IN -63 -7M/
91505
P-38

Modelling Default and Likelihood Reasoning as Probabilistic

Wray Buntine

AI Research Branch, Mail Stop 244-17
NASA Ames Research Center
Moffett Field, CA 94035

NASA Ames Research Center

Artificial Intelligence Research Branch

Modelling Default and Likelihood Reasoning as Probabilistic Reasoning

Wray Buntine*

RIACS†and AI Research Branch

NASA Ames Research Center, MS 244-17

Moffett Field, CA, 94025, USA

September 11, 1990

Abstract

This paper presents a probabilistic analysis of plausible reasoning about defaults and about likelihood. "Likely" and "by default" are in fact treated as duals in the same sense as "possibility" and "necessity". To model these four forms probabilistically, a logic QDP and its quantitative counterpart DP are derived that allow qualitative and corresponding quantitative reasoning. Consistency and consequence results for subsets of the logics are given that require at most a quadratic number of satisfiability tests in the underlying propositional logic. The quantitative logic shows how to track the propagation error inherent in these reasoning forms. The methodology and sound framework of the system highlights their approximate nature, the dualities, and the need for complementary reasoning about relevance.

Index Terms: default, likelihood, plausible reasoning, qualitative reasoning, subjective probability.

*Much of this research was done while at the University of Technology, Sydney, Broadway, NSW, Australia, and some at the Turing Institute, 36 Nth. Hanover Str., Glasgow, Scotland.

†Research Institute for Advanced Computer Science.

1 Introduction

Default reasoning is a form of non-monotonic reasoning which can be introduced by Delgrande [1] as follows:

> Many common sense assertions about the real world express default or prototypical properties of individuals or classes of individuals, rather than strict conditional relations. Thus, for example, "birds fly" attributes the property of flight to birds, even though birds with broken wings generally don't fly, and quite probably no penguin flies. The import of "birds fly" then certainly isn't that *all* birds fly, but rather is more along the lines of "typically birds fly".

This form of default reasoning then is concerned with drawing "typical" conclusions. There is a continuously growing and diverging variety of theoretical treatments on this and other forms of non-monotonic reasoning [2,3,1,4,5,6,7,8].

Likelihood reasoning, another form of plausible reasoning, is more concerned with drawing "likely" conclusions. For example, it is "likely" or reasonably possible that a coin tossed twice will land heads both times, although this certainly is not "typically" the case. It is not "likely", however, that a coin tossed twice will land on its side one of those times. Although the laws of physics might treat this as a "possible" outcome, for most practical purposes it is not. When one is considering possible outcomes, rather than looking at *all*, likelihood reasoning is intended to be applied to find only those outcomes that are reasonably possible. A historical perspective and further discussion for this form of reasoning can be found in [9].

The relationship between probability and plausible reasoning is best introduced by Polya [10, Chapter XV] in his work on reasoning in mathematics. Polya introduced a system of guides to the mathematician of the form:

> [Given a conjecture,] the verification of a consequence renders the conjecture more credible.
>
> Our confidence in a conjecture can only increase when an incompatible rival conjecture has been exploded.

These guides were based on belief about conjectures modelled as subjective probabilities. Plausible reasoning about default and likelihood, however, has more often been modelled in AI using purely logical formalisms [5,1,6,9] or non-standard probabilistic methods [11,12,13], although probability-motivated approaches exist [7,14]. Another form of reasoning seen in areas such as qualitative physics and model-based diagnosis systems is the qualitative and approximate reasoning about physical devises. In this paper we combine the two paradigms, probabilistic and qualitative/approximate, to model default and likelihood reasoning, and so take up Polya's theme more fully.

Surprising to some, it is controversial whether these plausible reasoning forms can be modelled with probabilities[1] [15,4,16]. Likelihood reasoning and some forms of default reasoning, however, will always remain problems of uncertainty or incomplete information. With some forms of non-monotonic reasoning, such as the closed-world assumption used in database systems and PROLOG, uncertainty does not exist because the default is actually a convention. These exceptions aside, there comes a time when something that is currently "typically" or "likely" to hold becomes known true or false. Until that time, we are in a state of uncertainty. However well logical systems may cope with modelling these reasoning forms, we

---

[1] Cheeseman has said [15, p1002]:

> Unfortunately, the logical style of reasoning is so prevalent in AI that many have attempted to force intrinsically probabilistic situations into a logical straight jacket with predictable limited success.

should at least see how they can be modelled by a theory of uncertainty like probability. Perhaps there is more to learn?

This paper follows the view that subjective Bayesian probability theory provides a benchmark against which methods for reasoning about uncertainty can be compared. The theory is a *normative* theory of reasoning about uncertainty, which means it gives a prescription for how uncertain reasoning should be done. The prescription itself has been derived from a set of fundamental axioms about belief (an introduction to this in the AI context is in [17]). One can model default and likelihood reasoning as either qualitative or quantitative approximations to full normative probabilistic reasoning. One can then argue that the resulting model seems to exhibit the required properties, and compare the model with some existing methods.

A logic QDP (a mnemomic for qualitative default probabilistic logic) is developed here from a suitable quantitative counterpart DP as a demonstration. This yields a probabilistic system as a canvas on which a number of more significant issues can be sketched. These issues are: (1) the interplay between quantitative and qualitative forms of plausible reasoning, (2) the duality between default and likelihood reasoning, (3) the approximate nature of these reasoning forms, for instance, the propagation of errors in reasoning, and (4) the need for complementary reasoning about, for instance, relevance.

The logic QDP, being probabilistically based, is easily able to express sentences[2] such as "most birds fly". This is using a "default" conditional style operator "$\Rightarrow$" as in: $Bird(x) \Rightarrow Flies(x)$. Similarly, "an Australian is likely to drink Foster's"[3] can be represented with a "likely" conditional style operator "$\approx\!\!\succ$" as in: $Australian \approx\!\!\succ Drinks\text{-}Foster's$. This operator also has iterated forms indicated by numeric superscripts, "$\approx\!\!\succ^2$", that express lesser degrees of likelihood, as in: $Australian \approx\!\!\succ^2 Drinks\text{-}another\text{-}Foster's$, which expresses the fact that, at least occasionly, an Australian will drink even more Foster's.

Surprisingly enough, QDP is also able to express sentences more in the spirit of autoepistemic [18] and default logics [2]. We can interpret the sentence "a professor has a Ph.D. unless known otherwise" two ways:

$$\diamond(Prof(x) \land Phd(x)) \longrightarrow (Prof(x) \Rightarrow Phd(x)) ,$$
$$\diamond(Prof(x) \land Phd(x)) \longrightarrow \Box(Prof(x) \rightarrow Phd(x)) ,$$

where the "$\Box$" operator represents necessity interpreted as "known with certainty", and the dual "$\diamond$" operator represents possibility interpreted as "the negation is not known with certainty". Read as "if it is possible that a particular professor has a PhD, then the professor *most likely* has a Ph.D.", and "if it is possible that a particular professor has a PhD, then the professor *definitely* has a Ph.D." respectively. The default logic representation, from $Prof(x) \land M\ Phd(x)$ infer $Phd(x)$, corresponds to the second reading. So the possibility operator, "$\diamond$", behaves rather like the M operator of default logic.

The default component of the logic QDP is a variant and extension of Adams' conditional logic [19], applied to default reasoning by Pearl [7]. The probabilistic semantics of QDP differs slightly from Adams' logic however, because QDP is developed as a qualitative model for order of magnitude reasoning about probabilities, rather than being based on infinitesimal arguments. Like Adams' logic, QDP can be combined with a notion of relevance or causality to resolve the so-called default paradoxes: the Yale shooting problem [8] and "can Joe read and write?" [7]. The logics also resolves the "vanishing subclasses" paradox [20]. These three paradoxes are discussed in Section 5. A fourth paradox is the lottery paradox [3], considered in Section 3. This has a version both in default and likelihood reasoning, and provides an example of the propagation of errors inherent in these reasoning forms.

The logics has been modelled after Delgrande's modal conditional logic NP which allowed reasoning *about* default rules. Likewise, reasoning about defaults and likelihood is an important feature of the

---

[2] Although propositional sentences are dealt with throughout, pseudo-first-order sentences will sometimes be used. They are effectively propositional if there are known to be a finite number of constants, no quantifiers are allowed, and a sentence with variables is intended to represent a sentence schema.

[3] For the record, many Australians don't. Some drink XXXX, others Swan, ...

3

approach here. For example, suppose you know friends have travelled to Australia. Then they are likely to have visited Sydney. Although any visitor to Sydney will typically see the Sydney Harbour Bridge, it is only likely that they will visit Bondi Beach. We can infer that your friend is likely to (rather than "typically") have seen the Harbour Bridge but is less likely to have visited Bondi. In QDP, this argument can be summed up as follows:

$$true \approx\!\!\!\!- Visit\text{-}Sydney(Bruce)$$
$$Visit\text{-}Sydney(x) \Rightarrow See\text{-}Harbour\text{-}Bridge(x)$$
$$Visit\text{-}Sydney(x) \approx\!\!\!\!- Visit\text{-}Bondi(x)$$
$$\models_{QDP} \quad (true \approx\!\!\!\!- See\text{-}Harbour\text{-}Bridge(Bruce)) \wedge (true \approx\!\!\!\!-^2 Visit\text{-}Bondi(Bruce)) .$$

Consistency and consequence tests developed for subsets of the default and likelihood components of the logics also show how this form of reasoning can be automated in a manner requiring at most a quadratic number of satisfiability tests in the underlying propositional logic. With a careful choice of the underlying propositional logic, the operation can then be quite efficient.

Perhaps most significantly, this reasoning can be easily complemented with error tracking facilities to indicate when the conclusions from a chain of such plausible reasoning may be coming doubtful. For instance, it is shown in some circumstances that error when reasoning about defaults can increase at most additively, while error when reasoning about likelihood can increase multiplicatively. It is not claimed, however, that these tracking facilities are a substitute for a more thorough probabilistic approach; they are merely an approximation.

The paper follows the following course. First, the philosophical problem of modelling default reasoning with probabilities is considered in Section 2. The corresponding discussion for likelihood reasoning is not given here, because the principle objections in AI to modelling likelihood reasoning with probabilities do not centre around the use of probability theory at all, but whether the modelling should be qualitative or quantitative [9], and both are done here. A basic probabilistic framework for plausible reasoning is then proposed in Section 3. Two logics, one with a probabilistic semantics, DP, and a qualitative version, QDP, are then introduced in Section 4. Here, the duality between default and likelihood is introduced, and the consistency and consequence results are developed. Section 5 demonstrates a methodology for applying the qualitative logic, using relevance, and Section 6 draws some comparisons with other probabilistic approaches.

2 On Modelling Default Reasoning with Subjective Bayesian Probability

Non-monotonic reasoning is generally considered to have three broad forms [4,18]: *autoepistemic reasoning* is reasoning about self-knowledge of beliefs [18], for instance, "if I had an older brother I would know about it"; *conventions* are used in the interpretation of natural language and with the closed-world assumption often made for database systems; and *typicality* or *default reasoning* is the form discussed in the Introduction.

To illustrate the use of convention in natural language, consider the sentence "birds lay eggs" [16], which is certainly not true for the male half of the bird population. The sentence is more accurately stated as "[female] birds lay eggs [to reproduce]". The parts in the square brackets are implicit. Most people realise that male birds cannot lay eggs, so in the interests of brevity, the speaker leaves "female" to be inferred from the remainder of the sentence. This implicit convention is handled in nonmonotonic systems using knowledge of the form "an X is a Y unless known otherwise". As illustrated in the introduction, this form can also be represented in a probabilistic framework using the probabilistic version of the possibility and necessity operators.

When modelling the third form, typicality or default reasoning, we are hampered by the fact that

4

there is little consensus as to its exact nature [20]. Hanks and McDermott [8] say,

> While it is not entirely clear *exactly* what constitutes default reasoning, the phenomenon commonly manifests itself when we know what conclusions should be drawn about *typical* situations or objects, ...

Neufeld, Poole and Aleliunas [20] make an even stronger statement. They say,

> What, then, does a default mean? Within the default logic camp, we know of no work which provides a semantics for defaults, in the sense that an experiment is described that can be performed in the semantic domain to verify the truth of a default.

However, there is general agreement that default reasoning is a form of "defeasible inference", or "plausible reasoning" [21,8], and that default conclusions have some (often small [21]) degree of uncertainty to them.

Given that default reasoning is an admittedly specialised form of reasoning under uncertainty, it is natural to pose the question: can probability theory model default reasoning (see also [7])? Critics of a Bayesian approach claim that probabilities are just not suited for describing "prototypical" knowledge. Must arguments, however, are based on some misunderstanding.

Nutter [16] gives the following argument:

> For instance: if ... the by now tormented example "Birds fly" really means "Most birds fly", then birds don't fly in spring. In the nesting season, baby birds outnumber adults. Baby birds don't fly. Hence in the nesting season, "Most birds fly" is false.

To the Bayesian, "Most birds fly" is interpreted as "if we know nothing else about a particular bird, then that bird most likely flies". Notice the "most likely" conclusion is conditioned on our current knowledge about the bird. In particular, if we know it is nesting season, we cannot conclude the bird most likely flies because we do now know some additional thing about the bird. Two rules are relevant to the situations Nutter gives: "Most birds fly" and "In the nesting season, most birds don't fly". If we do not know that it is the nesting season, then the first rule is applicable because it usually is not the nesting season. The importance of conditioning probabilistic statements with context or current knowledge is a key feature of probabilistic reasoning and the cornerstone of the subjective Bayesian approach.

McCarthy address a similar concern [4, p92].

> Note that the general probability that a bird can fly may be irrelevant, because we are interested in the facts that influence our opinion about whether a particular bird can fly in a particular situation.

Classical statistics, with its concern about long term frequencies and samples spaces, can have problems in adapting general knowledge to specific situations. The ability to adapt knowledge to particular situations, however, is a hallmark of Bayesian methods. In this case, suppose we know that the bird is a male yellow-bellied warbler, but we have no knowledge at all about this type of bird, or even what they may be similar to. The only relevant knowledge we have is the general probability statement that most birds fly. In the absense of information to the contrary, we assume that other details about the bird are irrelevant (this is the maximum entropy argument [7]), which leads us to the quite reasonable conclusion that most male yellow-bellied warblers fly. We can now reason about this particular bird.

There are, however, strong arguments that default reasoning should be modelled by probability *with caution*. In practice, an intelligent system may not be able to supply precise probabilities for its beliefs

and may not be able to perform all the exact calculations required to maintain its beliefs in accord with Bayesian principles as new evidence becomes available. People certainly cannot. It is of course not just the computation that causes problems but the communication required to prime and then update an intelligent system with an adequate set of beliefs.

The normative properties of Bayesian theory assures us that despite these problems, by trying to approximate the Bayesian approach our reasoning at least remains approximately rational. Essentially, it is the best we can do in an inherently imprecise and computationally complex world. This view has been supported in AI alone in a range of areas [22,23,24,25].

3 A Framework for Plausible Reasoning

In this section, a basic framework for default and likelihood reasoning is developed. These two forms of reasoning are referred to below as plausible reasoning. Before presenting the framework, we first consider some major features of plausible reasoning, and then infer properties that a plausible reasoning system should have.

3.1 Basic features of plausible reasoning

There are several basic features of plausible reasoning that must effect the design of a plausible reasoning system. While these can be derived from the probabilistic model presented in the next section, the features are presented here independently of any probabilistic analysis.

Plausible reasoning is non-monotonic

With standard logical reasoning, conclusions derivable from a set of sentences increase *monotonically* as the set of sentences is extended. That is, if S logically implies C, and we extend S with A, then $S \wedge A$ also logically implies C.

Default reasoning is known to be non-monotonic [3]; the above monotonicity property breaks down. So while you might well believe that birds fly, on discovering that a certain bird is a baby bird in nesting season, you would no longer believe that particular bird flies. So your set of beliefs have extended one way but contracted another. Similarly, something that initially seems likely can become, with changing circumstances, well nigh impossible.

Error combines along a chain of plausible reasoning

A second key feature of standard logical reasoning is that if the premises are known to be true, then the conclusion from a long chain of reasoning steps must also be true. With plausible reasoning, however, there is an inherent element of uncertainty involved, so it is natural to suspect this key feature might break down.

The famous lottery paradox [3] is an excellent example of this. For a single lottery entrant, Leslie say, one can conclude by default that Leslie will not win the lottery. But we can apply this sort of reasoning to every potential lottery entrant. There are two paradoxes here. First, why is it that someone actually wins the lottery. Second, why does Leslie bother to enter the lottery in the first place.

For a lottery with one million entrants, the default conclusion about Leslie has an obvious statistical error of one ten-thousandth of 1%, acceptable by most standards. If we make a logical deduction based on one million such default conclusions, the one million errors certainly combine to give a total error of 100% (after all, someone definitely wins the lottery). That Leslie would enter the lottery at all is as much irrational behaviour due to the effect of large sums of money, as it is the result of plausible reasoning. Perhaps it is because most people do not mind losing one dollar just to be given the remotest chance of winning one million dollars. In the former, their life is no different; in the latter, well ...

This last point anticipates the next basic feature of plausible reasoning.

6

Plausible reasoning is effected by the decision context

After a system performs plausible reasoning, it would typically decide some course of action. As a result of the action, the system might make some gain or incur some loss. For Leslie in the lottery situation above the potential loss is one dollar while the potential gain is one million minus one dollars. This feature of reasoning is referred to as the *decision context* and the losses and gains as the *utilities*.

Shoham provides the following illustration of how the decision context can effect plausible reasoning.

> ... think of making the default inference "people you'll meet on the street will not stab you in the back" in a city in which only 5% of the population are back stabbers. In this case the relatively small chance of being hurt seems to outweigh the computational resources needed to reason about individual people on the street, and the discomfort of wearing a steel-plated vest. Notice that if the 5% dropped to 0.00000000005%, we'd take off the armor and stop looking darkly at passers by.

Clearly, the decision context should be taken into consideration (see also [26]).

3.2 Basic properties of a plausible reasoning system

The above features can be used to argue that a method for plausible reasoning should have certain basic properties.

A first property is that *plausible reasoning needs to be sensitive both to the current knowledge of the system and to the decision context*. This is directly suggested by the features given in the previous subsection. Sensitivity to the decision context can be handled by targeting a default system for a single decision context.

Now the number of different states of knowledge is potentially exponential in the number of propositional symbols. So a system could not reasonably keep separate default rules for each possible state of knowledge and decision context. To get around this problem, a second property seems important: *it should be possible to reason about plausible rules and the relevance of different facts to the applicability of a plausible rule*. It may also be useful to give a system the ability to compile plausible rules from some more fundamental knowledge form.

Third, because of non-monotonicity and error propagation, *plausible conclusions need to be flagged as such, and should not be confused with the current knowledge*. In fact, because of the possible need for weighing up belief when combining error or considering the decision context, plausible conclusions may need to be tagged with some form of qualitative or quantitative measure of belief. Whether this is done and how surely depends on the application concerned; no single approach will be favoured in this paper.

3.3 A probabilistic framework

It is beyond the scope of this paper to cover the basic notions of probability and decision theory underlying subsequent sections. Suitable introductions from an AI perspective can be found in [26,27,7]. The problem of the decision context in plausible reasoning is side-stepped here by assuming that a default system is being prepared for a specific binary (yes/no) decision. In this simple case, a decision has to be made whether some condition, A say, is "true" or "false". Once utilities of the problem are taken into account, the problem invariably reduces to "is $Pr(A) \geq p$?" for some $p \in [0, 1]$. Given a particular decision context for a binary decision, we can therefore use approximate inequality reasoning to make decisions in a normative manner.

The notion of probability used here is *subjective probability*, which is a *measure of belief* prescribed to some proposition by an intelligent system. This is represented as $Pr(A|B) \in [0, 1]$, interpreted as follows: a particular intelligent system, on knowing just B, has a measure of belief $Pr(A|B)$ in A being true. The "|" operator is called the conditioning operator. Its left hand side is the proposition whose belief is being

considered and its right hand side specifies all current knowledge relevant to A of the intelligent system. A *probability distribution* is a particular function Pr consistent with the standard axioms of probability theory.

A probabilistic framework for plausible reasoning is based on the assumptions that (1) plausible statements that are uncertain should be interpreted in some way using subjective probability statements, and that (2) methods of plausible reasoning which deal with uncertainty should be interpreted as approximations to subjective probability or decision theory. We shall treat a default conclusion as a plausible proposition in which one has "sufficiently high belief". Similarly, a likely conclusion is a plausible proposition in which one has "belief that it is reasonably possible". In both cases, the belief is modelled as subjective probability and should be conditioned on current knowledge using the conditioning operator. Due to the decision theoretic argument above, both these types of plausible reasoning should, in many cases, be a good approximation to the normative probabilistic approach.

Notice that this rough probabilistic interpretation of defaults and likelihood automatically provides a framework which addresses the basic properties of plausible reasoning discussed in this section. Decision theory provides the basis for considering the decision context. The conditioning operator provides the mechanism for making plausible conclusions sensitive to a system's current knowledge and for keeping plausible conclusions (on the left hand side) separate from current knowledge (on the right hand side). Probability theory also provides the potential for developing ways of reasoning about plausible rules, and with the notion of independence, ways of reasoning about relevance. Some of these connections are explored more fully in the next section. Finally, probability theory provides a framework for both testing and developing default rules for a given application, for instance, by learning them from examples.

4 Default Probabilistic Logic

This section introduces two logics for default and likelihood reasoning: a probabilistic logic DP and its qualitative counterpart QDP. These are applicable in the broad framework given in Section 3 for reasoning about defaults and likelihood. Notation and semantics of these logics are first covered in Sections 4.2 and 4.3. Some basic properties of the logics are then outlined. One theme of this paper is the importance of reasoning about relevance; Section 4.5 motivates this and shows how relevance information can interface with default and likelihood reasoning. Another theme of the paper is the approximate nature of both these reasoning forms; Section 4.6 shows how, for small errors at least, the quantitative logic DP can be treated as a simple numeric extension of the qualitative logic QDP. This last section presents consistency and consequence results for fragments of both logics.

4.1 Introduction

DP is a propositional logic annotated with probability bounds, and has a probabilistic rather than a possible world semantics. This allows the sort of inequality reasoning found in Quinlan's INFERNO [28]. Inequality reasoning is an approximation to normative reasoning about point probabilities when a decision is binary, as explained in Section 3.3. So the justification for DP is approximation, rather than some fundamental principle about intervals or fuzzy sets for reasoning under uncertainty. In this sense, it differs in philosophy from Ginsberg's suggestion [12] or Dubois and Prade's treatment of syllogism's [13].

QDP has the annotations dropped, and the default component is almost identical to Geffner and Pearl's logic of defaults [7,29] borrowed from Adams' logic of conditionals [30,19]. QDP is also similar to Delgrande's conditional logic NP [1].

QDP is designed to be a qualitative counterpart of DP. It is intended to be an approximation to DP for reasoning about "small" but not infinitesimal probabilities. The semantics of QDP complements DP and is based on order of magnitude reasoning. Like NP, dynamic aspects of plausible reasoning (for instance, involving action and time) are not handled directly by either DP or QDP, although they can often be handled with a simple situation calculus, as is done in Section 5.3. In the general case, an

8

extension of the logic would be required.

4.2 Basic notation

A standard propositional language denoted L_P is used here. This is formed in the usual manner from a finite set of atomic propositions $P = \{p_1, \ldots, p_n\}$ together with *true* and *false*, the standard connectives, $\neg$ (negation), $\rightarrow$ (conditional), $\wedge$ (conjunction), $\vee$ (disjunction) and $\leftrightarrow$ (biconditional). "$\models A$" denotes that propositional formula A is a theorem of the usual propositional logic.

Probability distributions can be given over the language L_P as follows. An *event space* E_P, a mutually exclusive and exhaustive set of events, is readily constructed from a subset of L_P. Given n atomic propositions P as described above, this would have cardinality 2^n and one such set is given by

$$E_P = \{ L_1 \wedge \ldots \wedge L_n \mid \text{for } i = 1, \ldots, n, \ L_i = p_i \text{ or } \neg p_i \} . \tag{1}$$

A probability distribution $Pr : E_P \mapsto [0, 1]$ maps events to measures of belief. For $A, B \in L_P$

$$Pr(A) = \sum_{\substack{e \in E_P \\ \models e \rightarrow A}} Pr(e) ,$$

$$Pr(B|A) = \begin{cases} \frac{Pr(B \wedge A)}{Pr(A)} & \text{if } Pr(A) > 0 \\ 1 & \text{otherwise} \end{cases} .$$

In many probability texts, if $Pr(A) = 0$ then $Pr(B|A)$ is undefined. Instead we assert that if $Pr(A) = 0$ then $Pr(B|A) = 1$. This means we can reason about conditional probabilities even if the antecedent of the conditioning is false. A probability distribution like Pr above is termed a distribution over L_P.

The probabilistic logic DP describes constraints on probability distributions over the language L_P. It is built on the language D_P that is constructed from L_P together with four modal operators: the unary connectives $\Box$ (necessity), $\diamond$ (possibility), and the binary connectives $\Rightarrow$ (default with error bound) and $\leftrightsquigarrow$ (likelihood with lower bound). There is no nesting of these operators. Nesting would represent second and higher-order probability statements [31], as used in learning to reason about belief in probabilistic models [25], but is unnecessary for the initial treatment here. The operators can be interpreted as follows.

$\Box A$: A is necessarily true in any situation.

$\diamond A$:   Some situation can possibly arise in which A is true.

$A \Rightarrow_\epsilon B$:   Given that you know just A about the current situation, it is safe to infer B by default (with error in belief at most ϵ).

$A \leftrightsquigarrow_\epsilon B$:   Given that you know just A about the current situation, B is at least likely (with belief no less than ϵ).

In the language QD_P the subscripts are dropped. QD_P also has successively weaker forms of the likelihood operator. $A \leftrightsquigarrow B$ denotes "likely", whereas $A \leftrightsquigarrow^2 B$ would denote "barely likely", etc. This is related to the iterated likelihood operator found in [14].

$A \leftrightsquigarrow^n B$:   Given that you know just A about the current situation, B is at least likely to be ... to be likely (to order n).

Both the likelihood and default operators are conditional operators, in a similar sense to [1]. For instance, in the cases above each is conditioned on A. It will be shown later that it is unnecessary for the necessity and possibility operators to have conditional forms.

Definition 4.1 *The sentences or well formed formulae (wffs) of D_P comprise the least set such that*

9

1. *If $A \in L_P$ then $\Box A$ is a wff.*

2. *If $A, B \in L_P$ then $A \Rightarrow_\epsilon B$ is a wff for $0 \leq \epsilon < 1$.*

3. *If $D, E \in D_P$ then $\neg D$ and $D \to E$ are wffs.*

Conjunction ($\wedge$), disjunction ($\vee$), and biconditional ($\leftrightarrow$) on sentences in D_P, and possibility ($\diamond$) and likelihood ($\approx\!\!\!-$) on sentences in L_P are introduced by definition.

Definition 4.2 *The sentences or well formed formulae of QD_P consist of the sentences of D_P with the numeric subscripts dropped from "$\Rightarrow$" and "$\approx\!\!\!-$". The "$\approx\!\!\!-$" operator may have optional integer superscripts weakening the order of likelihood.*

Some examples of QDP sentences were given in the introduction. The four modal operators have operator precedence midway between disjunction and conditional/biconditional. So a disjunction binds before a default operator, and a default operator binds before a conditional. For instance, the sentence

$$A \vee B \wedge C \Rightarrow D \to \diamond E \wedge F$$

is identical to the sentence

$$((A \vee (B \wedge C)) \Rightarrow D) \to \diamond(E \wedge F) \,.$$

Although,

$$\diamond A \wedge B \Rightarrow C$$

is identical to the sentence

$$(\diamond A) \wedge (B \Rightarrow C) \,,$$

because otherwise the sentence does not parse.

4.3 Semantics

In DP, "$\models_{Pr} D$" denotes that $D \in D_P$ is true for the probability distribution Pr. Pr plays a role not unlike an interpretation in standard propositional logic.

Definition 4.3 *Given a probability distribution Pr on L_P, "$\models_{Pr}$" is defined on sentences from D_P as follows.*

1. *$\models_{Pr} \Box A$ if and only if $Pr(A) = 1$.*

2. *$\models_{Pr} A \Rightarrow_\epsilon B$ if and only if $Pr(B|A) \geq 1 - \epsilon$.*

3. *$\models_{Pr} \neg D$ if and only if not $\models_{Pr} D$.*

4. *$\models_{Pr} D \to E$ if and only if not $\models_{Pr} D$ or $\models_{Pr} E$.*

Possibility and likelihood are by definition dual operators for necessity and default respectively. "$\diamond A$" is defined as "$\neg \Box \neg A$", so $\models_{Pr} \diamond A$ if and only if $Pr(A) > 0$. "$A \approx\!\!\!-_e B$" is defined as "$\neg(A \Rightarrow_e \neg B)$", so $A \approx\!\!\!-_e B$ if and only if $Pr(B|A) > e$. In addition, "$\mapsto_e B$" is shorthand for "$true \Rightarrow_e B$", and likewise for "$\mapsto\!\!\!-$".

If the necessity operator were to have a conditional version, it would have the semantics $Pr(B|A) = 1$, but since this is equivalent to $Pr(A \to B) = 1$, a conditional form of necessity can be adequately constructed as $\Box(A \to B)$. Likewise, a conditional version of the possibility operator can be constructed as $\diamond A \to \diamond(A \wedge B)$.

A map translating probabilities to subsequent modal representation is given in Figure 1. By convention, "$\Rightarrow$" is subscripted by greek letters ϵ, δ, etc., which are intended to be small ($\ll 1$), whereas, "$\approx\!\!\!-$" is subscripted by the letters e, f, etc., which are intended to be not as small. This is no absolute restriction; it gives an indication of the intent of the sentences.

Definition 4.4 *A sentence $D \in D_P$ is a theorem of the probabilistic logic DP if $\models_{Pr} D$ for all possible probability distributions Pr. This is denoted "$\models_{DP} D$". D is a consequence of a set of sentences Γ if there are $D_1, \ldots, D_n \in \Gamma$ such that $\models_{DP} (D_1 \wedge \ldots \wedge D_n) \rightarrow D$. D is consistent if $\neg D$ is not a theorem of DP.*

To obtain qualitative rules about default and likelihood from the quantitative rules in DP, we can perform order of magnitude reasoning. We can consider a representative default error, ϵ, where ϵ might be less than 0.01, or whatever the decision context requires. Likewise, we can consider a representative default likelihood, e, where e might be greater than 0.05, say. The choice for modelling particular limits rather than some arbitrary infinitesimal is motivated by the decision theoretic argument at the beginning of Section 3.3. In order to approximate the behaviour of our reasoning with these particular limits in mind, we can parameterise the system by ϵ and e and consider only approximate calculations to $O(\epsilon)$ and $O(e)$. A map translating probabilities to these kinds of qualitative values is given in Figure 2. The hashed regions represent those fuzzy boundaries where the qualitative reasoning becomes most susceptible to error.

QD_P is defined in a manner such that ϵ and e are arbitrarily small, but ϵ is also arbitrarily smaller than e. Of course, it is unrealistic to expect arbitrarily small magnitudes for ϵ and e to be achieved, let alone the right relative magnitudes. This, however, is irrelevant, as far as the application of the logic is concerned. The "arbitrarily small" magnitudes are only being used as a theoretical device to investigate the *approximate* behaviour of notions like "$\Rightarrow$" and "$\leadsto$" for ϵ being small and e being not quite as small (see also [7, Section 10.2.4]). In addition, the choice of relative magnitude between ϵ and e is a particular design decision that might just as well have been made some other way. Applications of QDP should of course take this into account.

Definition 4.5 *A sentence $D \in QD_P$ is a theorem of the qualitative probabilistic logic QPD if there exists a theorem $D' \in D_P$ corresponding to D (that is, identical except for any super or subscripts), in which all subscripts to "$\Rightarrow$" and "$\leadsto$" are parameterised by some variables ϵ and e and each subscript to "$\Rightarrow$" is of order ϵ as ϵ approaches 0 and e remains finite, and each subscript in D' corresponding to "$\leadsto_n$" in D is of order e^n as e and $\frac{\epsilon}{e}$ approach 0. This is denoted "$\models_{QDP} D$". Consequence and consistency are defined as before.*

From the definition of "$\leadsto$" (for $e > \epsilon$) it follows that

$$\models_{DP} \neg(\leadsto_e A \wedge \Mapsto_\epsilon \neg A) \, .$$

Consequently,

$$\models_{QDP} \neg(\leadsto A \wedge \Mapsto \neg A) \, . \tag{2}$$

That is, if something is likely, its negation cannot be true by default. But the complementary sentence $(\leadsto A \vee \Mapsto \neg A)$, is not a theorem.

It follows directly from this definition that the set of theorems of QDP is closed under application of *modus ponens* and conjunction. That is,

$$\models_{DP} D \text{ and } \models_{DP} D \rightarrow E \qquad \text{implies} \qquad \models_{DP} E \, ,$$
$$\models_{DP} D \text{ and } \models_{DP} E \quad \text{if and only if} \quad \models_{DP} D \wedge E \, .$$

Because the definition of QDP is based on an order of magnitude argument, there are potential pitfalls with these closure properties. Order of magnitude arguments invariably give dubious results when the constant factors become too large. Suppose a lottery has 1,000,000 participants. The following sentence can be shown to be a theorem of DP.

$$\bigwedge_{i=1}^{1,000,000} \Mapsto_\epsilon \langle \text{person } i \text{ will not win the lottery} \rangle \longrightarrow \Mapsto_{1,000,000*\epsilon} \langle \text{no-one will win the lottery} \rangle \, . \tag{3}$$

11

Moreover, replacing the error bound $1,000,000 * \epsilon$ by $999,999 * \epsilon$ yields a sentence that is not a theorem of QDP. Without the error bounds, the sentence would seem to read "if, by default, any particular person will not win the lottery, then, by default, no-one will win the lottery at all". The illusory lottery paradox has reappeared. In DP this is not the correct reading because with the natural value for ϵ, $\frac{1}{1,000,000}$, the right hand side of rule (3) is impotent (its default error is 1). In QDP unfortunately, it is the correct reading: QDP drops the subscripts (both are of order ϵ as ϵ approaches 0) and loses the error information.

If we wish a purely qualitative default logic to be closed under conjunction and *modus ponens*, two seemingly intuitive properties, then we have no choice but to accept that the above kind of anomaly may occur. People get around this with an intuitive knowledge of where plausible reasoning is likely to break down, for instance, by not making default or likelihood inference to any great depth: "don't rest your argument on too many assumptions, something is bound to go wrong along the way!". Default and likelihood reasoning may well produce incorrect results when carried on indefinitely; they should, however, be "locally" correct. Imprecision is an *inherent* property of plausible reasoning; so knowledge of how to contain the imprecision is a prerequisite for safe plausible reasoning. Hence the importance of DP in understanding QDP.

4.4 Basic theorems

This section introduces a few basic theorem schemata, and discusses several notable but unrelated properties of the logics. Examples of using the logic QDP are given later in Section 5.

First, the default and likelihood operators can be broken down into two components, according to whether the antecedent is possible or impossible. This is done using

$$\models_{QDP} \quad A \Rightarrow B \;\leftrightarrow\; (\Box\neg A \;\vee\; \Diamond A \wedge A \Rightarrow B) . \tag{4}$$

The second component here, $\Diamond A \wedge (A \Rightarrow B)$, is referred to as the *proper* default operator, and likewise for the likelihood operator. This corresponds to Adams' notion of the conditional over "proper" distributions [19, p49], that is, distributions where the antecedent of the operator must be possible. The unmodified, improper version of the default, $A \Rightarrow B$, corresponds to Adams' original notion of the conditional [30]. While the mathematics of the improper default is generally easier, it is sometimes better to break down the default and likelihood operators into the two components, and then put the pieces back at the end.

Second, both DP and QDP can be seen as natural extensions to propositional logic. For instance, the theorems for "$\Box$" given later in Table 1 encode the provability relation in propositional logic. The following lemma further highlights the connection.

Lemma 4.1 *First, all substitution instances of the theorems and rules of inference of standard propositional logic that are sentences of D_P hold for DP. Second, in DP necessary equivalences can be substituted. That is,*

$$\models_{DP} \Box(A \leftrightarrow B) \;\rightarrow\; (D(A) \leftrightarrow D(B)) ,$$

where $D(A)$ denotes any sentence of D_P with an occurrence of the propositional formula A in a particular position. Corresponding results for QDP hold.

Third, some examples of theorem schemata of DP are given in Tables 1–3. These hold for d, e, ϵ and δ all less than $\frac{1}{2}$. Certain dual forms, either on "$\Box$" or on "$\Rightarrow$", are given in the third column. These are obtained by restructuring the formula and converting either "$\Box$" or "$\Rightarrow$" to their dual. In each case, either the original form or the dual form can be proven by the consistency or consequence theorems presented in Section 4.6. For each DP theorem in Tables 1–3, the QD_P sentence obtained be removing subscripts (and in the case of the duals for theorems T14 and T16, making the "$\rightsquigarrow_{ed}$" operator "$\rightsquigarrow^2$") is a theorem of QDP.

One important aspect of any DP theorem is the relationship between errors on the defaults and

likelihoods. For instance, we can rewrite theorem T17' as

$$(A \approx_e C) \wedge (B \approx_d C) \rightarrow A \vee B \approx_f C \,,$$

and note that this only holds for some values of f, and in particular holds for

$$f \leq \frac{ed}{e + d - ed} \leq \min(e, d) \,.$$

In this case, f represents an *error propagation function*, which relates the errors in the DP theorem. If we were to apply this theorem in some chain of reasoning to deduce $A \vee B \approx C$, then we could either choose to forget about the error f, as we implicitly do when using QDP, or we could use the error propagation function to compute a value for f from e and d. Bear in mind that an error propagation function only represents a worst-case bound on error. If we were to do a more precise probabilistic analysis, we may find that error has shrunk to nothing, however, the error propagation function represents an upper-bound on what error can be. In Section 4.6 it is shown that for small errors, DP behaves just like QDP, so a system for reasoning about defaults and likelihoods can be constructed using the qualitative logic QDP, and then optionally, error tracking facilities can be grafted on top with the use of error propagation functions to give approximate probabilistic reasoning.

Finally, theorems of the logics can be generalised by uniformly changing conditioning information.

Lemma 4.2 *Any theorem of DP (QDP) can be transformed to another by uniformly changing conditioning information. Given conditioning information C, a formula D is transformed by uniformly applying the following transformations to all non-propositional operators in "$\diamond$true $\rightarrow$ D:*

$$
\begin{aligned}
\Box A &\mapsto \Box(C \rightarrow A) \\
\diamond A &\mapsto \diamond(C \wedge A) \\
B \Rightarrow_e A &\mapsto (C \wedge B) \Rightarrow_e A \\
B \approx_e A &\mapsto (C \wedge B) \approx_e A
\end{aligned}
$$

Versions of some of the theorems extended using this transformation are given in Table 4. Notice that for theorems T11', T12', T13', and T14', the initial term "$\diamond C \rightarrow$" has been dropped: this is safe because the "$\Rightarrow$" and "$\approx$" operators are always true and false respectively if the conditioning part is necessarily equivalent to false. A similar situation holds for theorem T6'.

4.5 Relevance

The antecedent of a default or likelihood corresponds to the context in which the rule can be applied. So the rule $B \Rightarrow C$ can be applied when we know just B, nothing more or less. This feature is inherited from the semantics of the conditioning operator in probability theory. As a result, defaults and likelihoods cannot have their antecedents arbitrarily specialised. That is, the QD_P sentence

$$(B \Rightarrow C) \rightarrow (A \wedge B \Rightarrow C)$$

is not a theorem of QDP; so the context B cannot in general be specialised to include other information, in this case A.

A second related feature of the logics is that there is no transitive relation applying to defaults or likelihoods. The same holds for NP [1, Section 7]. That is, the QD_P sentence

$$(A \Rightarrow B) \wedge (B \Rightarrow C) \rightarrow A \Rightarrow C$$

is not a theorem of QDP. For instance, a counterexample to this transitive sentence is that penguins

13

are birds, most birds fly, and penguins do not fly. So we would not expect the sentence to be a theorem. However, if we are told that the yellow-bellied warbler is a bird, and know nothing else about it, it is quite plausible to us that the warbler should fly.

So for plausible reasoning in certain situations, we would like some form of transitive reasoning. Notice the QD_P sentence

$$(A \Rightarrow B) \wedge (A \wedge B \Rightarrow C) \rightarrow A \Rightarrow C$$

is a theorem of QDP (T12$'$ in fact). Suppose we can obtain some additional information that implies the rule $B \Rightarrow C$ is the same as $A \wedge B \Rightarrow C$, so the condition A in the antecedent is not relevant. Then this additional information together with theorem T12$'$ shows the original transitivity form above does hold.

This ability to modify the antecedent of a default or plausible rule requires reasoning about relevance, where a condition in the antecedent is irrelevant if it can be added or deleted and still maintain the correctness of the rule. In probability theory, such information can be obtained in a number of ways. We can represent this information using the notion of independence, and in a more limited sense, following Neufeld *et al.* [20], the notion of favouring.

Definition 4.6 *Proposition A is* independent of *proposition B given proposition C if*

$$Pr(B|C) = Pr(B|C \wedge A) \ .$$

Proposition A favours *proposition B given proposition C if*

$$Pr(B|C) < Pr(B|C \wedge A) \ .$$

Lemma 4.3 *If proposition A is independent of proposition B given proposition C then the following sentences of QD_P are true:*

$$C \Rightarrow B \ \longleftrightarrow \ (C \wedge A) \Rightarrow B \ ,$$
$$C \Rightarrow A \ \longleftrightarrow \ (C \wedge B) \Rightarrow A \ ,$$
$$C \not\Rightarrow A \ \longleftrightarrow \ (C \wedge B) \not\Rightarrow A \ ,$$
$$C \not\Rightarrow B \ \longleftrightarrow \ (C \wedge A) \not\Rightarrow B \ .$$

If proposition A favours proposition B given proposition C then the sentences above only hold for the forward direction, that is, replacing "$\leftrightarrow$" by "$\rightarrow$".

It should be clear from this lemma that methods for reasoning about relevance are vital in plausible reasoning in order to modify plausible rules so they can be applied to each particular context. Some examples are given in Section 5. Causal (or Bayesian) networks can be used for this form of reasoning, and the maximum entropy method provides a way of making independence assumptions "by default" [7].

4.6 Consistency and consequence

The question of whether a sentence from D_P is consistent can be converted to the question of whether one of a set of simplex problems in the 2^n variables $\{Pr(p)|p \in E_P\}$ has a solution. Consequently, DP is decidable (this is similar to Probabilistic Logic [32]). For the purposes of this paper, it is not worth obtaining axiom schemata and rules of inference for the whole of DP, since we are really only interested in the case where the errors are quite small. A system encompassing the whole of DP would most likely degenerate to the kind found in [33, p10], where the schemata is close to an enumeration of primitive operations in the simplex algorithm. Fortunately, a different approach is available. Adams [30,19] has developed tests for consistency and entailment in his conditional logic, which have been extended by Goldszmidt and Pearl [34]. Similar consistency and consequence tests are presented below for the default and likelihood components of DP, and are easily adapted to QDP. These results show that reasoning can

14

be performed using the qualitative system QDP, and the approximate error bounds of DP propagated concurrently.

Tests on consistency and consequence are presented below in terms of a clausal form. Consider the default component of QDP. An arbitrary sentence containing the default, possibility and necessity operators can be turned into a conjunction of clauses, where each clause has the form

$$\Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \Rightarrow B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \Rightarrow H_i \ ,$$

for some index sets I_V, I_A, and I_C. Notice that all necessity and possibility operators have been gathered in the antecedent of the clause, by converting $\neg \Box A$ to $\Diamond \neg A$ where necessary, and all the necessity operators have been combined into one using theorem T3.

It is also of interest, though not essential for the development of this section, to consider a more precise interpretation of what it means for a clause to be a theorem in QDP. Lemma 4.4 uses the above clausal form to reinterpret the definition of a QDP theorem.

Lemma 4.4

$$\models_{QDP} \Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \Rightarrow B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \Rightarrow H_i \ ,$$

if and only if there exists a δ and η such that for all $\epsilon < \eta$

$$\models_{DP} \Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \Rightarrow_\epsilon B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \Rightarrow_{\delta \epsilon} H_i \ .$$

For the D_P sentence in the lemma, δ is an *error propagation factor*, and $\delta \epsilon$ is an *error propagation function*, which in this case is linear. The larger the value of δ, the faster error can propagate when this particular clause is applied in some chain of reasoning. By comparison, Adams' notion of entailment corresponds to:

if and only if for all ϵ there exists a δ such that

$$\models_{DP} \Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \Rightarrow_\delta B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \Rightarrow_\epsilon H_i \ .$$

The difference between the two notions is that in QDP error is restricted to propagate linearly.

Likewise, we can convert sentences containing the likelihood, necessity and default operators in a clausal form. The corresponding notion of a QDP theorem is given in Lemma 4.5.

Lemma 4.5

$$\models_{QDP} \ \Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \approx^{n_i} B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \approx^{m_i} H_i \ ,$$

if and only if there exists a δ and η such that for all $\epsilon < \eta$

$$\models_{DP} \ \Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \approx_{\epsilon^{n_i}} B_i \ \longrightarrow \ \vee_{i \in I_C} G_i \approx_{\delta \epsilon^{m_i}} H_i \ .$$

In this case, the error propagation functions, $\delta \epsilon^{m_i}$, are polynomial, and δ is the error propagation factor. Since a smaller likelihood represents more room for error, the smaller the value of δ, the faster error will propagate when this particular clause is applied in some chain of reasoning.

Results below on consistency and consequence of clauses using the default operator are extensions of several theorems in [30,19], and similar extensions can be found in [34], although Adams' terminology is not used here. The extensions introduce necessity and possibility. Consistency turns out to be the operation on which the three kinds of consequence tests are based.

Logical tests for consistency and consequence are given in Theorem 4.6 for clauses containing the default operator. These are given for DP and, because the error propagation functions are linear, can be extended to QDP simply by dropping the error subscripts. While the consistency test in the theorem with its "there exists a subset J" looks fairly involved, the kind of trick used in [34, Sect. 4] can be applied to develop an algorithm that constructs the subset efficiently.

Theorem 4.6 *Consider the D_P sentence D given by*

$$\Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \Rightarrow_{\epsilon_i} B_i \ ,$$

where $\epsilon_i < \frac{1}{|I_A|}, \frac{1}{2}$ for $i \in I_A$.

1. *The sentence D is inconsistent if and only if there exists some $j \in I_V$ such that $U \wedge V_j$ is unsatisfiable or there exists some $J \subseteq I_A$ such that*

$$U \ \wedge \left(\vee_{i \in J} A_i \right) \ \wedge_{i \in J} \left(A_i \to B_i \right)$$

is unsatisfiable.

2. *The D_P sentence $C \Rightarrow_\delta B$ is a consequence of D for some $\delta < \frac{1}{2}$ if and only if $D \wedge (C \Rightarrow_\delta \neg B)$ is inconsistent. This holds if and only if D itself is inconsistent or there exists some $J \subseteq I_A$ such that*

$$U \ \wedge (C \vee_{i \in J} A_i) \ \wedge_{i \in J} \left(A_i \to B_i \right) \ \wedge (C \to B)$$

is unsatisfiable. If $C \Rightarrow_\delta B$ is a consequence for some δ and can be demonstrated so using J, then $\delta = \sum_{i \in J} \epsilon_i$ is a correct error propagation function.

3. *The D_P sentence $\Diamond C$ is a consequence of D if and only if $D \wedge \Box \neg C$ is inconsistent.*

4. *The D_P sentence $\Box C$ is a consequence of D if and only if D itself is inconsistent or $\models U \to C$.*

When determining the consequences of a consistent set, whether a possibility is a consequence may depend on all elements of the set including the defaults, whereas whether a necessity is a consequence depends on only the other necessities. In this case, necessity can only follow from other necessities or inconsistency. Also note that the theorem shows error propagates additively when reasoning with default rules. This is a clear warning against long chains of such reasoning.

The resultant algorithm for checking consistency of QDP sentences is given in Figure 3.

Corollary 4.6.1 *The defaults-consistency algorithm is correct and uses at most $|I_V| + |I_A|^2 / 2$ satisfiability tests on the underlying propositional logic.*

The first step of this algorithm also forms the basis of testing the consistency of sentences containing only the necessity and possibility operator. Any such sentence can be converted to a conjunctive normal form consisting of a disjunction of conjuncts of the form $\Box U \wedge_{i \in I_V} \Diamond V_i$. Each conjunct can be tested for consistency using the first step.

Corollary 4.6.2 *Let the QDP sentence D containing only the necessity and possibility operators be in conjunctive normal form, and let $|D|$ denote the number of modal operators in the sentence. Then the consistency of D can be determined using less than $|D|$ satisfiability tests on the underlying propositional logic.*

The drawback with this result, however, is that the size of the conjunctive normal form of a sentence can be exponential in the size of the original sentence.

Tests for consistency and consequence using the likelihood operator are given in Theorem 4.7.

Theorem 4.7 *Consider the D_P sentence D given by*

$$\Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \rightsquigarrow_{e_i} B_i \ ,$$

where $e_i < \frac{1}{|I_A|}$ for $i \in I_A$. Let I_{min} denote the least subset of I_A, I, such that $U \wedge_{i \in I} \neg A_i \wedge A_j \wedge B_j$ is satisfiable for all $j \in I_A - I$. Such a minimum set is unique.

16

1. The sentence D is inconsistent if and only if there exists some $j \in I_V$ such that $U \wedge_{i \in I_{min}} \neg A_i \wedge V_j$ is unsatisfiable.

2. The D_P sentence $C \approx_f B$ is a consequence of D for some $f < 1$ if and only if D is inconsistent or there exists an ordered subset of the indices in $I_A - I_{min}$, $i_1, i_2, \ldots, i_h$, possibly empty ($h = 0$), such that for $j = 1, \ldots, h$,

$$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \wedge A_{i_j} \wedge B_{i_j} \ \wedge_{k < j} \neg A_{i_k} \ \rightarrow (C \wedge B) \qquad , and \tag{5}$$

$$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \wedge_{k \leq h} \neg A_{i_k} \ \rightarrow \ (C \rightarrow B) \ . \tag{6}$$

If consequence holds, then a lower bound on f, the error propagation function, is given by

$$f \ \geq \ \left(\frac{e}{1+e}\right)^h , \qquad where \ e = \min_{1 \leq k \leq h} e_{i_k} \ ,$$

although the error propagation function can be linear in the e_i in some cases.

3. The D_P sentence $\diamond C$ is a consequence of D if and only if $D \wedge \square \neg C$ is inconsistent.

4. The D_P sentence $\square C$ is a consequence of D if and only if D is inconsistent or

$$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \rightarrow C \ .$$

There is also a special case of this theorem that applies to non-iterated versions of the likelihood operator.

Corollary 4.7.1 *Consider the QD_P sentence D given by*

$$\square U \ \wedge_{i \in I_V} \diamond V_i \ \wedge_{i \in I_A} A_i \approx B_i \ .$$

The QD_P sentence $C \approx B$ is a consequence of D if D is inconsistent, or there exists some $I \subset I_A - I_{min}$ such that for $j \in I$,

$$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \wedge A_j \wedge B_j \rightarrow (C \wedge B) \qquad , and \tag{7}$$

$$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \wedge_{j \in I} \neg A_j \ \rightarrow \ C \rightarrow B \ . \tag{8}$$

I conjecture that the converse of this theorem also holds. The dual form of the corollary, converted to apply to defaults, allows a disjunction of defaults to be the consequence of a single default. An example of this corollary is theorem T17 and its dual.

An algorithm for checking consistency of QDP sentences is given in Figure 4. Step 2(b) has been added to this to make the algorithm more efficient when some of the likelihood operators are proper.

Corollary 4.7.2 *The likelihood-consistency algorithm is correct and uses at most $|I_V| + |I_A|^2/2$ satisfiability tests on the underlying propositional logic.*

An algorithm for checking consequence is given in Figure 5. This algorithm assumes the consistency check has already been made. The error propagation function in this case can be taken from Theorem 4.7 part 2, and a tighter error propagation function is given in the proof of that theorem.

Corollary 4.7.3 *The likelihood-consequence algorithm is correct and uses at most $(|I_A| + 1)^2/2$ satisfiability tests on the underlying propositional logic.*

5   Applications

This sections demonstrates the use of the qualitative logic QDP on three anecdotal problems that reoccur in the default reasoning literature.

The first example resolves the paradox of the "vanishing subclasses". The second example demonstrates how reasoning about independence using causal networks can be integrated with the forms of plausible reasoning just developed. The final example is the classic Yale shooting problem [8]. This example highlights a subtle problem with the situation calculus when it is used for plausible reasoning.

5.1 The "vanishing" emus

Neufeld *et al.* have criticised the modelling of default reasoning based on infinitesimal probabilities [20, p123] on the grounds that it makes "subclasses vanish". Consider the following rules:

$$Emu \to Bird , \tag{9}$$

$$Emu \Rightarrow \neg Flies , \tag{10}$$

$$Bird \Rightarrow Flies . \tag{11}$$

The following are consequences.

$$Bird \Rightarrow \neg Emu$$

$$\Rightarrow \neg Emu$$

We can conclude that "typically, birds aren't emus" and "typically, things aren't emus". To show the first is a consequence using Theorem 4.6, notice $U = (Emu \to Bird)$, $I_V = \emptyset$, and try to show the rules together with $Bird \Rightarrow Emu$ is inconsistent. This follows because the rules themselves are consistent and

$$U \wedge (Emu \vee Bird) \wedge (Emu \to \neg Flies) \wedge (Bird \to Flies) \wedge (Bird \to Emu)$$

is unsatisfiable.

If we take the $O(\epsilon)$ semantics of the default operator literally then we could conclude, since ϵ is infinitesimal, that "no birds are emus", or "nothing is an emu". The real intent of the semantics, however, is about approximations for ϵ small. So instead we should conclude that the emu is just an uncommon or non-typical bird, which in reality is true of emus. The approximate probabilistic semantics does not cause subclasses to vanish; but it may cause you to deduce some subclasses must be non-typical.

5.2 Can Joe read and write

The importance of independence in default reasoning, and plausible reasoning generally, has been underlined by Pearl in his simple problem "can Joe read and write?" [7, Sct. 10.3]. This is a good example of why general transitivity should not hold for default reasoning. A twist is also given at the end to show how likelihood reasoning can complement default reasoning.

Pearl introduces the propositions (I have altered the symbols)

$$
\begin{aligned}
Over\text{-}7 &\equiv \quad \text{Joe is over 7 years old ,} \\
RdWr &\equiv \quad \text{Joe can read and write ,} \\
EngPrf &\equiv \quad \text{Joe's father is a Professor of English ,} \\
Shakes &\equiv \quad \text{Joe can recite passages from Shakespeare .}
\end{aligned}
$$

and the default rules (expressed in QDP)

$$
\begin{aligned}
RdWr &\Rightarrow Over\text{-}7 , \\
EngPrf &\Rightarrow RdWr , \\
Shakes &\Rightarrow RdWr .
\end{aligned} \tag{12}
$$

Let $\Delta_{literacy}$ denote ruleset (12). Pearl also assumes that Joe is over 6 years old and is not retarded, so that the default rules above seem reasonable.

Given, in addition, that Joe recites Shakespeare, Pearl argues that a reasonable conclusion is that Joe is over seven years old. That is, we want to be able to infer the default rule

$$Shakes \Rightarrow Over\text{-}7 . \tag{13}$$

On the other hand, given that Joe's father is a Professor of English, it is *not* a reasonable conclusion that Joe is over seven years old. An argument being that Joe's father's profession adequately explains Joe's literacy, so we don't need the more common explanation that Joe is over seven years old. We do not want to be able to infer the default rule

$$EngPrf \Rightarrow Over\text{-}7 . \tag{14}$$

The problem with the formulation at present is that the constraints on *Shakes* and *EngPrf* are syntactically identical, but we hope to infer conflicting default rules for them. In QDP (and in NP) it happens that *neither* default rule (13) nor (14) can be derived. We do get, however, that

$$\Delta_{literacy} \models_{QDP} EngPrf \Rightarrow Over\text{-}7 \longleftrightarrow (EngPrf \wedge RdWr) \Rightarrow Over\text{-}7 ,$$
$$\Delta_{literacy} \models_{QDP} Shakes \Rightarrow Over\text{-}7 \longleftrightarrow (Shakes \wedge RdWr) \Rightarrow Over\text{-}7 . \tag{15}$$

The problem as it stands is underconstrained. So, what information is missing?

Pearl's solution to the problem introduces the notion of causality. For instance, Joe's literacy is a partial cause (and the only direct one occurring in the formulation) of Joe being able to recite Shakespeare. What Pearl alludes to but never explicitly mentions is the causal network (a Directed-Acyclic Graph (DAG) [35]) given in Figure 6. In this network, arcs correspond to the intuitive notion of "can cause".

As Pearl and Verma explain, such a causal network provides information about independence [35, definition for DAGD, p376]. It should be pointed out that the notion of causality is merely incidental to their analysis: it serves as a useful, intuitive focus for acquiring knowledge about independence. We can subsequently apply the dependence information so obtained to default and likelihood reasoning using Lemma 4.3.

Applying Pearl and Verma's technique of deducing independence relations to Figure 6, we get that Joe's Shakespearean recital is independent of Joe being over seven, given he is literate. In QDP, it follows that

$$RdWr \Rightarrow Over\text{-}7 \longleftrightarrow (Shakes \wedge RdWr) \Rightarrow Over\text{-}7 .$$

Let us denote by Γ_1 the dependence information obtainable from Figure 6. Together with the default conclusion (15), we get

$$\Delta_{literacy} \cup \Gamma_1 \models_{QDP} Shakes \Rightarrow Over\text{-}7 .$$

The same does not hold for *EngPrf*, however, because in contrast we get that Joe's father's profession is not independent of Joe being over seven, given Joe is literate. Because of this, the truth or falsehood of default rule (14) is undetermined from $\Delta_{literacy}$ and Γ_1. But, if we were also told that it is likely for a child of a Professor of English to be under seven years old and literate, then default rule (14) becomes false as required. That is[4].

$$\Delta_{literacy} \cup \{EngPrf \rightsquigarrow (RdWr \wedge \neg Over\text{-}7)\} \models_{QDP} \neg(EngPrf \Rightarrow Over\text{-}7) .$$

---

[4] Derive this result as follows. From $EngPrf \rightsquigarrow (RdWr \wedge \neg Over\text{-}7)$, the conditioned version of theorem T10, and the conditioned version of the theorem given in Equation (2), infer $\neg(EngPrf \wedge RdWr \Rightarrow Over\text{-}7)$. Finally, combine this with default rule (15).

5.3 The Yale shooting problem

A second problem that needs to incorporate independence for a solution is the Yale shooting problem [8]. This problem has been the subject of considerable discussion in AI, and it is beyond the scope of this paper to give a reasonable survey. In this section, the specific solution of Delgrande [36, Section 6.2] is considered. In probabilistic reasoning it is important to differentiate between what is currently known, and what is not. However, the situation calculus, in which the Yale shooting problem is usually presented, allows the representation of knowledge about static properties of a state but represses the representation of knowledge about events. This causes problems in the subsequent representation of defaults, which we discuss below.

The Yale shooting problem can be presented briefly as follows: a gun is loaded; one waits for a moment; a shot is fired. We should conclude by default that the person is dead, assuming, of course, the gun was well aimed at the person, etc. Early default reasoning systems could not make this conclusion; during the wait, the gun would not stay loaded by default.

Delgrande [36, Section 6.2] initially suggested a situation calculus representation of this problem in NP that in QDP becomes:

$$\Box T(Alive, S_0) \, ,$$
$$\Box T(Loaded, Result(Load, s)) \, ,$$
$$T(Loaded, s) \Rightarrow T(Dead, Result(Shoot, s)) \, ,$$
$$T(f, s) \Rightarrow T(f, Result(e, s)) \, .$$

Variables are given by e, f and s, and state S_0 is some constant starting state. The first sentence reads "*Alive* is necessarily true in state S_0", the third "if *Loaded* is true in some state s then typically *Dead* will be true in the state resulting from a *Shoot* in state s", etc. Assume $Result(S_n)$ is denoted S_{n+1}. To adequately handle the shooting problem we now wish to infer that contingent on a certain sequence of events taking place, a death will occur.

$$T(Load, S_0) \wedge T(Wait, S_1) \wedge T(Shoot, S_2) \Rightarrow T(Dead, S_3) \, .$$

As Delgrande points out, this formulation cannot be correct. From the second sentence and theorem T6' we get

$$\neg T(Loaded, s) \Rightarrow T(Loaded, Result(Load, s)) \, ,$$

and together with an instance of the fourth sentence $(f = \neg Unloaded)$,

$$\neg T(Loaded, s) \Rightarrow \neg T(Loaded, Result(Load, s)) \, ,$$

from theorem T7' we get

$$\Box T(Loaded, s) \, .$$

That is, the gun is *always* loaded! If we added an *Unload* event to the above formulation that resulted in the gun being unloaded, we could similarly deduce that the gun is *always* unloaded!

Delgrande suggests repairing this conflicting state of affairs by changing the last sentence to (assuming that equality is introduced)

$$(f = Alive) \vee (T(f, s) \Rightarrow T(f, Result(e, s))) \, ,$$
$$(e = Shoot) \vee (T(Alive, s) \Rightarrow T(e, Result(Shoot, s))) \, ,$$

which together say people do not tend to remain alive if they are shot, or changing the second last to

$$(T(Alive, s) \wedge T(Loaded, s)) \Rightarrow T(Dead, Result(Shoot, s)) \, .$$

20

In addition, we will have to take this kind of evasive action for every event type. Adopting the first strategy, the simple concept "things tend to stay the same" is starting to look decidedly lengthy. We are required to explicitly detail all those exceptions default reasoning is supposed to circumvent. The second strategy seems to introduce an unnecessary complication: if you shoot a dead person they will remain dead, so why bother specifying they should be alive before the shooting.

The real problem lies with the representation of knowledge about events. *Without knowing* which event occurs at a state, we know things will tend to stay the same. Once we *know* which particular event occurs, however, we also know for sure that certain things will change. The antecedents in the conditionals in Delgrande's formulation need to be qualified with knowledge about events to block the conflict between the second and fourth sentences. We do this by modifying the sentences to allow explicit representation of knowledge about events:

$$\Box T(Alive, S_0) \, ,$$
$$\Box(T(Load, s) \rightarrow T(Loaded, Next(s))) \, ,$$
$$T(Loaded, s) \wedge T(Shoot, s) \Rightarrow T(Dead, Next(s)) \, ,$$
$$T(f, s) \Rightarrow T(f, Next(s)) \, ,$$

where $T(e, s)$ about an event e such as *Shoot* denotes that it is known that the event e occurred in situation s, and $Next(s)$ denotes the state after state s. Denote this set of sentences by $\Delta_{shooting}$.

But with the problem as formulated in $\Delta_{shooting}$, the required result is not forthcoming. Again we need information about relevance to show how the redrafted sentences can have their antecedents sufficiently specialised.

First, the following can be inferred from the third rule in $\Delta_{shooting}$ given that if a loaded gun is shot at someone, then events strictly prior to the shooting are independent of possible death,

$$T(Load, S_0) \wedge T(Wait, S_1) \wedge T(Shoot, S_2) \wedge T(Loaded, S_1) \wedge T(Loaded, S_2) \Rightarrow T(Dead, S_3) \, .$$

Second, the following can be inferred from the fourth rule in $\Delta_{shooting}$ given that whether a gun stays loaded is only dependent on prior *Unload* or *Shoot* events.

$$T(Load, S_0) \wedge T(Wait, S_1) \wedge T(Shoot, S_2) \wedge T(Loaded, S_1) \Rightarrow T(Loaded, S_2) \, .$$

This information about independence, call it Γ_2, is sufficient to yield the required result.

$$\Delta_{shooting} \cup \Gamma_2 \models_{QDP} T(Load, S_0) \wedge T(Wait, S_1) \wedge T(Shoot, S_2) \Rightarrow T(Dead, S_3) \, .$$

Notice that Γ_2 could have been obtained automatically using the "default" independent assumptions inherent in a maximum entropy approach [7].

The specification of Γ_2 can be seen to involve as much detail as Delgrande's earlier suggestion. So where is the advantage? The defaults remain in a simple form, and the exceptions are instead coded in the modular form of causal (independence) information about events.

6   Further Comparisons

This section compares the logics DP and QDP with some related approaches. Halpern and Rabin's and Halpern and McAllester's likelihood logics, and Neufeld *et al.* influence graphs are compared because they have also been motivated by probability. Comparisons with Adams' conditional logic have been sprinkled throughout Section 4, and are not reiterated here. The last comparison given here is with Delgrande's NP; this system had an historical influence on the logics DP and QDP.

6.1 Halpern and Rabin's likelihood logic

Halpern and Rabin propose the unary likelihood operator L with semantics [14, p386]

> Lp is best thought of as saying "p is reasonably likely to be a consistent hypothesis."

This should not be confused with "p is reasonably likely", the interpretation Halpern and McAllester give to Lp [9, p5].

For instance, suppose a lottery with $1,000,000$ tickets is being held, then the following can be deduced by applying their Axiom AX6 repeatedly:

$$L(\text{someone will win the lottery}) \longleftrightarrow \bigvee_{i=1}^{1,000,000} L(\text{person } i \text{ will win the lottery}) . \tag{16}$$

The right hand side of this equivalence reads, there exists a particular person who is likely to win the lottery. In the Oxford dictionary sense of the word "likely", this is certainly not true before the lottery is held. So in the Halpern-McAllester interpretation, the sentence (16) above can be interpreted as $true \leftrightarrow false$. This is a variant of the lottery "paradox". Because they assume that likelihood reasoning is precise, they conclude that the Halpern-Rabin interpretation must be more appropriate.

By contrast, in the framework proposed here it is taken for granted that likelihood reasoning may be imprecise. As explained after Definition 4.5, QDP suffers from the lottery "paradox" in a sense, but it is viewed as an anomaly, an inherent consequence of modelling imprecise reasoning with a precise logic. Of course, such anomalies can be avoided by either using heuristics about plausible reasoning ("don't make too many assumptions"), or by resorting to numeric methods which allow more careful tallying of degrees of imprecision.

Notice that interpreting Lp to mean "p is a consistent hypothesis" yields the following transformation to QPD: $Lp \mapsto \diamond p$, and $Gp \mapsto \Box p$. Indeed, their axioms on non-iterated modalities each have a corresponding theorem in QPD.

Halpern and Rabin propose instead that iterated modalities of the form $L^i Gp$ be used to model "p is reasonably likely", and they give soundness results to support their claim. There is, however, a serious methodological problem with this approach: knowledge expressed in the form they propose is non-modular and cumbersome. A sentence such as "P_1 is reasonably likely given P_3" is represented in QDP simply as $P_3 \leadsto P_1$. In their logic it translates to

$$\neg G \neg P_1 \wedge \neg G P_2 \wedge \neg G P_4 \wedge G P_3 \Rightarrow L G P_1$$

in one situation, and

$$\neg G P_1 \wedge \neg G \neg P_2 \wedge \neg G P_4 \wedge G P_3 \Rightarrow L G P_1$$

in another. These cumbersome translations occur because, as they explain [9, p7], their representation has no means of making likelihood contingent on what is currently known (for instance, by using conditioning, the role played by the left hand side of the "$\leadsto$" operator). Worse still, if the model (and consequently the atomic propositions used) becomes extended, the appropriate translation must be extended as well. In addition, Halpern and Rabin give no evidence that non-trivial theorems hold about iterated modalities of the form $L^i G$.

6.2 Neufeld and Poole's favouring formalism

Neufeld *et al.* present *influence graphs*, a qualitative system for reasoning about *favouring* [20] that is related to Suppes' causal algebra [37]. B favours A when $Pr(A|B) > Pr(A)$. It was argued in Section 4.5 that favouring provides an important complement to the logics presented here.

22

Favouring alone, however, is not sufficient information on which to base a decision. This stems from the fact that favouring is for reasoning about *shift* in belief and not current strength in belief. For instance, it is well known now that smoking favours cancer (that is, a smoker is more likely to have cancer than a non-smoker). But the knowledge that a person smokes is not sufficient evidence on which to base a conclusion that the person has cancer. It merely provides an additional degree of support for such a conclusion.

6.3 Delgrande's conditional logic NP

There is a strong correspondence between the theorems of QDP and Delgrande's NP [1]. The only axiom of NP that is not also a theorem of QDP is the CV axiom given by

$$\neg(A \Rightarrow B) \;\rightarrow\; (A \Rightarrow C) \;\rightarrow\; (A \wedge \neg B) \Rightarrow C \;,$$

although this is similar to the QDP theorem T14'. Notice, however, that, by adapting T14' we get that

$$\models_{DP} \neg(A \Rightarrow_{\epsilon\delta/2} B) \;\rightarrow\; (A \Rightarrow_{\epsilon} C) \;\rightarrow\; (A \wedge \neg B) \Rightarrow_{\delta} C \;.$$

This version of the CV axiom does not also become a theorem in QDP because the first default has an error that is a different order of magnitude to the second two defaults ($\epsilon\delta$ compared with ϵ and δ).

Also, necessity is introduced into QDP and NP in a very different manner. Nevertheless, theorems involving necessity in NP given in [1] are also theorems for QDP. Consequently, almost every theorem of NP that is a sentence of QD_P is also a theorem of QDP.

7 Conclusion

This paper has examined the problem of reasoning about defaults and likelihood from a probabilistic perspective. The presentation has been one of theoretical analysis, comparison with existing systems, and review of anecdotal examples. The approach developed has extended some existing systems [19,7, 1] and put some others in a clearer perspective [14]. This highlighted the approximate nature of the reasoning forms, the duality between them, and the need for complementary reasoning about relevance and error propagation. Algorithms have also been presented for determining some types of consistency and consequence for both logics, qualitative and quantitative.

The following research issues give some idea of how this area might be further developed.

Causality, independence [7] and favouring [20] play a complementary but vital role to default and likelihood reasoning. They help in the determination of relevance, for the derivation of plausible rules applicable to a system's current context. Suppose we have separate information about relevance and defaults. How might reasoning about both these forms be integrated? For instance, how can the consistency and consequence algorithms be interfaced with algorithms for reasoning about independence?

There is a remarkable similarity between Delgrande's conditional logic NP and the probabilistic logics presented here. With the necessity and possibility operators, the logics presented here have an ability to express sentences roughly in the realm of autoepistemic or default logics. What are the relationships to these other approaches?

How should the effect of the decision context be integrated? For instance, one would like to be able to obtain the default reasoning structure present in the layered control systems of Brooksian robots [38], where each layer is intended to handle a different class of decision problems. How might these layered systems be developed?

Given that defaults and likelihoods have been represented here as probabilistic rules, how might they be learned from data? Machine learning techniques for rule induction have been developed, but these only allowing one particular propositional symbol (or concept) in the consequence of the rule. Some methods are described in [39,40,25]. To learn a set of defaults and likelihoods, more general approaches are required

that simultaneously learn rules with a variety of different propositional symbols in the consequence, as found in [41].

At what point does qualitative reasoning of QDP have to be augmented with quantitative reasoning of DP to produce reliable results? Furthermore, when do the approximations inherent in DP break down so that a system needs to be developed using more thorough probabilistic reasoning? It may be necessary to reason about uncertainty using approximate numeric techniques, and to use the plausible logics developed here merely at the man-machine interface. For instance, one observable use of default and likelihood reasoning in people is explanation and presentation of results.

Implementation and application to real problems is clearly one important way to explore these plausible reasoning forms further.

Acknowledgements

References

[1] J. Delgrande, "A first-order conditional logic for prototypical properties—extended report," LCCR TR 86-15, Simon Fraser University, Burnaby, Canada, 1986.

[2] R. Reiter, "A logic for default reasoning," *Artificial Intelligence*, vol. 13, pp. 81–132, 1980.

[3] D. McDermott, "Nonmonotonic logic II: nonmonotonic modal theories," *JACM*, vol. 29, pp. 33–57, January 1982.

[4] J. McCarthy, "Applications of circumscription to formalising common-sense knowledge," *Artificial Intelligence*, vol. 28, pp. 89–116, 1986.

[5] M. Ginsberg, ed., *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann, 1987.

[6] Y. Shoham, *Reasoning about Change.* Cambridge, Massachusetts: MIT Press, 1988.

[7] J. Pearl, *Probabilistic Reasoning in Intelligent Systems.* Morgan and Kauffman, 1988.

[8] S. Hanks and D. McDermott, "Nonmonotonic logic and temporal projection," *Artificial Intelligence*, vol. 33, pp. 379–412, 1987.

[9] J. Halpern and D. McAllester, "Likelihood, probability, and knowledge," RJ 4313 (47141), IBM Research Laboratory, San Jose, CA, 1984.

[10] G. Polya, *Patterns of Plausible Inference.* Princeton University Press, 1954.

[11] R. Yager, "Using approximate reasoning to represent default knowledge," *Artificial Intelligence*, vol. 31, pp. 99–112, 1987.

[12] M. Ginsberg, "Non-monotonic reasoning using Dempster's rule," in *Fourth National Conference on Artificial Intelligence*, (Austin, Texas), pp. 126–129, 1984.

[13] D. Dubois and H. Prade, "On fuzzy syllogisms," *Computational Intelligence*, vol. 4, no. 2, pp. 171–179, 1988.

[14] J. Halpern and M. Rabin, "A logic to reason about likelihood," *Artificial Intelligence*, vol. 32, pp. 379–405, 1987.

[15] P. Cheeseman, "In defence of probability," in *Proceedings of IJCAI-85*, pp. 1002–9, Los Angeles, CA: Kaufmann, 1985.

[16] J. Nutter, "Uncertainty and probability," in *International Joint Conference on Artificial Intelligence*, (Milan), pp. 373–379, 1987.

[17] E. Horvitz, D. Heckerman, and C. Langlotz, "A framework for comparing alternative formalisms for plausible reasoning," in *Fifth National Conference on Artificial Intelligence*, (Philadelphia), pp. 210–214, 1986.

[18] R. Moore, "Semantical considerations on nonmonotonic logic," *Artificial Intelligence*, vol. 25, pp. 75–94, 1985.

[19] E. Adams, *The Logic of Conditionals*. Boston: Reidel, 1975.

[20] E. Neufeld, D. Poole, and R. Aleliunas, "Probabilistic semantics and defaults," in *Uncertainty in Artificial Intelligence 4* (R. Schachter, T. Levitt, L. Kanal, and J. Lemmer, eds.), North Holland, 1990.

[21] Y. Shoham, "Nonmonotonic logics: meaning and utility," in *International Joint Conference on Artificial Intelligence*, (Milan), pp. 388–393, 1987.

[22] M. Henrion and D. Cooley, "An experimental comparison of knowledge engineering for expert systems and for decision analysis," in *Sixth National Conference on Artificial Intelligence*, (Seattle), pp. 471–476, 1987.

[23] D. Heckerman, "Probabilistic interpretations for MYCIN's certainty factors," in *Uncertainty in Artificial Intelligence* (L. N. Kanal and J. F. Lemmer, eds.), Amsterdam: Elsevier Science Publishers, 1986.

[24] P. Cheeseman, M. Self, J. Kelly, W. Taylor, D. Freeman, and J. Stutz, "Bayesian classification," in *Seventh National Conference on Artificial Intelligence*, (Saint Paul, Minnesota), pp. 607–611, 1988.

[25] W. Buntine, *A Theory of Learning Classification Rules*. PhD thesis, University of Technology, Sydney, 1990. Forthcoming.

[26] C. Langlotz and E. Shortliffe, "Logical and decision theoretic methods for planning under uncertainty," *AI Magazine*, vol. 10, no. 1, pp. 39–48, 1989.

[27] E. Horvitz, J. Breeze, and M. Henrion, "Decision theory in expert systems and artificial intelligence," *International Journal of Approximate Reasoning*, vol. 2, pp. 247–302, 1988.

[28] J. Quinlan, "Inferno: A cautious approach to uncertain inference," *The Computer Journal*, vol. 26, no. 3, pp. 255–269, 1983.

[29] H. Geffner, "On the logic of defaults," in *Seventh National Conference on Artificial Intelligence*, (Saint Paul, Minnesota), pp. 449–454, 1988.

[30] E. Adams, "Probability and the logic of conditionals," in *Aspects of Inductive Logic* (J. Hintikka and P. Suppes, eds.), pp. 265–316, Amsterdam: North-Holland, 1966.

[31] P. Haddawy and A. Frisch, "Modal logics of higher-order probability," in *Uncertainty in Artificial Intelligence 4* (R. Schachter, T. Levitt, L. Kanal, and J. Lemmer, eds.), North Holland, 1990.

[32] N. Nilsson, "Probabilistic logic," *Artificial Intelligence*, vol. 28, pp. 71–87, 1986.

[33] R. Fagin and J. Halpern, "Reasoning about knowledge and probability," in *Second Conference on Theoretical Aspects of Reasoning about Knowledge*, (Vancouver), 1988.

[34] M. Goldszmidt and J. Pearl, "Deciding consistency of databases containing defeasible and strict information," in *Uncertainty in Artificial Intelligence 5* (M. Henrion, R. Schachter, L. Kanal, and J. Lemmer, eds.), Amsterdam: Elsevier Science Publishers, 1990. An extended version appears as UCLA Cognitive Systems Laboratory, Technical Report CSD-890034 (R-122).

[35] J. Pearl and T. Verma, "The logic of representing dependencies by graphs," in *Sixth National Conference on Artificial Intelligence*, (Seattle, Washington), pp. 374–379, 1986.

[36] J. Doyle, "An approach to default reasoning based on a first-order conditional logic: revised report," *Artificial Intelligence*, vol. 36, pp. 63–90, 1988.

[37] P. Suppes, *A Probabilistic Theory of Causality*. Amsterdam: North-Holland, 1970.

[38] R. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. 2, no. 1, 1986.

[39] J. Quinlan, "Generating production rules from decision trees," in *International Joint Conference on Artificial Intelligence*, (Milan), pp. 304–307, 1987.

[40] S. Weiss and I. Kapouleas, "An empirical comparison of pattern recognition, neural nets, and machine learning classification methods," in *International Joint Conference on Artificial Intelligence*, (Detroit), pp. 781–787, Morgan Kaufmann, 1989.

[41] P. Smyth and R. Goodman, "An information theoretic approach to rule induction from databases," *IEEE Trans. on Knowledge and Data Engineering.* to appear.

Appendix   Proofs of Lemmas and Theorems

Proof of Theorem 4.1   First, substitution instances of propositional logic hold for DP because the interpretation of theorems for DP is given in terms of propositional logic ("not", "and", etc.). Suppose $E \in QD_P$ is some substitution instance of a theorem of propositional logic. Consider $E' \in D_P$ obtained by transforming $\Rightarrow$ to $\Rightarrow_\epsilon$ and $\leadsto^n$ to $\leadsto_{\epsilon^n}$ for some ϵ and e. E' is also a substitution instance of propositional logic, so it is a theorem of DP. As this holds for arbitrary ϵ and e, E is a theorem of QDP.

Second, equivalences can be substituted. Suppose $Pr(A \leftrightarrow B) = 1$, then

$$Pr(C(A)) = Pr(C(A) \wedge (A \leftrightarrow B)) = Pr(C(B) \wedge (A \leftrightarrow B)) = Pr(C(B)),$$

and the result follows from the definition of a theorem in DP. A similar proof applies for QDP.   □

Proof of Lemma 4.2   Any sentence that holds for an arbitrary probability distribution must hold for an arbitrary probability distribution conditioned on some C given that $Pr(C) > 0$. That is, given that $Pr(C) > 0$, we can make the transformations

$$Pr(A) \;\mapsto\; Pr(A|C)$$
$$Pr(A|B) \;\mapsto\; Pr(A|B \wedge C)$$

and the sentence must still hold for any arbitrary probability distribution. This corresponds to the transformations given in the lemma. Notice there is no confusion in applying the transformations because the operators do not nest.   □

Proof of Lemma 4.4   First, notice that the order of magnitude definition of Definition 4.5 applies if and only if there exists constants c_i for $i \in I_A$ and d_i for $i \in I_C$ and η such that for all $\epsilon < \eta$ and probability distributions Pr,

$$\models_{Pr} \quad \Box U \; \wedge_{i \in I_V} \diamond V_i \; \wedge_{i \in I_A} A_i \Rightarrow_{c_i \epsilon} B_i \quad \longrightarrow \quad \vee_{i \in I_C} G_i \Rightarrow_{d_i \epsilon} H_i \; . \tag{17}$$

To show the only if part of the theorem, assume the δ, ϵ condition in the lemma holds. Then let $c_i = 1$ for $i \in I_A$ and $d_i = \delta$ for $i \in I_C$, so by above, the clause is a theorem of QDP.

To show the if part of the theorem, assume the clause is a theorem of QDP so constants c_i for $i \in I_A$ and d_i for $i \in I_C$ and η exist as above. Let $\alpha = min_{i \in I_A} c_i$, and $\delta = max_{i \in I_C} d_i / \alpha$, and $\eta' = \eta \alpha$. Pick any $\epsilon' < \eta'$ and note $\epsilon = \epsilon'/\alpha < \eta$. We now have that $1 - c_i \epsilon \leq 1 - \epsilon'$, and $1 - d_i \epsilon \geq 1 - \delta \epsilon'$. So if a probability distribution Pr satisfies the clause (17) using η and ϵ, then it also satisfies the D_P clause in the theorem using η' and ϵ'. So this clause is satisfied for every distribution.   □

Proof of Lemma 4.5   The proof proceeds as for Lemma 4.4 but using

$$\models_{Pr} \quad \Box U \; \wedge_{i \in I_V} \diamond V_i \; \wedge_{i \in I_A} A_i \leadsto_{c_i \epsilon^{n_i}} B_i \quad \longrightarrow \quad \vee_{i \in I_C} G_i \Rightarrow_{d_i \epsilon^{m_i}} H_i \; ,$$

instead of formula 17. There is a difference in showing the if part of the theorem; δ is now constructed in an inverse manner. Let

$$\alpha = max_{i \in I_C} \sqrt[n_i]{c_i} \qquad \text{and} \qquad \delta = min_{i \in I_C} \frac{d_i}{\alpha^{m_i}} \; ,$$

and proceed as before, noticing that we are dealing with quantities such as $\delta(\epsilon')^{n_i}$ rather than $1 - \delta \epsilon'$.   □

Proof of Theorem 4.6 part 1 First, assume D is consistent and we shall prove both unsatisfiability conditions fail. With D consistent, $Pr(U \wedge V_j) > 0$, so $U \wedge V_j$ is satisfiable, and the first unsatisfiability condition fails. To show the second unsatisfiability condition fails for each J, it is sufficient to show it fails for $J = I_A$, because if D is consistent, then any subset of D is consistent, so, correspondingly, the same reasoning applies for any subset J of I_A. We shall show this failure by contradiction. Assume the second unsatisfiability condition holds for $J = I_A$, then for any probability distribution such that $Pr(U) = 1$, $Pr(\vee_{i \in J} A_i \wedge \neg B_i) = Pr(\vee_{i \in J} A_i)$. Therefore,

$$max_{i \in J} \, Pr(A_i \wedge \neg B_i) \geq \frac{\sum_{i \in J} Pr(A_i \wedge \neg B_i)}{|J|} \geq \frac{Pr(\vee_{i \in J} A_i \wedge \neg B_i)}{|J|} \geq \frac{Pr(\vee_{i \in J} A_i)}{|J|} \geq \frac{Pr(A_j)}{|J|}$$

for any j. Therefore either $Pr(A_j) = 0$ for all $j \in I_A$, or for at least one $j \in I_A$, $Pr(\neg B_j | A_j) \geq \frac{1}{|J|}$. This gives the contradiction since $\epsilon_j < \frac{1}{|J|}$.

Second, assume both unsatisfiability conditions fail and we shall prove D is consistent. Notice because the second unsatisfiability condition fails for every $J \subseteq I_A$, there exists an ordered subset of I_A given by $i_1, \ldots, i_m$ (where $m = |I_A|$), such that $U \wedge A_{i_j} \wedge_{k \geq j} (A_{i_k} \rightarrow B_{i_k})$ is satisfiable for $j = 1, \ldots, m$. Let truth assignment t_j demonstrate this satisfiability for $j = 1, \ldots m$. Also let truth assignment t'_j demonstrate the satisfiability of $U \wedge V_j$ for $j \in I_V$. These second assignments exist because the first unsatisfiability condition fails. Now define the probability distribution Pr as

$$Pr(C) = \sum_{k=1,\ldots,m} (1 - \epsilon_{i_k}) t_k(C) \prod_{l=1,\ldots,k-1} \epsilon_{i_l} + \frac{\prod_{l=1,\ldots,m} \epsilon_{i_l}}{|I_V|} \sum_{j \in I_V} t'_j(C) \, ,$$

where the truth assignment $t(C)$ takes the value 1 if C is satisfied by t, and 0 otherwise. By construction, this is a well-defined probability distribution that satisfies all the right inequalities for arbitrary $\epsilon_i < 1$.
$\square$

Proof of Theorem 4.6 part 2 To show the only if part of the theorem, assume $D \wedge (C \Rightarrow_\delta \neg B)$ is consistent. If $C \Rightarrow_\delta \neg B$ with $\delta < \frac{1}{2}$ then clearly $\neg(C \Rightarrow_\delta B)$, so $D \wedge \neg(C \Rightarrow_\delta B)$ is consistent, so it must be false that $\models_{DP} D \rightarrow (C \Rightarrow_\delta B)$.

To show the if part of the theorem, assume $D \wedge (C \Rightarrow_\delta \neg B)$ is inconsistent. If D is inconsistent, then clearly $\models_{DP} D \rightarrow (C \Rightarrow_\delta B)$. If D is consistent, by part 1 of the theorem and the consistency assumptions just made, it must follow that

$$\models \neg(U \wedge (\vee_{i \in J} A_i \vee C) \wedge_{i \in J} (A_i \rightarrow B_i) \wedge (C \rightarrow \neg B)) \, ,$$

for some $J \subseteq I_A$. From the second half of the proof for Theorem 4.6 part 1, this follows for any $\delta < 1$, not just $\delta < \frac{1}{|I_A|}$. Noting that $(\vee_{i \in J} A_i \vee C)$ is equivalent to $(\vee_{i \in J} A_i \vee C \wedge_{i \in J} \neg A_i)$ and taking this disjunction out through the negation, it follows that

$$\models \quad U \wedge (\vee_{i \in J} A_i) \wedge_{i \in J} (A_i \rightarrow B_i) \rightarrow C \wedge B \, ,$$
$$\models \quad U \wedge C \wedge \neg B \rightarrow \vee_{i \in J} A_i \wedge \neg B_i \, .$$

Notice that if $\models E \rightarrow F$ then $Pr(F) \geq Pr(E)$ for any probability distribution Pr. So for any distribution Pr such that $Pr(U) = 1$,

$$Pr(C \wedge B) \geq Pr((\vee_{i \in J} A_i) \wedge_{i \in J} (A_i \rightarrow B_i)) = Pr(\vee_{i \in J} A_i) - Pr(\vee_{i \in J} A_i \wedge \neg B_i) \, ,$$

and $Pr(\vee_{i \in J} A_i \wedge \neg B_i) \geq Pr(C \wedge \neg B)$. Let Pr be any probability distribution satisfying $\square U \wedge_{i \in I_V} \diamond V_i \wedge_{i \in I_A}$

28

$(A_i \Rightarrow_{\epsilon_i} B_i)$. So $Pr(A_j \wedge \neg B_j) \le \epsilon_j Pr(A_j) \le \epsilon_j Pr(\vee_{i \in J} A_i)$ for any j even if $Pr(A_j) = 0$. Consequently,

$$Pr(\vee_{i \in J} A_i \wedge \neg B_i) \le \sum_{i \in J} Pr(A_i \wedge \neg B_i) \le \left(\sum_{i \in J} \epsilon_i\right) Pr(\vee_{i \in J} A_i) .$$

Let $\delta = \left(\sum_{i \in J} \epsilon_i\right)$. So

$$Pr(C \wedge B) \ge (\frac{1}{\delta} - 1) Pr(\vee_{i \in J} A_i \wedge \neg B_i) \ge \frac{1 - \delta}{\delta} Pr(C \wedge \neg B) ,$$

which is the required inequality to show $C \Rightarrow_\delta B$. □

Proof of Theorem 4.6 parts 3 and 4 The sentence $\diamond C$ is a consequence of D if and only if $D \wedge \neg \diamond C$ is inconsistent, by definition of consequence and inconsistency. Replacing $\neg \diamond C$ by $\square \neg C$ shows part 3 holds. Similarly, part 4 holds but this time we can simplify the inconsistency of $D \wedge \diamond \neg C$ by using part 1 of the theorem. □

Proof of Corollary 4.6.1 We shall show that the algorithm reaches to the end if and only if the sentence is consistent. By the theorem, it is sufficient to show it reaches to the end if and only if for every $j \in I_V$, $U \wedge V_j$ is satisfiable and for every $J \subseteq I_A$, $U \wedge (\vee_{i \in J} A_i) \wedge_{i \in J} (A_i \to B_i)$ is satisfiable.

Clearly Step 1 handles the first case correctly.

Consider step 3(a). Notice that if for $j \in J$, $U \wedge A_j \wedge_{i \in J}(A_i \to B_i)$ is satisfiable, then $U \wedge (\vee_{i \in J'} A_i) \wedge_{i \in J'} (A_i \to B_i)$ is satisfiable for every $J' \subseteq J$ containing j. So we now only need to consider subsets not containing j. The repeat loop in Step 3 simply uses this fact iteratively to eliminate each possible index j from the original set J. So the loop terminates short if and only if this satisfiability fails, which means the original sentence was inconsistent. □

Proof of Theorem 4.7 part 1 First prove I_{min} exists and has a unique minimum. Notice I_A is an upperbound on I_{min}, so some (but not necessarily unique) I_{min} exists. Suppose a set I' exists which is a subset of every possible I_{min}, and that $U \wedge_{i \in I'} \neg A_i \wedge A_j \wedge B_j$ is unsatisfiable for some $j \in I_A - I'$. Then this unsatisfiability will also hold for any I_{min}, so j must also be in I_{min}. So we can place j in I' too. If we start with $I' = \emptyset$ and iterate this operation to a fixed point, we clearly obtain the unique $I' = I_{min}$ because an invariant of the operation is "any I_{min} must be a superset of I'".

Suppose the unsatisfiability condition fails, that is, for each $j \in I_V$ that $U \wedge_{i \in I_{min}} \neg A_i \wedge V_j$ is satisfiable. So there exist truth assignments demonstrating the satisfiable of these. There also exist truth assignments satisfying $U \wedge_{i \in I_{min}} \neg A_i \wedge A_j \wedge B_j$ for $j \in I_A - I_{min}$, by the definition of I_{min}. Take a probability distribution that makes each assignment in the first set infinitesimally small, each assignment in the second set equiprobable, and any other truth assignments probability zero. So $Pr(V_j) > 0$, and for $j \in I_A - I_{min}$, $Pr(B_j|A_j)$ is greater than or arbitrarily close to $\frac{1}{|I_A| - |I_{min}|}$, etc. This distribution demonstrates D is consistent.

Suppose D is consistent. Consider any probability distribution Pr that demonstrates this. Let $I = \{i \in I_A : Pr(A_i) = 0\}$. So $Pr(U \wedge_{i \in I} \neg A_i) = 1$. Since $Pr(V_j) > 0$ for each $j \in V_j$, it follows that $Pr(U \wedge_{i \in I} \neg A_i \wedge V_j) > 0$ as well, so the corresponding propositional sentence must be consistent. Also, $Pr(A_j) > 0$ for $j \in I_A - I$, so since D is consistent $Pr(A_j \wedge B_j) > 0$ and $Pr(U \wedge_{i \in I} \neg A_i \wedge A_j \wedge B_j) > 0$, so the corresponding propositional sentence must be consistent. A side effect of this is that $I_{min} \subseteq I$, therefore the above satisfiability conditions holding for I also hold for I_{min}, as required for the theorem. □

Proof of Theorem 4.7 part 2 First prove the only if part of the theorem. So assume $C \approx_\delta B$ is a consequence of D for some δ. It is sufficient to prove that if D is consistent and $U \wedge_{i \in I_{min}} \neg A_i \wedge C \wedge \neg B$

is satisfiable, then there exists an ordered subset of the indices in $I_A - I_{min}$, $i_1, i_2, \ldots, i_h$ such that formulas (5) and (6) are true. Do this by contradiction. Suppose there does not exist such an ordered set of indices. Then there exists an ordered subset of the indices in $I_A - I_{min}$, $I = i_1, i_2, \ldots, i_h$ such that formulas (5) are true, but formula (6) fails and there does not exist an index i_{h+1} such that formula (5) applies for that index. Note that this occurs only if

$$U \wedge_{i \in I_{min}} \neg A_i \wedge A_j \wedge B_j \wedge_{i \in I} \neg A_i \wedge \neg (C \wedge B)$$

is satisfiable for every $j \in I_A - I_{min} - I$, and

$$U \wedge_{i \in I_{min}} \neg A_i \wedge_{i \in I} \neg A_i \wedge C \wedge \neg B$$

is satisfiable. Call the set of $|I_A - I_{min} - I|$ truth assignments satisfying the first form above T_1, and the truth assignment satisfying the second form T_2. By the definition of I_{min}, we have that $U \wedge_{i \in I_{min}} \neg A_i \wedge A_j \wedge B_j$ is satisfiable for $j \in I$. Call the set of $|I|$ truth assignments satisfying this T_3. Finally, since D is consistent, we also have that $U \wedge_{i \in I_{min}} \neg A_i \wedge V_j$ is satisfiable for $j \in I_V$. Call the set of $|I_V|$ truth assignments satisfying this T_4. Now for η vanishingly small, consider the probability distribution Pr that makes truth assignments in T_4 have probability $\frac{\eta^3}{|I_V|}$, those in T_3 have probability $\frac{(1-\eta)\eta^2}{|I|}$, the one in T_2 have probability $\eta(1-\eta)$, those in T_1 have probability $\frac{1-\eta}{|I_A - I_{min} - I|}$, and any other truth assignment have probability 0. This makes $Pr(B_j | A_j) \geq \frac{1-\eta}{|I_A - I_{min} - I|}$ for $j \in I_A - I_{min} - I$, $Pr(\neg B | C) > 1 - \eta$, etc. Clearly, Pr with a suitable value of η can be used to demonstrate $D \wedge \neg (C \approx_f B)$ is consistent for any $f < 1$. So we have proven the contradiction.

Next prove the if part of the theorem. Clearly, consequence holds if D is inconsistent. So assume it is consistent, and assume without loss of generality that $i_j = j$ for notational convenience. Consider any probability distribution Pr such that

$$\models_{Pr} \Box U \wedge_{i \in I_V} \Diamond V_i \wedge_{i \in I_A} A_i \approx_{e_i} B_i .$$

So for each $j \in I_A - I_{min}$,

$$Pr(A_j \wedge B_j) \geq \frac{e_j}{1 - e_j} Pr(A_j \wedge \neg B_j) , \tag{18}$$

even if $Pr(A_j) = 0$. From part 1 of the theorem we also know that $Pr(U \wedge_{i \in I_{min}} \neg A_i) = 1$, so this term can be effectively ignored in probability statements that follow. If $h = 0$, then from formula (6) it follows that $Pr(C \rightarrow B) = 1$, which implies $Pr(B \mid C) = 1$, so the if part holds. Otherwise, $h \geq 1$. From formulas (5) and (6), we have that

$$\models U \wedge_{i \in I_{min}} \neg A_i \wedge_{k \leq h} (A_k \rightarrow B_k) \rightarrow (C \rightarrow B) ,$$

so

$$\sum_{k \leq h} Pr(A_k \wedge \neg B_k) \geq Pr(\vee_{k \leq h} A_k \wedge \neg B_k) \geq Pr(C \wedge \neg B) . \tag{19}$$

Also, there must exist sets $P_j \subseteq \{1, \ldots, j-1\}$ such that formulas (5) can be replaced by

$$\models U \wedge_{i \in I_{min}} \neg A_i \wedge A_j \wedge B_j \wedge_{k \in P_j} \neg A_k \rightarrow (C \wedge B) ,$$

for $j = 1, \ldots, h$. Then

$$\sum_{k \in P_j} Pr(A_k) \geq Pr(\vee_{k \in P_j} A_k) \geq Pr(A_j \wedge B_j \wedge \neg (C \wedge B)) . \tag{20}$$

These inequalities are strung together below to produce the desired result.

30

For $j \in 1, \ldots, h$, define

$$\beta_j = \min_{k \in P_j} e_k \, ,$$

$$\gamma_j = \frac{1 - e_j}{e_j} + \sum_{k \, : \, j \in P_k} \frac{\gamma_k}{\beta_k} \, .$$

We shall prove by induction on j that

$$\sum_{k=j}^{h} \gamma_k Pr(A_k \wedge B_k \wedge C \wedge B) + \sum_{k=j}^{h} \frac{\gamma_k}{\beta_k} \sum_{l \, : \, l \in P_k, l < j} Pr(A_l \wedge B_l) \geq \sum_{k=j}^{h} Pr(A_k \wedge \neg B_k) \, . \qquad (21)$$

Assume it is true for $j + 1$. Consider the case for j. Notice that by formula (18)

$$\frac{1 - e_j}{e_j} Pr(A_j \wedge B_j) + \frac{\gamma_j}{\beta_j} \sum_{l \in P_j} Pr(A_l \wedge B_l)$$

$$\geq Pr(A_j \wedge \neg B_j) + \frac{\gamma_j}{\beta_j} \sum_{l \in P_j} e_l Pr(A_l)$$

$$\geq Pr(A_j \wedge \neg B_j) + \gamma_j \sum_{l \in P_j} Pr(A_l) \, .$$

Adding this to the inequality for the induction hypothesis, formula (21) with $j + 1$, we get that

$$\gamma_j Pr(A_j \wedge B_j) + \sum_{k=j+1}^{h} \gamma_k Pr(A_k \wedge B_k \wedge C \wedge B) + \sum_{k=j}^{h} \frac{\gamma_k}{\beta_k} \sum_{l \, : \, l \in P_k, l < j} Pr(A_l \wedge B_l)$$

$$\geq \sum_{k=j}^{h} Pr(A_k \wedge \neg B_k) + \gamma_j \sum_{l \in P_j} Pr(A_l) \, .$$

So

$$\sum_{k=j}^{h} \gamma_k Pr(A_k \wedge B_k \wedge C \wedge B) + \sum_{k=j}^{h} \frac{\gamma_k}{\beta_k} \sum_{l \, : \, l \in P_k, l < j} Pr(A_l \wedge B_l)$$

$$\geq \sum_{k=j}^{h} Pr(A_k \wedge \neg B_k) + \gamma_j \left(\sum_{l \in P_j} Pr(A_l) - Pr(A_j \wedge B_j \wedge \neg(C \wedge B)) \right) \, .$$

By formula (20), the induction step is proven. Notice that this same argument works for the base case of the induction proof, if we start at $j = h$ using $0 \geq 0$, so the induction proof is complete.

Finally, for $j = 1$ in formula (21), we have that

$$\sum_{k=1}^{h} \gamma_k Pr(A_k \wedge B_k \wedge C \wedge B) \geq \sum_{k=1}^{h} Pr(A_k \wedge \neg B_k) \, .$$

By formula (19), it follows that

$$\left(\sum_{k=1}^{h} \gamma_k \right) Pr(C \wedge B) \geq Pr(C \wedge \neg B) \, .$$

So

$$Pr(B|C) \geq \frac{1}{1 + \sum_{k=1}^{h} \gamma_k} \cdot$$

The right-hand side of this inequality gives the error propagation function f for this consequence. This can be evaluated using the definitions of γ_j, β_j and P_j given previously. Clearly, the smallest this error propagation function can be is when $P_j = \{1, \ldots, j-1\}$ and $\beta_j = e_1 = e$ for $j \leq h$. In this case, a simple induction proof shows that

$$\gamma_j = \frac{1-e}{e}\left(1 + \frac{1}{e}\right)^{h-j} .$$

Summing these over j and simplifying gives

$$f \geq \frac{1}{e + (1-e)\left(1 + \frac{1}{e}\right)^h} \geq \left(\frac{e}{1+e}\right)^h .$$

In contrast, if $P_j = \emptyset$, and $\beta_j = e$, then $f \geq \frac{e}{h}$. $\quad\square$

Proof of Theorem 4.7 parts 3 and 4 The same as for Theorem 4.6 parts 3 and 4. Notice, also, that I_{min} remains the same if a possibility is added to D. $\quad\square$

Proof of Corollary 4.7.1 The if part of the corollary follows from Theorem 4.7 part 2. Notice that if $P_j = \emptyset$ for each j, then the error propagation function developed in the proof of Theorem 4.7 part 2 becomes

$$\frac{1}{1 + \sum_{k=1}^{h} \frac{1-e_k}{e_k}} \cdot$$

This behaves linearly for small e_k. If any $P_j \neq \emptyset$, however, this linear behaviour no longer exists. $\quad\square$

Proof of Corollary 4.7.2 The repeat loop in Step 2 simply performs the construction described in the proof of Theorem 4.7 part 1. This iteratively builds up I_{min}. The repeat loop terminates when for all $j \in I_A - I$, $U \wedge_{i \in I} \neg A_i \wedge A_j \wedge B_j$ is satisfiable. Otherwise, the algorithm is a direct implementation of Theorem 4.7 part 1. $\quad\square$

Proof of Corollary 4.7.3 The algorithm builds the ordered set of indices in turn. Clearly, if it fails at step 3(c), then no such ordered set can exist. $\quad\square$

Figure 1: Quantitative measures of beliefs in A

Figure 2: Qualitative measures of beliefs in A

T1	$\Box A$ (when $\models A$)	$\neg \Diamond A$ (when $\models \neg A$)
T2	$\Box(A \to B) \to (\Box A \to \Box B)$	$(\neg \Diamond A \wedge \Diamond B) \to \Diamond(\neg A \wedge B)$
T3	$(\Box A \wedge \Box B) \leftrightarrow \Box(A \wedge B)$	$(\Diamond A \vee \Diamond B) \leftrightarrow \Diamond(A \vee B)$
T4	$\Box A \to \Diamond A$	$\Box A \to \Diamond A$

Table 1: Theorem schemata (and duals) on "$\Box$" and "$\Diamond$"

T5	$\Box A \leftrightarrow \models_0 A$	$\Diamond A \leftrightarrow \not\models_0 A$
T6	$\Box(A \to B) \to (\models_\epsilon A \to \models_\epsilon B)$	$\Box(A \to B) \to (\not\models_e A \to \not\models_e B)$
T7	$\neg(\models_\epsilon A \wedge \models_\epsilon \neg A)$	$(\not\models_e A \vee \not\models_e \neg A)$
T8	$\Diamond A \wedge (A \Rightarrow_\epsilon B) \to \Diamond B$	$(\Diamond A \wedge \Box B) \to A \not\Rightarrow_e B$

Table 2: Theorem schemata (and duals) relating "$\Rightarrow$" and "$\rightsquigarrow$" to "$\Box$" and "$\Diamond$"

T9	$A \Rightarrow_\epsilon A$	$A \not\rightsquigarrow_e A$
T10	$A \Rightarrow_\epsilon B \to \models_\epsilon (A \to B)$	$\not\rightsquigarrow_e (A \wedge B) \to A \not\rightsquigarrow_e B$
T11	$(\models_\epsilon A \wedge \models_\delta B) \to \models_{\epsilon+\delta} (A \wedge B)$	$\not\rightsquigarrow_{e+d} (A \vee B) \to (\not\rightsquigarrow_e A \vee \not\rightsquigarrow_d B)$
T12	$\models_\epsilon A \to (A \Rightarrow_\delta B \to \models_{\epsilon+\delta} B)$	$\models_e A \to (\not\rightsquigarrow_d B \to A \not\rightsquigarrow_{d-e} B)$
T13	$\models_\epsilon A \to (\models_\delta B \to A \Rightarrow_{\epsilon+\delta} B)$	$\models_e A \to (A \not\rightsquigarrow_d B \to \not\rightsquigarrow_{d-e} B)$
T14	$\not\rightsquigarrow_e A \to (\models_\epsilon B \to A \Rightarrow_{\frac{e}{\epsilon}} B)$	$\not\rightsquigarrow_e A \to (A \not\rightsquigarrow_d B \to \not\rightsquigarrow_{\frac{d}{e}} B)$
T15	$(A \Rightarrow_\epsilon C) \wedge (B \Rightarrow_\delta C) \to (A \vee B) \Rightarrow_{\epsilon+\delta} C$	$(A \vee B) \not\rightsquigarrow_{e+d} C \to (A \not\rightsquigarrow_e C \vee B \not\rightsquigarrow_d C)$
T16	$\not\rightsquigarrow_e (A \wedge B) \to (\neg A \Rightarrow_\epsilon \neg B \to B \Rightarrow_{\frac{e}{\epsilon}} A)$	$\not\rightsquigarrow_e (A \wedge B) \to (B \not\rightsquigarrow_d \neg A \to \neg A \not\rightsquigarrow_{\frac{e}{d}} B)$
T17	$(A \vee B \Rightarrow_{\frac{\epsilon\delta}{\epsilon+\delta-\epsilon\delta}} C) \to (A \Rightarrow_\epsilon C \vee B \Rightarrow_\delta C)$	$(A \not\rightsquigarrow_e C) \wedge (B \not\rightsquigarrow_d C) \to A \vee B \not\rightsquigarrow_{\frac{ed}{e+d-ed}} C$

Table 3: Theorem schemata (and duals) on "$\Rightarrow$" and "$\rightsquigarrow$"

33

T6′	$\Box(C \wedge A \to B) \; \to \; (C \Rightarrow_\epsilon A \; \to \; C \Rightarrow_\epsilon B)$
T7′	$(C \Rightarrow_\epsilon A) \; \wedge \; (C \Rightarrow_\epsilon \neg A) \; \to \; \neg \Diamond C$
T11′	$(C \Rightarrow_\epsilon A) \; \wedge \; (C \Rightarrow_\delta B) \; \to \; C \Rightarrow_{\epsilon+\delta} (A \wedge B)$
T12′	$C \Rightarrow_\epsilon A \; \to \; (C \wedge A \Rightarrow_\delta B \; \to \; C \Rightarrow_{\epsilon+\delta} B)$
T13′	$C \Rightarrow_\epsilon A \; \to \; (C \Rightarrow_\delta B \; \to \; C \wedge A \Rightarrow_{\epsilon+\delta} B)$
T14′	$C \not\Rightarrow_\epsilon A \; \to \; (C \Rightarrow_\epsilon B \; \to \; C \wedge A \Rightarrow_{2\epsilon} B)$

Table 4: Conditioned theorem schemata

Input: A QD_P sentence

$$\Box U \; \bigwedge_{i \in I_V} \Diamond V_i \; \bigwedge_{i \in I_A} A_i \Rightarrow B_i \; .$$

Output: The consistency or inconsistency of the sentence.

Algorithm: Check the two sets of satisfiability conditions in turn.

1. If for some $j \in I_V$, $U \wedge V_j$ is unsatisfiable, return *inconsistent*.
2. Let $J = I_A$.
3. Repeat,
   (a) Find a $j \in J$ such that $U \wedge A_j \bigwedge_{i \in J} (A_i \to B_i)$ is satisfiable.
   (b) If no such j can be found, return *inconsistent*.
   (c) Else, $J = J - \{j\}$.
   Until $J = \emptyset$.
4. Return *consistent*.

Figure 3: The defaults-consistency algorithm

Input: A QD_P sentence

$$\Box U \ \wedge_{i \in I_V} \Diamond V_i \ \wedge_{i \in I_A} A_i \approx\!\!\!\!\approx^{n_i} B_i \ .$$

Output: The consistency or inconsistency of the sentence.

Algorithm: First construct I_{min}, then check the possibility conditions.

1. Let $I = \emptyset$.

2. Repeat,

   (a) Find some $j \in I_A - I$ such that $U \wedge_{i \in I} \neg A_i \wedge A_j \wedge B_j$ is unsatisfiable.

   (b) If some j found and $\Diamond A_j$ is in the possibilities in the input sentence, then return *inconsistent*.

   (c) Else, if some j found, $I = I \cup \{j\}$.

   Until no j found.

3. I is now equal to I_{min}. If for some $j \in I_V$, $U \wedge_{i \in I} \neg A_i \wedge V_j$ is unsatisfiable, return *inconsistent*.

4. Else return *consistent*.

Figure 4: The likelihood-consistency algorithm

Input: A likelihood $C \approx\!\!\!\!\!-^m D$, a consistent QD_P sentence

$$\Box U \ \wedge_{i \in I_V} \diamond V_i \ \wedge_{i \in I_A} A_i \approx\!\!\!\!\!-^{n_i} B_i \ ,$$

and its index set I_{min}.

Output: Whether the likelihood is a consequence of the sentence for some value of m.

Algorithm: Build up the ordered subset of I_A iteratively.

1. If $U \ \wedge_{i \in I_{min}} \neg A_i \wedge C \wedge \neg B$ is unsatisfiable, return *is a consequence* for any m.
2. Set $I = \emptyset$.
3. Repeat,
    (a) Find some $j \in I_A - I_{min} - I$ such that
    $$\models \ U \ \wedge_{i \in I_{min}} \neg A_i \ \wedge A_j \wedge B_j \ \wedge_{i \in I} (A_i \rightarrow B_i) \ \rightarrow \ C \wedge B \ .$$
    (b) If some j found, $I = I \cup \{j\}$.
    (c) Else return *not a consequence*.
    Until $U \ \wedge_{i \in I_{min}} \neg A_i \wedge C \wedge \neg B \wedge_{i \in I} A_i$ is unsatisfiable or $I = I_A - I_{min}$.
4. If the loop terminated only because $I = I_A - I_{min}$, return *not a consequence*.
5. Else return *is a consequence* for some m.

Figure 5: The likelihood-consequence algorithm

Figure 6: Dependency network for "can Joe read and write?"