neg Kp$ (5). The modal operator K models KD45 belief if it satisfies the axiom schemata $Kp \to \neg K\neg p$ (D), $Kp \to KKp$ (4), and $\neg Kp \to K\neg Kp$ (5). Both operators also satisfy the schema $K(p \to q) \to (Kp \to Kq)$ (K). These schemata actually contain formula variables, not propositional variables; see section 4. In computer science the system S 5 is well-accepted, but the negative introspection axiom (5) has been heavily debated among philosophers.

$$K(p \land \neg Kp)$$

$$\Rightarrow Kp \land K \neg Kp$$

$$\Rightarrow positive introspection$$

$$KKp \land K \neg Kp$$

$$\Rightarrow K(Kp \land \neg Kp)$$

$$\Rightarrow K\bot$$

$$\Rightarrow consistency of belief$$

Another way to derive the inconsistency is as follows:

$$K(p \land \neg Kp)$$
 \Rightarrow $Kp \land K \neg Kp$ \Rightarrow truth / property of belief $Kp \land \neg Kp$ \Rightarrow \bot

It *seems* as if the proof depends on the property that known propositions are true, and therefore only applies to knowledge. However, for the modal operator satisfying KD45, 'knowledge of ignorance' is equivalent to 'ignorance' (Meyer and van der Hoek, 1995): $K \neg Kp \leftrightarrow \neg Kp$. So also the second proof only depends on the properties of belief.

These proofs have gone around in the community. Hintikka (1962, p. 69) already mentions both and it also reappears in the recent literature, e.g., it is mentioned again by Linsky in Salerno's knowability volume (2009, p. 165). Our experience is that people seem unaware of proofs not based on the essential property $Kp \rightarrow p$ of knowledge, so the reader will excuse us for refreshing their memory.

2. Knowable - the Historical Record

The knowability paradox was clearly and recognizably formulated in a 1945 referee report on a submission by Fitch to the *Journal of Symbolic Logic* (see Figure 1):

... there is always a true proposition which it is empirically impossible for a to know at time t. For let k be a true proposition which is unknown to a at time t, and let k' be the proposition that k is true but unknown to a at time t. Then k' is true. But it would seem that if a knows k' at time t, then a must know k at time t, and must also know that he does not know k at time t. (Church, 1945)

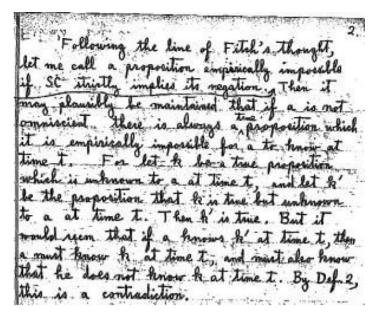


Figure 1: Fragment of an anonymous referee report of Fitch's paper, now attributed to Church

For 'empirically impossible', read 'inconsistent'. This citation from a handwritten note was uncovered by Salerno's archival efforts and after some further effort and handwriting comparison indisputably attributed to Church. Fitch's "A Logical Analysis of Some Value Concepts" (1963), based on the rejected 1945 version, only appeared 18 years later, namely in 1963. Fitch (1963) writes, very similar to the cited (Church, 1945) (and attributed to the reviewer, who was anonymous for him):

If there is some true proposition which nobody knows (or has known, or will know) to be true, then there is some true proposition that nobody can know to be true. (Fitch, 1963, p. 139)

Fitch's inconsistency proof uses the property that known formulas are true – it is therefore not surprising that this paradox has become associated with knowledge and not with belief. (For Fitch, 'knowledge' is a notion satisfying 'conjunction elimination', meaning $K(p \land q) \rightarrow (Kp \land Kq)$, 'conjunction introduction', meaning the converse implication; and the set of known propositions is a 'truth class', meaning that $Kp \rightarrow p$.) However, we have shown above that it is sufficient to assume properties of belief in order to derive a contradiction.

By now, 'knowability' does the round among philosophers as $\Diamond Kp$ (p is knowable) or $p \to \Diamond Kp$ (every truth can be known). The additional modal diamond \Diamond slipped in, to give meaning to the word 'can' in 'can know to be true'. Fitch suggests some implicit temporal connotation for 'can', as he mentions:

... the element of time will be ignored in dealing with these various concepts. (Fitch, 1963, p. 135)

Indeed, Fitch *does* ignore it and makes no difference between known and th-knowable, between 'knows to be true' and 'can know to be true'. Fitch (1963) does not distinguish two distinct modalities \diamondsuit and K.

A standard analysis of the Fitch paradox is as follows. We base our exposition on the excellent review of the literature on Fitch's paradox in the *Stanford Encyclopedia of Philosophy* (Brogaard and Salerno, 2004), an analysis one can see repeated many times in Salerno (2009). The existence of unknown truths is formalized as $\exists p$ $(p \land \neg Kp)$. The requirement that all truths are th-knowable is formalized as $\forall p$ $(p \to \Diamond Kp)$, where \Diamond formalizes the existence of some process after which p is known, or an accessible world in which p is known. Fitch's paradox is that the existence of unknown truths is inconsistent with the requirement that all truths are knowable.

The Moore-sentence $p \land \neg Kp$ witnesses the existential statement $\exists p(p \land \neg Kp)$. Assume that it is true. From $\forall p(p \to \Diamond Kp)$ follows the truth of its instance $(p \land \neg Kp) \to \Diamond K(p \land \neg Kp)$, and from that and $p \land \neg Kp$ follows $\Diamond K(p \land \neg Kp)$. Whatever the interpretation of \Diamond , it results in having to evaluate $K(p \land \neg Kp)$. But this is inconsistent for knowledge and belief.

Moore's paradox is traditionally more associated with the notion of belief, whereas Fitch's paradox is traditionally more associated with the notion of knowledge. The former is not often mentioned in the same breath as the latter. That is not surprising, as Moore talks about belief, and as Fitch talks about knowledge and derives the inconsistency of $K(p \land \neg Kp)$ with a property of knowledge. As we have seen, $p \land \neg Kp$ is not just unknowable, it is *unbelievable*.

Church's 1945 report and Moore's 1942 edited volume are very close in time. It makes one wonder if Moore, Church and Fitch were in contact with each other, and if Church, more a mathematician, was aware of the work of Moore, who was more a philosopher. We do not know. We think that Moore-sentences and Fitch-type paradoxes are closely related and deserve a combined treatment. We can achieve the integration by taking into account the dynamic turn in logic, which became eminent from the 1980s onward.

2.1 Why should everything be knowable?

The topic of *knowability* has done the rounds of philosophical communities (Salerno, 2009; Tennant, 1997; Dummett, 2001) since Fitch's 1963 publication. The knowability paradox answered a question posed in analytical philosophy: it is relevant in verificationism and in non-realism. Verificationism was for example proposed by A.J. Ayer (1936) in *Language, Truth and Logic*. The verification principle requires a non-analytic, meaningful true sentence to be empirically verifiable. Replace 'empirically verifiable' with 'knowable' (or recall 'empirically

impossible for *a* to know', above) and we are there. Anti-realism or non-realism is the philosophy that denies the existence of an objective reality of entities. In other words, there are no true unknowable propositions: a true proposition about the objective reality that has no counterpart in a knowing subject would be such an unknowable proposition. A contemporary influential proponent is Michael Dummett (1982) with his influential paper 'Realism'.

3. Successful - the Dynamic Turn

The further development of the Moore-sentence firstly gives a *multi-agent* perspective of announcements of the form "(I tell you that:) p is true and you don't know that", and, secondly, gives a *dynamic* perspective namely that such announcements cannot be believed after being announced. Both are quite different from Moore's original analysis that $p \land \neg Kp$ cannot be sincerely announced/uttered!

Unlike the single-agent version, the multi-agent version of the Moore-sentence is not problematic. If I tell you "You don't know that I play the cello", this has the conversational implicature "You don't know that I play the cello and it is true that I play the cello", and again we have the form $p \land \neg Kp$. However, this is not believed by you, but by me. (The announcement can be assumed to be made by an outsider not modelled in the logic with an epistemic operator. But in principle we can model both the speaker and the listener and we would get K_{me} ($p \land \neg K_{you}$ p) for different epistemic modalities K_{me} and K_{you} , as in the logic APAL presented in section 5.)

But we are now facing another problem. Suppose I were tell you again "You don't know that I play the cello". Then you can respond: "You're lying. You just told me that you play the cello." We can analyse what is going on here in modal logic. We model your uncertainty, for which a single epistemic modality suffices. Initially, there are two possible worlds, one in which p is true and another one in which p is false, and that you cannot distinguish from one another. Although in fact p is true, you don't know that: $p \wedge \neg Kp$. In this logic, we can also model the informative consequences of announcements. On the assumption that such announcements are public (all agents know that they are being informed, and know this about one another, etc.) and truthful (the announcements are assumed to be true), an announcement can be interpreted as a model restriction: the announcement of $p \wedge \neg Kp$ results in a restriction of these two possibilities to those where the announcement is true: in the p-world, $p \land \neg Kp$ is true, but in the $\neg p$ -world, $p \land \neg Kp$ is false. In the model restriction consisting of the single world where p is true, p is known: Kp. Given that Kp is true, so is $\neg p \lor Kp$, and $\neg p \lor Kp$ is equivalent to $\neg (p \land \neg Kp)$, the negation of the announced formula. So, announcement of $p \land \neg Kp$ makes it false! Gerbrandy (1999, 2007) calls this phenomenon an unsuccessful

update; the matter is also taken up in van Ditmarsch and Kooi (2006) and Qian (2002) and more recently in Holliday and Icard (2010). We will formally define public announcements in section 5.

The word 'unsuccessful update' is not coincidental. Another philosophical root of the dynamic turn in logics of belief and knowledge is the notion of *success*. In the area of belief revision (Alchourrón, Gärdenfors and Makinson, 1985) a well-known postulate describes that if you revise a theory (set of formulas) with novel information described in a proposition p, then p should after that revision process form part of the theory, it should be believed! This postulate is called the *success postulate*. Initially, belief revision had nothing to do with modal logic and with explicit knowledge K. We review how this came about. A theory K consists of a set of believed propositions, in propositional or first-order logic – let us assume this is in propositional logic, and let p be such a proposition. There are different theory change operators, modelling expansion, contraction and revision. For the purpose of explaining unsuccessful updates, it is sufficient to look at expansion. For the typical expansion we have that $p \notin K$, for the theory expanded with p we write $K \oplus p$, and the success postulate is the requirement that $p \in K \oplus p$.

Here, $p \notin \mathcal{K}$ means that p is initially not believed and $p \in \mathcal{K} \oplus p$ means that p is believed after expansion with p.

The AGM framework has been redescribed and expanded in modal logic. In retrospect, one could say that this required three steps.

The first step made it possible to have belief revision operators in the logical language, by formalizing these (meta-logical) operations as dynamic modal operators. In the case of belief expansions, we can let $[\oplus p]q$ express that after revision with p, q holds – where $[\oplus p]$ means 'perform belief expansion with p', a dynamic modal operator. This approach was suggested by van Benthem (1989) and further developed by de Rijke (1994).

The next step allowed for explicit modelling of belief and knowledge with K (or B) operators, where these operators bind propositional logical formulas. For example, $\neg Kp \land [\oplus p]Kp$ means that p is not believed ('known') and after revision with p, p is believed. This approach was followed in work by Segerberg and collaborators (1999), and a partial generalization was proposed to lift this to belief of modal propositions, such as $K(p \land \neg Kp)$ (Lindström and Rabinowicz, 1999), so-called higher-order belief.

The final step (although chronologically this took place independently of the second step) is to allow unrestricted belief revision with higher-order beliefs (truly 'unlimited DDL'). One might say that this was achieved for belief expansion in Plaza's public announcement logic (Plaza, 1989; van Ditmarsch, van der Hoek and Kooi, 2007), wherein we can say in the logical language that it may be true in some given Kripke model that $\neg Kp \land [p]Kp$ (the agent does not know p but after

announcement of p he knows that p), and also that $[p \land \neg Kp] \neg (p \land \neg Kp)$ is a validity: after announcing that p is true and that you don't know that, it is always false that (p) is true and that you don't know that). Like before, [p] stands for 'public announcement of p'. As public announcements are interpreted as Kripke model restrictions, this can be seen as a form of belief expansion. Here the parallel stops. Which AGM postulates for belief expansion are satisfied in public announcement logic, depends on your point of view (see, e.g., Bonanno, 2005). At least, the 'unsuccessful update' demonstrates that the success postulate is clearly not satisfied.

4. Knowable – the Dynamic Turn

The suggestion to interpret 'knowable' as 'known after an announcement' was made by van Benthem (2004), and Balbiani *et al.* (2008) propose a logic where ' φ is knowable' is interpreted in that way. In this setting, Φp stands for 'there is an announcement after which p (is true)', so that ΦKp stands for 'there is an announcement after which p is known', which is a form of 'proposition p is knowable'. To distinguish this specific interpretation of knowability from the more general Fitch setting we have written ΦKp instead of $\Diamond Kp$.

Before we present the logic in detail, let us first explore an example. Consider the proposition p for 'it rains in Liverpool'. Suppose you are ignorant about $p: \neg(Kp \lor K\neg p)$. First, suppose that p is true. I can announce to you here and now that it is raining in Liverpool (according to your expectations, maybe . . .), after which you know that: $\langle p \rangle Kp$ stands for 'p is true and after announcing p, p is known'. Now, suppose that p is false. In a similar way, after I announce of that, you know that; so that we have $\langle \neg p \rangle K \neg p$. If you already knew whether p, having its value announced does not have any informative consequence for you. Therefore, $\langle p \rangle Kp \lor \langle \neg p \rangle K \neg p$ is a validity: either the atom p holds and you can get to know that it is true, or it is false and you can get to know that it is false.

Let us now generalize the statement 'there is a proposition p such that after its announcement, p is known', to 'there exists a proposition q, such that after its announcement, p is known', where q is not necessarily the same as p. Then we have informally captured the meaning of $\blacklozenge Kp$. In other words, this operator is a

In public announcement logic, the 'box'-form [p]q stands for 'if p is true, then after (every) announcement p,q is true'; whereas the 'diamond'-form $\langle p \rangle q$ stands for 'p is true, and (there is an announcement of p such that) after announcement of p, q is true'. Of course, there is only one way to make an announcement of p: it is a functional operation. This is formalized by the principle $\langle p \rangle q \rightarrow [p]q$.

quantification over announcements. We have just proved that $\Phi(Kp \lor K \neg p)$ is a validity: given a model, if p is true, announce that it is true, and if p is false, announce that it is false. This schema captures the meaning of getting to know whether p, i.e., getting to know whether p is true or false.

Announcing the value of p is not the only possible announcement that I can make. Consider again your state of initial ignorance about p. Were I to make the trivial announcement \top , you would remain ignorant, $\langle \top \rangle \neg Kp$, so that, with negative introspection, we also have $\langle \top \rangle K \neg Kp$. Ignorance of p is knowable. But if you already knew p, not only does announcing p then make no difference, but announcing p would also not make a difference: $\langle \top \rangle Kp$, so, with positive introspection, $\langle \top \rangle KKp$: knowledge of p is knowable. On the other hand, after announcing $p \wedge \neg Kp$, this is not known, as $\langle p \wedge \neg Kp \rangle \neg K(p \wedge \neg Kp)$. And no other announcement can achieve that either: $\blacklozenge K(p \wedge \neg Kp)$ is not valid.

In the presentation so far, mainly of historical interest, we have been treading cautiously in order to avoid the crucial distinction in modal logics between propositional variables and formula variables. A propositional variable, or propositional letter, p cannot at will be replaced by what in modal logic tends to be called a proposition. In the tautology $p \vee \neg p$ we can replace p by $\neg p$ and $\neg p \vee \neg \neg p$ is still a tautology, and we can replace p by $p \wedge \neg Kp$ and $(p \wedge \neg Kp) \vee \neg (p \wedge \neg Kp)$ is a validity in the modal logic of knowledge. (To distinguish p from $p \land \neg Kp$ we call the former an atomic proposition.) Modal logics that satisfy this substitution property (and some other properties) are called normal modal logics. Multi-agent epistemic modal logic is a normal modal logic. However, public announcement logic and arbitrary public announcement logic are not normal modal logics: [p]Kpis valid (after announcing atomic proposition p, p is known). But substitute $p \wedge \neg Kp$ for p and disaster strikes, as $[p \wedge \neg Kp]K(p \wedge \neg Kp)$ is invalid. Similarly, $p \to \Phi Kp$ is valid but $(p \land \neg Kp) \to \Phi K(p \land \neg Kp)$ is invalid. We write φ, ψ, \ldots for (modal) formula variables, instead of p, q, \ldots , for propositional variables. Public announcement logic is not a normal modal logic, but there are many validities that can be formulated in terms of formula variables, such as $\phi \vee \neg \phi$, and $K\varphi \to KK\varphi$.

We now continue with the overview of arbitrary public announcement logic, followed by the investigation of knowable and successful in that logic.

5. Arbitrary Public Announcement Logic

Arbitrary public announcement logic is an extension of public announcement logic (Plaza, 1989). Let a finite set of agents Ag and a countable set of propositional variables At be given. These parameters can remain implicit.

Definition 1 (Language) The language $\mathcal{L}(K, [\cdot], \blacksquare)$ is defined as

$$\varphi := p |\neg \varphi| (\varphi \land \varphi) |K_a \varphi| [\varphi] \varphi | \blacksquare \varphi$$

where $p \in At$ and $a \in Ag$.

Disjunction and implication are defined as usual. A formula that only contains atoms from At and Boolean connectives is called *objective*. The language of public announcement logic $\mathcal{L}(K, [\cdot])$ is the fragment of $\mathcal{L}(K, [\cdot], \blacksquare)$ without \blacksquare . Likewise, we define the language of (multi-agent) epistemic logic $\mathcal{L}(K)$, the language of knowability logic $\mathcal{L}(K, \blacksquare)$, and the language \mathcal{L} of propositional logic. For the dual of $[\psi]\varphi$ we write $\langle \psi \rangle \varphi$, the dual of $K_a\varphi$ is written $\hat{K}_a\varphi$ and the dual of $\blacksquare \varphi$ is $\blacklozenge \varphi$. Formula $K_a\varphi$ stands for 'agent a knows φ ', $[\varphi]\psi$ stands for 'after announcement of φ , ψ ', and $\square \varphi$ stands for 'after every announcement, φ '. We have chosen the symbols \spadesuit and \blacksquare to contrast them with the operators used for knowability in section 2, but beware that our \spadesuit and \blacksquare are the same as \diamondsuit and \square in APAL as defined in Balbiani et al. (2008).

Definition 2 (Epistemic model) An epistemic model M is a tuple $M = (S, \sim, V)$ such that

- S is a non-empty set of possible worlds,
- $\sim : Ag \rightarrow \wp(S \times S)$ assigns an equivalence relation to each agent,
- $V: At \to \wp(W)$ assigns a set of possible worlds to each propositional variable.

If $M = (S, \sim, V)$, rather than $s \in S$, we will also write $s \in M$. For $\sim(a)$ we write \sim_a . A *pointed model* is a pair (M, s) where $s \in M$.

Definition 3 (Submodel) *Let two epistemic models* $M = (S, \sim, V)$ and $M' = (S', \sim', V')$ *be given. The pointed model* (M, s) *is a* submodel *of the pointed model* (M', s') *if*

- 1. $S' \subset S$,
- 2. s=s',
- 3. $\sim_a' = \sim_a \cap (S' \times S'),$
- 4. $V(p') = V(p) \cap S'.$

Note that for each non-empty subset X of S there is a unique submodel: the model M restricted to X, notation M|X. If X is the denotation of a formula φ , we write $M|\varphi$. In other words, $M|\varphi$ is the model M restricted to those worlds where φ is true (φ may not be valid on $M|\varphi$, as in the case for $p \land \neg K_a p$).

Definition 4 (Semantics)

$M, s \models p$	iff	$s \in V(p)$
$M, s \models \neg \varphi$	iff	$M, s \not\models \varphi$
$M, s \models \varphi \land \psi$	iff	$M, s \models \varphi \text{ and } M, s \models \psi$
$M, s \models K_a \varphi$	iff	$M, t \models \varphi \text{ for every } t \text{ such that } s \sim_a t$
$M, s \models [\varphi] \psi$	iff	if $M, s \models \varphi$, then $M \varphi, s \models \psi$
$M, s \models \blacksquare \psi$	iff	for all $\varphi \in \mathcal{L}(K)$, $M, s \models [\varphi] \psi$

If M, $s \models \varphi$ for all M and s, we write $\models \varphi$, for ' φ is valid'. The restriction to multi-agent epistemic formulas $\varphi \in \mathcal{L}(K)$ in the semantics of $\blacksquare \varphi$ is for technical reasons; if $\varphi \in \mathcal{L}(K, [\cdot], \blacksquare)$ were allowed, the semantics would be a circular definition, as this would quantify over the precise $\blacksquare \psi$ we are trying to determine. The restriction to epistemic formulas amounts to a restriction to 'Box-free' formulas, as public announcement logic is equally expressive as multi-agent epistemic logic.

Arbitrary public announcement logic has a complete axiomatization, for at least two agents it is strictly more expressive than multi-agent epistemic logic, it is non-compact and it is undecidable. For details, see Balbiani *et al.* (2008). Valid principles of the logic include:

• $\blacksquare \varphi \rightarrow \varphi$

If φ holds after every announcement, then it holds also after the trivial announcement of \top , so it was already true.

• $\blacksquare \varphi \rightarrow \blacksquare \blacksquare \varphi$

The composition of two announcements is again an announcement. The dual version $\Phi \Phi \varphi \to \Phi \varphi$ more clearly corresponds to that intuition: if there are ψ and χ such that $\langle \psi \rangle \langle \chi \rangle \varphi$, then we also have, using a property of public announcement logic, $\langle \psi \wedge [\psi] \chi \rangle \varphi$ (which is also equivalent to $\langle \langle \psi \rangle \chi \rangle \varphi$), and therefore $\Phi \varphi$.

- Church-Rosser: $\blacklozenge \blacksquare \varphi \rightarrow \blacksquare \blacklozenge \varphi$
- McKinsey: $\blacksquare \blacklozenge \varphi \rightarrow \blacklozenge \blacksquare \varphi$

5.1 Example

Consider model $M = (S, \sim, V)$ of Figure 2, modelling the uncertainty of two agents 1 and 2, where $S = \{s_1, s_2, s_3, s_4\}$, where \sim_1 is the reflexive closure of $\{(s_3, s_4), (s_4, s_3), (s_1, s_2), (s_2, s_1)\}$, and where, similarly, agent 2 cannot distinguish s_2 from s_3 nor s_1 from s_4 . Also stipulate $V(p) = \{s_1, s_2\}$ and $V(q) = \{s_1, s_4\}$. Then

$$M, s_1 \vDash p \land q \land \neg K_1 q \land \neg K_2 p \land \hat{K}_1 \hat{K}_2 (\neg p \land q)$$

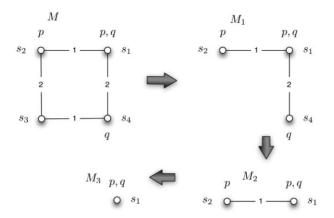


Figure 2: A model M and three consecutive announcements

Now consider the announcement $p \vee q$: this transforms M into $M_1 = M|(p \vee q)$. The following is true in (M, s_1) since (1) $p \vee q$ is true in s_1 , and (2) the formula bound by $\langle p \vee q \rangle$ is true in (M_1, s_1) :

$$M, s_1 \models \langle p \lor q \rangle \Big(\neg K_1 q \land K_2 p \land \neg \hat{K}_1 \hat{K}_2 (\neg p \land \neg q) \Big)$$

Suppose that in (M_1, s_1) agent 1 now publicly announces the true statement that he does not know q. Since in (M_1, s_4) agent 1 *does* know q, this state gets eliminated from the model, resulting in (M_2, s_1) . We have:

$$M, s_1 \vDash \langle p \lor q \rangle \langle \neg K_1 q \rangle (\neg K_1 q \land K_2 p)$$

In other words, one effect of agent 1 announcing he does not know that q is that agent 2 comes to know that p! Finally, if in (M_2, s_1) , agent 2 announces the true proposition K_2q , we end up in model $M_3 = M_2|K_2q$. So we have

$$M, s_1 \vDash \langle p \lor q \rangle \langle \neg K_1 q \rangle \langle K_2 q \rangle (K_1(p \land q) \land K_2(p \land q))$$

Now for some examples involving 'arbitrary announcement' operators. The previous establishes that M, $s_1 \models \Phi(K_1(p \land q) \land K_2(p \land q))$; the three announcements can be made into one, using the property of the logic that $\langle \phi \rangle \langle \psi \rangle \chi$ is equivalent to $\langle \phi \land \langle \phi \rangle \psi \rangle \chi$. We also have M, $s_1 \models \Phi(K_1K_2(p \lor \neg q) \land K_2K_1(p \lor \neg q))$: there is an announcement such that the agents have mutual knowledge that $(p \lor \neg q)$; the announcement could in fact be $p \lor \neg q$. Now take $\phi = p \land \neg K_1 p$. Although M, $s_1 \models \varphi$, there is no announcement that can reveal φ to agent 1: M, $s_1 \models \neg \Phi K_1 \varphi$.

However, it *is* possible to do an announcement with $K_1 \neg \varphi$ as effect: M; $s_1 = \Phi K_1 \neg \varphi$. This is true because, e.g., M, $s_1 \models \langle p \rangle K_1 \neg \varphi$. So for that φ we have $\Phi(K_1 \varphi \vee K_1 \neg \varphi)$.

6. Successful and Knowable

Given some formula φ , an intuitive way in which it can be said to be 'successful' or 'knowable' is relative to a pointed model (M, s) in which φ is true. An update with a formula φ is successful in a pointed model (M, s) (for an agent a) if it is true and if after announcing it (i.e., after the update), it is known:

$$M, s \models \varphi \land \langle \varphi \rangle K_a \varphi$$
.

Similarly, a true formula can be called *knowable in a pointed model* (M, s) (to agent a) if it is indeed true and if there is a way to make it known, i.e., if there is an announcement after which it is known:

$$M, s \models \varphi \land \spadesuit K_a \varphi$$
.

Note that $\varphi \wedge \langle \varphi \rangle K_a \varphi$ is equivalent to $\langle \varphi \rangle K_a \varphi$, but the former makes the relation to knowability clearer. Clearly, the notions are related. It seems possible that announcing a knowable formula may not result in knowing that formula – we will indeed give a counterexample.

It is not obvious what the most natural definition is of a successful and knowable *formula*, independent from a specific model. First we deal with successful, and then with knowable.

6.1 Successful

Definition 5 (Successful) A formula $\varphi \in \mathcal{L}(K, [\cdot], \blacksquare)$ is successful (for agent a) iff $\models [\varphi] \varphi$. A formula is unsuccessful if it is not successful.

The definition is *global*: it refers to a validity of the logic. A formula φ is successful if in any model and any state, announcing φ in that state, would result in a state where φ is true. Objective formulas do not change the state, so they are successful. A formula like Kp is also successful. Note also that any contradiction is successful: there is no model that is the result of an announcement with \bot , so in any such model, anything holds. Examples of formulas that are *not* successful are $p \land \neg Kp$, or, in a multi-agent setting, $K_a p \land K_b \neg K_c p$.

Definition 5 of successful entails that successful for any agent means successful for all agents. Proposition 1 below makes clear why this is reasonable.

Proposition 1 (Different views on successful (van Ditmarsch and Kooi, 2006)) All the following descriptions of successful are equivalent:

- $\models [\varphi]\varphi$
- $\models \varphi \rightarrow \langle \varphi \rangle K_a \varphi$
- $\models \varphi \rightarrow \langle \varphi \rangle \varphi$

The equivalence between $[\varphi]\varphi$ and $\varphi \to \langle \varphi \rangle K_a \varphi$ validates the former as a definition of successful, because the latter is the exact paraphrase of 'if φ is true, then after announcement of φ it is known', modulo the conditional this is the notion of successful update given above. This is similar to typical belief expansion in an AGM belief revision setting: if a theory is expanded with consistent information φ (read: if the believing agent has decided to accept φ as true information, and does not believe the opposite), then φ is believed in the expansion.

We list some results for successful formulas from van Ditmarsch and Kooi (2006) and Holliday and Icard (2010), to which we add some novel ones of our own.

Proposition 2 (Successful)

- 1. Not all formulas are successful.
- 2. There are φ such that both φ and $\neg \varphi$ are successful.
- 3. There are φ such that both φ and $\neg \varphi$ are unsuccessful.

Proof

- 1. $p \wedge \neg K_a p$ is not learnable and not successful.
- 2. p and $\neg p$ are both successful.
- 3. Consider $(p \land \neg K_a p) \land \neg (q \land \neg K_a q)$. Clearly, this formula is not successful, as after announcing it, p is known by a, so $(p \land \neg K_a p)$ is false. But its negation is also not successful. Consider the pointed four-state model with maximal uncertainty about (universal access between) the value of two atoms p and q, and where these are both true. After announcing the negation $\neg p \lor K_a p \lor (q \land \neg K_a q)$ of the formula above, three states remain (the formula is only false when p is true and q is false). In the state where p and q are both true, this formula is now false.

There are more results of this kind, for example, formulas φ and ψ may be successful, but not $\varphi \wedge \psi$, or not $\varphi \rightarrow \psi$, or not $\neg \varphi$ (van Ditmarsch and Kooi, 2006), or not $\varphi \vee \psi$ (Holliday and Icard, 2010). The recent investigation by Holliday and Icard (2010) characterizes the successful formulas for single-agent $\mathcal{L}(K, [\cdot])$, a remarkable result. They also distinguish further notions, such as the *supersuccessful formulas* φ that are always known after being announced but additionally

© 2011 Stiftelsen Theoria.

remain true in every model in between the initial model M and the model restriction $M|\varphi$. The characterization of the successful formulas for multi-agent $\mathcal{L}(K, [\cdot])$ is unknown.

We can contrast this 'global' notion of successful, a formula property, with the more intuitive 'local' notion at the beginning of this section: a relation between a pointed state and a formula. An update with a formula φ is *successful in a pointed model* (M, s) (for an agent a) if it is true and if after announcing it, it is known, $M, s \models \varphi \land \langle \varphi \rangle K_a \varphi$; and an update is *unsuccessful in a pointed model* (M, s) (for an agent a) if it is true and if after announcing it, it is known to be false. Clearly, a formula is successful if it is a successful update in *all* models, whereas an unsuccessful formula may well be successful in some models. The typical example is 'not stepping forward' ('nobody knows whether he is muddy') in the Muddy Children problem (Fagin and Vardi, 1986): this formula is only unsuccessful when the muddy children finally step forward; otherwise, it is successful: they still don't know it!

We continue with the investigation of 'knowable'.

6.2 Knowable

We recall the relative notion of knowability. A true formula is knowable in a pointed model (M, s) to agent a if it is indeed true and if there is a way to make it known: $M, s \models \varphi \land \blacklozenge K_a \varphi$. Given that, it might be tempting to call a formula knowable if $igspace K_a \varphi$ is satisfiable, but that amounts to the same as requiring that $K_a \varphi$ is satisfiable: let (M, s) be a pointed model such that $M, s \models \Phi K_a \varphi$, then there is a $\psi \in \mathcal{L}(K)$ such that $M, s \models \langle \psi \rangle K_a \varphi$, and therefore $M | \psi, s \models K_a \varphi$. So, $K_a \varphi$ is satisfiable. To require the even stronger validity of $\Phi K_a \varphi$ is also doomed, as now even propositional variables would not be knowable: for a simple formula as $\oint K_a p$ to be valid, p has to be true in all models; but of course, it is sometimes true and sometimes false. Given the popular requirement $\forall p(p \to \Diamond Kp)$ in the literature on the Fitch paradox, as discussed in section 2, our next best option is to require validity of $\varphi \to \Phi K_a \varphi$, for 'all true formulas are knowable' (for agent a). This has also been called 'learnability' in the dynamic epistemic logic literature (Balbiani et al., 2008; Holliday and Icard, 2010). Now, indeed, we can rightfully call a propositional variable p knowable, as $p \to \Phi K_a p$ is valid (see section 4). We call this th-knowability, for knowing that a formula is true. One further option down the road, slightly weaker, is to require validity of $\Phi(K_a \varphi \vee K_a \neg \varphi)$; we call that wh-knowability, for knowing whether a formula is true.

Definition 6 (Knowable) *Let* $\varphi \in \mathcal{L}(K, [\cdot], \blacksquare)$ *be a formula.*

- φ is th-knowable (for agent a) iff $\models \varphi \rightarrow \Phi K_a \varphi$;
- φ is wh-knowable (for agent a) iff $\models \Phi(K_a \varphi \vee K_a \neg \varphi)$.

Th-knowable formulas are those that, if true, are known after some announcement, whereas wh-knowable formulas are those that are known to be true after some announcement or are known to be false after some announcement. The conditional flavour of the definition makes any contradiction th-knowable. Theorem 1 will show that all formulas are wh-knowable, so that wh-knowable for any agent means wh-knowable for all agents. This then trivially entails that th-knowable implies wh-knowable. We anticipate that result in Proposition 3, in an interesting variation. Proposition 4 will show that a formula may be th-knowable for one agent but not th-knowable for another agent.

Examples of formulas that are th-knowable for agent a are: p, $\neg p$, $K_a \varphi$ and $\neg K_a \varphi$ for all φ (use introspection, and the trivial announcement). Whereas $K_a(p \land \neg K_b p)$ is th-knowable for a, but not for b (see Proposition 4).

The schema $\Phi(K_a \varphi \vee K_a \neg \varphi)$ we have not encountered before in the literature. A wh-knowable formula may be true now, but known to be false after an announcement. For a pregnant example, $p \wedge \neg K_a p$ is wh-knowable, because after its own announcement it is known to be false.

Proposition 3 *Th-knowable implies wh-knowable.*

Proof Consider the following equivalences:

Clearly, th-knowable implies the weaker wh-knowable.

Proposition 4 (Th-knowable and successful)

1. Let $\varphi \in \mathcal{L}(K, [\cdot])$. If φ is successful, then φ is th-knowable (van Ditmarsch and Kooi, 2006).

- 2. There are th-knowable formulas that are not successful.
- 3. Th-knowable for a given agent does not imply th-knowable for all agents.

Proof

1. This follows from the observation that the validity of $[\varphi]\varphi$ is equivalent to the validity of $\varphi \to \langle \varphi \rangle K_a \varphi$, and that $\varphi \to \langle \varphi \rangle K_a \varphi$ implies $\varphi \to \Phi K_a \varphi$; if φ is

known after announcement of φ , then there is an announcement after which it is known. Given the semantics of \blacksquare , the announcement witnessing it should be \blacksquare -free.

- 2. Take $\varphi = K_a(p \land \neg K_b p)$. Take an M, s where φ is true. Then we have M, $s \models \blacklozenge K_a \varphi$: after the trivial announcement, the formula is still true. But M, $s \not\models \langle \varphi \rangle K_a \varphi$: after agent a announces her knowledge, b knows p as well so φ is now false.
- 3. The formula $K_a(p \wedge \neg K_b p)$ used in the previous item is clearly not th-knowable for agent *b*.

Our proposals to define th-knowable and wh-knowable as the validity of $\varphi \to \Phi K \varphi$ and $\Phi (K \varphi \lor K \neg \varphi)$, respectively, are tentative in the sense that there are yet other ways to pin down syntactic fragments of the logical language $\mathcal{L}(K, \lceil \cdot \rceil, \blacksquare)$.

We already saw that $\blacklozenge(K\varphi \lor K\neg \varphi)$ is equivalent to $(\varphi \to \blacklozenge K\varphi) \lor (\neg \varphi \to \blacklozenge K\neg \varphi)$, and that this is obviously weaker than $\varphi \to \blacklozenge K\varphi$. What formulas satisfy the stronger

$$(\varphi \to \bigstar K\varphi) \land (\neg \varphi \to \bigstar K \neg \varphi)$$
?

We note that the negation $\neg(p \land \neg Kp)$ of the Moore-sentence satisfies th-knowability but, obviously, not $(\varphi \to \blacklozenge K\varphi) \land (\neg \varphi \to \blacklozenge K\neg \varphi)$; as its negation is the Moore-sentence again, so the second conjunct is not satisfied.

If, in order to avoid that inconsistencies are th-knowable, we were to require that $\varphi \to K\varphi$ is valid and φ satisfiable, how close does that come to requiring that $\varphi \land \blacklozenge K\varphi$ is satisfiable? Closer than mere th-knowability indeed, but not close enough. For example, given two agents i and j, let φ be $p \land \neg K_j p$. We then have that $\varphi \land \blacklozenge K_i \varphi$ is satisfiable, namely in a model consisting of a p-state and a $\neg p$ -state and wherein j is uncertain about p but i is not. On the other hand, this formula is not th-knowable for agent i, because not every model satisfying φ also satisfies $\blacklozenge K_i \varphi$. For a counterexample, take again a model consisting of a p-state and a $\neg p$ -state, but now such that neither i nor j can distinguish between those states.

6.3 Everything is knowable

In order to derive our main result that all formulas are wh-knowable, we first repeat the following lemma.

Lemma 1 ([3, Lemma 3.2]) Let $\varphi \in \mathcal{L}(K, [\cdot], \blacksquare)$. Consider the set P_{φ} of atoms occurring in φ . Let M be a model where all states correspond on the valuation of P_{φ} (i.e., $\forall p \in P_{\varphi}(V(p) = S \text{ or } V(p) = \varnothing)$). Then $M \models \varphi \text{ or } M \models \neg \varphi$, i.e., either φ or its negation is a model validity.

© 2011 Stiftelsen Theoria.

П

Theorem 1 Every formula φ is wh-knowable:

$$\models \spadesuit(K_a \varphi \lor K_a \neg \varphi)$$

Proof Given a formula φ and a pointed model M, s, define δ_s^{φ} as the characteristic formula of the atoms P_{φ} in s:

$$\delta_s^{\varphi} = \bigwedge \{ p | p \in P_{\varphi} \text{ and } M, s \models p \} \land \bigwedge \{ \neg p | p \in P_{\varphi} \text{ and } M, s \models \neg p \}$$

From Lemma 1 we immediately obtain

$$M|\delta_s^{\varphi} \vDash \varphi \text{ or } M|\delta_s^{\varphi} \vDash \neg \varphi$$

Take $M | \delta_s^{\varphi}, s$: this is nothing else than the model obtained when δ_s^{φ} is announced in M, s, so for any ψ with $M | \delta_s^{\varphi} \models \psi$ we have M, $s \models \spadesuit \psi$. It follows immediately from (1) that $M | \delta_s^{\varphi} \models K_a \varphi$ or $M | \delta_s^{\varphi} \models K_a \neg \varphi$ (for an *arbitrary* agent a). Hence $M | \delta_s^{\varphi} \models K_a \varphi \vee K_a \neg \varphi$, and hence M, $s \models \spadesuit (K_a \varphi \vee K_a \neg \varphi)$.

Our result says that for every formula, either that formula or its negation can be known, where 'can be known' means 'known after some announcement'. The result is *not* that for every formula, if currently true it can be known to be true, and if currently false it can be known to be false. In other words, we cannot get to know for every formula that it *was* true or that it *was* false; only that it *is* true or that it *is* false. The value of the formula may change as a result of the announcement, as in the case of $p \land \neg K_a p$. This formula, when true, can be known to be false after its announcement.

The proof of Theorem 1 is constructive, in the sense that we know which announcement leads to the knowledge of either φ or its negation: announce the current truth value of all atoms involved. This is in some sense disappointing: the agents do not learn what multi-agent uncertainty about factual information actually was the case, but the world is *manipulated* for them. For example, suppose a pointed model wherein $\varphi = \hat{K}_b q \wedge \hat{K}_b \neg q$ is true, and wherein (obviously) a considers that possible, but wherein also $\hat{K}_a K_b q$ and $\hat{K}_a K_b \neg q$ are true, and suppose that currently q holds. Is φ knowable? Yes, a can get to know that it is false: announce q, and a knows that b knows a. But if a did not already know that $\hat{K}_b q \wedge \hat{K}_b \neg q$ was true, a does not learn from the announcement of a that a was ignorant about a before that announcement. The formula a0 cannot be said to find out the truth about a0.

In that sense, it is not very meaningful to say that everything is knowable. It does not mean that everything true now can be known to be true in future, and everything false now can be known to be false in future. (And of course since (Fitch, 1963) we know that we don't *want* that, as in the case of $p \land \neg Kp$.) It means that we can find out the value of every proposition in the future, not necessarily the value it currently has, but the truth-value it will have at the point we have found it out, possibly different from its current truth-value.

In another sense, it is meaningful: it is still not the case that an agent always knows in advance what will be known or not. Sometimes he does: if the proposition is $p \land \neg K_a p$, i.e., if it is self-refuting, clearly the only value that can be known later is that it is false. Also for the formula $\hat{K}_b q \land \hat{K}_b \neg q$ in the above example, a knows that there is an announcement after which he knows it to be false. He just does not know what the announcement is! If q is true, announcing that makes it false, otherwise, announcing $\neg q$ makes it false. But sometimes the agent does not know in advance what will be known, as in the above case for the formula K_bq : a knows that the truth about q can be announced, if q is true then after announcing that $K_a \neg K_bq$ is true and if q is false then after announcing that $K_a \neg K_bq$ is true (because $K_a \neg q$ entails $K_a \neg K_bq$). In that case, agent a truly only can get to know $whether K_bq$ is true.

6.4 Back to Fitch

Let us summarize the results for the \blacklozenge operator, where $\blacklozenge K\varphi$, for ' φ is knowable', means ' φ is known after an announcement'.

- For every true proposition we can get to know that it is true. False
- For every true proposition we can get to know that it was true. False
- For every proposition we can get to know whether it is true. *True*
- For every proposition we can get to know whether it was true. False

How does this bear on Fitch knowability, with $\Diamond K\varphi$ instead of $\bigstar K\varphi$? We have enforced a concrete interpretation of 'getting to know' and we should ask ourselves how far we have strayed from the trodden knowability-path while doing that.

The verificationists and non-realists will not be satisfied by our result. For every proposition they want to get to know whether it *was* true, not whether it *is* true. As known since Fitch, this is not possible for higher-order knowledge. In epistemic systems, as in experimental physics, observing the system may change the properties of the system. An agent who is being informed about the truth of a given proposition is like an experimenter observing the value of a system parameter. Just as information may no longer be true because you are being informed, performing a measurement may change the value(s) of the measured parameter.

We hope that our contribution may further the philosophical investigation of the schemata we proposed for other logics of knowability: given some other interpretation of \diamondsuit (than \spadesuit), for what φ is $\diamondsuit(K\varphi \lor K\neg \varphi)$ valid? Or, to mention another

schema we discussed: for what φ is $\varphi \to \diamondsuit K \varphi$ valid and φ satisfiable? For which φ is both φ and $\neg \varphi$ th-knowable, i.e., $(\varphi \to \diamondsuit K \varphi) \land (\neg \varphi \to \diamondsuit K \neg \varphi)$ valid? Or, after all, given a dynamic epistemic logic with history operators, for which φ can we get to know whether they *were* true? And the schema $\varphi \to \diamondsuit K \varphi$ also appears in other modal logical settings, such as topologic (Parikh, Moss and Steinsvold, 2007). Are there parallels with yet other modal logics for spatial reasoning?

7. Further Research

There are a number of topics for further research in the technical setting of the logic where ◆ means 'there is an announcement after which'.

- A syntactic characterization of the th-knowable formulas in arbitrary public announcement logic is unknown, and also how this would relate to the successful formulas.
- To investigate the wh-knowability construct $\bigstar K \varphi$ in a dynamic epistemic logic, it is convenient to have the announcement operator, but the announcement does not occur in the construct, so a fair question seems the investigation of knowability in the logical language $\mathcal{L}(K, \blacksquare)$ defined as $\varphi := p \mid \neg \varphi \mid (\varphi \land \varphi) \mid K_a \varphi \mid \blacksquare \varphi$. In that language, the semantics of \blacksquare is: $M, s \models \blacksquare \psi$ iff for all $\varphi \in \mathcal{L}(K)$, $M \mid \varphi$, $s \models \psi$. The axiomatization of that logic is unknown, and it seems non-trivial.
- The multi-agent setting of knowability allows for different generalizations. Consider that $\Phi \varphi$ does not stand for 'there is an announcement after which φ ' but for 'there is an informative action after which φ '. An informative action may be a private announcement, or any other complex but not public action. Are there propositions that are only knowable in that setting?
- Onto a different track comes knowability with group epistemic operators, e.g., which propositions are commonly knowable, or distributedly knowable, or transferable between agents? These questions come with the schemata $\Phi C_{Ag} \varphi$, $\Phi D_{Ag} \varphi$, and $K_a \varphi \to \Phi K_b \varphi$, respectively.

Acknowledgements

We thank the anonymous reviewers of *Theoria* for their comments. Part of this work evolved over various presentations – we thank Paul Egré, Wes Holliday, Thomas Icard, Stephen Read and Joe Salerno for their comments. Hans van Ditmarsch is also affiliated as associated researcher to the Institute of Mathematical Sciences, Chennai, India.

References

- ALCHOURRÓN, C.E., GÄRDENFORS, P. and MAKINSON, D. (1985) "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." *Journal of Symbolic Logic* 50: 510–530.
- AYER, A.J. (1936) Language, Truth and Logic. London: Victor Gollancz Ltd.
- BALBIANI, P., BALTAG, A., VAN DITMARSCH, H., HERZIG, A., HOSHI, T. and DE LIMA, T. (2008) "'Knowable' as 'Known after an Announcement'." *Review of Symbolic Logic* 1(3): 305–334.
- BONANNO, G. (2005) "A Simple Modal Logic for Belief Revision." Synthese (Knowledge, Rationality & Action) 147(2): 193–228.
- BROGAARD, B. and SALERNO, J. (2004) *Fitch's Paradox of Knowability*. http://plato.stanford.edu/archives/sum2004/entries/fitch-paradox/.
- CHURCH, A. (1945) First anonymous referee report on Fitch's "a definition of value". Sent to Ernest Nagel, co-editor of the *Journal of Symbolic Logic*.
- DE RIJKE, M. (1994) "Meeting Some Neighbours." In J. van Eijck and A. Visser (eds), *Logic and Information Flow*, pp. 170–195. Cambridge MA: MIT Press.
- DUMMETT, M. (1982) "Realism." Synthese 52(1): 55-112.
- DUMMETT, M. (2001) "Victor's Error." Synthesis 61: 1-2.
- FAGIN, R. and VARDI, M.Y. (1986) "Knowledge and Implicit Knowledge in a Distributed Environment." In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning about Knowledge*, pp. 187–206. San Francisco, CA: Morgan Kaufmann.
- FITCH, F.B. (1963) "A Logical Analysis of Some Value Concepts." *Journal of Symbolic Logic* 28(2): 135–142.
- GERBRANDY, J.D. (1999) "Bisimulations on Planet Kripke." PhD thesis, University of Amsterdam. ILLC Dissertation Series DS-1999-01.
- GERBRANDY, J.D. (2007) "The Surprise Examination." Synthese 155(1): 21-33.
- HINTIKKA, J. (1962) Knowledge and Belief. Ithaca, NY: Cornell University Press.
- HOLLIDAY, W. and ICARD, T. (2010) "Moorean Phenomena in Epistemic Logic." In L. Beklemishev, V. Goranko and V. Shehtman (eds), *Advances in Modal Logic 8*, pp. 178–199. London: College Publications.
- LINDSTRÖM, S. and RABINOWICZ, W. (1999) "DDL Unlimited: Dynamic Doxastic Logic for Introspective Agents." *Erkenntnis*, 50: 353–385.
- MEYER, J.-J.Ch. and VAN DER HOEK, W. (1995) *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41. Cambridge: Cambridge University Press.
- MOORE, G.E. (1912) *Ethics*. Oxford University Press. Consulted edition: Home University Library of Modern Knowledge, Vol. 54, OUP, 1947.
- MOORE, G.E. (1942) "A Reply to My Critics." In P.A. Schilpp (ed.), *The Philosophy of G.E. Moore*, pp. 535–677. The Library of Living Philosophers, Vol. 4. Evanston, IL: Northwestern University.
- MOORE, G.E. (1944) "Russell's 'Theory of Descriptions'." In P.A. Schilpp (ed.), *The Philosophy of Bertrand Russell*, pp. 175–225. The Library of Living Philosophers, Vol. 5. Evanston, IL: Northwestern University.
- PARIKH, R., Moss, L.S. and Steinsvold, C. (2007) "Topology and Epistemic Logic." In M. Aiello, I. Pratt-Hartmann and J. van Benthe (eds), *Handbook of Spatial Logics*, pp. 299–341. Berlin: Springer Verlag.

- PLAZA, J.A. (1989) "Logics of Public Communications." In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic and Z.W. Ras (eds), *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, pp. 201–216. Oak Ridge National Laboratory.
- QIAN, L. (2002) "Sentences True after Being Announced." In *Proceedings of the Student Conference of NASSLLI 2002*. Stanford University.
- SALERNO, J. (ed.) (2009) New Essays on the Knowability Paradox. Oxford: Oxford University Press
- SEGERBERG, K. (1999) "Two Traditions in the Logic of Belief: Bringing Them Together." In H.J. Ohlbach and U. Reyle (eds), *Logic, Language, and Reasoning*, pp. 135–147. Dordrecht: Kluwer Academic Publishers.
- TENNANT, N. (1997) The Taming of the True. Oxford: Oxford University Press.
- VAN BENTHEM, J. (1989) "Semantic Parallels in Natural Language and Computation." In *Logic Colloquium* '87. Amsterdam: North-Holland.
- VAN BENTHEM, J. (2004) "What One May Come to Know." Analysis 64(2): 95-105.
- VAN DITMARSCH, H. and KOOI, B. (2006) "The Secret of My Success." Synthese 151: 201–232.
- VAN DITMARSCH, H., VAN DER HOEK, W. and KOOI, B. (2007) *Dynamic Epistemic Logic*. Synthese Library Vol. 337. Berlin: Springer.

Note: Some small typographical errors in section 6.2 have been corrected in this version of the article on 20 February 2012 after first publication online on 21 October 2011.