5.1 Doxastic and Epistemic Logic

Hans van Ditmarsch

Summary

We present modal-logical semantics for knowledge and belief, alternative semantics for knowledge, and also systems with interaction between knowledge and belief. We then present common knowledge and common belief, as well as distributed knowledge. We conclude with topics involving change of knowledge and belief: public announcements, unsuccessful updates, knowability, growth of common knowledge, and resolving distributed knowledge.

1. Knowledge

1.1 Knowledge and Ignorance

Let p stand for the proposition "Paneer is rose water." This is true. Amir, however, is uncertain whether it is true (maybe paneer is soft cheese?). If it is true, he considers it possible that it is true and that it is false. If false, he considers it possible that it is false and that it is true. "Considers possible" induces a binary *indistinguishability* relation between states of the world that is therefore obviously reflexive and symmetric and that is also transitive (see below for exceptions involving vagueness). It is an equivalence relation. Figure 5.1.1b depicts the situation.

If Amir and his friend Bala are both ignorant whether p, and they know this about each other (and so on), then a simple model for that situation again consists of two states for the different truth values of p and is such that both Amir and Bala cannot distinguish them: figure 5.1.1b.

In that situation, both know what the other is uncertain about. It is easy to conceive of scenarios where this is not the case, for example, when Bala considers it possible that Amir knows p, that he knows $\neg p$, and that he is ignorant about p: figure 5.1.1d. This is known as *higher-order uncertainty*.

Models for higher-order uncertainty can be arbitrarily large. Suppose that Amir and Bala both know a number and that the other's number is one more or one less: the *consecutive numbers riddle* (Littlewood, 1953). An infinite chain of possibilities results: figure 5.1.1f (p now stands for "the number is even").

1.2 The Modal Logic of Knowledge

The language of epistemic logic is that of propositional logic expanded with inductive clauses $K_a p$ for "Agent a knows proposition p." The dual modality $\langle K \rangle_a p$ can be defined by abbreviation as $\neg K_a \neg p$ and stands for "Agent a considers p possible." The language of epistemic logic is interpreted on pointed Kripke models. A *Kripke model* consists of a domain of abstract objects called "states" or "worlds," a set of binary relations R_a (one for each agent a) on this domain, and valuations V_p for atomic propositions p, specifying for each state whether atom p is true or false there. (In this simplified presentation, the names of states are their value of atom p, and we also use p for other than atomic propositions.) A *pointed* Kripke model has a designated "actual" state. Formula $K_a p$ (for any proposition p) is true in state s of a model iff p is true in all states t of the model that are R_a-accessible from s. On models interpreting knowledge, the accessibility relations are assumed to be equivalence relations. We will list and discuss the principles of knowledge below in view of how rational they are. The logic of knowledge is called **S5**.

1.3 Belief

A difference between knowledge and belief is that knowledge is supposed to be correct, whereas belief may be mistaken. But a belief's being true does not make it knowledge. The difference between knowledge and belief is a notoriously difficult question. It also involves evidence and justification, as well as Gettier examples (Gettier, 1963). It will not be addressed here.

The correctness of knowledge is formalized in the axiom $K_a p \rightarrow p$ (**T**), which characterizes the reflexivity of R_a. Reflexivity guarantees nonemptyness of the relation. Removing this axiom from the logic allows an accessibility relation R_a to be empty, in which case beliefs are

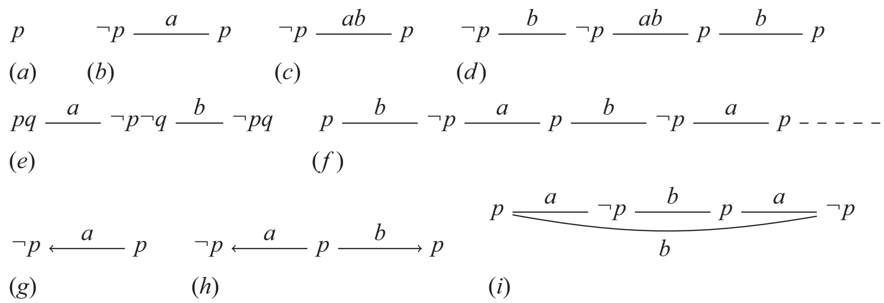

Figure 5.1.1
Models for knowledge and belief. States are named with the valuations of atoms. Visual conventions for relations R_a are transitivity, symmetry *except for a-labeled edges with arrows at end* and reflexivity *except for states from which lead a-labeled edges with arrows at end*. These conventions are sufficient to display **S5** models and **KD45** models unambiguously for a given number of agents.

inconsistent. A more austere way of preserving consistency but removing correctness is merely to require that *you consider possible what you believe*. In the logical language, instead of an inductive clause for knowledge, we now have an inductive clause $B_a p$ for "Agent a believes p." Formally, the above requirement is then $B_a p \rightarrow \langle B \rangle_a p$ (**D**). The logic of belief is called **KD45**.

Typical examples are when p is true but agent a believes it to be false (figure 5.1.1g), and where p is true, a mistakenly believes that a and b believe it to be false, whereas b believes that a and b correctly believe it to be true (figure 5.1.1h).

Knowledge and belief in these modal **S5** and **KD45** incarnations have become standard since Hintikka (1962) but have roots going back to medieval and classical times.

1.4 Rationality and Knowledge

We now present the principles of knowledge by way of questioning their rationality. Except for the correctness of knowledge, they are also the principles of belief. There are no principles for multiagent interaction. We therefore omit the index a in K_a.

T $Kp \rightarrow p$.

What you know is true (factivity or correctness of knowledge).

If this fails, why not call the notion "belief"? Some approaches still call it knowledge: if the set of counterexamples, which may or may not include what is actually the case, is sufficiently *small* (according to some formal notion of small and large, such as measure theory; Ben-David & Ben-Eliyahu, 1994), in other words, if the counterexamples are insignificant. Or it is the *fallible knowledge* of Özgün (2017, chapter 5), which ties the notion of knowledge only to justification and evidence.

4 $Kp \rightarrow KKp$.

You know what you know (positive introspection).

You have to make pairs of matching color out of the heap of socks coming out of the washing. Two socks may have the same color, and the next pair, and the next pair . . . But you are likely to end up with a last pair of socks of different color. The "same-color" relation is not transitive. Knowledge of socks' colors is *vague* (van Deemter, 2010). A different kind of attack emphasizes that rationality is bounded. Birds know how to fly, but they do not *know* that they know how to fly. Higher-order cognition in animals is (also) addressed in Verbrugge (2009).

5 $\langle K \rangle p \rightarrow K \langle K \rangle p$.

You know what you do not know (negative introspection).

Negative introspection is most often attacked: I am unaware of much that I do not know. So I do not know what I do not know. Negative introspection is justifiable if the relevant facts that agents may be uncertain about are known: the closed-world assumption. The **S5** notion of knowledge is considered suitable for describing the cognitive architecture of artificial agents. Well known is the *interpreted system* (Fagin, Halpern, Moses, & Vardi, 1995) consisting of a set of global states composed in turn of the local states of individual agents/processors. An agent cannot distinguish global states with the same local state value: **S5** knowledge.

Various alternative modal logics of knowledge weaken negative introspection. The logic **S4** is the system without **5**. Building a ladder from **S4** to **S5** are the systems **S4.2**, **S4.3**, and **S4.4**, obtained by adding the principles below to **S4**, for which (including **5**) we also give corresponding relational properties. The **.4** property is known as "remote symmetry."

.2 $\langle K\rangle Kp \rightarrow K\langle K\rangle p$ $\forall xyz(Rxy \wedge Rxz \rightarrow \exists w(Ryw \wedge Rzw))$

.3 $K(p \rightarrow \langle K\rangle q) \vee K(q \rightarrow \langle K\rangle p)$ $\forall xyz\ (Rxy \wedge Rxz \rightarrow Ryx \vee Rzy)$

.4 $p \rightarrow (\langle K\rangle p \rightarrow K\langle K\rangle p)$ $\forall xyz\ (Rxy \wedge Ryz \rightarrow Rzy)$

5 $\langle K\rangle p \rightarrow K\langle K\rangle p$ $\forall xyz\ (Rxy \wedge Rxz \rightarrow Ryz)$

Negative introspection has also been criticized because in combination with factivity, it implies $p \rightarrow K\langle K\rangle p$ (this characterizes symmetry; Hintikka, 1962; Stalnaker, 2005). This direct connection between objective truth and subjective truth is considered unacceptable: the agent should not be "allowed" to create knowledge out of thin air.

In logics combining knowledge and awareness (Fagin & Halpern, 1988; such logics are often known as logics for bounded rationality), the failure of negative introspection may even *define* unawareness. You are *unaware of p* if you don't know *p* and you also don't *know* that you don't know *p*: $\neg Kp \wedge \neg K \neg Kp$. This is known as the Modica–Rustichini definition of unawareness (Modica & Rustichini, 1994).

$K(p \rightarrow q) \rightarrow (Kp \rightarrow Kq)$
Deductive closure under the scope of knowledge.

This principle and the next principle of "necessitation" are together often known as those embodying logical *omniscience* and thus *full rationality*. Logics satisfying weaker notions, of which we will give two examples, are associated with *bounded rationality*.

In neighborhood structures (Pacuit, 2017), a formula is known in world *w* if there is a *neighborhood* of *w* (a designated subset typically containing *w*) that is the extension of the formula (i.e., the set of worlds where the formula holds). Let $p \rightarrow q$ and *p* have such extensions. Their intersection is contained in the extension of *q*. If the latter is not a neighborhood, $p \rightarrow q$ and *p* are both known, but not *q*.

In certain logics of knowledge and awareness (Fagin & Halpern, 1988; Halpern & Pucella, 2011), you know *p* if *p* is a member of the (fixed) set of "knowable" formulas in every accessible state; $p \rightarrow q$ and *p* may always be in that set, but not *q*. In that case, again, $p \rightarrow q$ and *p* are both known, but not *q*.

p implies K*p*.
Necessitation of knowledge (i.e., if *p* is *valid*, then K*p* is valid).

In neighborhood semantics, instead of necessitation, the weaker rule of *monotony* may hold: "$p \rightarrow q$ implies $Kp \rightarrow Kq$." In logics of knowledge and awareness, necessitation may not hold either: $p \vee \neg p$ is valid, but $K(p \vee \neg p)$ is invalid if the agent is unaware of *p*.

1.5 Combining Knowledge and Belief

Principles that are typically part of systems combining knowledge and belief, and where the main difference between the two is that beliefs, but not knowledge, may be incorrect, are $Kp \rightarrow Bp$ (knowledge implies belief) and $Bp \rightarrow BKp$ (you believe to know what you believe).

The logics **KD45** and **S5** do not combine very well. If apart from $Bp \rightarrow BKp$ one also has $Bp \rightarrow KBp$, then belief trivializes to knowledge: $Bp \leftrightarrow Kp$ (van der Hoek, 1993).

But **KD45** and **S4.4** combine well. In **S4.4**, belief is definable in terms of knowledge, namely, as $Bp \leftrightarrow \langle K\rangle Kp$, as was observed by Lenzen (1978). In this combination, knowledge is also defined by belief, namely, as $Kp \leftrightarrow p \wedge Bp$. Halpern, Samet, and Segev (2009) call this *explicit* definability. They also show that in the **S5–KD45** combination, knowledge is not explicitly definable in terms of belief, but only implicitly.

1.6 Knowing Whether

You know *whether p*, formally Kw *p*, if you know that *p* or you know that $\neg p$—more generally, if all accessible states have the same value for *p*. The principles of "knowing whether" are different from the principles of knowledge, for example, $Kw\, p \leftrightarrow Kw\, \neg p$. For the expressivity of such epistemic logics, this does not matter, as "knowing that" and "knowing whether" are interdefinable: $Kp \leftrightarrow p \wedge Kw\, p$, and $Kw\, p \leftrightarrow Kp \vee K\neg p$. In modal logics without the principle $Kp \rightarrow p$, such as **KD45**, knowledge may no longer be definable in terms of knowing whether. Indeed, the models of figure 5.1.1a and figure 5.1.1g are indistinguishable in the *p* states; in the former, Amir believes that *p*, and in the latter, Amir believes that $\neg p$; therefore, in either case, Amir "believes whether *p*" (in natural language this is infelicitous). Knowing-whether logics are founded in logics of contingency (Montgomery & Routley, 1966). Related notions such as knowing value, knowing how, and so on are already discussed in Hintikka (1962). A recent survey is Wang (2018).

2. Common Knowledge and Belief, and Distributed Knowledge

When celebrating Christmas with your young cousin, you should not spoil the fun by telling her that Santa Claus is actually dressed-up Uncle Ben. In fact, she knows this. So you both know this. But it is not common knowledge. Only when she yells, "But it is Uncle Ben," does it become common knowledge. Lewis (1969) gave the example of knowing that you have to drive on the left-hand side of the road: you know it, the car approaching you on this narrow road knows it, but you would feel very unsafe if it were all you know. However, it is common knowledge that you have to drive on the left-hand side. But how can you have obtained common knowledge with the total stranger in the other car? This is absurd. This sort of common knowledge is a *convention*: you are supposed to know. If you cause an accident, you cannot claim ignorance.

2.1 The Modal Logic of Common Knowledge

To the logic of knowledge, we add operators Cp for "The agents have common knowledge of p." The dual operator is $\langle \mathrm{C} \rangle p$. The formula $\langle \mathrm{C} \rangle p$ is true in a state of a Kripke model iff p is true at the end of any finite path of accessibility/indistinguishability links involving any agent. This accessibility is described by the arbitrary iteration of the union of all accessibility relations: let R_{C} be the transitive closure of the union of all relations R_a, then Cp is true in state s if p is true in all states t that are R_{C}-accessible. Similarly, we add to the logic of belief modalities CB p for common belief of p.

Common knowledge has the properties of knowledge, but common belief does not have the properties of belief. It is easy to see that $\mathrm{C}p \rightarrow p$, $\mathrm{C}p \rightarrow \mathrm{CC}p$, and $\langle \mathrm{C} \rangle p \rightarrow \mathrm{C}\langle \mathrm{C} \rangle p$. But for common belief, negative introspection does not hold. Figure 5.1.1g demonstrates that $\neg \mathrm{CB}\, p \rightarrow \mathrm{CB}\, \neg \mathrm{CB}\, p$ is false.

For example, in figure 5.1.1f, it is common knowledge that Bala does not know whether p. The logic of common knowledge is indeed more expressive than the logic of knowledge. This is proved by comparing the set of all finite-length chains (with leftmost actual state) consisting of iterated paired *b–a* links with that same set plus the infinite chain as in figure 5.1.1f. The formula $\langle \mathrm{C} \rangle \langle \mathrm{K} \rangle_b p$ can distinguish the former from the latter set, but no formula in the logic of knowledge can distinguish these sets, as it cannot look into the chains beyond the finite modal depth of that formula.

The history of common knowledge comes with the names Friedell (1969) and Lewis (1969). Van Ditmarsch, van Eijck, and Verbrugge (2009) give a historical survey in the form of a Socratic dialogue.

2.2 Relativized Common Knowledge

Temporal logic with the binary "until" modality is more expressive than temporal logic with the unary "going to be" (always in the future) modality. Similarly, epistemic logic with the binary so-called *relativized common knowledge* modality is more expressive than that with the unary common knowledge modality (Kooi & van Benthem, 2004). In a given state s, the agents have common knowledge of p relative to q (notation $\mathrm{C}_q p$) iff p is true in all states t that are reachable by finite R_{C}-paths satisfying q. Figure 5.1.1e illustrates the difference between relativization and a "mere" antecedent in an implication: in the pq state, $\mathrm{C}_q p$ is true, but $\mathrm{C}(q \rightarrow p)$ is false.

2.3 Distributed Knowledge

Suppose that Amir knows that p and that Bala knows that $p \rightarrow q$. Then together they know that q: if they were to communicate, they could make q common knowledge. Seen from a static perspective: if they were the same agent, then that agent could deduce q from p and $p \rightarrow q$. That agent only considers a world possible if Amir and Bala both consider it possible. We write Dp for "The agents have distributed knowledge of p"; Dp is true if p is true in all states accessible by the intersection of all accessibility relations R_a. The notion of distributed belief goes back to Hayek (1945) and Halpern and Moses (1990).

The logic with distributed knowledge is more expressive than the logic without: figure 5.1.1c and figure 5.1.1i cannot be distinguished in the logic with merely K_a and K_b (the models are bisimilar), but in the former, Dp is false in the p state, whereas in the latter, it is true in the p states.

A survey of logics of knowledge, belief, common knowledge, and distributed knowledge is van Ditmarsch, Halpern, van der Hoek, and Kooi (2015, chapter 1). Goldblatt (2003) presents modal logics in their historical context.

3. Change of Knowledge and Belief

The formation of knowledge and belief is crucial to give meaning to the concepts of knowledge and belief, and even more so for common knowledge and belief (see also chapter 2.6 by van Benthem, Liu, & Smets, this handbook).

3.1 Unsuccessful Update

Moore (1942) observed the incoherence of the statement "I went to the movies last Tuesday and I don't believe that." This form, $p \wedge \neg \mathrm{K}p$ (belief can be treated similarly), became later known as the *Moore sentence*. Assuming that you believe what you say, we get $\mathrm{K}(p \wedge \neg \mathrm{K}p)$. Using the discussed properties of knowledge, this is equivalent to $\mathrm{K}p \wedge \mathrm{K}\neg \mathrm{K}p$ and thus to $\mathrm{K}p \wedge \neg \mathrm{K}p$, which is contradictory. Now consider Bala informing Amir that Amir does not know that p is true; by conversational implicature, this is formalized as $p \wedge \neg \mathrm{K}_a p$, which is not contradictory. (Hintikka [1962] devotes a whole chapter to such phenomena.) However, Bala cannot truthfully inform Amir *twice* in this way. After informing Amir of $p \wedge \neg \mathrm{K}_a p$, true at the moment of utterance, Amir has learnt that p, that is, $\mathrm{K}_a p$, so $p \wedge \neg \mathrm{K}_a p$ is now false.

We can semantically verify this in *public announcement logic* (Plaza, 1989). A public announcement is a novel piece of information that is considered reliable (it is assumed true) and is simultaneously observed/incorporated by all agents. The result of the public announcement of p is the restriction of the Kripke model to the subdomain satisfying p. Now consider figure 5.1.1b: in the p state, $p \wedge \neg K_a p$ is true. In the $\neg p$ state, it is not. This announcement therefore induces the transition of figure 5.1.1b to figure 5.1.1a, the singleton-p state, in which obviously $K_a p$ is true. The dynamic interpretation of the Moore sentence as an *unsuccessful update* is due to Gerbrandy (2007; see also van Ditmarsch & Kooi, 2006). Furthermore, the Moore sentence demonstrates that announcing a proposition need not make it common knowledge.

3.2 Growth of Common Knowledge

The other problem with common knowledge is that it does not exist. This is very clear in the case of complete strangers in approaching cars having "common knowledge" of the traffic rules but less clear in other cases. In order to obtain common knowledge, special observation conditions need to apply. The agents must simultaneously observe the event and must also have joint awareness of their observing this event. Such conditions are fulfilled in a group of people facing each other wherein one of them is addressing the group: the observation is of a sound under conditions of mutual eye contact. This explains the term "public announcement." Loopholes can always be found: you may not have been paying attention, you could not hear what was being said, and so on. But you would then later need an excuse if you did not act upon the information: you were *supposed* to have heard.

The directness of speech in company should be contrasted with separating the sending and receiving of messages, after which growth of common knowledge is impossible (Halpern & Moses, 1990).

3.3 Knowability

In figure 5.1.1b, we can make two announcements, assuming the truth of p: the trivial announcement or the announcement of p. So, there is an announcement after which $K_a p$, and there is an announcement after which (still) $\neg K_a p$. Consider a logic with *quantification* over announcements: let $\blacklozenge Kp$ mean that there is an announcement after which p is known. Inspired by Fine (1970), such a logic was proposed in Balbiani et al. (2008). We can also read $\blacklozenge Kp$ as "p is knowable," such that $p \rightarrow \blacklozenge Kp$ means that all truths are knowable. This is not valid. We have already seen that $p \wedge \neg Kp$ is not knowable. Knowability (Fitch, 1963) is indeed interpreted in such dynamic terms in van Benthem (2004) and Balbiani et al. (2008). A surprising result is that $\blacklozenge(Kp \vee K\neg p)$ (i.e., $\blacklozenge Kwp$) is valid (van Ditmarsch, van der Hoek, & Iliev, 2012).

3.4 Changing Knowledge and Belief Simultaneously

We can also change **S5** knowledge and **KD45** belief simultaneously. Plausibility models are Kripke models with equivalence relations R_a to encode knowledge, but where the states in each epistemic equivalence class are ordered according to plausibility. The agent believes p iff p is true in all most plausible states (the order has to be a well-preorder, which guarantees that most plausible states do exist). This induces a doxastic relation R'_a to interpret belief, contained in R_a. A public announcement restricts the epistemic relations to the novel domain. It is then determined anew what the most plausible states are, and this determines the new doxastic relation.

Figure 5.1.1b and figure 5.1.1g jointly represent that Amir does not know whether p but believes that $\neg p$. After public announcement of p, the single remaining state is p: figure 5.1.1a. This p state has now become the most plausible state. So figure 5.1.1a also depicts the doxastic relation. Amir now knows and believes that p. He has revised his knowledge and his belief (see Baltag & Smets, 2008; van Benthem, 2007; van Ditmarsch, 2005; a precursor is Spohn, 1988. These matters are also discussed in chapter 5.2 by Rott and chapter 5.3 by Kern-Isberner, Skovgaard-Olsen, & Spohn, both in this handbook).

3.5 Resolving Distributed Knowledge

Not all distributed knowledge can become common knowledge. Let us give two very different examples. In the p state of figure 5.1.1b, it is distributed knowledge that p is true, and agent a does not know this: we have that $K_b(p \wedge \neg K_a p)$, and thus $D(p \wedge \neg K_a p)$. But this cannot become common knowledge: if b informs a of p, p is common knowledge, not $p \wedge \neg K_a p$. Second, consider figure 5.1.1i: as we have already seen, p is distributed knowledge between a and b. But communication between a and b cannot make this common knowledge, because figure 5.1.1i and figure 5.1.1b encode the same information in epistemic logic (they are bisimilar). Making distributed knowledge common knowledge is called *resolving* distributed knowledge by Ågotnes and Wáng (2017).

Acknowledgments

Hans van Ditmarsch is also affiliated with IMSc (Chennai, India) and IRISA (Rennes, France). Reviewers are kindly thanked for their comments.

References

Ågotnes, T., & Wáng, Y. N. (2017). Resolving distributed knowledge. *Artificial Intelligence*, *252*, 1–21.

Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., & de Lima, T. (2008). "Knowable" as "known after an announcement." *Review of Symbolic Logic*, *1*(3), 305–334.

Baltag, A., & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, & M. Wooldridge (Eds.), *Logic and the foundations of game and decision theory (LOFT 7)* (Texts in Logic and Games, Vol. 3, pp. 13–60). Amsterdam, Netherlands: Amsterdam University Press.

Ben-David, S., & Ben-Eliyahu, R. (1994). A modal logic for subjective default reasoning. In *Proceedings, Ninth Annual IEEE Symposium on Logic in Computer Science* (pp. 477–486). IEEE Computer Society Press.

Fagin, R., & Halpern, J. Y. (1988). Belief, awareness, and limited reasoning. *Artificial Intelligence*, *34*(1), 39–76.

Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about knowledge*. Cambridge, MA: MIT Press.

Fine, K. (1970). Propositional quantifiers in modal logic. *Theoria*, *36*(3), 336–346.

Fitch, F. B. (1963). A logical analysis of some value concepts. *Journal of Symbolic Logic*, *28*(2), 135–142.

Friedell, M. F. (1969). On the structure of shared awareness. *Behavioral Science*, *14*, 28–39.

Gerbrandy, J. (2007). The surprise examination in dynamic epistemic logic. *Synthese*, *155*(1), 21–33.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, *23*, 121–123.

Goldblatt, R. (2003). Mathematical modal logic: A view of its evolution. *Journal of Applied Logic*, *1*(5–6), 309–392.

Halpern, J. Y., & Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, *37*(3), 549–587.

Halpern, J. Y., & Pucella, R. (2011). Dealing with logical omniscience: Expressiveness and pragmatics. *Artificial Intelligence*, *175*(1), 220–235.

Halpern, J. Y., Samet, D., & Segev, E. (2009). Defining knowledge in terms of belief: The modal logic perspective. *Review of Symbolic Logic*, *2*(3), 469–487.

Hayek, F. (1945). The use of knowledge in society. *American Economic Review*, *35*, 519–530.

Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, NY: Cornell University Press.

Kooi, B., & van Benthem, J. (2004). Reduction axioms for epistemic actions. In R. A. Schmidt, I. Pratt-Hartmann, M. Reynolds, & H. Wansing (Eds.), *AiML-2004: Advances in modal logic* (Technical Report Series No. UMCS-04-9-1, pp. 197–211). Manchester, England: University of Manchester.

Lenzen, W. (1978). Recent work in epistemic logic. *Acta Philosophica Fennica*, *30*(2), 1–219.

Lewis, D. K. (1969). *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.

Littlewood, J. E. (1953). *A mathematician's miscellany*. London, England: Methuen.

Modica, S., & Rustichini, A. (1994). Awareness and partitional information structures. *Theory and Decision*, *37*, 107–124.

Montgomery, H., & Routley, R. (1966). Contingency and noncontingency bases for normal modal logics. *Logique et Analyse*, *9*, 318–328.

Moore, G. E. (1942). A reply to my critics. In P. Schilpp (Ed.), *The philosophy of G. E. Moore* (pp. 535–677). Evanston, IL: Northwestern University.

Özgün, A. (2017). *Evidence in epistemic logic: A topological perspective* (Doctoral thesis, University of Lorraine, Nancy, France).

Pacuit, E. (2017). *Neighborhood semantics for modal logic*. Cham, Switzerland: Springer.

Plaza, J. (1989). Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, & Z. W. Ras (Eds.), *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems (ISMIS 1989): Poster session program* (pp. 201–216). Charlotte, NC: Oak Ridge National Laboratory.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper & B. Skyrms (Eds.), *Causation in decision, belief change, and statistics* (Vol. II, pp. 105–134). Dordrecht, Netherlands: Kluwer.

Stalnaker, R. (2005). On logics of knowledge and belief. *Philosophical Studies*, *128*(1), 169–199.

van Benthem, J. (2004). What one may come to know. *Analysis*, *64*(2), 95–105.

van Benthem, J. (2007). Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics*, *17*(2), 129–155.

van Deemter, K. (2010). *Not exactly: In praise of vagueness*. Oxford, England: Oxford University Press.

van der Hoek, W. (1993). Systems for knowledge and beliefs. *Journal of Logic and Computation*, *3*(2), 173–195.

van Ditmarsch, H. (2005). Prolegomena to dynamic logic for belief revision. *Synthese*, *147*, 229–275.

van Ditmarsch, H., Halpern, J. Y., van der Hoek, W., & Kooi, B. (Eds.). (2015). *Handbook of epistemic logic*. London, England: College Publications.

van Ditmarsch, H., & Kooi, B. (2006). The secret of my success. *Synthese*, 151, 201–232.

van Ditmarsch, H., van der Hoek, W., & Iliev, P. (2012). Everything is knowable—how to get to know *whether* a proposition is true. *Theoria*, *78*(2), 93–114.

van Ditmarsch, H., van Eijck, J., & Verbrugge, R. (2009). Common knowledge and common belief. In J. van Eijck & R. Verbrugge (Eds.), *Discourses on social software* (pp. 99–122). Amsterdam, Netherlands: Amsterdam University Press.

Verbrugge, R. (2009). Logic and social cognition: The facts matter, and so do computational models. *Journal of Philosophical Logic*, *38*(6), 649–680.

Wang, Y. (2018). Beyond knowing that: A new generation of epistemic logics. In H. van Ditmarsch & G. Sandu (Eds.), *Jaakko Hintikka on knowledge and game-theoretical semantics* (Outstanding Contributions to Logic, Vol. 12, pp. 499–533). Cham, Switzerland: Springer.

This is a section of doi:10.7551/mitpress/11252.001.0001

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn

Citation:

The Handbook of Rationality
Edited by: Markus Knauff, Wolfgang Spohn
DOI: 10.7551/mitpress/11252.001.0001
ISBN (electronic): 9780262366175
Publisher: The MIT Press
Published: 2021

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Knauff, Markus, editor. | Spohn, Wolfgang, editor.
Title: The handbook of rationality / edited by Markus Knauff and Wolfgang Spohn.
Description: Cambridge : The MIT Press, 2021. | Includes bibliographical references and index.
Identifiers: LCCN 2020048455 | ISBN 9780262045070 (hardcover)
Subjects: LCSH: Reasoning (Psychology) | Reason. | Cognitive psychology. | Logic. | Philosophy of mind.
Classification: LCC BF442 .H36 2021 | DDC 153.4/3—dc23
LC record available at https://lccn.loc.gov/2020048455