Bayesian Inductive Logic

Doctoral Dissertation for Philosophy
Bayesian Inductive Logic
Jan-Willem Romeyn

University of Groningen
Facullty of Philosophy
Oude Boteringestraat 52
9712 GL Groningen

Research supported by NWO, RuG and BCN.

 Published 2005.
Printed in the Netherlands by
Haveka BV Alblasserdam.
ISBN 90-367-2279-9

Explanation of cover The cover shows the physical models of two Bernoulli processes. One model consists of a pair of identical cubic dice, both with six possible outcomes. The other model consists of two different dice, a tetrahedron and an octahedron with four and eight possible outcomes respectively. If one of the pairs is selected at random and rolled repeatedly, Bayes' rule can be used to infer which of the two pairs is selected on the basis of the observations of the added outcomes. The two models can thus be told apart even though they have the same range of possible observations, namely 2 to 12.

RIJKSUNIVERSITEIT GRONINGEN

BAYESIAN INDUCTIVE LOGIC

INDUCTIVE PREDICTIONS

FROM

STATISTICAL HYPOTHESES

PROEFSCHRIFT

ter verkrijging van het doctoraat in de
Wijsbegeerte
aan de Rijksuniversiteit Groningen
op gezag van de
Rector Magnificus, dr. F. Zwarts,
in het openbaar te verdedigen op
donderdag 2 juni 2005
om 16.15 uur

door

Johannes Wilhelmus Romeijn
geboren op 13 augustus 1975
te Delft

Promotor:	Prof. dr. T. A. F. Kuipers
Copromotores:	Dr. I. Douven
	Dr. J. Peijnenburg
Beoordelingscommissie:	Prof. Dr. C. Howson
	Prof. Dr. E. C. W. Krabbe
	Prof. Dr. M. van Lambalgen

ISBN 90-367-2279-9

Preface

Writing a thesis is like gardening. Every spring of green ideas must be followed by a careful process of pruning and tripping. Accordingly, the garden that is forever fixed in this book shows only one of the many collections of ideas that I have been cultivating over the years. But I believe it contains ideas that may grow into trees one day, and perhaps even carry fruit.

When I started with the project 'Inductive Rules and the Structure of Evidence', I knew very little to nothing of statistical inference. This has proved to be an enormous advantage. Already after a few months, it became apparent that the criticisms against Bayes' rule that I had envisioned were unsound. The reorientation of the project that resulted from this insight has laid the basis for the research that I carried out in the remaining three and a half years.

The circumstances for gardening have been almost ideal. It has been a time of great intellectual freedom. This is for a large part due to the setting in which I found myself as a PhD in Groningen, and for this I am very grateful. It has made the return to academic life from an occupation in consultancy into one of the best decisions I have ever made. The intellectual freedom may further explain the fact that next to some promising plants and trees, I seem to have nurtured all sorts of weeds over the past years. There is no need for complaining. I have come to the conviction that creative development is often the result of a fearless engagement in error and confusion.

I invite the reader to wander around in the result of some four years of gardening. Many ideas in it will still appear to have only just begun growing, some others may look rather dull and dry. But luckily, as I am writing these last words, spring is in the air.

J. W. R.
April 2005

Acknowledgements

In this thesis I argue that knowledge is always a co-production of an observer and an observed world. In the same way this thesis is itself a co-production of its author and his environment. It is a pleasure to be able to thank some people in this environment for producing this thesis with me.

First of all, I thank Theo Kuipers and Jeanne Peijnenburg for their knowledge and critical reflection during all the stages of the project. Many ideas in this thesis made their first appearance in one of our discussions. Another wonderful stage for sharing ideas was provided by Igor Douven. I am grateful for his comments, his sharp mind and his good advice, from which the whole thesis has benefitted. I am further grateful to the members of the reading committee. I thank Colin Howson for his encouraging words, and for his hospitality during my visit of the London School of Economics. I thank Erik Krabbe, especially for his careful reading of the manuscript. And I thank Michiel van Lambalgen for helpful discussions on my ever diverging research plans.

Many people have been involved in writing this thesis on a more occasional basis, by looking at drafts or discussing a common interest. Alan Hajek, Alexander van den Bosch, Allard Tamminga, Atocha Aliseda, Barteld Kooi, Beerend Winkelman, Boudewijn de Bruin, Branden Fitelson, Conor Dolan, Daan Franken, David Atkinson, David Makinson, Denny Borsboom, Elliot Sober, Frank Hindriks, Franz Huber, Fred Muller, Fred Keijzer, Gabriele Contessa, Hans Rott, Hasok Chang, Hauke de Vries, Hedde Zeijlstra, Herman Philipse, Ioannis Votsis, James McAllister, Jan Albert van Laar, Jeff Seidman, Jeroen van Maanen, Jim Joyce, Johan van Benthem, John Woods, Jon Williamson, Jos Uffink, Julian Reiss, Karin de Boer, Kim van Gennip, Luc Bovens, Maarten van Dyck, Marc van Duijn, Marian Counihan, Martin van Hees, Melle Lyklema, László Pólos, Martine Prange, Menno Rol, Michael Esfeld, Miklos Redei, Nils-Eric Sahlin, Patrick Maher, Peter Grunwald, Peter Lipton, Peter Milne, Pieter Sjoerd Hasper, Prasanta Bandyopadhyay, Richard Bradley, Robert Bishop, Roberto Festa, Roman Frigg, Ruurik Holm, Stephan Hartmann, Talal Debs, and Wouter Meijs, thanks a lot.

During the project I have travelled quite a bit. Firstly, I visited the London School of Economics, which has been an absolutely wonderful experience. I want to thank all those involved in making my visit possible, and for making it such fun. Further, I have spent a short time in Edinburgh, where Peter Milne made me feel very welcome.

Apart from these two visits I travelled to give talks. I have benefitted from discussions during and after the talks I gave for the GROningen LOGicians, for the Promotion Club Cognitive Patterns, for the ILLC Logic Tea, for the Foundations of Physics Group in Utrecht, for the PhD-seminar of the Erasmus Institute of Philosophy and Economics, for the Philosophy and Probabilistic Modelling Group in Konstanz, for the Belief Revision seminar and the Popper seminar at the LSE, for the Work in Progress seminar in Edinburgh, and for audiences at the Dutch Conference on Philosophy of Science in Leusden, the CENSS conference in Gent, the Workshop on Bayesian Epistemology at the Wittgenstein conference, the Popper Centennial conference in Vienna, the LMPS conference in Oviedo, the VlaPoLo 9 in Gent, the BSPS conference in Kent, the PSA conference in Austin Texas, and the Wiener Kreis lecture of Michael Friedman. I thank all the audiences and academic communities that allowed me to share my ideas, and thus, finally, I thank the reader of these words.

Contents

Outline 1

Introduction 3

I Bayesian Inductive Inference 15

1 Carnapian and Bayesian Inductive Predictions 17
1.1 Introduction 17
1.2 Inductive predictions 20
1.2.1 Observation framework 21
1.2.2 Carnapian scheme 23
1.3 Bayesian scheme 26
1.3.1 Predictions from statistical hypotheses 26
1.3.2 Probability assignments over hypotheses 28
1.4 A logical picture 32
1.4.1 Conclusions and inference rules 32
1.4.2 Probability models as premises 35
1.4.3 Multiple probability models 37
1.5 Normative aims 39
1.6 Conclusion 42

2 A Frequentist Semantics of Hypotheses 45
2.1 Statistical hypotheses 45
2.2 A restricted class of hypotheses 48
2.3 Hypotheses as elements of $\mathcal{Q}$ 52
2.3.1 Definition of the elements $H_{w\theta}$ 52
2.3.2 Collections of Kollektivs 55
2.3.3 Partitions 57
2.4 Objective probability models 59
2.4.1 Objective probability 59
2.4.2 Deriving direct probabilities 60
2.4.3 Physical models 63
2.5 Comparing inductive schemes 65

3 Hypotheses as Inductive Assumptions 71
3.1 Introduction 71
3.2 Carnap-Hintikka inductive logic 72
3.3 Examples on crash data 77
3.3.1 Bernoulli partition 77
3.3.2 Crash hypotheses 80
3.4 The use of partitions 83
3.4.1 Immediate insights 84
3.4.2 Partitions as inductive assumptions 85
3.4.3 Advantages of the Bayesian scheme 87
3.5 Conclusion 89

II Analogy and Independence 93

4 Analogical Predictions for Explicit Similarity 95
4.1 Analogy within Carnapian rules 95
4.2 Analogical predictions 98
4.3 A model for explicit similarity 102
4.4 Problematic aspects 105
4.5 Statistical underpinning for $\lambda\gamma$ rules 108
4.6 Analogy partition 112
4.6.1 Defining the analogy partition 112
4.6.2 Transforming partitions 114
4.7 Generalizing the analogy partition 117
4.8 Conclusion 121

5 A General Model for Analogical Predictions 123
5.1 Introduction 123
5.1.1 Inductive relevance 123
5.1.2 Aim of this chapter 125
5.2 Hypotheses schemes for Q-predicates 127
5.3 Schemes using underlying predicates 129
5.4 Transformations between partitions 132
5.4.1 Coordinate transformations 133
5.4.2 Transformed probability models 135
5.5 General model 137
5.5.1 A more modest aim 137
5.5.2 Towards a general model 138

5.5.3 Tentative solution 141
5.6 Qualitative and numerical characterisations 143
5.6.1 The form of the analogy priors 143
5.6.2 Numerical examples 145
5.7 Other models of analogical predictions 147
5.8 Conclusion 151

6 Inductive Inference for Bayesian Networks 155
6.1 Induction and Bayesian networks 155
6.2 Bayesian scheme 160
6.2.1 Observational algebra 160
6.2.2 Hypotheses and predictions 162
6.3 Hypotheses on nodes 165
6.3.1 Relative frequencies 165
6.3.2 Unconnected nodes 166
6.4 Tools for Bayesian networks 169
6.4.1 Marginal and conditional densities 169
6.4.2 Predicting and updating 171
6.4.3 A system of $\lambda\gamma$ rules 173
6.5 Capturing the links 175
6.5.1 Hypotheses schemes 175
6.5.2 Prediction rules 177
6.5.3 Illustration 179
6.6 Inductive dependence 180
6.6.1 General idea 181
6.6.2 Varying conditional distributions 182
6.6.3 Approximated predictions 184
6.7 Generalising to multiple links 186

III Philosophy of Science 189

7 Induction in the Bayesian Scheme 191
7.1 The problem of induction 191
7.2 Hypotheses as tools 194
7.2.1 Indispensable assumptions 195
7.2.2 General uniformity 196
7.2.3 Logical solution 197
7.3 Induction as co-production 198

7.3.1 The role of observations 199
7.3.2 The theoretical level 200
7.3.3 Externalism 202
7.4 Suppositions on substance 202
7.5 Summary and conclusion 205

8 Bayesian Theory Change 207
8.1 Introduction 207
8.2 Hypotheses, conditioning and predictions 208
8.3 Contaminated cows 211
8.4 Careless vaccination 213
8.5 A framework for changing partitions 216
8.6 Distance between partitions 218
8.7 Concluding remarks 221

9 Abducted by Bayesians? 223
9.1 Statistical inference using partitions 223
9.2 Duplicate partitions 225
9.3 The use of underdetermination 229

Conclusion 233

Summary in Dutch 239

Word of Thanks in Dutch 247

References 249

Outline

It is probably not the intention of the reader to go through this thesis from cover to cover. To accommodate this eclecticism, chapters are set up more or less independently. The technical parts of the chapters show a considerable overlap, but each chapter emphasises other aspects of the technical material relating to the current theme. On the other hand, since complete independence of the chapters leads to too many repetitions, especially in the first part of the thesis, preceding sections or chapters are sometimes presupposed.

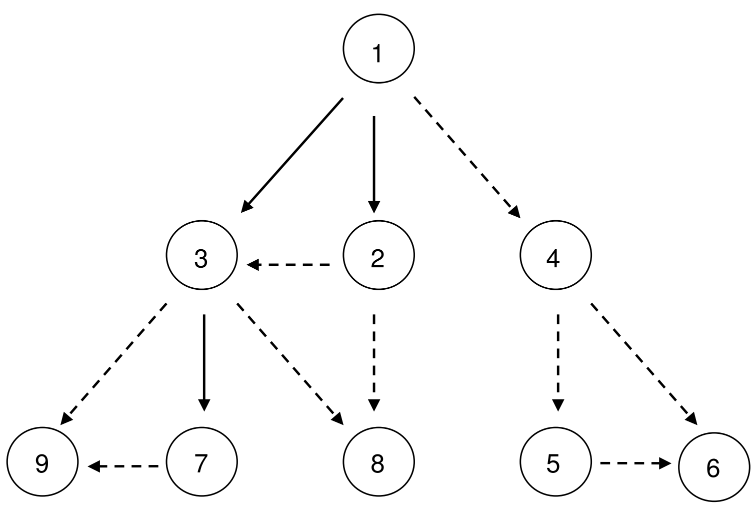

The diagram summarises the relations of presupposition by means of arrows. For example, the arrow from chapter 1 to chapter 2 indicates that the latter presupposes the former. If a chapter presupposes only certain sections, references to these sections are included in the chapter itself. The dotted arrows, as for example from chapter 2 to chapter 3, mean that the former may be helpful for a better understanding of the latter, while it is not necessary for understanding the main line.

Two further comments are in order. Firstly, it may be noted from the diagram that the second part of this thesis is quite independent from the first and third part. Even though it may be read without any specialist knowledge, the second part concerns a rather specialist subject within inductive logic. The thesis has been organised so as to allow the reader to leave aside the second part completely, and jump from chapter 3 to chapter 7. Secondly, the introduction and conclusion are not included in the diagram. They provide a general understanding of the position of the chapters within this thesis, and of this thesis as a whole in the debate on inductive logic and its relation to the philosophy of science.

Introduction

> "You will reply that reality has not the least obligation to be interesting. And I will answer you that reality may avoid this obligation, but that hypotheses may not." – J.L. Borges in Death and the Compass

Inductive Inference. This thesis concerns inductive inferences, that is, inferences running from given observations to as yet unknown observations and general observational statements. Inductive inferences are almost everywhere, and come so naturally to us that they easily escape philosophical attention. We drink water because it has always refreshed us, we stride forward confidently because the earth has always attracted and carried us, and we give way to sleep because we have always woken up to renewed presence. In all these cases, the trust that we put in the stability of a pattern can be seen as a trust in the inherent inductive inference. A bit further removed from the backbone of life, these inferences perhaps become more easily recognisable. If the post has been delivered every morning until now, we expect it to be delivered on future mornings too, but if delivery is late on Saturday and Sunday, it will not take long before we only expect it in the morning on weekdays. Here again, the trust in the stability of the pattern is eventually a trust in an inductive inference.

The basic assumption of inductive inferences is that the world is a boring place, and that the same pattern in the observations will keep repeating itself. Usually the sameness is taken as the result of some structure in the world, which is supposed to underpin the patterns in the observations. Fortunately, the world is boring in a rather interesting way. Many observations show patterns that are occasionally violated. For example, it may be that of two equally expensive stalls at the market the fruit of one is usually better than that of the other. But this need not be the case every week. These so-called weak patterns suggest that the world contains a certain structure, but that this structure is perturbed by other structures or effects, which may be deemed noise for the occasion. The one merchant may be better in spotting good fruit at the auctions than the other, but the trade at such auctions always contains an element of chance. Nevertheless we may come to know that the fruit of the one stall is usually better than that of the other. So inductive practice is able to pick up on weak patterns as well.

Inductive inferences are abundant in daily life, and no less so in the daily life of scientists. Much of experimental science concerns the identification of, possibly weak, patterns in the observations. The interdependence between electrodynamic and magnetic forces is an example of a strong, exceptionless pattern. The correlation between vaccination and disease in a population of cows exemplifies a weak pattern, because it only shows in large herds. In general, in experimental science there is a basic trust in the stability of specific patterns, and the typical activity of theoretical science is to motivate this trust by providing a picture or story on the structure behind the stable patterns. Of course, the discussion on structures may get far removed from the observations, but this must not distract from the original intention of experimental science to select stable patterns, and to strip them of noise.

The problem of induction. Once we have realised how deeply both common sense and science are permeated with inductive inference, the destructive force of the problem of induction becomes apparent. David Hume, in his 'Treatise on Human Nature' of 1739, was the first to put his finger on the sore spot. In book 1, part 4, section 2 he writes:

> "Any degree, therefore, of regularity in our perception, can never be a foundation for us to infer a greater degree of regularity in some objects which were not perceived, since this supposes a contradiction, viz., a habit acquired by what was never present to the mind."

So a perceived regularity never allows us to infer a greater regularity, that is, a regularity that includes objects or facts that were not yet perceived. In other words, observations alone can never justify inductive inferences. Any conclusion that transcends observations is not properly inferred from these observations alone, but invokes an additional component, namely the assumption of the stability of some pattern in the observations. Moreover, according to Hume this assumption of stability is a component that cannot itself be derived from the observations. Common sense and science are both resting on nothing more than sheepish habit.

Not surprisingly, this destructive conclusion invited a lively discussion, to which the present thesis is yet another contribution. The essential characteristic of this contribution is its focus on the logical part of the problem of induction, that is, the part that concerns the inferences themselves. I claim that this thesis solves the logical problem, by presenting a scheme for valid inductive inference. It further clarifies the role of the input components of this scheme, and employs the scheme in clearing up some more specific problems concerning inductive inference. Note that the logical part is strictly separated from the epistemological part of the problem of induction, which concerns the input components of inductive inference. The latter part is dismissed as irrelevant to the task of the inductive logician. This perspective resembles the perspective of Howson (2000), but it is more explicitly focused on a specific scheme for inductive inference. The general aim is a revival of inductive logic, and a better control over inductive inference in science.

The remainder of this introduction provides a sketch of the framework within which the inductive logic of this thesis finds its place. It then introduces Bayesian inference, as it is used in this thesis, and discusses the relevance of this thesis for scientific method and the philosophy of science more generally. The introduction ends with an overview of the chapters.

Logical empiricist framework. Let me make the general perspective of this thesis precise. First, as indicated in the foregoing, it is concerned with inductive inference and the problem of induction. But in this context it has a specific aim, and it assumes a particular framework and a particular position therein. As for the aim, it is strictly normative. Inductive practice is only briefly discussed, and only to contrast practice with the norms that this thesis is concerned with. As for the framework, it is that of logical empiricism, as represented by Reichenbach (1935) and most notably Carnap (1950, 1952). In this framework, observations are considered to be clear-cut packages of information, which may be expressed in a formal observation language. Further, inductive inferences are cast in the form of probability judgements over this language. And finally, the inductive inferences are primarily concerned with predictions of single observations. On these three points, the present thesis adopts the framework of Carnap.

Before highlighting the differences with Carnap when it comes to the position of this thesis within the logical empiricist framework, it may be noted that this framework already pushes a number of philosophical positions on the problem of induction out of the picture. Some of those positions stress the theoretical content of observations over and above their empirical content. This surplus value can then be used to derive more from the observations than is warranted by their strict empirical content. The conclusion of this thesis picks up on this line of argument, but considerations on the nature of observations are not part of this thesis itself. Other positions on the problem of induction, such as the so-called structuralist position, rely not so much on the theoretical content of the observations, but on their structural aspects, which may then be connected to a theory on structures behind the observations. The present thesis relates to this possibility only indirectly. Still other positions do not employ the notion of probability, or eschew formal means altogether. Such alternative positions will not be dealt with in this thesis at all, but it may be remarked that certain forms of eliminative induction present a limiting case of the frameworks studied in this thesis.

Positioning this thesis. While the framework of this thesis is basically the one of Carnap, the position that will be developed is very different from the Carnapian position. The next few paragraphs highlight the main point of departure. For Carnapian inductive logic, valid inductive inferences are basically determined by the choice of an observation language. More specifically, probability judgements on future observations, or predictions for short, are derived by means of the notion of logical probability, where logical probability comes down to applying a principle of indifference to the observation language. It is supposed that before obtaining any observations, all the exhaustive descriptions of some system have the same epistemic status, and must therefore be assigned equal probability. On the assumption of this logical probability, both the initial predictions and the effect of accumulating observations on further predictions are determined by the structure of the observation language. The Carnapian idea is thus that the probabilistic predictions are analytic: they follow logically, namely according to logical probability, from the observation language and the preceding observations.

The inductive logic of Carnap was dealt a severe blow when Goodman (1954) proposed a new version of the problem of induction, calling it the new riddle of induction. There is hardly any need to reiterate the famed puzzle for its own sake. However, it provides a convenient way to make explicit the differences between Carnapian logic and the treatment of inductive inference in this thesis. For Carnap, all the work of induction is done by choosing the right predicates for the language, which are in the words of Goodman the projectable predicates. These projectable predicates select the weak or strong patterns that the inductive inferences focus on. If, for example, we choose to employ the predicate 'green' in a study on emeralds, we can derive predictions of green emeralds in the future, but if we employ 'grue', we can derive predictions of blue emeralds with equal force. Now the new riddle is not damaging for Carnapian inductive logic because the logic allows for crazy predictions, such as those on 'grue' and thus blue emeralds. Nobody has ever blamed deductive logic for generating crazy conclusions, since the responsibility for such conclusions lies in the premises. The damaging aspect is rather that Carnapian logic can only start working after a choice of language, and thus of projectability assumptions, has been made. It cannot itself express the choice of projectability assumptions as part of the inductive inference.

This is where the present treatment deviates strongly from Carnapian logic. Generally speaking, logic is concerned with the validity of arguments and not with the truth of conclusions of the arguments: if the premises are true, then so is the logically inferred conclusion, but there is no guarantee to truth if some of the premises are false, or perhaps not even well-formed statements. However, as indicated above, some substantial assumptions of inductive inferences cannot be expressed in Carnapian logic, simply because they are inherent to the observation language and its logical probability assignment. It thus seems that Carnapian logic provides not just valid inferences, but a number of implicit premises as well. Certainly, from the point of view of Carnap these premises are tautological, and therefore do not present implicit premises at all. But the new riddle makes perfectly clear that in fact they do. It is to resolve this seeming conflation of premise and inference that the present thesis presents an alternative logical scheme, following Ramsey (1921), De Finetti (1937) and Jeffrey (1984). It turns out to be perfectly possible to express projectability assumptions in an inductive logic, and thus to separate the part on valid inductive inference from the part on true inductive premises. The failure to disentangle these two aspects, so clearly separated in deductive logic, has obstructed a comparable development of inductive logic.

Other perspectives. At this point it is illustrative to consider an alternative approach to the problem of induction, which has only been touched upon implicitly so far. It is that the problem of induction presupposes a sceptical starting point that need not be accepted. It seems that the destructive conclusion of the problem is immediate once a bare language of observations is put in place: if set apart in that way, it is hardly controversial that the observations do not entail anything about each other. The answer of Carnap, if viewed from this angle, is to deny the sceptical starting point by employing a notion of logical probability over the language, thus creating an inherent dependence between observations. But note that he thereby reacts to the problem of inductive scepticism, or in other words, he is solving a problem in epistemology. Put more dramatically, Carnap seems to attack a problem in epistemology and one in logic at the same time.

Against this, I propose to consider inductive scepticism not as a problem in epistemology, but rather as a philosophical tool in logic. The tool allows us to analyze inductive knowledge in terms of observations and projectability assumptions, which must both be given a place in a scheme for inductive inferences. Moreover, after having settled the issue of valid inductive inference in a logical scheme, there is also a natural way to resolve the epistemological problem of inductive scepticism, by using an externalist theory in which inductive knowledge ultimately rests on the truth of inductive assumptions. I come back to this latter point in the conclusion.

The logical perspective of this thesis must not be mistaken for a rather different view on inductive logic, as developed in Maher (2004) and Fitelson (2005), which is in a sense closer to the initial intentions of Carnap. In this view inductive logic concerns an explication of the strength of the argument running from evidence to a hypothesis, or, in other words, the degree of confirmation that the evidence gives to the hypothesis. The position of Fitelson is that this degree of confirmation is objectively given, but further that it is a three-place function: next to evidence and hypothesis, it must include the probability model on which the confirmation relation supervenes. Unlike Maher, I agree that the probability model must be seen as a separate input component to inductive logic. However, the logic contained in this thesis does not assess the strength of arguments. Instead it simply classifies arguments as valid or invalid, and in this sense it may even be considered deductivist.

The function of Bayesian inference. Let me return to a sketch of the logical perspective of this thesis. In addition to the need for an expression of projectability assumptions as part of inductive inference, a truly logical view on induction is in need of one more thing: an innocent, or epistemically neutral, inference rule. By this I mean a rule that combines projectability assumptions and given observations to produce valid inductive predictions, or more generally, valid probability assignments, without entailing any substantial or synthetic assumptions itself. In a sense, asking for such an inference rule is equivalent to asking for a scheme that brings out all the assumptions that underlie inductive inference. The conclusions of inductive inference derive completely from the input components and the inference rule, so anything that is not implicit to the rule is driven back into the corner of the input components, and anything made explicit as input cannot hide away in the rule anymore. Employing an innocent inference rule seems to provide a natural insight into all input components of inductive inference, as conceived from within the chosen empiricist framework.

At this point, the Bayesian theory of probabilistic inference enters. Bayes' rule, or strict conditioning, prescribes how observations can be incorporated in a probability assignment over an observation language. Many arguments suggest that this rule is innocent in the required way, as long as we assign full certainty to the observations that we have made. The specific aspect of Bayes' rule that is significant here is that its use in incorporating observations induces minimal changes to the probability assignment. In other words, Bayes' rule is maximally conservative. If used to incorporate a specific observation, it takes care that no other change in the probability assignment is induced than those effected by deeming the observation itself certain. Furthermore, along the same lines, the Bayesian theory indeed determines the location of the projectability assumptions. While it is not yet clear in what form these assumptions can be stated, the assumptions must be implicit in the prior probability assignment over the observation language. Thus the Carnapian decision to choose a particular language and use logical probability is in the Bayesian scheme replaced by the decision to adopt a particular prior probability, which encodes the projectability assumptions.

Numerous Bayesianisms. It must be stressed that there is no unique Bayesian theory of inference. There are some common roots and standard texts on inductive inference and Bayesian statistics, most notably De Finetti (1936), Jeffreys (1951), Savage (1956) and more recently Howson and Urbach (1996). But there is certainly not a shared view on what it means to be a Bayesian.

Many of the quarrels among Bayesians come down to three related issues, to wit, the interpretation of epistemic probability, the origin of priors, and the basic form of the axioms of probability. Subjectivists take epistemic probability to be the expression of free personal opinion, and declare this to be the origin of all probability assignments. This point of view is associated with a further defence of the Bayesian theory and its inference rule, based on the relation between probabilities and betting contracts. Objective Bayesians, on the other hand, feel that there are certain rationality constraints on epistemic probability assignments, which may derive from physical probability or some principle of indifference. As for quarrels on the axioms of probability, subjectivism is sometimes associated with empiricist worries concerning probability assignments to opinions that cannot be expressed in finite form, while some objectivists have proposed to replace basic probability assignments with conditional probability assignments.

Bayes' Bayesianism. This thesis falls between all these positions. It employs inferences that have most in common with the inferences first put forward by the reverend Bayes himself. Statistical hypotheses occupy a central place in the original form of Bayesian inference. In the inferences, observations are first reflected in a probability assignment over statistical hypotheses, from which predictions on further observations can be derived. It may be noted that choosing a collection of such hypotheses restricts the probability assignment over the observation algebra. But more importantly, and as will be argued below, the hypotheses are related directly to the patterns in the observations that are considered to be of interest. In this way the hypotheses provide direct access to the projectability assumptions inherent in the prior probability assignment over the observations. The replacement of language in the Carnapian logic with a prior probability can therefore be made more precise. The choice of a range of projectable predicates in the Carnapian scheme can be replaced by the choice of a range of statistical hypotheses in a Bayesian scheme.

All this leads more or less to a middle position in between the above forms of Bayesianism. Statistical hypotheses are taken as so-called tail events in the observational algebra, and are defined by means of limiting relative frequencies. As for the interpretation and origin of priors, it may be noted that the use of hypotheses is connected to the dual nature of probability. On the one hand, the hypotheses pertain to weak patterns in the observations and thus to physical probability, and the restriction on the epistemic probability imposed by the hypotheses thus points to objective probability. On the other hand, the probability assignment over hypotheses is entirely free and reflects personal opinion, so it must somehow be interpreted subjectively. This is so even while it is difficult to connect the probability assignments over hypotheses to betting contracts, simply because statistical hypotheses cannot be tested with finite means. In sum, the Bayesian scheme presented in this thesis leads to a mixture of physical, epistemic, objectivist and subjectivist views on both the interpretation and origin of probability.

This blend of Bayesianism is much more natural than it may now seem. Bayes' original idea is precisely that epistemic and physical probability may be used in the very same inference, and that these two probabilities can coexist peacefully. It is a small step from this to the position that some epistemic probabilities are subjective, whereas others are restricted by physical probabilities and are thus objective, as in Jeffreys' principle of direct probability. Apart from that, the discussion on interpretation and origin loses some of its relevance once we recall that in the present thesis a prior probability assignment is an expression of a premise in an inductive inference. In view of this, both the hypotheses and the priors over them are instruments to express premises. Note that classical deductive logic does not set itself the task to clarify the exact world picture or conceptual interpretation that lays behind a premise or truth assignment. The task of logic only starts after the truth assignment has been given. Similarly, inductive logic need not fix conceptual categories for the probability assignments either. The main task is to investigate the inductive inferences themselves, and the instruments for expressing premises in them. Conceptual categories and interpretations are useful only insofar as they promote that task.

The use of hypotheses. As indicated, the central element in the Bayesian inductive schemes sketched is their use of statistical hypotheses. On this point the present treatment deviates most strongly from the empiricist and subjectivist views of respectively Carnap and De Finetti. Where Carnap localised the projectability assumptions in the choice of an observation language, this thesis makes the projectability assumptions explicit in the statistical hypotheses. Moreover, the hypotheses are introduced as an extension of the Carnapian observation language, as they are defined by means of limiting relative frequencies. Now the representation theorem of De Finetti revealed that statistical hypotheses are redundant in inductive schemes: they can be replaced by exchangeability requirements over the subjective inductive predictions. The present treatment takes the opposite view. It shows that hypotheses, even while they are redundant, are useful tools in expressing inductive assumptions and prior information. They provide a grip on a number of issues in the philosophy of science.

Philosophical import of this thesis. With these remarks on the position of the thesis and its relation to the Carnapian and Bayesian traditions in place, we can zoom out again and look at the overall relevance of this thesis. I first discuss its philosophical import.

With respect to the internal task of clarifying inductive inference, the Bayesian scheme presents a number of advantages over Carnapian inductive logic. A large part of this thesis is dedicated to making these advantages clear. As suggested, the Bayesian scheme provides a way of expressing and controlling the assumptions on the relevance of patterns in the observations. This ability allows us to solve a number of problems in Carnapian logic. A whole package of such problems relates to analogical reasoning. Here the Bayesian scheme allows us to take the package apart, and then to solve part of it. As it appears, this package is intimately related to another, seemingly different package of problems, namely that of encoding relations of probabilistic independence into inductive predictions. Furthermore, the Bayesian scheme is more readily applicable to current themes in the philosophy of science. In particular, it suggests a specific view on the problem of induction, it offers space to model dynamic changes in the projectability assumptions, and it sheds light on the role of theoretical notions in inductive inference.

This brings us to the relevance of this thesis for the philosophical discussion on induction and inductive knowledge. It is most easy to enter this discussion at the point of tension between Carnapian inductivism on the one hand, and the searchlight theory of Selz (1913) and Popper (1959) on the other. The former states that knowledge may be built up by observation alone, while the latter emphasises the importance of conjectures or theoretical starting points before collecting information. Hintikka (1966) made clear that this methodological distinction is not strict. With the perspective and scheme of this thesis, it becomes apparent that there is no methodological tension at all. The notion of conjecture may be combined with the inductivist point of view in a logical scheme, and in this scheme it is even seen to be indispensable. Put in more popular terms, Popper can finally be accepted as a member of the Vienna circle.

From this insight concerning inductivism we can move to the consequences for inductive knowledge, and the related theme of scientific realism. The proposed scheme can be used here to formalise a view that finds its roots already in Kant, and that connects my position with that of Kuipers (2000). It is the view that knowledge can only emerge on the intersection of observation, presented by a mind-independent world, and a conceptual framework, devised by, partly world-independent, minds. I hope that both radical constructivists and hardcore realists take this constructive realist message to heart.

Relevance of this thesis for science. Regarding the relevance of this thesis for science, first note that the empiricist framework accords well with the statistical inductive inferences of experimental science. In almost all experimental cases, weak patterns in observations are dealt with by means of statistics, and it is on these kind of inferences in experimental science that this thesis has its bearing. In this context, the aim of the thesis is not so much descriptive but normative, and more precisely, passively normative. The claim is not that scientists must, in all their investigations, follow the scheme laid down in this thesis. There may be practical reasons for using other procedures. However, scientists must eventually be clear on the exact inductive inference that they are making, and this they can find out by writing their procedures down in terms of the scheme provided here. In other words, they can check the validity of their procedures by writing them down in a Bayesian form.

This passive form of normativity indicates how the thesis relates to inductive inferences performed by means of classical statistics, as presented in Cramèr

(1946), Mood and Graybill (1973), Barnett (1999) and numerous other textbooks. It is well-known that classical statistics faces a number of paradoxes, some owing to the base-rate fallacy, others owing to the failure to respect the likelihood principle. But it is also well-known that procedures from classical statistics sometimes provide practical solutions where Bayesian statistics remains silent. Moreover, as far as the procedures of classical statistics are indeed inferences, they do not necessarily lead to false conclusions. It would therefore be misguided to advise scientists not to use classical statistics. On the other hand, the inferential steps in the classical procedures are often elliptic, or in other words, incomplete, and experimental scientists may not always be aware of the things they are presupposing when using these procedures. Classical statistics provides inferential shortcuts, whose applicability simply varies from case to case. The Bayesian reformulation of classical procedures can help to determine their applicability.

There is an enormous amount of literature on statistics, and many of the points made in this thesis have in some form or other been made elsewhere. Apart from the benefit of repeating the truth from time to time, the reader may wonder what innovations this thesis offers in the field of statistics. One innovative aspect is the connection of Bayesian inference with the Carnapian programme, and specifically, the use of statistical hypotheses in solving problems on analogy and inductive dependence. Another innovative aspect concerns the relation between Bayesian inference and problems with theory change and underdetermination. But perhaps the most important innovation is the use of frequentist chances in the Bayesian scheme. Statistical hypotheses can therefore be seen as part of an extended observation language, which allows for the integration of empiricist, subjectivist and frequentist views on probability.

Overview of chapters. I will now briefly run through the chapters, and indicate how their contents link up with the topics discussed in this introduction. The first three chapters form the first part of the thesis. This part concerns the reformulation of Carnapian inductive logic in terms of Bayesian logic and the improvement of the latter logic by the explicit use of hypotheses. Chapter 1 presents Carnapian inductive predictions as the conclusion of valid inductive inferences, and contrasts these with predictions deriving from the Bayesian scheme. Chapter 2 then deals with the nature and use of statistical hypotheses, and in particular with a frequentist semantics for hypotheses and the fundamental change they present to the Carnapian scheme. Chapter 3 argues that statistical hypotheses provide access to the projectability assumptions in the inductive inferences.

Chapters 4, 5 and 6 form the second part of the thesis. They concern the use of the Bayesian scheme in solving two problems in inductive logic, namely that of analogical predictions and that of causal relations or correlations between predicates. The general idea is that hypotheses provide a convenient handle on inductive dependencies between predicates, which prove hard to capture in terms of direct prediction rules. Chapter 4 shows that a natural system of prediction rules for capturing so-called explicit analogical predictions can be understood as the result of transforming a certain space of hypotheses. In chapter 5 these transformations are employed further to include analogical predictions of any kind, but unfortunately an exact match between the resulting predictions and the classification of analogical effects cannot be derived. Finally, chapter 6 employs the very same techniques to tackle the seemingly different problem of inductive inference for Bayesian networks. The mathematical structure of the problem turns out to be exactly the same.

The last part of the thesis is much smaller. It contains three short chapters on the Bayesian scheme in relation to venerable themes in the philosophy of science: the problem of induction, the problem of new theories and theory change, and the problem of underdetermination, which relates to abduction. It will be shown that these problems in methodology can be elucidated with the Bayesian scheme. In particular, chapter 7 investigates in what sense the Bayesian scheme solves the problem of induction. Chapter 8 proposes an addition to the Bayesian scheme that enables us to incorporate changes in the partition of hypotheses that are used in an inductive inference. Chapter 9 concerns the use of theoretical hypotheses in inductive inferences, and in this way provides a first sketch of a Bayesian model for abductive inference. The thesis ends with some general conclusions, and a perspective on further research.