

Naturalising Normativity

MARK COLYVAN*

Abstract

In this paper I discuss the problem of providing an account of the normative force of theories of rationality. The theories considered are theories of rational inference, rational belief and rational decision—logic, probability theory and decision theory, respectively. I provide a naturalistic account of the normativity of these theories that is not viciously circular. The account offered does have its limitations though: it delivers a defeasible account of rationality. On this view, theories of rational inference, belief and decision are not *a priori*. Rather, they are *a posteriori* and may change over time. Finally, I compare this approach with another that emerges from the Ramsey-Lewis approach to defining theoretical terms.

1 Introduction

Normativity is built into the very notion of rationality. For example, when we describe some set of beliefs or actions as rational, we mean that these beliefs or actions are sanctioned by the relevant (normative) theories (Bayesian belief theory and decision theory, respectively). We might add logic, the study of rational inferences, to the mix. Indeed, logic, belief theory and decision theory are quite unlike other scientific theories in this regard. These three theories do not purport to simply describe or systematise the way real-world agents reason and act; these theories purport to describe and systematise how idealised rational agents reason and act. Or, to explicitly invoke the normativity here, they purport to prescribe how real-world agents *ought* to reason and act. But this raises an important question about where such theories derive their normative force.¹

The problem is that normative operators such as “ought” are rather odd modal operators—they are not truth functional² and, as David Hume pointed out (1975, p. 469), one cannot derive normative claims from matters

*Philosophy Program, University of Queensland, Brisbane, Queensland, 4072, Australia. *Email:* mcolyvan@uq.edu.au.

¹There is, of course, a place for purely descriptive theories of human reasoning, but these are not the focus of this paper.

²For example, the truth of ‘you ought not affirm the consequent’ cannot be determined purely from the truth of ‘you don’t affirm the consequent’.

of fact. Yet, “one should ensure that one’s beliefs accord with Kolmogorov’s axioms” was not given to us by divine decree. In this paper, I propose one way to understand the origin and force of the normativity of theories of rationality. I do this from a broadly naturalistic point of view. Although from this point of view, normativity is returned to its rightful place in centre stage, the advice delivered about rationality is not the usual advice. I argue that the price we pay for normativity is a kind of fallibilism, and this opens the door on alternatives to classical logic, Kolmogorov probability theory, and standard decision theory. Finally, I compare this approach with another naturalistic approach to normativity proposed by Robert Nola (2003) and Frank Jackson (1998).

2 From the Non-Modal to the Modal

It might seem clear that one cannot get a normative claim from a descriptive claim. For instance, people use *modus ponens* in every day reasoning, but from this it certainly does not follow that people *ought* to use modus ponens. Indeed, it would seem that it is not just normative notions that are problematic; there is a broader class of modal notions that resist derivation from matters of fact. After all, modal notions like necessity seem to suffer similar problems: the modal proposition $\Box P$ does not follow from the non-modal matter of fact P .³ But in the more general modal case, at least, it would be rash to conclude that modal notions cannot be derived from descriptive claims. It’s true that you can’t get $\Box P$ from P , but from a descriptive claim about how all the relevant possible worlds are, we can most definitely derive modal claims. So, for example, from the descriptive claim ‘for all worlds w , P is true in w ’, we can legitimately derive the modal claim $\Box P$.⁴

Invoking possible worlds semantics for modal claims to move between the descriptive and the modal might be seen as something of a cheat and irrelevant to the main issue. After all, it might be argued that unless one is a realist about possible worlds, claims about possible worlds other than the actual are not really descriptive claims at all, but, rather, are just disguised modal claims. I won’t take a stand on this issue here, for the simple reason that it is not my purpose to defend the more general claim that all modal notions can be derived from matters of fact; I will be content to demonstrate that various normative claims concerning rationality can be derived from matters of fact. Why then raise the issue of the modal semantics? Because it demonstrates that often we are quick to think of modal claims as being of a quite different kind to non-modal claims and a case can be made, at least,

³Although the modal $\Diamond P$ *does* follow from the descriptive P , at least under some interpretations of ‘ \Diamond ’.

⁴See Hintikka (1962) for a classic formal account of knowledge and belief in terms of modal operators.

that this is mistaken. Once one looks for the *right kind* of non-modal facts, the apparent gap between modal facts and non-modal facts disappears. So too, I will argue, is the case with claims about rationality.

3 From the Descriptive to the Prescriptive

Consider some typical prescriptive claim concerning rationality:

- (1) It is rational to use *modus ponens*.

Now consider some descriptive claim such as

- (2) Agent *A* employs *modus ponens*.

Clearly (1) does not follow from (2), because, apart from anything else, *A* might be a poor reasoner, and quite unjustified in her use of *modus ponens*. But (2) is not the only relevant descriptive claim. Let's look further afield (as we did in the last section when looking for descriptive counterparts for modal claims). Consider

- (3) Most agents employ *modus ponens*.

Does this help in deriving (1)? No, for much the same reason: most agents could be bad reasoners.⁵ But what if we restrict attention to good reasoners? Consider

- (4) Good reasoners employ *modus ponens*.

There does not seem to be such a gap between (1) and (4), but now another problem emerges, namely, circularity. The phrase 'good reasoners' in (4) is very suggestive of 'rational agent', so we are a small step away from providing a circular definition of rationality, in terms of what rational agents do. It would seem that we've closed the normative–descriptive gap, but at the price of rendering the account circular, and thus useless.

But if we can give an independent account of what a good reasoner is, we can avoid the charge of circularity. I claim that we can do this, but to see how, it will be useful to look at a related example. Consider the concept of 'a good chess move' in a given setup. Obviously a good chess move must be a legal move, but it must be more than this. We might define such a move in terms of what good chess players would do in such circumstances. Our definition of 'a good chess move' also seems to be circular—at least until we can spell out what 'a good chess player' is. But the latter is easy. It does not have to be spelled out in terms of someone who makes good chess

⁵Indeed, anyone who has ever taught introductory formal or informal logic will know that a large number of incoming students take various fallacies (e.g., affirming the consequent) as valid argument forms. Wason and Johnson-Laird (1972) explore these issues.

moves; we can spell it out in terms of someone who wins games of chess. Or better still, we can spell out what a good chess player is in terms of their moves according well with the relevant bits of chess theory. Now, with this example in mind, let's return to the question of rationality.

One way to answer our question of what a good reasoner is, is to provide a teleological account, according to which a good reasoner is one who is better adapted to achieving certain goals. There are a number of candidates for the appropriate goals. The most popular of these is survival. While such evolutionary accounts of rationality undoubtedly have something to recommend them,⁶ typically such accounts are too indiscriminating. After all, not all poor reasoners have been selected against. A better answer to the question of what a good reasoner is, is that it's someone whose reasoning is well supported by our best theories of reasoning. Again, it might seem that we are dangerously close to circularity, but we're not. Our best theories of rationality—formal logic, belief theory, and decision theory—are not defined purely in terms of what good reasoners do. These theories earn their place as our best theories of rationality in much the same way that general relativity earns its place as our best scientific theory of space and time. I'll elaborate on this shortly.

There is another feature of theories of rationality that deserves mention. Theories of rationality must employ a great deal of what they set out to explain and systematise. Take, for example, a debate in logic—the debate over the validity of excluded middle, say. In order to settle the issue of whether we ought to employ a logic in which excluded middle holds (such as classical logic) or one in which it fails (such as Kleene strong K_3) we need to reason and, in particular, we need to make, and assess the validity of, various logical inferences. But which logic do we use for this task? Philosophical logic, like other subjects that use the very theories under study, is difficult, but not impossible.⁷ Indeed, it would seem that the three theories of rationality that I have in mind for this discussion—formal logic, belief theory, and decision theory—all have this reflexive aspect. This makes their study difficult, but it does not undermine their value or preclude progress on the study of rationality.⁸

But putting aside the worries about the reflexive nature of the rational study of rationality, another more pressing problem emerges. So far, I've argued that we can get from descriptive statements to normative ones by

⁶See, for example, Nozick (1993) and Papineau (2003).

⁷I won't go into the details here. Elsewhere, with Otávio Bueno (2004), I've argued that even apparently fundamental principles such as the law of non-contradiction can be sensibly debated without begging questions.

⁸After all, psychology manages to use the human mind to study the human mind and even physics employs physics to study physics. We should not become what Adam Elga calls 'reflexophobic'. That is, we should not be so struck with the various problems and paradoxes of self reference that we shy away from self reference wherever it occurs. Much self reference is benign.

invoking a certain respect for scientific theories. But on what grounds can this be justified? Indeed, this move, once spelled out, just smuggles normativity in through the back door. For surely the principle I'm appealing to is some version of the doctrine of naturalism⁹ that counsels us to believe our best scientific theories. Or, to make the normative assumption manifest, it's rational to believe or accept our best scientific theories.¹⁰ I agree that naturalism (of the kind I subscribe to¹¹) is normative, but I deny that this undermines the account of the normativity of rationality I've given so far. I discuss this in the next section.

4 The Role of Naturalism

As we saw in the previous section, my justification of the normativity of various theories of rationality hangs crucially on the doctrine of naturalism. We've also seen that this doctrine is normative—it tells us that it is rational to believe (only) our best scientific theories. So the cogency of my argument hangs on a defence of naturalism. Unfortunately, a defence of such a fundamental doctrine as naturalism is hard to come by. One finds different positions on this issue. Some suggest that naturalism is just an attitude one takes towards investigating the world. It's just a basic commitment to attempt to describe and explain the world without resorting to supernatural or otherwise spooky entities and forces. But I think we can do better than simply taking a stand. I suggest that the naturalistic approach has many benefits and we should judge the doctrine by its fruits, not by arguing for it from more fundamental principles.¹²

So what are the fruits of naturalism? First, the scientific enterprise has a remarkably successful history, and naturalism is little more than a statement of our continued support for that enterprise. After all, rejecting naturalism amounts to claiming that sometimes we ought not accept our best scientific theories. Let's get clear what this amounts to in the current context. With Quine, I'm understanding science very broadly here, to include all theoretically and empirically well-supported areas of study (including philosophy). In short, our best scientific theories are simply our best theories. To reject naturalism is to deny that we ought to accept our best theory of some domain. But what are the other options? Accept the second best theory? Accept no theory at all?¹³ Once put this way, naturalism, if not self ev-

⁹See Quine (1969).

¹⁰I hedge by including both 'accept' and 'believe' here so as not to beg any questions against constructive empiricists such as van Fraassen (1980).

¹¹See Colyvan (2001, chap. 2)

¹²This is the same way modern set theorists argue for the acceptance or rejection of new set-theoretic axioms. See Russell (1973) and Gödel (1946) for early advocates of this modern view of the foundations of mathematics.

¹³Naturalism, properly construed, doesn't rule against an occasional bout of agnosti-

ident, is at least a rather compelling doctrine. It's not trivial though. It does rule against certain mystical and religious world views, for instance—at least when there are better (scientific) theories of the same phenomena. Other benefits of naturalism are that it enables a rather plausible reply to the sceptic¹⁴ and it provides a satisfying account of the relationship between philosophy and science.

Once again, I stress that although the above justifications of naturalism *are* circular, they are not viciously so. All chains of justification ultimately result in either infinite regress or circularity. All other things being equal, circularity seems the more palatable option, so long as the circles are not too tight (i.e., vicious).¹⁵ If all I've said so far is correct, we have a sketch of an account of the normativity of rationality theories. This account is based on the acceptance of the naturalistic point of view. There are, however, no free lunches. Let's consider the consequences of this particular defence of normativity.

The account of normativity I've been defending depends on naturalism's respect for the (current) best scientific theories of rationality. It is thus a defeasible account and typically canons of rationality will change depending on the state of science at the time. What counts as a valid inference at one time may not count as a valid inference at some other time. Those who would have classical logic, standard Bayesian belief theory, and standard decision theory as *the* normative theories of rationality may be uncomfortable with this result.¹⁶ For, on the view I'm advancing here, these theories are only *defeasibly* the normative theories of rational inference, rational belief, and rational decision making, respectively.¹⁷ But seen in the right light, this defeasibility is both natural and desirable. Our current best scientific theories of rationality should not be any different from other scientific theo-

cism. For instance, naturalism may license agnosticism when there is no best theory or when the best theory is clearly inadequate.

¹⁴See Quine (1974, p. 3).

¹⁵It has been claimed by Katz (2000) that the kind of naturalism under consideration here is worse than circular—it leads to paradox. Elsewhere (Colyvan, 2005a), I've argued that Katz is wrong and that no paradox can be generated—at least not without making some rather implausible assumptions about the logic of belief revision.

¹⁶For example, a great deal of work on rationality in psychology takes these classical theories of rationality to be the final word on normative rational theory. Kahneman and Tversky (see many of the essays in Kahneman, Slovic, and Tversky (1982)) take standard Kolmogorov probability theory to be (uncontroversially) *the* theory of belief and belief updating. And Wason and Johnson-Laird (1972) take classical logic as *the* theory of deductive inference. See Hájek (2003a and b) for some of the shortcomings of classical (conditional) probability theory and Shafer (1976) and Walley (1991) for examples of non-classical belief theories.

¹⁷In fact, I think classical logic has already passed its use-by date. And in so far as standard Bayesian belief theory, and standard decision theory presuppose classical logic—it's the tautologies of classical logic that are defined to have maximal probability in Kolmogorov probability theory—those theories too are questionable. See Colyvan (2004) for more on this issue.

ries. All these theories are known a posteriori, they are defeasible, and their acceptance or rejection is sensitive to new evidence and new developments.

5 The Ramsey-Lewis Approach to Normativity

Let me finish up by saying a little about the points of contact with, and the differences between, the account I've presented and the Ramsey-Lewis approach to defining theoretical terms (otherwise known as “the Canberra Plan”).¹⁸ The Ramsey-Lewis approach provides a rigorous means of providing implicit definitions of theoretical terms by appealing to their role in the theories in which they occur. Suppose, for ease of exposition, that there is only one theoretical term, τ , we are hoping to define. We separate the theory in question, Γ , into theoretical terms (which in this case is just τ) and non-theoretical terms, d_i . The resulting theory can be written as

$$\Gamma(\tau, d_1, d_2, d_3, \dots d_n).$$

We then existentially generalise to get the Ramsey sentence

$$(\exists x)\Gamma(x, d_1, d_2, d_3, \dots d_n).$$

Finally, we replace the existential quantifier in the Ramsey sentence with the definite description operator (ιx) to obtain our implicit definition of τ .

$$\tau =_{\text{Defn}} (\iota x)\Gamma(x, d_1, d_2, d_3, \dots d_n).$$

This can be generalised to cases of m theoretical terms, $\tau_1, \tau_2, \dots, \tau_j, \dots, \tau_m$, with the corresponding Ramsey sentence:

$$(\exists x_1, x_2 \dots x_j \dots x_m)\Gamma(x_1, x_2, \dots x_j, \dots x_m, d_1, d_2, \dots d_n).$$

We can then provide implicit definitions for each τ_j in terms of the non-theoretical terms and other theoretical terms.¹⁹

$$\begin{aligned} \tau_j =_{\text{Defn}} & (\iota x_j)(\exists x_1, \dots x_{j-1}, x_{j+1}, \dots x_m)(\forall y_1, \dots y_j, \dots y_m) \\ & (\Gamma(y_1, y_2, \dots y_j, \dots y_m, d_1, d_2, \dots d_n) \equiv \bigwedge_{i=1}^n y_i = x_i). \end{aligned}$$

In the case of theories of rationality, the relevant theories will be Bayesian belief theory, deductive logic, and rational choice theory. Jackson (1998) and Nola (2003) have shown how to extend the Ramsey-Lewis approach to

¹⁸See Ramsey (1990) and Lewis (1983).

¹⁹The complications are a result of defining one theoretical term by way of other theoretical terms, where the latter are to be defined in the same manner. These complications need not concern us here. See Lewis (1983) for details.

normative theories such as ethics and rationality (respectively). As Nola (2003, p. 162) points out, all we need to do is separate the normative terms (such as ‘rational’) and non-normative or descriptive terms of the theory. The resulting implicit definition of the normative terms in question appeals only to their functional role in the theory and this is articulated only in non-normative vocabulary. There is a sense in which this bridges the gap between the normative and the descriptive in a naturalistically respectable way. More on this shortly.

It might be thought that the Ramsey-Lewis account is simply a formalisation of the approach I’ve argued for in the earlier sections of this paper. There are certainly some similarities between the two approaches. Both take naturalistic approaches to dealing with normativity and both focus attention on theories of rationality. And they both agree that there is something unsettling about normative notions. The Ramsey-Lewis approach’s response is to eliminate the normative in favour of the non-normative. The approach I’m suggesting reduces all normativity to just one—the normativity implicit in naturalism. But both accounts are naturalistic and they both take it that normativity deserves to come in for some special treatment. That is, they both set out to address the Humean worry about deriving normative claims from descriptive claims.

Despite these initial similarities, there are also some important differences between the two approaches. Take rational choice theory, for instance. Here the normativity—what an agent ought to do—is spelled out in non-normative terms—the maximising of expected utility. Both approaches seem to be on the same page so far, but notice a couple of important differences. In the Ramsey-Lewis approach (at least the version advanced by Nola (2003)) the normative–descriptive gap is bridged by providing a *definition* of normative terms in terms of purely descriptive terms of the relevant theory. The approach I’m advocating doesn’t claim to provide anything so strong as a definition of the normative terms in question. After all, nowhere is it claimed that the meaning of “what you ought to do” is “maximise expected utility”. So there’s one significant difference between the two approaches: the Ramsey-Lewis approach invokes implicit definitions to bridge the normative–descriptive gap; the approach I’m defending does not offer definitions at all.

There is a second, more significant difference between the two approaches. The Ramsey-Lewis approach defines the normative terms via their functional role in the relevant scientific or folk theories. My approach, on the other hand, uses the scientific theory to back up the normative claims in question: you ought to do what your best theory says to do. Why? because you ought to believe your best scientific theories. And the threatening regress is blocked by invoking the doctrine of naturalism. In effect, on this approach, all normativity arises from the normativity built into the doctrine of naturalism. Of course the Ramsey-Lewis approach is also naturalistic.

But it is naturalistic in a quite different way. The naturalism doesn't provide the source of normativity; the Ramsey-Lewis approach is naturalistic in the sense that normative terms are defined via naturalistically-respectable descriptive terms. In the Ramsey-Lewis approach the aim is to arrive at a naturalistically respectable theory—one without suspect terms—but naturalism does no real work in the account. Naturalism provides only the motivation. On my account naturalism does some serious work, for it is naturalism that is the ultimate source of the normativity in question.

The final difference between the two approaches concerns the scope of their applicability. It is not clear that the approach I've argued for can be extended to areas where there is no well-defined, consistent, and unique formal theory. My approach can only be applied to well-systematised (indeed, scientific) theories of rationality. The account works well for the theories I've been explicitly discussing—formal logic, belief theory, and decision theory—but it's not at all clear that the account presented here can be used for less rigorous theories of rationality (such as common-sense judgements and informal logic). For similar reasons, extending this account of normativity to ethics may also prove problematic; it's not clear that there is a suitably rigorous and systematic theory of ethics on which the account can rely.²⁰ The strict version of the Ramsey-Lewis account may too suffer such limitations, but it is generally taken to be a core part of the Canberra plan that the Ramsey-Lewis approach to defining theoretical terms can be extended to *folk* theories.²¹ But in any case, I see no prospect for the approach I've argued for to be extended in this way. This would thus seem to mark a significant difference between the two approaches.

This brings us to the question of which is the better approach. Indeed, in light of the last point, the wider scope of the Ramsey-Lewis approach would suggest that the latter has the edge. But I think that it is a mistake to compare the two like this. There is a sense in which these are two completely different games. The Ramsey-Lewis approach, motivated by a naturalistic suspicion of normativity, aims to dispense with normative terms in favour of non-normative terms. As I've already mentioned, though, in this approach, naturalism plays no role beyond motivation. The Ramsey-Lewis approach seeks to satisfy the exacting standards of a certain conception of naturalistic philosophy. The approach that I'm advocating, on the other hand, employs naturalism to defend the normativity so central to theories of rationality. The Ramsey-Lewis approach is apologetic for normativity and, in the name of naturalism, seeks to do away with normativity. My approach embraces

²⁰Though see Colyvan *et al.* (to appear) for some steps towards formalising ethical theories.

²¹Although I won't enter that debate here, there is some reason to be suspicious of such extensions. For the Ramsey-Lewis approach to work, we need to have something that resembles a theory, and folk "theories" are often inconsistent, or at least non-unique and often imprecise.

normativity and employs naturalism to allow normativity to take its rightful place in theories of rationality. Once put this way, the question of which approach you should prefer boils down to a question of what you're trying to do and how you conceive of the naturalistic project of philosophy. Indeed, it may well turn out that these two approaches to normativity are not competitors at all, just different parts of, or different takes on, a broader naturalistic project of naturalising normativity.²²

References

- [1] Bueno, O. and Colyvan, M. (2004) 'Logical Non-Apriorism and the Law of Non-Contradiction', in G. Priest, JC Beall, and B. Armour-Garb (eds.), *The Law of Non-Contradiction: New Philosophical Essays*, Oxford University Press, Oxford, pp. 156–175.
- [2] Colyvan, M. (2001) *The Indispensability of Mathematics*, Oxford University Press, New York.
- [3] Colyvan, M. (2004) 'The Philosophical Significance of Cox's Theorem', *International Journal of Approximate Reasoning*, Vol. 37, No. 1, 71–85.
- [4] Colyvan, M. (2005a) 'Naturalism and the Paradox of Revisability', *Pacific Philosophical Quarterly*, forthcoming.
- [5] Colyvan, M. (2005b) 'No Expectations', *Mind*, forthcoming.
- [6] Colyvan, M., Cox, D. and Steele, K. (to appear) 'Modelling the Moral Dimension of Decisions'.
- [7] Gödel, K. (1947). What is Cantor's Continuum Problem? reprinted (revised and expanded) in P. Benacerraf and H. Putnam (eds.), *Philosophy of Mathematics Selected Readings*, second edition. Cambridge University Press, Cambridge, 1983.
- [8] Hájek, A. (2003a) 'Conditional Probability Is the Very Guide of Life', in H.E. Kyburg and M. Thalos (eds.), *Probability Is the Very Guide of Life: The Philosophical Uses of Chance*, Open Court, La Salle, Ill.
- [9] Hájek, A. (2003b) 'What Conditional Probability Could Not Be', *Synthese*, Vol. 137, No. 3 (December 2003), 273–323.
- [10] Hintikka, J. (1962) *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell University Press, Ithaca, NY.
- [11] Hume, D. (1975) (edited by L.A. Selby-Bigge) *A Treatise of Human Nature*, Clarendon Press, Oxford.

²²An earlier version of this paper was presented at the joint 4th International Conference on Cognitive Science and the 7th Australasian Society for Cognitive Science Conference at the University of New South Wales in July 2003 and in a joint Philosophy and Science Studies Seminar at the University of South Carolina. I'd like to thank the audiences at those presentations for their contributions. I am also very grateful to Carol Booth, Damian Cox, and Juliana Weingaertner for helpful discussions, and Alan Hájek and Carol Booth for comments on an earlier draft.

- [12] Jackson, F. (1998) *From Metaphysics to Ethics*, Oxford University Press, Oxford.
- [13] Katz, J.J. (2000) *Realistic Rationalism*, MIT Press, Cambridge, Ma.
- [14] Kahneman, D., Slovic, P., and Tversky, A. (1982) *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.
- [15] Lewis, D. (1983) 'How to Define Theoretical Terms', in D. Lewis, *Philosophical Papers Volume I*, Oxford University Press, Oxford, pp. 78–95. (Originally published in 1970.)
- [16] Nola, R. (2003) *Rescuing Reason: A Critique of Anti-Rationalist Views of Science and Knowledge*, Kluwer, Dordrecht.
- [17] Nover, H. and Hájek, A. (2004) 'Vexing Expectations', *Mind*, 113, 237–249.
- [18] Nozick, R. (1993) *The Nature of Rationality*, Princeton University Press, Princeton.
- [19] Papineau, D. (2003) *The Roots of Rationality: Philosophical Essays on Rationality, Evolution, and Probability*, Oxford University Press, Oxford.
- [20] Quine, W.V. (1969) 'Epistemology Naturalized', in *Ontological Relativity and Other Essays*, Columbia University Press, New York, pp. 69–90.
- [21] Quine, W.V. (1974) *The Roots of Reference*, Open Court, La Salle, Ill.
- [22] Ramsey, F.P. (1990) 'Theories', in F.P. Ramsey (ed. D.H. Mellor), *Philosophical Papers*, Cambridge University Press, Cambridge, pp. 112–136. (Originally published in 1931.)
- [23] Russell, B. (1973) 'The Regressive Method of Discovering the Premises of Mathematics' reprinted in D. Lackey (ed.) *Essays in Analysis*, George Allen and Unwin Ltd. (Originally published in 1907.)
- [24] Shafer, G. (1976) *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ.
- [25] van Fraassen, B. (1980) *The Scientific Image*, Clarendon, Oxford.
- [26] Walley, P. (1991) *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, London.
- [27] Wason, P.C. and Johnson-Laird, P.N. (1972) *Psychology of Reasoning: Structure and Content*, Harvard University Press, Cambridge, Ma.