UvA-DARE (Digital Academic Repository)

The logical response to a noisy world

Stenning, K.; van Lambalgen, M.

DOI

[10.1093/acprof:oso/9780199233298.003.0005](https://doi.org/10.1093/acprof:oso/9780199233298.003.0005)

Publication date

2010

Document Version

Author accepted manuscript

Published in

Cognition and Conditionals

[Link to publication](#)

Citation for published version (APA):

Stenning, K., & van Lambalgen, M. (2010). The logical response to a noisy world. In M. Oaksford, & N. Chater (Eds.), *Cognition and Conditionals: Probability and Logic in Human Thinking* (pp. 85-102). Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199233298.003.0005>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

To appear in: *The Psychology of Conditionals* edited by
Mike Oaksford. OUP.

The logical response to a noisy world

Keith Stenning

Human Communication Research Centre, Edinburgh University

Michiel van Lambalgen

Department of Philosophy, Amsterdam University

January 5, 2012

1 Introduction

We live in a noisy world and our premises are subject to doubt. There is a thriving school of opinion in the psychology of reasoning that this means that it is both desirable, and inevitable for realism's sake, that theories of human reasoning should be framed in probability theory rather than mental process theories in logical frameworks; a prominent recent example is Oaksford and Chater [15]. A forceful earlier statement of this position is

... it is, in fact, rational by the highest standards to take proper account of the probability of one's premises in deductive reasoning ... performing inferences from statements treated as absolutely certain is uncommon in ordinary reasoning. We are mainly interested in what subjects will infer from statements in ordinary discourse that they may not believe with certainty and may even have serious doubts about ... [26, p. 615]

This approach assumes that reasoners come with their interpretations of materials already fixed. How otherwise could they have any ideas about the probabilities

of component propositions? When subjects hear, to take a famous example we will return to, “If she has an essay, she is in the library” how could they have any idea about a probability without knowing anything about who ‘she’ might be or anything else about the situation? As Cummins [6] has shown, there is a strong correlation between subjects’ willingness to draw for example a modus ponens inference from such conditionals, and the ease with which they can retrieve ‘additional’ and ‘alternative’ conditionals, such as “If the library is shut, she isn’t in the library”, or “If she has a reference book to read, she’s in the library”. Does this mean that subjects come to such experiments with an estimate of the probability of all such propositions which they merely need to retrieve? Or does it mean that there is a process of discourse interpretation which can explain how discourses composed of these conditionals give rise to interpretations on the basis of which subjects are willing to hazard likelihoods? A process of interpretation, moreover, which is endlessly sensitive to subtleties of the kind of discourse subjects believe they should engage in? Does Cummins’ result suggest, rather than an archive of probabilities, that subjects estimate likelihoods from the ease of retrieving schematic knowledge about alternatives and defeaters? At least in the situations where they construe the task as interpreting a kind of fiction about typical students and typical libraries? Our most basic claim is that what subjects cannot do is to engage in *no* process of interpretation. We cannot study some ‘central’ process of inference, without studying its ‘front-end’ process of interpretation.

So, in contrast to the probabilists, we believe that the loudness of the world is one important reason why human reasoning is a two component process [25]. We impose interpretations on discourse before we can reason from those interpretations, and we generally start by reasoning to interpretations as if their assumptions were perfectly certain *within those interpretations*. That is, our guiding question as discourse comprehenders is “What would it be like for the speaker’s discourse to be absolutely true?” Only when we have constructed an intended interpretation, can we derive any consequences *from* the interpretation, or decide how it relates to the world. Needless to say, we believe the twin theories of reasoning *to* interpretations and reasoning *from* them must be couched in logic, though in very different logical frameworks than have hitherto figured in the psychology of reasoning. Our aim in this paper is to motivate such a logical account, by showing that probability cannot provide the kind of defeasible framework within which interpretations can be established. Our positive logical account has appeared elsewhere [23, 25]; here we concentrate on the shortcomings of probabilistic analyses of defeasible reasoning. We start by discussing a number of faulty assumptions about logic that have hindered its application to actual human reasoning.

1.1 Dusting off logic

The probabilistic critique of ‘logician’ accounts of the psychology of reasoning trades on several assumptions about logic that may perhaps be reasonably applied to logic’s applications in some psychological theories (mental logics, mental models, ...), but are quite foreign to a modern logical approach to human reasoning and discourse. First there is an assumption that logic is to be identified with classical logic whose logical forms can be read directly from natural language sentence forms. Second, the interpretative component of logic is ignored at the expense of total concentration on the derivational component. Third, it is assumed that using probability theory as a basis for rational analysis does not involve using a logic, and so avoids the criticism levelled at ‘logician’ cognitive science.

The first assumption, that logic must be identified with classical logic, is unfortunately still rather common. On the formal side, it overlooks the great wealth of logical systems currently explored. More to the point, it unjustifiably takes classical logic to be the only logic with normative force [25, Chapters 2, 11]. Once one drops this assumption, we must assume instead that the process of interpretation of natural language sentences is a substantial one, and is best viewed as involving the setting of parameters not only for the syntactic form of natural language sentences, but for their semantics, and for the concept of validity inherent in the reasoning task at hand. The end result of parameter setting is to fix on what we call a *logical form* [25, Chapters 2]. Here we will concentrate on the contrast between defeasible logics of interpretation and classical logics of derivation, but many other logics are often involved. A notable further example is deontic logic as involved in some selection task variants [22].

The second assumption was that logic is to be identified with its formal, derivational component. The probabilistic approach identifies ‘logician’ approaches with certainty of conclusions given premises and sees the central advantage of probabilistic approaches as their capacity to handle uncertainty. But even traditional logic always assumed that uncertainty was to be accommodated in its interpretative component. Traditional logic (that is, logic as it was practised and taught before the era of formalisation) had no formal account of interpretation but, as we shall see, it shares this property with current probabilistic approaches, and besides traditional logic had a substantial informal account of the criteria for coherent interpretation. This problem perhaps arises from identifying logic with its applications in the foundations of mathematics, a late 19C employment. Traditional logic was developed as a theory of the discourse of argumentation – a subject matter as prone to uncertainty as it is possible to find.

On the third assumption, that the relevant properties of logic do not apply to prob-

abilistic approaches, one must counter that probability theory simply is a (family of) logic(s) – probability logics – and as such qualifies as a target of all the criticisms levelled at ‘logician’ approaches. In fact probability logic is a deductive theory which vastly extends the range of *certainty* of inference. For example, classical probability theory assures us that if the probabilities of X and Y are both .5 and are independent, then the probability of $X \wedge Y$ is .25—exactly. Even when probability theories only licence inferences of probability intervals rather than point probabilities, they often give exact bounds on the confidence intervals of our conclusions.¹ Everywhere uncertainty retreats before a tidal wave of exactitude. We agree entirely that handling uncertainty is of great importance, but we feel that this militates strongly against probabilistic frameworks. We agree entirely that we must give accounts of how content affects reasoning, but for precisely this reason we propose that a ‘defeasible logic of interpretation’ is required to explain reasoning over long-term knowledge *to* an interpretation.

In the next section, we will provide a brief sketch of our own positive logical approach through its account of the suppression task in the next section [5, 23, 25]. The last reference gives an overview of the reanalysis of several psychology of reasoning tasks in this interpretative framework: the selection and suppression tasks, categorial syllogisms and some non-linguistic executive function tasks. In the middle section we will examine the Oaksford & Chater probabilistic account of the suppression task and show, first that they too require an additional initial interpretative component which cannot be couched in probability theory, and second, that even when this interpretative component is added, the resulting probabilistic treatment cannot cover all the argument patterns which people readily employ.

2 A logic of discourse interpretation

Whilst traditional logic had a substantial but informal understanding of interpretation it had no formal account of the reasoning involved. Its formal accounts, such as they were, were of derivations within an interpretation – examples of which are known and loved by every introductory student. Interpretation figured mostly in the injunction that it must be held constant throughout an argument – no equivocation (surreptitious shifting of interpretation) was to be allowed. Since, with the advent of symbolic logic, teachers were often content to teach the formal manipulation of the symbols within systems, the process of interpretation was backgrounded in introductory courses, often to the point of disappearance.

¹There are important differences among probabilistic approaches (for example, [18] and [15] differ significantly), but they share the property of assuming prior interpretation.

Since the 1970s, artificial intelligence researchers have provided formal accounts of defeasible logics of interpretation [13, 20]. Defeasibility, the capacity to withdraw earlier conclusions on the arrival of new evidence, is clearly a criterial requirement for any process of interpretation, even though it is equally obviously anathema to derivational processes. Unfortunately, despite their invention by those interested in computer implementation, the motivating belief that these ‘natural’ logics must be easier for people to use turned out to be disappointed. These logics are provably seriously computationally intractable – more so than classical logic which was already known to be highly intractable. Out of the frying pan into the fire. This impasse had a baleful effect on the application of defeasible logics in psychology.

Fortunately, in the eighties and early nineties, the more tractable alternative logics of interpretation which we will use here became generally available. It is of some interest that they originated obliquely from the very practical research program of adapting logic as a computer programming language, issuing in languages such as PROLOG. Whilst the practitioners saw themselves as using fragments of classical logic to achieve efficient computations, later logical study of the resulting systems revealed that they could more perspicuously be conceptualised as computing in defeasible logics. The particular logic we will introduce here is known as ‘logic programming with negation as failure’.

Needless to say, the tractability of these new defeasible logics is purchased at a certain price, in this case, as one should expect, in expressiveness of the languages concerned. Logic programming [7] restricts conditionals to ones of the form $L_1 \wedge \dots \wedge L_n \rightarrow A$, where A is atomic, and the L_i are either atoms or negated atoms.² However, we believe this loss of expressiveness is well worth the price. The fact that the resulting languages can be used in programming is a good indication that they are expressive enough to get lots done, and some of the restrictions (such as that against some syntactic iteration of the implication) also arguably apply to natural language conditionals. We think this is a good compromise balance between expressiveness and tractability. This is a stark contrast to the intractability of probability theory which is acknowledged also by those advocating the probabilistic approach. This acknowledgment forces them back to a position where probability is a distant competence theory whose relation to the data we will argue is hard to understand, and whose implementation they have little to say about. Our defeasible logic can be shown to have a direct neural implementation and so brings competence and performance dramatically closer together [23], [25, Chapter 8].

²The restrictions on the antecedent can be liberalised somewhat, but it is not allowed to have conditionals in the antecedent.

These defeasible logics specify valid patterns of reasoning over *databases* of default conditional rules. We conceive of these rules as representing long term general knowledge of environmental regularities, and of connections between the agent's beliefs and his actions.³ Since we are concerned with modelling discourses, we conceive of the new statements of the discourse arriving sequentially and being incorporated into the 'hearer's' working memory (WM) which is conceived of as the activated part of long term memory, rather in the style in which ACT-R would model this relation (Anderson and Lebiere [2]).

We model conditionals whose surface form is 'if p then q ' as *default rules*, which are defined as logic programming clauses of the form: $p \wedge \neg ab \rightarrow q$ read as "If p and nothing is abnormal, then q ".⁴ The abnormality (and possibly other propositions as well) is governed by *closed world reasoning* (CWR), which says that if a proposition is not known to be true, it may be assumed false. Thus, if there is no information that an abnormality will occur, we may assume that none will occur, and the above conditional reduces to $p \rightarrow q$. CWR is the logic that one often applies in planning and prediction. There is an unlimited number of events – abnormalities – which could interfere with our goal to be in London on August 28, 2007 (ranging from major natural or man-made disasters to strikes, to waking up too late and missing the plane). The vast majority of these events are not accounted for in our plan to go to London at said date – that is, they are treated as if they will not occur. CWR leads to a *non-monotonic* logic. New information may become available which forces us to consider an event as an abnormality which was not previously classified as such.

With this much logical background we can introduce the main reasoning task to be studied in this chapter, the suppression task (Byrne [5]). When subjects are presented with the modus ponens material "If she has an essay she studies late in the library. She has an essay." they almost universally draw the conclusion she studies late. When instead they are presented with the same premises *plus* the premise: "If the library is open she studies late in the library", about half of them withdraw the inference. Byrne concludes that inference rules cannot be used to explain the former performance since they evidently cannot explain the latter non-inference. Rules, if they are to be invoked must be invoked universally and uniformly. So much the worse for mental logics – roll on mental models, concludes Byrne.

³Stenning and van Lambalgen [24] discusses the role of default conditionals in 'executive function', an umbrella term for processes involving planning of actions to achieve a given goal.

⁴Strictly speaking we would have to write $p \wedge \neg ab(p, q) \rightarrow q$, a notation that emphasises that the abnormality is specific to the conditional. We ask the reader to remember that this is the intended interpretation and we will continue using the less cumbersome notation.

Stenning and van Lambalgen [23] analyse this task in terms of CWR. In the condition of the suppression task where we are given that ‘she has an essay’ and ‘if she has an essay she’s in the library’, the argument proceeds as follows. The underlying logical form of the premise set is $\{p, p \wedge \neg ab \rightarrow q\}$. By itself this does not justify the modus ponens inference, but now we invoke CWR, in the guise of the assumption that since no abnormalities are mentioned, to conclude that $\neg ab$ is in fact true; whence q follows.

When, in the second condition of the experiment we get the extra premise that ‘if the library is open she is in the library’, it is conceived of as triggering the relevance of an abnormality – namely that the library may be closed (along with the relevance of other general knowledge conditionals such as ‘if the library is closed, readers are not inside’ etc. etc.). Now we do have information that something may be abnormal, and from $p \wedge \neg ab \rightarrow q$ and ab we can no longer conclude q , though we, along with lots of subjects, may be happy to make the weaker conclusion that ‘if she has an essay *and* the library is open she’ll be in the library’.

Because of closed world reasoning over abnormalities, some patterns of reasoning are valid in this logic which are invalid in classical logic, notably affirmation of the consequent and denial of the antecedent. So from ‘If she has an essay she is in the library’ and ‘She doesn’t have an essay’, it is valid to conclude that ‘She isn’t in the library’, but again only by closed world reasoning. This is a valid move if this is all there is in the database concerning essays or libraries, because we can conclude that this is the only conditional with the relevant consequent. If we then add the ‘alternative’ premise ‘If she has a textbook to read she’s in the library’ then this defeats the conclusion because the world is now large enough to contain another explanation for her presence in the library. Similarly, affirmation of the consequent is valid in this logic in the right context. From ‘If she has an essay then she’s in the library’, and ‘She is in the library’ we can conclude that ‘she does have an essay’, but only providing the world is closed against abnormalities such as the library being closed. For a fuller exposition of the suppression task treated in this logic see [23].

As an aside which we return to below, it is worth raising the question “What is worth modelling in the suppression task?” Stenning and van Lambalgen [23] show that each of the four conditional inference forms can be valid in the proposed logic, depending on context, and therefore why nonmonotonic revisions of conclusions should or should not occur with the arrival of different kinds of premises. It is not unreasonable to fit this to group data such as Byrne’s, showing that indeed the changes in proportion of subjects making or retracting conclusions fits well with the logical model. But it is important we don’t forget that this is a very weak sort of data. What we would really like is longitudinal data on the

same subject's revisions of belief in each of the conditions studied – the model is about belief revision, not about the development of group belief. Of course there are problems of interference between conditions measures which forced the original design. Our point is that other methods of investigation need to be brought to bear. Lechler [11] used Socratic dialogue to show that the range of interpretations subjects adopted was much wider than allowed for in Byrne's analysis, and that the classification of premises as 'alternative' or 'additional' was far from reliable.⁵ So the logical model as it stands only models a particular small range of the interpretations people adopt, but it is also clear, even from Byrne's data, that different subjects are doing different things and it makes little sense to model the details of their average mental processes.

We mention here in passing that there is a highly significant difference in performance between autists and normal controls, in that the former do *not* engage in closed world reasoning with respect to abnormalities, although they are capable of other forms of CWR. Thus autists suppress MP and MT much less, although they are indistinguishable from normals with respect to AC and DA. See Stenning and van Lambalgen [24] for more details.

Defeasibility is achieved in this logic by closed world reasoning, and this is central to why the logic is tractable. In fact the semantics assigned to the logic guarantees that a unique minimal 'intended' model (valuation of all the atomic propositions) for the discourse is derivable at each step of incorporation of a new discourse sentence. The semantics is three-valued (true, false, so-far-undecided) but not truth functional and the concept of validity is not the classical 'truth of conclusion in *all* models of the premises' but rather 'truth of conclusion in the intended model of the premises'. The non-truth functionality of the system is clear for all to see. The defaults in long term memory cannot be shown to be false because of their robustness to exceptions invoked as abnormality conditions. Genuine counterexamples would require relearning – repair of the database which is a process outside this logic. None of this is surprising as long as we keep in mind that this is a logic of *credulous* discourse processing in which our goal as hearer is to make the speaker's statements true. that is, to find a model that makes them true. The three-valued semantics is required to allow negation-as-failure ($\neg p$ is adopted as true if we fail to prove p).

Interestingly, this three-valued semantics allows a neural network implementation of the logic in which a feedforward spreading activation network computes the intended model of a discourse in linear time [25, Chapter 8]. It should be obvious that this is not a logic of cogitation – slow deliberate derivation of conclusions. In psychological terms it is rather a logic of System 1 processes [8] – fast automatic

⁵Stenning and van Lambalgen [25, Chapter 7] features extracts of Lechler's data.

processes of inference operating over general knowledge, often well outside of conscious control, or indeed awareness – exactly appropriate for the processes of discourse understanding for simple well-wrought discourse.

2.1 So why is discourse processing still hard?

Computational linguistics tells us that discourse processing is hard because discourses are massively ambiguous and require continual backtracking in their processing. How can a defeasible logic deliver a unique minimal intended model at each stage? We do not pretend to have solved the computational linguistic problem. What we are claiming is that the human problem is solved by the mobilisation of general knowledge by the human processor. Human processors do not normally experience these ambiguities because they can apply commonsense knowledge to avoid them. The logical theory shows how this can happen in example cases but does not explain how LTM can be organised on a global scale to achieve this in a computer. Baggio, van Lambalgen and Hagoort [3] show how ERP data can provide some evidence that the invocation of abnormalities in the logic leaves traces in brain activity in psycholinguistic experiments.

Let us take an example that is used by Oaksford and Chater to argue against the applicability of defeasible logics:

Birds fly.
One-minute-old birds don't fly.
Can the bird Tweety at one-minute-old fly?

Here there is a conflict between defeasible rules. Our response would be that defeasible logical databases allow the mobilisation of general knowledge ‘theories’ to resolve such impasses and this gives a plausible explanation as to how humans resolve them. The database also contains conditionals such as *Birds hatch from eggs*, *Helpless chicks take weeks to mature*, and lots of other linking materials ... which are sufficient to arrive at the conclusion *So Tweety can't fly, yet*. *Contra* Oaksford and Chater, this tractable defeasible logic really does solve some versions of the frame problem (Shanahan [21]).

This problem of conflict between two applicable rules is only one kind of impasse. The other kind of impasse the credulous discourse processor may hit is where no models can be found i.e. where the speaker's utterances appear to contradict each other. *Socrates is mortal*. *Socrates lives forever in the dialogues of Plato*. Given a small amount of inference about living for ever and immortality,

we have a contradiction. No model is available. Here, we propose, the credulous processor interrupts its activity while classical reasoning attempts to repair: “Something is wrong. There must be equivocation in the interpretation. We must find an interpretation in which these two statements are consistent.” The intended (unique) model of the discourse thus far, now has to be inflated into a set of classical models representing various possibilities. The single truth valuation of the atoms of the discourse fragment define a set of all possible valuations. Now we have to find some modification of this interpretation which will allow an intended model in which the rogue statements are both true. Is the first Socrates perhaps the Brazilian football player? Or are there two senses of immortality involved? The reasoning is only classical in the sense that we *may* have to explore all possible classical models in order to repair the impasse by finding some one that will do. Failure to do so will turn the discourse from credulous cooperative to classical adversarial, at least until some mutually acceptable repair is found.

Our claim is that we have provided a logical framework for interpretation whose known properties make it plausible that these problems can be solved and will scale up. Much of the empirical psychology remains to be done. Note that in this account of ‘suppression’, nothing is suppressed, and the account is completely neutral as to whether the reasoning is done over syntactic rules or over models. Both proof-theoretic and equivalent model-theoretic expositions can be given. The task cannot provide any evidence whatsoever for resolving debates about rules versus models.

This concludes our survey of what we believe is the proper setting of the suppression task: defeasible reasoning on discourse interpretations. We now turn to a critical examination of a very different probabilistic analysis of the same task, that of Oaksford and Chater [14] (reprinted in [15]). Our purpose here is general however: we believe the potential of probability theory to model defeasible reasoning has been overestimated, and the suppression task allows us to show why.

3 A probabilistic model for the suppression task

Oaksford and Chater [14] attempt to show that a probabilistic model can account for the observed suppression effects in conditional reasoning. We will carefully analyse the assumptions underlying the model, because they are fairly typical of applications of probability to human reasoning. These assumptions can be divided into four broad categories:

- I. philosophical assumptions forcing a close connection between uncertainty and probability (while denying a role for logic)
- II. largely implicit assumptions about what is and what is not to be included in the model
- III. coordinating definitions linking the probabilistic model to the task at hand
- IV. assumptions concerning what counts as validation of the model.

In our discussion below we indicate the category of assumptions at issue by the corresponding Roman numeral.

3.1 Probability and logic

The following quote from [14] provides much material for reflection on I.:

[M]uch of our reasoning with conditionals is uncertain, and may be overturned by future information; that is, they are non-monotonic. But logic based approaches to inference are typically monotonic, and hence are unable to deal with this uncertainty. Moreover, to the extent that formal logical approaches embrace non-monotonicity, they appear to be unable to cope with the fact that it is the content of the rules, rather than their logical form, which appears to determine the inferences that people draw. We now argue that perhaps by encoding more of the content of people's knowledge, by probability theory, we may more adequately capture the nature of everyday human inference. This seems to make intuitive sense, because the problems that we have identified concern how uncertainty is handled in human inference, and probability is the calculus of uncertainty [14, p. 100].

The argument sketched here is that logic is not useful as a model of reasoning, because it is either monotonic, in which case it flies in the face of data such as suppression, or non-monotonic, in which case it is closer to the data, but still far off, because it cannot account for the difference between additional and alternative conditionals, which is supposed to be one of content not form. Probability fares much better in this respect, because it can incorporate changes in content as changes in the values of probabilities. Furthermore, and most importantly, there are *a priori* arguments to show that probability theory is *the* normatively justified calculus of uncertainty. We will discuss this issue first.

3.1.1 Probability and uncertainty

The standard justification for probability – interpreted subjectively – as representation of uncertainty proceeds via so-called ‘Dutch Book arguments’, a simplified version of which goes as follows.⁶ Assume given a sample space X and a set of events $\mathcal{A}$ on X , which has a Boolean structure. For ease of exposition we take $\mathcal{A}$ to be the powerset of X . What is important to note, however, is that the Boolean structure of the set of events is given not derived; in other words, one must assume that the events satisfy classical logic. This severely restricts the kinds of uncertainty to which probability theory can be applied.

In the next step it is assumed that one’s degree of belief in the occurrence of an event E in $\mathcal{A}$ can be determined via betting. In very rough outline, the agent’s degree of belief in E equals the price of a promise to pay 1 if E occurs, and 0 otherwise. The phrase ‘a promise to pay’ refers to a subtlety in the wager: the bookmaker decides *after you have set the price* which side of the bet will be the agent’s. Thus, the agent knows that the bookmaker will either buy such a promise from him at the price set, or will require the agent to buy such a promise, again at the price set by the agent. The price set by the agent is the subjective probability assigned to E . A *Dutch Book* is a betting scheme which guarantees a gain for the bookmaker. A famous foundational result in subjective probability then says that the following are equivalent

1. no Dutch Book can be made against the agent
2. the agent’s degrees of belief satisfy the axioms of probability

Since it is supposedly irrational to engage in a bet in which you are sure to lose, the result is often glossed as: a rational assignment of degrees of belief is one which satisfies the axioms of probability theory. This is a so-called *synchronic* Dutch Book theorem; the *diachronic* Dutch Book theorem is concerned with Bayesian updates of beliefs and justifies the rule of Bayesian conditionalisation

(BaCo) The absolute subjective probability of event D given that evidence E has been observed is equal to the conditional probability $P(D \mid E)$. Here E should encompass *all* available evidence.

(BaCo) will be important in constructing probabilistic analogues for the standard propositional inference patterns MP, MT, AC, DA.

⁶For an excellent exposition, see Paris [17].

An objection against justifications of this type is that the agent is all-knowing and ‘all-willing’ as well: he must be willing to bet on any event in $\mathcal{A}$ and he must fix a unique price for the promise, instead of, say, an upper and a lower bound. While these idealisations may be alright in some philosophical contexts, they are definitely out of order when it comes to cognition. It is not excluded that agents reason with subjective probabilities in some situations, but there is no guarantee that they assign probabilities to every event of interest.

There exists a way out of these problems, at least technically, by considering finite approximations to the space of all events. To fix ideas, suppose the space of all events corresponds to the denumerable set of propositions $p_0, p_1, p_2, p_3, \dots$. An agent is aware of only finitely many propositions at any one time, but over time this set may increase. For the sake of argument we may suppose that the above betting procedure assigns a joint distribution P_n over $\mathcal{A}_n = \{p_0, \dots, p_n\}$, for all $n \geq 1$. In principle, however, the $P_n(p_i)$ can be different for all n , and hence there is no guarantee that there exists a $\lim_{n \rightarrow \infty} P_n$ which defines a joint distribution over the whole space $\{p_0, p_1, p_2, p_3, \dots\}$. The reader may wish to consider the special case where the P_n are actually truth value assignments to the $p_0, \dots, p_n$; in this case $\lim_{n \rightarrow \infty} P_n(p_i)$ exists only if the truth value of p_i changes finitely many times. That is, incoming information about propositions not considered so far may non-monotonically change the truth value of p_i , but only for finitely many propositions – the last change must be monotonic.

The general case in which the P_n are probability distributions is a difficult mathematical problem (Rao [19]), much studied in the Bayesian literature, for example in connection with ‘hierarchies of beliefs’ (see Brandenburger and Dekel [4] for an interesting example). In order for the limit $\lim_{n \rightarrow \infty} P_n$ to exist for all events, a complicated coherence condition must be satisfied. A special case of this condition will be of interest later when discussing the suppression task. Suppose P_2 is the joint distribution over p, q and P_3 the joint distribution over p, q, r . The coherence condition then implies the following relation among conditional probabilities:

$$P_2(q \mid p) = P_3(q \mid p \wedge r)P(r \mid p) + P_3(q \mid p \wedge \neg r)P(\neg r \mid p). \quad (1)$$

This property is the probabilistic analogue of monotonicity.

Summing up the discussion so far: we have seen that many assumptions are necessary to conclude that probability theory is the uniquely designated formalism to deal with every form of uncertainty. In fact the assumptions can be glossed by saying that the application of probability assumes quite a bit of certainty: the structure of events must be of a very specific type described by classical logic, the

agent must be certain about the price he wants to set for any bet, and the agent must also be certain that his future probabilities conform to the coherence condition. This is too much to ask for a real cognitive agent. Of course, we do not deny that the brain uses *frequency* information about the environment – e.g. the frequency with which a given word occurs – in processing. We just doubt that, where frequency information is lacking, subjective probability can always do duty for frequencies.

It is therefore unclear why subjective probability can claim the status of a computational model (in the sense of Marr) for dealing with uncertainty [14, p. 103]. Note that Marr did not intend his computational level to be an idealised model only. This level specifies the inputs and outputs – and their relation – of a given cognitive function, and thus functions as a declarative specification of the algorithm at the next level. Its character as specification of an algorithm means that the computational level is not just idealisation and is connected in a lawlike manner to the algorithm.

3.2 Setting up the model

Of course, for the purposes of the argument we grant Oaksford and Chater the assumption that subjects do indeed assign degrees of belief to the events of interest in the suppression task. This assumption (of type III.) can be divided into two components:

1. reasoning with a conditional $p \rightarrow q$ is *de facto* reasoning with the conditional probability⁷ $P(q \mid p)$ – which is assumed to be known exactly
2. subjects assign precise probabilities to atomic propositions

Three parameters then suffice to determine a joint probability on $\{p, q\}$: $a := P(p)$, $b := P(q)$, $\epsilon := P(\neg q \mid p)$. This last parameter is the ‘exception parameter’, which models the defeasibility of the conditional $p \rightarrow q$. We will see below that suppression of MP and MT in the presence of an additional conditional is explained via modulation of ϵ , but we first discuss the two premise inference patterns. MP is viewed as (BaCo) applied to $P(q \mid p)$. That is, if p is given as the categorical premise, it is assumed this premise constitutes all the available evidence, so that (BaCo) can be applied. This way of justifying probabilistic MP gives it a nice non-monotonic flavour, for if in addition to p the categorical

⁷As a consequence, some reasoning with iterated conditionals cannot be represented, as is also the case in closed world reasoning.

premise r is given, (BaCo) can no longer be applied to $P(q \mid p)$. Thus, it would seem that a probabilistic analysis is in principle suitable to model the suppression task.⁸ In a similar way, MT is viewed as (BaCo) applied to $P(\neg p \mid \neg q)$, AC as (BaCo) applied to $P(p \mid q)$, and lastly DA as (BaCo) applied to $P(\neg q \mid \neg p)$. The three parameters a, b, ϵ suffice to generate the required conditional probabilities via the probabilities of conjunctions. For example, $P(p \wedge \neg q)$ is computed as $P(p)P(\neg q \mid p) = a\epsilon$, $P(p \wedge q) = P(p)P(q \mid p) = a(1 - \epsilon)$, $P(\neg p \wedge q) = P(q) - P(p \wedge q) = b - a(1 - \epsilon)$. It follows that e.g. $P(p \mid q) = \frac{P(p \wedge q)}{P(q)} = \frac{a(1 - \epsilon)}{b}$, and so on for the other conditional probabilities.

3.2.1 Validating the model for two premises

Why is it plausible to assume that subjects reason according to such a probabilistic model? Oaksford and Chater attempt to show this by fitting the model to characteristic data on two premise inferences (here taken from Byrne [5]). The fitting⁹ proceeds as follows: ϵ was set to .1, and values of a, b were sampled¹⁰ in the interval [.1, .9].¹¹ The values of the four conditional probabilities were averaged over the a, b in the sample. Just to give an example, the value of the average for $P(q \mid p)$ obtained in this way was 90.00. As a number, this is not dramatically different from the 97.36% rate of endorsement in Byrne's data, which are also averaged over subjects. This is a validation of sorts,¹² but it would be much stronger if the model were fitted to individual subjects – after all it's they who are supposed to reason with subjective probabilities. Further principled objections to this procedure will be discussed in section 3.3.3.

⁸This is not the justification given in [14], where it is simply observed that subjects, when given p , seem to rate the probability of q as proportional to $P(q \mid p)$. The present exposition makes it clearer why a probabilistic analysis of the suppression task is *prima facie* plausible.

⁹The assumptions used here are of type IV.

¹⁰If the probabilities a, b were frequencies, there would be no need to sample, because they would be uniquely determined by the environment.

¹¹Actually the sampling space was also restricted by the requirement that $a < b$. This is an assumption which should also count as a parameter of the model, so that we now have four parameters and four data points. The requirement is somewhat peculiar in a context where $P(q \mid p) < 1$. Note also that the rarity assumption prominent in Oaksford and Chater [16] is now dropped, since $a, b > 0.1$, because '[conditional] inferences are specific to context' [14, p. 103].

¹²If the problem indicated in footnote 11 can be solved.

3.3 Incorporating additional and alternative conditionals

There are two motivations for using probability theory as a model of reasoning in the suppression task. One is that Bayesian conditionalisation is itself a non-monotonic principle: an extension of the evidence p may invalidate a previous posterior probability for q derived by (BaCo) applied to $P(q \mid p)$.¹³ The second motivation is that the conditional probability $P(q \mid p \wedge r)$ may differ from $P(q \mid p)$, so that strengthening the antecedent of a conditional has a real effect. Thus there is face validity in the attempt to model suppression via change in conditional probability. As we will see, however, it is not straightforward to give a probabilistic account of the processes involved in changing the conditional probabilities.

There is also a weak probabilistic explanation of the suppression task, which is content with observing that because (BaCo) is a non-monotonic principle, expansion of the original evidence p with an additional conditional invalidates the probability assignment to q , so that suppression of MP is expected. Such an explanation would need to show how (the antecedents of) additional conditionals differ from (the antecedents of) alternative conditionals, and in any case requires an account of how a conditional, i.e. a conditional probability, can be incorporated into the evidence. Since Oaksford and Chater do not take this road, we shall not explore it either.

3.3.1 Additional conditionals

It is worth quoting Oaksford and Chater in full on suppression of MP:

Additional antecedents (or exceptions), for example that there is petrol in the tank with respect to the rule *if you turn the key, the car starts*, concern the probability of the car's not starting even though the key has been turned – that is, they concern ϵ . If you do not know there is petrol in the tank, you cannot unequivocally infer that the car will start (MP). Moreover, bringing to mind additional factors that need to be in place to infer that the car starts – for example the battery must be charged – will increase this probability [i.e. ϵ]. [...] [I]f there are many additional antecedents, that is, ϵ is high, the probability that the MP inference will be drawn is low [14, p. 104].

¹³This assumes that the additional conditional can itself be regarded as (probabilistic) evidence on which it is possible to conditionalise.

In more formal terms, the argument for the suppression of MP seems to be the following. Initially the subject works with the conditional probability $P(q | p)$ derived from a joint distribution on $\{p, q\}$. The antecedent of the additional conditional enlarges this space to $\{p, q, r\}$ (where r stands for ‘there is petrol in the tank’), and this leads to a new representation of $P(q | p)$. One may write

$$\begin{aligned} P(q | p) &= \frac{P(p \wedge q)}{P(p)} = \frac{P(p \wedge q \wedge r)}{P(p)} + \frac{P(p \wedge q \wedge \neg r)}{P(p)} = \\ &= \frac{P(p \wedge q \wedge r)}{P(p)} = \frac{P(p \wedge q \wedge r)}{P(p \wedge r)} \frac{P(p \wedge r)}{P(p)} = P(q | p \wedge r)P(r) \end{aligned} \quad (2)$$

where the last equality follows under the plausible assumption that p and r are independent.

In an orthodox Bayesian approach, $P(q | p \wedge r)P(r)$ is simply a different representation of $P(q | p)$, but the values of these expressions must be the same, since it is assumed that the subject had access to the joint distribution over $\{p, q, r\}$ in computing $P(q | p)$. But this is not how Oaksford and Chater use probability. They assume that the subject assigns a lower probability to $P(q | p)$ in the enlarged representation $P(q | p \wedge r)P(r)$; the quote suggests that this is because the subject lowers $P(r)$ from 1 to a value smaller than 1 when becoming aware of possible exceptions. This requires the kind of transition between algebras of events that we studied in section 3.1, but now we are in great difficulties, because such transitions must be governed by equation 1, which in the case at hand boils down to

$$P_2(q | p) = P_3(q | p \wedge r)P(r). \quad (3)$$

Thus, in Oaksford and Chater’s model the subject must change his probabilities in a manner that *conflicts* with the Bayesian desideratum of striving toward a coherent probability distribution over all events. As a consequence, no Bayesian explanation of the transition $P_2(q | p)$ and $P_3(q | p \wedge r)P(r)$ can be given – this transition must remain outside the model (a type II. assumption). In order to account for the phenomena the model therefore needs to be supplemented with a theory about non-monotonic changes in degrees of belief.

A technical aside The reader might be tempted to object that the above representation of the increase of information is obviously not what is intended. More to the point would be a construction in which the probability spaces remain the same, but the probability distributions change. For our running example this would mean

that the probability space is in both cases $\{p, q, r\}$, but that the probability distribution first assigns probability 0 to $\neg r$, and upon becoming aware of the additional conditional $r \rightarrow q$, a non-zero probability. The trouble with such a suggestion is twofold. Firstly, from a Bayesian point of view, the transition from the *a priori* probability $P(\neg r) = 0$ to the *a posteriori* $P(\neg r) > 0$ is not allowed, since this cannot be achieved via (BaCo): conditionalising on more evidence cannot make a null probability positive. Secondly, one would like to have the assurance that incoming information generates a stable probability distribution in the limit. However, the convergence theorems that guarantee this,¹⁴ also require a null probability to remain null. Clearly it does not help to assume that the initial probability of $\neg r$ is very small, because this would bring us back to the situation where the probability is essentially defined on the set of all propositions, and not on a finite subset.

Again the conclusion must be that Bayesian probability has too much monotonicity built in to account for non-monotonic belief change as witnessed in the suppression task. A form of closed world reasoning with probabilities must be developed which allows an agent to set conditional probabilities to 0 by default, and to change these to a positive value if relevant new information comes in.

3.3.2 Alternative conditionals

Let us now look at distinction between additional and alternatives: additional are related to $P(\neg q \mid p)$, alternatives to $P(q \mid \neg p)$ – note that this is determined by a, b, ϵ . The effect of incorporation of an alternative conditional is that $P(q \mid \neg p)$ must increase, but ϵ must remain constant to account for MP with alternative conditional. This could mean that b increases – more alternative antecedents ‘means’ there are more possibilities for the antecedent to occur, and the probability of the given antecedent a decreases.

We find the same mixture of probabilistic and non-probabilistic modelling assumptions in Oaksford and Chater’s account of the suppression of DA:

Alternative antecedents, such as information that the car can also be started by hot-wiring, with respect to the rule *if you turn the key, the car starts*, concern the probability of the car starting even though the key has not been turned; that is $P(q \mid \neg p)$. If you know that a car can be started by other means, you cannot unequivocally infer that the car will start even though the key has not been turned. Moreover, bringing

¹⁴The so-called ‘martingale convergence theorems’, see [27, Ch. 10].

to mind other alternative ways of starting cars, such as bump-starting, will increase this probability. [...] It is therefore an immediate consequence of our model that if there are many alternative antecedents, that is, $P(q \mid \neg p)$ is high, the probability that the DA inference should be drawn is low [14, p. 104].

Bearing in mind the analysis we gave above of the suppression of MP, it is now easy to see that what the Bayesian model *prima facie* predicts (via (BaCo)) is the following:

if the conditional probability $P(q \mid \neg p)$ increases, then the probability of DA decreases.

The essential question is, however, whether the Bayesian model also covers the first stage of the subject's reasoning, in which the presence of alternative antecedents increases the initial $P(q \mid \neg p)$. The argument seems to be this, using the same notation as in section 3.3.1: the subject makes a distinction between $P_2(q \mid \neg p)$ and $P_3(q \mid \neg p)$ given by

$$P_3(q \mid \neg p) = P_3(q \mid r \wedge \neg p)P_3(r \mid \neg p) + P_3(q \mid \neg r \wedge \neg p)P_3(\neg r \mid \neg p) \quad (4)$$

which, in the example discussed though not always, can be simplified to

$$P_3(q \mid \neg p) = P_3(q \mid r)P_3(r \mid \neg p) + P_3(q \mid \neg r \wedge \neg p)P_3(\neg r \mid \neg p) \quad (5)$$

since $r \rightarrow \neg p$. In this formula, the term $P_3(q \mid r)$ is large, since it represents the conditional $r \rightarrow q$, and the term $P_3(q \mid \neg r \wedge \neg p) = P_3(q \mid \neg(r \vee p))$ is much smaller, especially when $P_3(\neg r \mid \neg p) = 1$, since it then represents $P_3(q \mid \neg p)$.

Oaksford and Chater then appear to argue as follows. In the case of one conditional premise $p \rightarrow q$ only, the subject works with the conditional probability $P_2(q \mid \neg p)$ on the probability space $\{p, q\}$. This probability is then identified with $P_3(q \mid \neg p)$ on the probability space $\{p, q, r\}$ under the assumption that $P_3(\neg r \mid \neg p) = 1$. Consideration of the alternative antecedent $r \rightarrow q$ then leads to a decrease in $P_3(\neg r \mid \neg p)$, so that a larger proportion of the large term $P_3(q \mid r)$ contributes to $P_3(q \mid \neg p)$, whence $P_3(\neg q \mid \neg p)$, i.e. the probability that $\neg q$ is concluded, decreases. But we can now see that this reasoning is not Bayesian, since it requires the subject to update $P_3(r \mid \neg p)$ from zero to strictly positive. This update must necessarily remain outside the model.

3.3.3 Validating the model for three premises

The hypothesis to be tested is that subjects reason with subjective probabilities in order to solve the suppression task. It is in the nature of this hypothesis that it does not concern the exact values of the subject's degrees of belief in concrete conditionals and categorical propositions. The hypothesis is that whatever these values are, they obey the Bayesian inference rules. This makes it difficult to test the hypothesis directly.

Above we have seen that Oaksford and Chater take the revision of conditional probabilities in going from simple MP to MP with additional premise to lie outside the scope of the probabilistic model. Thus the model provides no constraints on the values of the probabilities in the simple condition versus the three premises conditions. To fit the model it is therefore sufficient to estimate a, b, ϵ from the data in each condition of a suppression task; Oaksford and Chater [14, p. 105] take Byrne's [5] for this purpose. Under the assumption that subjects apply (BaCo) when making an inference, the data supply information¹⁵ about the conditional probabilities $P(q | p)$ (MP), $P(\neg p | \neg q)$ (MT), $P(\neg q | \neg p)$ (DA) and $P(p | q)$ (AC). Given these conditional probabilities, the parameters of interest, a, b and ϵ , can be computed. The exception parameter ϵ is obtained from the rate of endorsement of MP (which is equal to $P(q | p)$). The parameter $a = P(p)$ is obtained from the rates of endorsement of MP, AC and DA via the following computation

$$x := \frac{P(q | p)}{P(p | q)} = \frac{P(q)}{P(p)}$$

$$y := \frac{P(q | \neg p)}{P(\neg p | q)} = \frac{P(q)}{1 - P(p)}$$

$$\text{hence } xP(p) = y(1 - P(p)) \text{ and so } P(p) = \frac{y}{x + y},$$

where the right hand side is supplied by the data. The computation for b is similar and yields

$$P(q) = \frac{xy}{x + y}.$$

Observe that the parameters are computed using only three conditional probabilities: $P(\neg p | \neg q)$ is nowhere used, and can be computed from the remaining

¹⁵It will soon become apparent why we use a deliberately vague expression here.

conditional probabilities. This puts a consistency requirement on the four data points, thus allowing the experimenter to check whether subjects are coherent in the probabilistic sense.

The result is that in every condition the three parameters can be fitted from the four data points MP, MT, DA, AC, where one is looking for a fit at .01 significance level. The question arises however whether the estimates of the parameters a, b, ϵ can be interpreted as *probabilities*: what is the event which is supposed to have probability $a = P(p)$? Recall the definition of $P(p)$: p does not stand for any specific proposition, but it refers to a particular role: the antecedent of the main conditional, the content of which differs from one conditional to the next. Similarly for the conditional probabilities used to estimate $P(p)$: these are in fact averages of the conditional probabilities hypothesised to be used by the subjects with the various experimental materials. Now suppose that the parameters were estimated using the data for a single subject. What would the estimated value $P(p)$ (or $P(q | p)$) mean in this situation? It would be an average over the very different contents occurring in the antecedent (and consequent) of the main conditional, not a degree of belief assigned by the subject to a particular event. What the model fitting procedure does, therefore, is construct a joint probability distribution (from a, b, ϵ) of which we can be certain that it represents the degrees of belief of no subject at all. This seems a weak justification for the model.

To summarise, we have not claimed that Bayesian probability plays no role at all in subjects' reasoning in the suppression task. Especially when the formulation of the task explicitly introduces qualitative probabilistic expressions like 'almost always' or 'rarely', as in Stevenson and Over's graded suppression task [26], a probabilistic model may be appropriate.¹⁶ We do claim however, that it cannot be the whole story in those reasoning tasks in which non-monotonic (degree of) belief revision plays an important part. If one does not incorporate this form of belief revision into the model, one is in effect saying it falls outside the scope of explanation. This is giving up too soon. If one is convinced that the last part of subjects' reasoning is indeed Bayesian conditionalisation, there arises the challenge of combining closed world reasoning with probabilistic reasoning.

More generally, the discussion points to the necessity of distinguishing between 'reasoning to an interpretation' and 'reasoning from an interpretation', as we have done in [25]. Once the subject fixes the interpretation of the task at hand as 'reasoning with uncertainty, more particularly probability', and assigns probabilities to the propositions of interest, reasoning may proceed entirely within (Bayesian) probability theory. The reasoning process that leads up to fixing the interpretation

¹⁶We do not take performance in the graded suppression task to indicate that subjects *always* reason according to a probabilistic model in a suppression task.

as probabilistic, and to fixing the probabilities, may well be of a very different nature. Here we have argued that this reasoning to an interpretation is best seen as a form of closed world reasoning, perhaps applied to probabilities. In any case it seems more fruitful to investigate how logic and probability must interact than to view them as rival approaches.

3.4 System P?

The reader acquainted with the literature might think that a fruitful interaction between non-monotonic logic and probability theory already exists, in the form of **System P** [10, 18]. This is a deductive logic for reasoning with exception-tolerant conditionals written as $\alpha \sim \beta$, and read as ‘if α , then *normally* β ’. A highly appealing feature of **System P** is that it allows probabilistic semantics, which are of two kinds:

1. ‘ $\alpha \sim \beta$ is true’ is interpreted as $P(\beta \mid \alpha) > \theta$, where $\theta < 1$ but $1 - \theta$ is infinitesimal [1].
2. ‘ $\alpha \sim \beta$ is true’ if $P(\beta \mid \alpha)$ takes a value in a suitable interval [9]. The details need not concern us here. It suffices to say that the semantics introduces a new kind of object, ‘conditional events’ $B \mid A$, corresponding to $\alpha \sim \beta$, allowing the possibility that the (absolute) probability of A is 0; and that the inference rules of **System P** in this semantics correspond to rules for transforming probability intervals.

It thus seems that **System P**, together with the interval semantics, is ideally suited to address the issues we raised above: it seems to allow a rational treatment of events with probability 0, and it is not wedded to the assumption that probabilities must be known precisely. This is not the case however, as can be seen when we look at **System P**’s rules, presented as sentences of the form ‘premises $\implies$ conclusion’:

$$\implies \alpha \sim \alpha \quad \text{REFLEXIVITY AXIOM} \quad (6)$$

$$\models \alpha \leftrightarrow \beta, \gamma \sim \alpha \implies \gamma \sim \beta \quad \text{LEFT LOGICAL EQUIVALENCE} \quad (7)$$

$$\models \alpha \rightarrow \beta, \gamma \sim \alpha \implies \gamma \sim \beta \quad \text{RIGHT WEAKENING} \quad (8)$$

$$\alpha \sim \gamma, \beta \sim \gamma \implies \alpha \vee \beta \sim \gamma \quad \text{OR} \quad (9)$$

$$\alpha \wedge \beta \sim \gamma, \alpha \sim \beta \implies \alpha \sim \gamma \quad \text{CUT} \quad (10)$$

$$\alpha \sim \beta, \alpha \sim \gamma \implies \alpha \wedge \beta \sim \gamma \quad \text{CAUTIOUS MONOTONICITY} \quad (11)$$

It is now immediately obvious that one rule of **System P** dashes all hopes of treating the suppression task: the rule OR which forces one to treat additional and alternative premises on the same footing.¹⁷ **System P** therefore reinforces the point made earlier, that probability is too much tied to classical logic to model the truly non-monotonic reasoning that occurs in the suppression task.

4 Conclusions

Interpretation processes are necessary, whether one then applies probability theory or some logic in reasoning from the resulting interpretations. In the case of suppression, understood in probabilistic terms, interpretation shows up as the necessity to change one's probabilities in ways not sanctioned by Bayesianism. We claim that a computational level analysis in the sense of Marr must also incorporate the interpretation process, not only the reasoning once the interpretation is chosen. This is not to deny the role of Bayesian probability in a characterisation of the computational level. If a subject construes the task as involving uncertain conditionals, in the sense of positive probability of exceptions, principles like Bayesian conditionalisation may well form part of the computational level. In this case competence theory is needed of how judgements of probabilities can change in non-Bayesian ways. This we regard as one of the most interesting technical challenges issuing from the present analysis. We do deny that the entire computational level can be characterised in this manner, and also that the same computational level analysis applies to all subjects engaged in this task.

Classical logic forced a single competence model upon reasoning tasks, and one of the merits of the Bayesian approach is to have loosened the grip of classical logic. But the danger exists that Bayesian probability is itself promoted to the absolute standard of competence. The arguments purporting to show that probability is *the* calculus of uncertainty are too weak to establish this, hence absoluteness of the Bayesian standard cannot be defended in this way. On the empirical side it is clear that different subjects are interpreting tasks very differently, and although Bayesianism has some built-in mechanisms to deal with individual differences via variation in assignments of subjective probabilities, this does not capture the full range of differences. For instance, closed world reasoning, for which there exists evidence in our tutorial data [25, Chapter 7], cannot be modelled in this way.

Oaksford and Chater view their work as an instance of 'rational analysis' and consider that giving an account of the economics of information processing tasks

¹⁷If one drops OR from **System P** one gets Makinson's 'cumulative logic' [12], which however has no close relationship to probability.

should occupy centre stage. Our own work adopts a very different computational level analysis of reasoning tasks, assimilating them to discourse understanding. The familiar laboratory deductive reasoning tasks are interesting, important and potentially psychologically insightful precisely because subjects assimilate them to this overall process of discourse processing. Because the normal cues as to how to make this assimilation are removed, subjects make it in a variety of ways, many of which are not the classical logical one the experimenter expects. So we absolutely agree with Oaksford and Chater that the classical logical computational model is generally not a reasonable choice for data analysis. Where we disagree is that we believe discourse processing has to be understood as at least a two-component process, and the rational analyses of these two processes cannot be the same. Trying to make them the same leads to the invocation of probability theory as an overall framework, and that makes the computational theory distant from the data, and inappropriate for dealing with the kinds of uncertainty that permeate interpretation.

References

- [1] E. Adams. *The logic of conditionals*. Reidel, Dordrecht, Netherlands, 1975.
- [2] J.R. Anderson and C. Lebiere. *The Atomic Components of Thought*. Lawrence Erlbaum, Mahwah, N.J., 1998.
- [3] G. Baggio, M. van Lambalgen, and P. Hagoort. Language, linguistics and cognition. In N Asher, R. Kempson, and T. Fernando, editors, *Handbook of the Philosophy of Linguistics*. Elsevier, Amsterdam, To appear.
- [4] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189–198, 1993.
- [5] R.M.J. Byrne. Suppressing valid inferences with conditionals. *Cognition*, 31:61–83, 1989.
- [6] D.D. Cummins. Naive theories and causal deduction. *Memory and Cognition*, 23(5):646–58, 1995.
- [7] K. Doets. *From Logic to Logic Programming*. The MIT Press, Cambridge, MA, 1994.
- [8] J.St.B.T. Evans. In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences*, 7(10):454–459, 2003.

- [9] A. Gilio. Probabilistic reasoning under coherence in system p. *Annals of Mathematics in Artificial Intelligence*, 34:5–34, 2002.
- [10] S. Kraus, D. Lehman, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [11] A. Lechler. Interpretation of conditionals in the suppression task. Msc thesis, HCRC, University of Edinburgh., 2004.
- [12] D. Makinson. General theory of cumulative reasoning. In M. Reinfrank, editor, *Proceedings Second International Workshop on Non-Monotonic Reasoning*, volume 346 of *Lecture Notes in Computer Science*. Springer, 1989.
- [13] J. McCarthy. Circumscription – a form of non–monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.
- [14] M. Oaksford and N. Chater. Probabilities and pragmatics in conditional inference: suppression and order effects. In D. Hardman and L. Macchi, editors, *Thinking: psychological perspectives on reasoning, judgment and decision making*, chapter 6, pages 95–122. John Wiley & Sons, Chichester, 2003.
- [15] M. Oaksford and N. Chater. *Bayesian rationality*. Oxford University Press, Oxford, 2007.
- [16] M.R. Oaksford and N.C. Chater. A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101:608–631, 1994.
- [17] J.B. Paris. *The uncertain reasoner’s companion*. Cambridge University Press, Cambridge, 1994.
- [18] N. Pfeiffer and G. Kleiter. Coherence and nonmonotonicity in human reasoning. *Synthese*, 146:93–109, 2005.
- [19] M.M. Rao. Projective limits of probability spaces. *Journal of Multivariate Analysis*, 1:28–57, 1971.
- [20] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [21] Murray P. Shanahan. *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press, 1997.
- [22] K. Stenning and M. van Lambalgen. A little logic goes a long way: basing experiment on semantic theory in the cognitive science of conditional reasoning. *Cognitive Science*, 28(4):481–530, 2004.

- [23] K. Stenning and M. van Lambalgen. Semantic interpretation as reasoning in nonmonotonic logic: the real meaning of the suppression task. *Cognitive Science*, 29(6):919–960, 2005.
- [24] K. Stenning and M. van Lambalgen. Logic in the study of psychiatric disorders: executive function and rule-following. *Topoi*, 26(1):97–114, 2007. Special issue on Logic and Cognitive Science.
- [25] K. Stenning and M. van Lambalgen. *Human reasoning and cognitive science*. MIT Press, Cambridge, MA., 2008.
- [26] R. Stevenson and D. Over. Deduction from uncertain premisses. *Quarterly Journal of Experimental Psychology A*, 48(3):613–643, 1995.
- [27] D. Williams. *Probability with martingales*. Cambridge University Press, Cambridge, 1995⁴.