

This is a section of [doi:10.7551/mitpress/11252.001.0001](https://doi.org/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

11.1 Deontic Logic

John Horty and Olivier Roy

Summary

Deontic logic is a general area encompassing many specific logical approaches united by a common concern with normative concepts—concepts, that is, involving norms, such as moral, legal, or rational norms. Based on normative information of this kind, deontic logics, most typically, provide judgments or conclusions about what we ought to do and what we are permitted to do, or about what ought to be the case. In this chapter, we sketch the basic ideas underlying the standard system of deontic logic. Next, just to highlight the interdisciplinary reach of deontic logic, we mention some connections with game theory that we believe may be of particular interest to readers of this volume. Finally, we survey some important variations on standard deontic logic.

1. What Is Deontic Logic?

Deontic logic is a general area encompassing many specific logical approaches united by a common concern with normative concepts—concepts, that is, involving norms, such as moral, legal, or rational norms; the norms of etiquette or aesthetics; or the rules of a game. Based on normative information of this kind, what deontic logics provide, most typically, is judgments or conclusions about what we ought to do and what we are permitted to do, or about what ought to be the case.

Like many areas of logic, deontic logic was studied in a fragmentary way by medieval theorists, but it was not until the 20th century, with the introduction of mathematical techniques into logic, that the subject developed into a systematic discipline. This development can be divided into three stages. In the first, beginning with the work of Mally (1926) early in the 20th century, a number of philosophers and logicians explored ways of extending the axiomatic approach to logic, in particular that of Whitehead and Russell's *Principia Mathematica* (1910–1913), to deontic concepts as well. Although these early axiomatizations were soon shown to be flawed, they

opened up the area for investigation and inspired later researchers. The second stage began near the midpoint of the 20th century with von Wright's (1951) interpretation of deontic logic as a species of modal logic, allowing for the subject to be studied using semantic techniques, and leading, with only slight modifications, to the system of modal logic now known as "standard deontic logic." Although this system stimulated some interesting technical developments, it also served as the primary focus of a philosophical debate concerning the relation between deontic principles and moral concepts that became, over time, increasingly arid and scholastic.

The third stage in the development of deontic logic, which began in the final decade of the 20th century and continues to this day, has been characterized by increasing interdisciplinary connection with many other fields, including artificial intelligence, multiagent systems and computer science more generally, legal theory, linguistics, political science, organizational theory, and economics. These interdisciplinary connections have revitalized the field, leading not only to new areas of application for ideas from deontic logic but also to the incorporation of new techniques into the study of deontic logic itself.

We cannot attempt here either to trace the historical development of deontic logic in any detail or to survey its current range of interdisciplinary connections; an authoritative historical treatment is presented by Hilpinen and McNamara (2014). Our discussion is organized as follows: first, we sketch the basic ideas underlying the standard system of deontic logic. Next, just to highlight the interdisciplinary reach of deontic logic, we mention some connections with game theory that we believe may be of particular interest to readers of this volume. Finally, we survey some important variations on standard deontic logic.

2. Standard Deontic Logic

Standard deontic logic is a species of modal logic, a form of logic in which the truth or falsity of a formula in a

particular state of affairs can depend on the truth value of different formulas in different states of affairs. These states of affairs are often referred to, somewhat poetically, as “possible worlds,” although the logic itself does not involve any metaphysical commitments, and the possible worlds can be provided with more concrete interpretations, as we shall see.

Let us suppose, then, that W is a set of possible worlds and that v is a valuation function mapping each atomic sentence into the set of possible worlds where it is true, so that, where p is an atomic sentence and w is a possible world, w is in $v(p)$ if p is true in w , and w is not in $v(p)$ if p is false in w . In addition to W and v , the structures against which formulas of standard deontic logic are interpreted contain a further, distinctive component: a function f mapping each possible world w into a non-empty set $f(w)$ of worlds that can be thought of as ideal from the standpoint of w . The basic idea underlying standard deontic logic is that what *ought* to be the case in the world w can be identified with what *is* the case in those worlds that are ideal from the standpoint of w .

To develop the theory formally, we first introduce models of the form $M = \langle W, f, v \rangle$, with W , f , and v defined as above. Next, letting $M, w \models A$ indicate that the formula A is true at the world w from the set W of worlds belonging to the model M , and where v is the valuation function from this model, we define truth conditions for atomic sentences through the clause

$$M, w \models p \text{ just in case } w \in v(p),$$

according to which the atomic sentence p is true at the world w just in case the valuation function says it is. Once the notion of truth at a world is settled for atomic sentences, it can be lifted to complex sentences formed from the connectives \neg , \wedge , \vee , and \rightarrow —representing negation, conjunction, disjunction, and implication, respectively—through the clauses

$$\begin{aligned} M, w \models \neg A & \text{ just in case it is not the case that } M, w \models A, \\ M, w \models A \wedge B & \text{ just in case } M, w \models A \text{ and } M, w \models B, \\ M, w \models A \vee B & \text{ just in case } M, w \models A \text{ or } M, w \models B, \\ M, w \models A \rightarrow B & \text{ just in case either not } M, w \models A, \text{ or} \\ & M, w \models B. \end{aligned}$$

What these further clauses tell us is simply that the truth value assigned to a complex sentence at a possible world respects the meaning of the connective through which that complex sentence is formed: according to the second clause, for example, a conjunction is true at a world just in case each of its conjuncts is; according to the fourth, an implication is true just in case the consequent of that implication is true whenever its antecedent is.

Finally, we turn to the characteristic deontic connective O , allowing for the construction of sentences of the form OA , to be read “It ought to be the case that A .” The clause governing this connective is

$$\begin{aligned} M, w \models OA & \text{ just in case} \\ M, w' \models A & \text{ for all } w' \text{ belonging to } f(w), \end{aligned}$$

where of course f is the function from M mapping each world w into the set of worlds that can be considered as ideal from the standpoint of w .

The clause governing the truth value assigned to a deontic statement captures in a formal way the idea that OA holds at a world w just in case A holds at all the worlds w' that are ideal from the standpoint of w . For a concrete illustration, let A represent the statement that there is no war, and suppose we can agree that, from the standpoint of our actual world, any ideal world would have to be one in which there is no war, so that A holds in each of these ideal worlds. Then, what our logic tells us is that, because A holds in each of these ideal worlds, OA is true in the actual world—it ought to be that there is no war.

One natural question to ask at this point is why the worlds considered to be ideal should be relativized to an initial world: why isn't there just a single set of ideal worlds? The answer is that, in general, ought-statements can be contingent, varying from world to world. Consider, for example, the proposition that Jo promises to meet her friend Jack for dinner. Presumably this is a contingent proposition, holding at some worlds but not at others. So imagine that w is a world at which Jo makes this promise and w' is a world at which she does not, and let A stand for the proposition that Jo, in fact, meets Jack for dinner—that she keeps her promise. If we assume that, ideally, promises are to be kept, then it follows that all the worlds ideal from the standpoint of w are worlds in which A holds, so that OA holds at w . However, since there is no particular reason to suppose that A holds at all the worlds ideal from the standpoint of w' , where Jo made no promise, it follows that OA fails at that possible world.

Once we have defined truth at a possible world in a model, we can, following the standard pattern in modal logic, define logical implication by stipulating that A logically implies B if B is true at every world from every model at which A is true. To illustrate, we can see that OA and $O(A \rightarrow B)$ implies OB , as follows. Suppose OA and $O(A \rightarrow B)$ are true at w . Then both A and $A \rightarrow B$ must be true at each world from $f(w)$. But it follows from ordinary logic that B is true whenever A and $A \rightarrow B$ are both true, so that B must likewise be true at each world from $f(w)$. From this, we can conclude that OB is true at w .

3. An Interlude: Connections with Game Theory

Although the talk of possible worlds in deontic logic, and in modal logic more generally, can seem hopelessly abstract, it is important to realize that this language can be provided with concrete interpretations in a number of application domains. We cannot describe these various application domains here, of course, but just for the sake of illustration, we briefly sketch how the machinery of deontic logic can be applied in the domain of game theory, where the norms at work are not moral or legal norms but norms of rationality. Our discussion proceeds by way of example; see chapter 9.1 by Albert and Kliemt (this handbook) for formal definitions.

Let us thus start with a simple example of a game in so-called normal or strategic form:

Ann \ Bob	L	R
U	1, 1	0, 0
D	0, 1	1, 0

In this game there are two players, Ann and Bob. Ann can choose the upper (*U*) or the lower (*D*) row and Bob either the left (*L*) or the (*R*) column. Each combination of these choices, for instance, (*U,L*) or (*D,R*), corresponds to one possible outcome of the game. Ann and Bob have preferences over these outcomes, which are expressed by the utility numbers in each cell of the matrix, with 1 preferred to 0. The left-hand numbers are Ann's utilities and the right-hand ones Bob's.

What is rational for Ann and for Bob to do in this game? What Ann should rationally do depends on what Bob does. If he plays *L*, then her best response is to play *U*. If, however, he plays *R*, then her best response is *D*. What is the best response for Bob, however, does *not* depend on what Ann does. If she plays *U*, he gets 1 by playing *L* and 0 by playing *R*, and the same holds if Ann plays *D*. Game theorists would say that *R* is strictly dominated by *L*. Ruling *R* out as a rational choice for Bob, we conclude that he ought to play *L*. But then, as we have seen already, if we only consider *L* as a possible action for Bob, the only rational choice for Ann is to play *U*, leaving (*U,L*) as the only rational solution of this game. In other words, Ann ought to play *U*, and Bob *L*. This combination of choices also happens to be a Nash equilibrium, meaning that neither Ann nor Bob have an incentive to unilaterally deviate from it. Bob has no incentive whatsoever to play *R* instead of *L*. Ann, given that Bob plays *L*, should not switch to *D* either.

Let us now look at how deontic logicians would analyze this situation. First, we need to define the set *W* of

possible worlds. In our example, we can take *W* to consist of the four possible outcomes of the game, that is, $W = \{(U,L), (U,R), (D,L), (D,R)\}$. Second, we need to define the valuation function *v*. The atomic sentences that we are interested in describe the strategies that Ann and Bob can choose. Let us use *u*, *d*, *l*, and *r* for that. The valuation is then defined as the natural one: $v(u) = \{(U,L), (U,R)\}$, $v(l) = \{(U,L), (D,L)\}$, and so on. The final component that we need to define is the function *f* mapping each possible world *w* onto its set of ideal worlds. Here, again, there is simple, although of course not a unique, way to define this function *f*: we can assign the set of rational solutions, $\{(U,L)\}$, as the set of ideal worlds for all the four possible worlds in *W*. That is, for all $w \in W$, $f(w) = \{(U,L)\}$.

With this in hand, we can see what kind of deontic statements are true in the model that we have just constructed. Recall that *u* is the statement that Ann plays *U*. Now, (*U,L*) is the only ideal world from the standpoint of any world in the model that we constructed above. In that world, *u* is true, so *Ou* is true everywhere in the model that we have constructed. Similarly, we have *Ol* true everywhere, reflecting the fact that in all ideal worlds, Bob plays *L*.

Now, as we already observed, obligations and permissions might be contingent, varying from world to world. This is *not* the case in the model we just built. By setting $f(w) = \{(U,L)\}$ everywhere, we have identified the set of ideal worlds with the unique outcome that results from first ruling out *R* as a rational choice for Bob and then identifying Ann's best response *given that Bob does not play R*. As we have seen, however, Ann's best response would be different if Bob would play *R* after all. To reflect that, we could instead define *f* as follows: $f(U,L) = \{(U,L)\}$, $f(U,R) = \{(D,R)\}$, $f(D,L) = \{(U,L)\}$, $f(D,R) = \{(D,L)\}$. With this modification, *Ou* becomes a contingent formula, true only in the possible worlds where Bob plays his only rational strategy *L* but false in the nonideal case where Bob plays *R*.

4. Variations

The standard deontic logic sketched earlier is a starting point for the study of deontic logic, but only a starting point—this standard theory has been elaborated along various dimensions, sometimes in response to difficulties or perceived difficulties, and a variety of alternative frameworks have been explored. We cannot hope even to mention all of these developments in this brief chapter but will merely list a few of the most important.

First, there are clear connections between deontic logics and logics of time, or temporal logics. Suppose, as in our earlier example, that Jo has promised to meet Jack

for dinner and so ought to have dinner with Jack. But of course, this event must occur at some particular point in the future, relative to the moment of the promise—it will do no good if Jo had dinner with Jack sometime in the past. Further, if Jo's promise is to function as a real constraint on her behavior, we must at least be able to envision alternative futures, or alternative ways in which the world might unfold, in which Jo does *not* meet Jack for dinner. These issues concerning the temporal aspects of ought-statements have been explored in a number of papers beginning with seminal work in the early 1980s by Åqvist and Hoepelman (1981), Thomason (1981), and van Eck (1982); a useful survey can be found in Thomason (1984). Interestingly, similar issues arise in the study of games as well, as illustrated, again, by chapter 9.1 by Albert and Kliemt (this handbook).

Second, deontic logics interact with issues of agency. Consider again our example, leading to the conclusion that Jo ought to dine with Jack. What we naturally mean by this is not simply that the event of Jo dining with Jack should occur but that Jo should bring it about, or see to it, that this event occurs—that this event should occur through Jo's agency. The study of agency in the context of deontic logic began with work by Kanger (1957/1971) and was explored in later work by Pörn (1962/1970) and Lindahl (1977) on the theory of "normative positions," an attempt to survey the various possible normative relations between individuals; this work has now been extended and solidified by Sergot (2014). More recently, Horty (2001) has explored a deontic logic developed within the *stit* analysis of agency introduced by Belnap, Perloff, and Xu (2001). This also has interesting connections with game theory, as illustrated, for instance, in earlier work by Apostel (1960) and van Benthem (1979) as well as more recent work by Tamminga (2013).

Third, although deontic logic typically focuses on the ought-connective introduced above, a connective representing permission can be defined as well, as a dual of the "ought." More precisely, where P represents permission, the statement PA can be defined as $\neg O\neg A$ —the idea is that A is permissible if it is not required that A not be the case. This standard treatment of permission is sensible and applies naturally in many situations, but there are also cases, known as cases of "free choice permission," in which this standard treatment seems to give the wrong results. Suppose the waiter tells you that, as part of your fixed-price meal, you can have either soup or salad. If we now let A represent the proposition that you have soup and B the proposition that you have salad, then it is reasonable to formalize the waiter's statement as $P(A \vee B)$ —it is permitted that you have soup or you

have salad. Of course, we would normally conclude from the waiter's statement that $PA \wedge PB$ —you are permitted to have soup and also permitted to have salad. The trouble is that, on the standard definition of permission as the dual of obligation, $P(A \vee B)$ does not entail $PA \wedge PB$ —we cannot conclude what we want to from what the waiter said. The problem of free-choice permission is the problem of finding another definition of permission that supports the desired inference; a contemporary survey of literature on the problem can be found in Hansson (2014). Again, this problem has also attracted attention from a game-theoretic perspective, as illustrated in work by Dong and Roy (2015) and Anglberger, Gratzl, and Roy (2015).

Fourth, there is the problem of contrary-to-duty oughts. To illustrate, return to our example in which, as the result of a promise, Jo ought to have dinner with Jack. Now suppose the time for dinner comes and goes and, for one reason or another, Jo has failed to meet Jack for dinner. What then? It seems sensible to conclude that Jo should try to make amends to Jack for breaking her promise—perhaps she should call up Jack to apologize. The trouble is that since standard deontic logic is concerned only with ideal situations, in which no obligations are violated, it has nothing to say about what Jo ought to do in the various subideal situations in which she has failed to meet some of her obligations. The problem of contrary-to-duty oughts is the problem of designing a deontic logic to deal with these subideal situations, telling us what we ought to do once we have violated our primary oughts. A very readable introduction to the problem, as well as an interesting proposal for a solution, is presented by Prakken and Sergot (1996).

Fifth, and finally, there is the problem of normative conflicts, sometimes known as "moral dilemmas." Let us say that a situation gives rise to a normative conflict if it presents each of two conflicting propositions as obligatory—if, for example, it supports the truth of both OA and OB , where A and B are inconsistent and so cannot hold jointly. We often seem to face conflicts like this in everyday life, and there are a number of vivid examples in philosophy, literature, and law. Perhaps the best known of these is Sartre's (1946) description of a student during World War II who felt for reasons of patriotism that he ought to leave home to join the Free French but who felt also, for reasons of sympathy and personal devotion, that he ought to stay at home to care for his aged mother.

Sartre presents this student's situation in a compelling way that does make it seem as if he had been confronted with conflicting, and perhaps irreconcilable, moral principles. However, if standard deontic logic is correct,

Sartre is mistaken: the student did not face a moral conflict—no one ever does, because according to standard deontic logic, such a conflict is impossible. This is easy to see. For the two statements OA and OB to be supported at a world w , both A and B must hold at each world that is ideal from the standpoint of w . It is a constraint on standard deontic models that the set $f(w)$ of worlds ideal from the standpoint of w must be nonempty, so there must be some world w' belonging to this set. Both A and B , therefore, must hold at w' , but this is impossible since, by assumption, A and B cannot hold jointly.

This feature of standard deontic logic—that it rules out normative conflicts—has received extensive discussion in the philosophical literature. There is currently no consensus among moral theorists on the question whether an ideal ethical theory could be structured in such a way that moral dilemmas might arise. The issue has been addressed, for example, by Donagan (1984), Foot (1983), Lemmon (1962), Marcus (1980), and Williams (1965). Still, it can seem like an objectionable feature of standard deontic logic that it rules out this possibility. Because the question is open, and the possibility of moral dilemmas is a matter for substantive ethical discussion, it seems to be inappropriate for a position on this issue to be built into the logic of the subject. And even if it does turn out, ultimately, that research in ethics can exclude the possibility of conflicts in a correct moral theory, it is useful all the same to have a deontic logic that is able to tolerate normative conflicts, since they arise in so many other domains.

The first intuitively adequate account of reasoning in the presence of normative conflicts was presented in van Fraassen (1973). The logic presented there was interpreted in Reiter's (1980) default logic by Horty (1994), establishing the first clear connection between existing deontic and nonmonotonic, or defeasible, logics; this interpretation was later developed from a philosophical perspective in Horty (2012). The connection between deontic and defeasible logics has now been explored in some detail. A useful collection of early papers is found in Nute (1997). Much of the most interesting recent work, summarized by Parent and van der Torre (2014), revolves around the framework of input/output logic, rather than older frameworks for defeasible logic. The most authoritative current treatment of normative conflicts in deontic logic is presented by Goble (2014).

5. Conclusion

The field of deontic logic began early in the 20th century with the introduction of a number of formal systems,

many of them flawed, aimed at analyzing normative concepts in philosophy and law. The field was solidified toward the middle of the century with the isolation of the system now known as standard deontic logic, which is sensible and well understood at least from a logical perspective. There followed a period of confusion, and eventual stagnation, centered on the failure of this standard deontic logic to yield what many considered the correct results in a variety of puzzle cases, sometimes known as “deontic paradoxes.” As a result of an infusion of new techniques, primarily from computer science, deontic logic is once again active and vibrant, with applications not only in computer science but also in a number of other fields, such as linguistics, political science, organizational theory, and economics.

Our goal here has been to present the central ideas underlying the system of standard deontic logic and to survey some of its more interesting variations. It is impossible, in a brief chapter such as this, to cover the range of current applications for deontic logic, but just to give the reader a sense of these connections, we have sketched a few of the relations between deontic logic and game theory.

References

- Anglberger, A. J. J., Gratzl, N., & Roy, O. (2015). Obligation, free choice, and the logic of weakest permissions. *Review of Symbolic Logic*, 8(4), 807–827.
- Apostel, L. (1960). Game theory and the interpretation of deontic logic. *Logique et Analyse*, 3(10), 70–90.
- Åqvist, L., & Hoepelman, J. (1981). Some theorems about a “tree” system of deontic tense logic. In R. Hilpinen (Ed.), *New studies in deontic logic: Norms, actions, and the foundations of ethics* (pp. 187–221). Dordrecht, Netherlands: Reidel.
- Belnap, N., Perloff, M., & Xu, M. (2001). *Facing the future: Agents and choices in our indeterministic world*. Oxford, England: Oxford University Press.
- Donagan, A. (1984). Consistency in rationalist moral systems. *Journal of Philosophy*, 81, 291–309.
- Dong, H., & Roy, O. (2015). Three deontic logics for rational agency in games. *Studies in Logic*, 8(4), 7–31.
- Foot, P. (1983). Moral realism and moral dilemmas. *Journal of Philosophy*, 80, 379–398.
- Goble, L. (2014). Prima facie norms, normative conflicts, and dilemmas. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 241–352). London, England: College Publications.

- Hansson, S. O. (2014). The varieties of permission. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 195–240). London, England: College Publications.
- Hilpinen, R., & McNamara, P. (2014). Deontic logic: A historical survey and introduction. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 3–136). London, England: College Publications.
- Horty, J. F. (1994). Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, 23, 35–65.
- Horty, J. F. (2001). *Agency and deontic logic*. Oxford, England: Oxford University Press.
- Horty, J. F. (2012). *Reasons as defaults*. Oxford, England: Oxford University Press.
- Kanger, S. (1971). New foundations for ethical theory. In R. Hilpinen (Ed.), *Deontic logic: Introductory and systematic readings* (pp. 36–58). Dordrecht, Netherlands: Reidel. (Original work privately distributed in 1957)
- Lemmon, E. J. (1962). Moral dilemmas. *Philosophical Review*, 70, 139–158.
- Lindahl, L. (1977). *Position and change: A study in law and logic*. Dordrecht, Netherlands: Reidel.
- Mally, E. (1926). *Grundgesetze des Sollens: Elemente der Logik des Willens* [Basic laws of ought: Elements of the logic of the will]. Graz, Austria: Leuschner & Lubensky.
- Marcus, R. B. (1980). Moral dilemmas and consistency. *Journal of Philosophy*, 77, 121–136.
- Nute, D. (Ed.). (1997). *Defeasible deontic logic*. Dordrecht, Netherlands: Kluwer.
- Parent, X., & van der Torre, L. (2014). Input/output logic. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 409–544). London, England: College Publications.
- Pörn, I. (1970). *The logic of power* (2nd ed.). Oxford, England: Blackwell. (Original work published 1962)
- Prakken, H., & Sergot, M. (1996). Contrary-to-duty obligations. *Studia Logica*, 57, 91–115.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–132.
- Sartre, J.-P. (1946). *L'existentialisme est un humanisme* [Existentialism is a humanism]. Paris, France: Nagel.
- Sergot, M. (2014). Normative positions. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 353–406). London, England: College Publications.
- Tamminga, A. (2013). Deontic logic for strategic games. *Erkenntnis*, 78, 183–200.
- Thomason, R. (1981). Deontic logic as founded on tense logic. In R. Hilpinen (Ed.), *New studies in deontic logic: Norms, actions, and the foundations of ethics* (pp. 165–176). Dordrecht, Netherlands: Reidel.
- Thomason, R. (1984). Combinations of tense and modality. In D. M. Gabbay & F. Guenther (Eds.), *Handbook of philosophical logic: Vol. II. Extensions of classical logic* (pp. 135–165). Dordrecht, Netherlands: Reidel.
- van Benthem, J. (1979). Minimal deontic logics. *Bulletin of the Section of Logic*, 8(1), 36–42.
- van Eck, J. A. (1982). A system of temporally relative modal and deontic predicate logic and its philosophical applications 2. *Logique et Analyse*, 25, 339–381.
- van Fraassen, B. C. (1973). Values and the heart's command. *Journal of Philosophy*, 70, 5–19.
- von Wright, G. H. (1951). Deontic logic. *Mind*, 60, 1–15.
- Whitehead, A. N., & Russell, B. (1910–1913). *Principia mathematica* (3 vols.). Cambridge, England: Cambridge University Press.
- Williams, B. A. O. (1965). Ethical consistency. *Proceedings of the Aristotelian Society*, 39(Suppl.), 103–124.

© 2021 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

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://lcn.loc.gov/2020048455>