International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems
Vol. 00, No. 00 (0000) 000—000

ON PROOFS IN SYSTEM P

SIMON PARSONS[†] and RACHEL A. BOURNE
Department of Electronic Engineering,
Queen Mary and Westfield College, University of London,
London, E1 4NS, United Kingdom.
{S.D.Parsons, R.A.Bourne}@elec.qmw.ac.uk

Received (September 1999)
Revised (January 2000)

This paper investigates how the rules of System P might be used in order to construct proofs for default consequences which take into account the bounds on the probabilities of the consequents of the defaults. Using a knowledge base of default rules which are considered to be constraints on a probability distribution, the result of applying the rules of P gives us new constraints that were implicit in the knowledge base and their associated lower bounds. The paper defines a proof system for such constraints, shows that it is sound, and then discusses at length the completeness of the system and the kind of proofs that it can generate.

Keywords: System P; Qualitative probability; Proof theory.

1. Introduction

Default reasoning has been widely studied in artificial intelligence, and a number of formalisms have been proposed as a means of capturing such reasoning[1], most prominent among which are default logic[2] and circumscription[3]. Many of these systems, including default logic and circumscription, have proposed particular mechanisms for default reasoning, and might therefore be considered quite specialised. However, there has also been work on more general approaches which attempt to analyse in broader terms what default reasoning involves. An early attempt to do this was Shoham's[4] proposal that all non-monotonic systems could be characterised in terms of the preference order over their models. A more proof-theoretic strand of this research has investigated the formalisation of the underlying requirements for any non-monotonic consequence relation. Perhaps the most influential piece of work within this area is that of Kraus *et al.*[5].

Kraus *et al.* investigated the properties of different sets of Gentzen-style proof rules for non-monotonic consequence relations, and related these sets of rules to the model-theoretic properties of the associated logics. Their major result was that a

[†]Current address: Department of Computer Science, University of Liverpool, Chadwick Building, Peach Street, Liverpool L69 7ZF, United Kingdom, S.D.Parsons@csc.liv.ac.uk.

particular set of proof rules had the same model-theoretic properties that Shoham had identified for logics in which there is a preference order over models. This system of proof rules was termed System P by Kraus *et al.*, the P standing for "preferential". System P has been the subject of much research, and is now widely accepted as the weakest interesting non-monotonic system; it sanctions the smallest acceptable set of conclusions from a set of default statements.

The reason that we are interested in the rules of System P is that, in addition to a semantics in terms of a preference order over models, they also have a probabilistic semantics. In particular, Pearl[6], following work by Adams[7], showed that a semantics based on infinitesimal probabilities satisfies the rules of System P. While the use of infinitesimal probabilities is theoretically interesting, it lacks something in practical terms. If we are to use System P to reason about the real world we will have to write defaults which summarise our knowledge about it, and we may well be unhappy making statements whose validity depends upon infinitesimal values. To overcome this difficulty, we suggest using real probabilities along with the rules of System P, giving each default statement a lower bounded probability, and showing that proofs in the System P can be used to propagate these bounds to find out something concrete about the probability of the derived results.

This paper is organised as follows. Section 2 sets the background to the paper by discussing the notion of entailment in System P. Then, Section 3 brings up the problem of using real ϵ-values, and shows how they affect the conclusions drawn by System P. At this point, in Section 4, we introduce our system $\mathcal{SP}$, which captures one possible proof theory for System P, we show that it is sound and complete, and in Section 5 detail the kinds of conclusions which can be drawn by $\mathcal{SP}$. Section 6 then gives some examples of the use of $\mathcal{SP}$, Section 7 discusses related work, and Section 8 concludes.

2. Entailment in System P

The rules of inference for the System P (see Figure 1) may be applied to a knowledge base made up of conditional assertions of the form $\alpha \mid\!\sim \beta$. In this context α and β are well-formed formulae of classical propositional logic, and $\mid\!\sim$ is a binary relation between pairs of formulae. The connectives $\wedge$, $\vee$, $\rightarrow$ and $\leftrightarrow$ have their usual meanings. The inference rules are written in the usual Gentzen style, with antecedents above the line and consequents below it. Thus the rule And says that if it is possible to derive $\alpha \mid\!\sim \beta$ and it is possible to derive $\alpha \mid\!\sim \gamma$, then it is possible to derive $\alpha \mid\!\sim \beta \wedge \gamma$. The inference rules can thus be viewed as a means of obtaining new conclusions from current knowledge; from an initial set of conditional assertions, further conditional assertions may be obtained by applying the rules.

Two things should be noted about the set of rules in Figure 1. Firstly, they only tell us how to derive new conditional assertions. If we wish to know whether we are justified in inferring a new fact, say γ, given that we currently know some other fact, say α, *and this is all we know*, it is necessary to determine whether $\alpha \mid\!\sim \gamma$ is derivable from our knowledge base of conditional assertions. Secondly, the proof

$$\alpha \mid\!\sim \alpha \qquad \text{Reflexivity}$$

$$\frac{\models \alpha \leftrightarrow \beta, \alpha \mid\!\sim \gamma}{\beta \mid\!\sim \gamma} \qquad \text{Left Logical Equivalence}$$

$$\frac{\models \beta \rightarrow \gamma, \alpha \mid\!\sim \beta}{\alpha \mid\!\sim \gamma} \qquad \text{Right Weakening}$$

$$\frac{\alpha \mid\!\sim \beta, \alpha \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \wedge \gamma} \qquad \text{And}$$

$$\frac{\alpha \mid\!\sim \beta, \alpha \mid\!\sim \gamma}{\alpha \wedge \beta \mid\!\sim \gamma} \qquad \text{Cautious Monotonicity}$$

$$\frac{\alpha \mid\!\sim \gamma, \beta \mid\!\sim \gamma}{\alpha \vee \beta \mid\!\sim \gamma} \qquad \text{Or}$$

Fig. 1. Rules of System P

rules in Figure 1 form a minimal set sufficient to characterise System P. Other rules may be derived from them in much the same way that new conditional assertions are derived. Two such rules are given in Figure 2—Cut which allows the elimination of a conjunct from the antecedent side, and S which allows the derivation of a material implication. Both of these (as we shall see later in the paper) may be derived directly by the application of the basic rules.

The semantics for System P introduced by Adams makes the assumption that the propositional variables are the basis of an unspecified joint probability distribution which is constrained by the conditional assertions. These conditionals are taken to represent conditional probabilities of the consequent given the antecedent being greater than or equal to $1 - \epsilon$ for any $\epsilon > 0$, that is:

Definition 1 *The conditional assertion* $\alpha \mid\!\sim \beta$ *denotes the fact that* $\Pr(\beta|\alpha) \geq 1 - \epsilon$ *for all* $\epsilon > 0$.

Given this interpretation, we can define the notion of the probabilistic consistency of a set of these conditional assertions[7]:

Definition 2 *A set of conditional assertions* Δ *is* p-consistent *if there is at least one probability distribution which satisfies the constraints imposed by the conditional assertions in* Δ.

Probabilistic entailment of a further conditional is then defined as probabilistic inconsistency of its counterpart, that is:

Definition 3 $\alpha \mid\!\sim \beta$ *is* p-entailed *by* Δ *iff* $\Delta \cup \{\alpha \mid\!\sim \neg\beta\}$ *is not p-consistent.*

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma, \alpha \mid\!\sim \beta}{\alpha \mid\!\sim \gamma} \quad \text{Cut}$$

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \rightarrow \gamma} \quad \text{S}$$

Fig. 2. Two derived rules of System P

This implies that all probability distributions that satisfy Δ also satisfy $\alpha \mid\!\sim \beta$. However this result may only be achieved by using infinitesimal analysis so that the derived conditional will be constrained to be greater than $1 - \delta$ for any $\delta > 0$ if the ϵ of the original conditionals is made small enough. This can be paraphrased as saying that System P allows us to make our conclusions as close to certainty as we like, provided the conditional probabilities associated with the conditional assertions are sufficiently close to certainty. In the literature this is used to justify the conclusions drawn using System P; if we are sure of the conditional assertions, and so are willing to give them high conditional probabilities, then the conditional assertions derived from them will also have high probabilities.

However, using this interpretation of the rules means assuming that we are able to give the conditional assertions arbitrarily high conditional probabilities. This is fine in the case that the assertions are pieces of default knowledge which are felt to hold almost all of the time. However, with less reliable information, for which ϵ is not infinitesimal, it seems less justifiable to accept the infinitesimal analysis. In particular, if a set of conditional assertions are used to derive new assertions and these new assertions are themselves used as the basis for new deductions, then it seems likely that some ϵ values will be far from infinitesimal. Because of this concern, the next section investigates the impact of non-infinitesimal ϵ values by considering what happens to values of ϵ and δ when the rules of P are applied. The result is twofold. First it is possible to track the effect of non-infinitesimal values, and second it becomes possible to identify bounds on the actual conditional probability of derived assertions.

3. Using real ϵ-values

We associate with each conditional assertion an ϵ-value which represents, for $\alpha \mid\!\sim \beta$, an upper bound on the conditional probability $\Pr(\neg\beta|\alpha)$. We demonstrate how using these values for each original conditional, we can generate δ values for the output conclusions. This enables us to calculate the lower bound on the probability of a conclusion based on the proof steps used to derive it. We consider first the six basic rules of System P, and then use the results obtained for those rules to obtain results for S and Cut.

Reflexivity: A reflexive conditional assertion may be introduced at any stage in a proof, and, since $\Pr(\alpha|\alpha) = 1$ for all formulae α, any such conditional will have an ϵ-value of zero.

Left Logical Equivalence: This rule means that we may take any conditional assertion and replace its antecedent with a logically equivalent expression. Clearly, the derived conditional will have the same ϵ-value as the original one.

Right Weakening: Right Weakening involves replacing the consequent of a conditional with any expression classically derivable from it. Now, $\beta \to \gamma$ means the models of β are a subset of the models of γ and hence:

$$\Pr(\gamma, \alpha) \geq \Pr(\beta, \alpha)$$

Now, since:

$$\begin{aligned}
\Pr(\gamma|\alpha) &= \frac{\Pr(\gamma, \alpha)}{\Pr(\alpha)} \\
\Pr(\beta|\alpha) &= \frac{\Pr(\beta, \alpha)}{\Pr(\alpha)}
\end{aligned}$$

it follows that:

$$\Pr(\gamma|\alpha) \geq \Pr(\beta|\alpha) \tag{1}$$

and therefore the ϵ-value of a rule obtained by Right Weakening will not be larger than the ϵ-value of the rule from which it was obtained. Since we are dealing with lower bounds, we may use the same value for the derived rule.

We now turn to the three basic rules of P which generate new conditionals from two known ones. Given we know the ϵ-values for the two known rules we obtain simple expressions which are functions of these for the derived conditional.

Cautious Monotonicity: The rule is as follows:

$$\frac{\alpha \mid\!\sim \beta, \alpha \mid\!\sim \gamma}{\alpha \wedge \beta \mid\!\sim \gamma}$$

hence we are interested in the value of $\Pr(\gamma|\alpha, \beta)$. Now:

$$\Pr(\gamma|\alpha) = \Pr(\gamma|\alpha, \beta)\Pr(\beta|\alpha) + \Pr(\gamma|\alpha, \neg\beta)\Pr(\neg\beta|\alpha) \tag{2}$$

Substituting $1 - \Pr(\beta|\alpha)$ for $\Pr(\neg\beta|\alpha)$ and rearranging, we obtain:

$$\Pr(\gamma|\alpha, \beta) = \frac{\Pr(\gamma|\alpha) - (1 - \Pr(\beta|\alpha))\Pr(\gamma|\alpha, \neg\beta)}{\Pr(\beta|\alpha)} \tag{3}$$

We are required to minimize this expression subject to the constraints:

$$\begin{array}{rcccl}
1 - \epsilon_1 & \leq & \Pr(\beta|\alpha) & \leq & 1 \\
1 - \epsilon_2 & \leq & \Pr(\gamma|\alpha) & \leq & 1 \\
0 & \leq & \Pr(\gamma|\alpha, \neg\beta) & \leq & 1
\end{array}$$

Equation (3) is linear in $\Pr(\gamma|\alpha)$ and $\Pr(\gamma|\alpha,\neg\beta)$ and will therefore attain its minimum when $\Pr(\gamma|\alpha)$ is minimum and $\Pr(\gamma|\alpha,\neg\beta)$ is maximum. This gives us:

$$\begin{aligned}\Pr(\gamma|\alpha,\beta) &\geq \frac{(1-\epsilon_1)-(1-\Pr(\beta|\alpha))}{\Pr(\beta|\alpha)}\\ &\geq 1-\frac{\epsilon_1}{\Pr(\beta|\alpha)}\end{aligned}$$

which will be minimum when $\Pr(\beta|\alpha)$ is minimum. This gives us an ϵ-value for the derived rule $\alpha\wedge\beta \mathrel{|\!\sim} \gamma$ of:

$$\frac{\epsilon_2}{1-\epsilon_1}$$

And: This time the rule is:

$$\frac{\alpha \mathrel{|\!\sim} \beta, \alpha \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \beta\wedge\gamma}$$

So we are interested in $\Pr(\beta,\gamma|\alpha)$. Consider:

$$\begin{aligned}\Pr(\beta,\gamma|\alpha) &= \frac{\Pr(\alpha,\beta,\gamma)}{\Pr(\alpha)}\\ &= \frac{\Pr(\alpha,\beta,\gamma)\Pr(\alpha,\beta)}{\Pr(\alpha,\beta)\Pr(\alpha)}\\ &= \Pr(\gamma|\alpha,\beta)\Pr(\beta|\alpha) \qquad (4)\end{aligned}$$

We are required to minimize this expression subject to the constraints:

$$\begin{array}{ccccc}1-\epsilon_1 & \leq & \Pr(\beta|\alpha) & \leq & 1\\ 1-\epsilon_2 & \leq & \Pr(\gamma|\alpha) & \leq & 1\end{array}$$

and in the previous case we saw that these constraints imply that:

$$1-\frac{\epsilon_2}{1-\epsilon_1} \quad \leq \quad \Pr(\gamma|\alpha,\beta) \quad \leq \quad 1$$

Equation (4) will be minimum when both factors in the product on the right-hand side are, so that

$$\begin{aligned}\Pr(\beta,\gamma|\alpha) &\geq \left(1-\frac{\epsilon_2}{1-\epsilon_1}\right)(1-\epsilon_1)\\ &= 1-(\epsilon_1+\epsilon_2) \qquad (5)\end{aligned}$$

which, as we would expect, is symmetrical in ϵ_1 and ϵ_2. This gives us an ϵ-value for the derived rule $\alpha \mathrel{|\!\sim} \beta\wedge\gamma$ of $\epsilon_1+\epsilon_2$.

Or: Here the rule is:

$$\frac{\alpha \mathrel{|\!\sim} \gamma, \beta \mathrel{|\!\sim} \gamma}{\alpha\vee\beta \mathrel{|\!\sim} \gamma}$$

$$
\begin{array}{rcccl}
0 & \leq & \Pr(\alpha) & \leq & 1 \\
0 & \leq & \Pr(\beta) & \leq & 1 \\
\Pr(\alpha)(1-\epsilon_1) & \leq & \Pr(\alpha \wedge \gamma) & \leq & \Pr(\alpha) \\
\Pr(\beta)(1-\epsilon_2) & \leq & \Pr(\beta \wedge \gamma) & \leq & \Pr(\beta) \\
\max\{0, \Pr(\alpha) + \Pr(\beta) - 1\} & \leq & \Pr(\alpha \wedge \beta) & & \\
\min\{\Pr(\alpha), \Pr(\beta)\} & \geq & \Pr(\alpha \wedge \beta) & & \\
\max\{0, \Pr(\alpha \wedge \gamma) + \Pr(\beta \wedge \gamma) - \Pr(\alpha \wedge \beta)\} & \leq & \Pr(\alpha \wedge \beta \wedge \gamma) & & \\
\min\{\Pr(\alpha \wedge \gamma), \Pr(\beta \wedge \gamma), \Pr(\alpha \wedge \beta)\} & \geq & \Pr(\alpha \wedge \beta \wedge \gamma) & &
\end{array}
$$

Fig. 3. The constraints for Or.

and we are interested in the value of $\Pr(\gamma|\alpha \vee \beta)$. Let $\Pr(\neg\gamma|\alpha) = \epsilon_1$ and $\Pr(\neg\gamma|\beta) = \epsilon_2$ and consider the following:

$$
\begin{aligned}
\Pr(\neg\gamma|\alpha \vee \beta) &= \frac{\Pr(\alpha \wedge \neg\gamma) + \Pr(\beta \wedge \neg\gamma) - \Pr(\alpha \wedge \beta \wedge \neg\gamma)}{\Pr(\alpha) + \Pr(\beta) - \Pr(\alpha \wedge \beta)} \\
&= \frac{\Pr(\alpha)\Pr(\neg\gamma|\alpha) + \Pr(\beta)\Pr(\neg\gamma|\beta) - \Pr(\alpha \wedge \beta \wedge \neg\gamma)}{\Pr(\alpha) + \Pr(\beta) - \Pr(\alpha \wedge \beta)} \\
&= \frac{\epsilon_1 \Pr(\alpha) + \epsilon_2 \Pr(\beta) - \Pr(\alpha \wedge \beta \wedge \neg\gamma)}{\Pr(\alpha) + \Pr(\beta) - \Pr(\alpha \wedge \beta)} \qquad (6)
\end{aligned}
$$

To find the maximum value of this expression, we note that $\Pr(\alpha), \Pr(\beta) \leq \Pr(\alpha \vee \beta)$ and we ignore the last term of the numerator since it is negative and could be zero. Maximizing this subject to the constraints in Figure 3 gives us

$$\Pr(\neg\gamma|\alpha \vee \beta) \leq \epsilon_1 + \epsilon_2$$

as an upper bound. This gives us an ϵ-value for the derived assertion of $\epsilon_1 + \epsilon_2$.

For completeness sake, we examine the derived rules Cut and S since they are the most useful rules when proving things. To make the presentation clearer, we write a conditional with ϵ-value of ϵ_1 as $\mid\!\sim_{\epsilon_1}$.

S: For S we need to derive $\alpha \mid\!\sim_{\epsilon_{new}} \beta \rightarrow \gamma$ and the value of ϵ_{new} from $\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \gamma$ just using the basic rules. This can be done as follows. First apply Right Weakening to $\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \gamma$ to get:

$$\frac{\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \gamma, \models \gamma \rightarrow (\beta \rightarrow \gamma)}{\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \beta \rightarrow \gamma} \qquad (7)$$

We then apply Reflexivity followed by Right Weakening (twice) to $\alpha \wedge \neg\beta$ to get:

$$\frac{\alpha \wedge \neg\beta \mid\!\sim_0 \alpha \wedge \neg\beta, \models \neg\beta \rightarrow (\beta \rightarrow \gamma)}{\alpha \wedge \neg\beta \mid\!\sim_0 \beta \rightarrow \gamma} \qquad (8)$$

Then we combine (7) and (8) using Or and apply Left Logical Equivalence to get:

$$\frac{\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \beta \rightarrow \gamma, \alpha \wedge \neg\beta \mid\!\sim_0 \beta \rightarrow \gamma}{\alpha \mid\!\sim_{\epsilon_1} \beta \rightarrow \gamma}$$

$$\alpha \mathrel{|\!\sim_0} \alpha \qquad \text{Reflexivity}$$

$$\frac{\models \alpha \leftrightarrow \beta, \alpha \mathrel{|\!\sim_{\epsilon_1}} \gamma}{\beta \mathrel{|\!\sim_{\epsilon_1}} \gamma} \qquad \text{Left Logical Equivalence}$$

$$\frac{\models \beta \rightarrow \gamma, \alpha \mathrel{|\!\sim_{\epsilon_1}} \beta}{\alpha \mathrel{|\!\sim_{\epsilon_1}} \gamma} \qquad \text{Right Weakening}$$

$$\frac{\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta, \alpha \mathrel{|\!\sim_{\epsilon_2}} \gamma}{\alpha \mathrel{|\!\sim_{\epsilon_1+\epsilon_2}} \beta \wedge \gamma} \qquad \text{And}$$

$$\frac{\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta, \alpha \mathrel{|\!\sim_{\epsilon_2}} \gamma}{\alpha \wedge \beta \mathrel{|\!\sim_{\frac{\epsilon_2}{1-\epsilon_1}}} \gamma} \qquad \text{Cautious Monotonicity}$$

$$\frac{\alpha \mathrel{|\!\sim_{\epsilon_1}} \gamma, \beta \mathrel{|\!\sim_{\epsilon_2}} \gamma}{\alpha \vee \beta \mathrel{|\!\sim_{\epsilon_1+\epsilon_2}} \gamma} \qquad \text{Or}$$

$$\frac{\alpha \wedge \beta \mathrel{|\!\sim_{\epsilon_1}} \gamma, \alpha \mathrel{|\!\sim_{\epsilon_2}} \beta}{\alpha \mathrel{|\!\sim_{\epsilon_1+\epsilon_2}} \gamma} \qquad \text{Cut}$$

$$\frac{\alpha \wedge \beta \mathrel{|\!\sim_{\epsilon_1}} \gamma}{\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta \rightarrow \gamma} \qquad \text{S}$$

Fig. 4. The extended rules of System P.

The consequent of this last derivation is the consequent of S, and comparing this with the antecedent, we can see that applying S has no effect on the ϵ-value; the value for the derived conditional assertion is the same as for the original assertion.

Cut: For Cut, we need to discover how $\alpha \mathrel{|\!\sim_{\epsilon_{new}}} \gamma$ may be derived from $\alpha \wedge \beta \mathrel{|\!\sim_{\epsilon_1}} \gamma$ and $\alpha \mathrel{|\!\sim_{\epsilon_2}} \beta$. This turns out to be easy given the result for S. S tells us that the ϵ-value of $\alpha \mathrel{|\!\sim} \beta \rightarrow \gamma$ is the same as that of $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$, so we have $\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta \rightarrow \gamma$ and applying And to $\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta \rightarrow \gamma$ and $\alpha \mathrel{|\!\sim_{\epsilon_2}} \beta$, followed by Right Weakening gives:

$$\frac{\alpha \mathrel{|\!\sim_{\epsilon_1}} \beta \rightarrow \gamma, \alpha \mathrel{|\!\sim_{\epsilon_2}} \beta}{\alpha \mathrel{|\!\sim_{\epsilon_1+\epsilon_2}} \gamma}$$

Cut is thus proved, and the ϵ-value for its consequent established. We will refer to this set of rules, the basic rules of System P plus Cut and S, augmented with details of how the ϵ-values of the conditional assertions are propagated, as the *extended rules* of System P. The extended rules are summarised in Figure 4. From the previous discussion we can state the following theorem:

Theorem 1 *The extended rules of System P are sound with respect to probability theory.*

In obtaining these results, we have shown that using each of the rules of P, and hence any derived rules, we can obtain lower bounds on the conditional probability of the conclusion given those of the antecedents. Figure 4 shows the basic rules plus S and Cut annotated with lower probability bounds on antecedents and consequents. It is clear that these lower bounds never improve. Using rules And and Or, or rules derived from these, means adding the ϵ-values so that after only a few proof steps our conclusions may attain high ϵ-values. A high ϵ-value means that the lower bound on the associated conditional probability is low and if this becomes too low then we don't have much information about the probability since the upper bound is always 1. Clearly, therefore, our input values must either be extremely small, or our proofs short, in order to obtain useful results. However, as our example in Section 7 shows, these conditions can be met without too much imagination.

Another point worth noting at this juncture is the fact that the following two rules are not included, although they might be expected to follow for systems with a probabilistic semantics:

$$\frac{\alpha \mathrel{|\!\!\sim}_{\epsilon} \beta}{\alpha \mathrel{|\!\!\sim}_{\epsilon'} \neg\beta}$$

and

$$\frac{\alpha \mathrel{|\!\!\sim}_{\epsilon} \beta}{\neg\alpha \mathrel{|\!\!\sim}_{\epsilon''} \neg\beta}$$

The antecedent first of these implies:

$$\Pr(\beta \mid \alpha) \geq 1 - \epsilon$$

which is only sufficient to set an upper bound on the conditional probability of the consequent:

$$\Pr(\neg\beta \mid \alpha) \leq \epsilon$$

so no useful value for ϵ' can be determined—the value we can determine cannot be used in further inferences because it is a upper, not an lower, bound. A similar thing occurs with the second rule. To find ϵ'' it is necessary to find the value of:

$$\begin{aligned}\Pr(\neg\alpha \mid \neg\beta) &= 1 - \Pr(\alpha \mid \neg\beta) \\ &= \frac{\Pr(\neg\beta \mid \alpha)\Pr(\alpha)}{\Pr(\neg\beta)}\end{aligned}$$

which, while it can be related to $\Pr(\beta \mid \alpha)$, and hence to ϵ, does not have a useful lower bound.

While the work described so far has solved the problem of determining the impact of the non-infinitesimal values, it falls short of providing a practical reasoning system. The problem is that although in System P we can tell whether or not $\alpha \mathrel{|\!\!\sim} \beta$ follows from the initial set of defaults, the procedures for determining this do not permit the propagation of the ϵ-values. Thus we can tell if $\alpha \mathrel{|\!\!\sim} \beta$ follows, and so we can find out if a proof exists, but we can't determine the associated ϵ-value. What we need is a proof theory which allows the ϵ-values to be propagated through the proof so that every inferred default has its ϵ-value determined, and providing such a proof theory is the subject of the remainder of this paper.

4. A proof theory for System P

Normally in generating a proof theory for some logical system the procedure[8] is to establish two rules for each connective in the underlying language. One rule relates to introducing the connective into a formula, and one relates to eliminating the connective from a formula. The set of rules then define all the legal transformations between formulae, and thus define what may be proved from some initial set of formulae. The process of defining a proof theory thus proceeds from the underlying language to the proof rules.

The situation here is a little different. System P already has a set of proof rules defined. However, these rules do not include introduction and elimination rules for all the connectives in the underlying language, and so do not support a conventional proof theory. However, it is possible to use the existing rules to define a proof theory for a significant part of the underlying language of System P, and this is the approach we adopt.

We start with a set of propositions $\mathcal{S}$, a set of connectives, $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow, \Rightarrow\}$, and the following rules for building well-formed formulae in this language:

1. If $\alpha \in \mathcal{S}$, then α is a basic well-formed formula (*bwff*).
2. If α and β are *bwff*s then $\neg\alpha$, $\alpha \wedge \beta$, $\alpha \vee \beta$, $\alpha \rightarrow \beta$, $\alpha \leftrightarrow \beta$ are *bwff*s.
3. If γ and δ are *bwff*s, then $\gamma \Rightarrow_\epsilon \delta$ is a default well-formed formula (*dwff*).
4. Nothing else is a *bwff* or a *dwff*.

Together all these formulae constitute a language $\mathcal{L}_\mathcal{S}$. The denotation of basic well-formed formulae is as in propositional logic, while the meaning of *dwff*s is the following:

Definition 4 *The default $\gamma \Rightarrow_\epsilon \delta$ denotes the fact that* $\Pr(\delta \mid \gamma) \geq 1 - \epsilon$.

Comparing Definitions 1 and 4 it is clear that the defaults of $\mathcal{L}_\mathcal{S}$ are exactly the conditional assertions of System P for a particular finite ϵ. Two things follow from this. The first is that this change from the general to the particular both allows us to have some notion of strength of the defaults in terms of the conditional probability associated with them and forces us to propagate the values every time we apply one of the rules of inference. The second thing which follows is that there is a duality between assertions and defaults. We formalise this as follows:

Definition 5 *The default $\alpha \Rightarrow_{\epsilon_i} \beta$ is the* default dual *of the conditional assertion $\alpha \mathrel{|\!\sim} \beta$ and the conditional assertion $\alpha \mathrel{|\!\sim} \beta$ is the* assertion dual *of $\alpha \Rightarrow_{\epsilon_i} \beta$.*

Similarly, by extension of this notion of duality, any set of conditional assertions $\Delta = \bigcup_i \{\alpha_i \mathrel{|\!\sim} \beta_i\}$ will have a corresponding set of *dwff*s $\Delta' = \bigcup_i \{\alpha_i \Rightarrow_{\epsilon_i} \beta_i\}$. More formally:

$$\frac{\alpha \Rightarrow_{\epsilon} \beta \in \Delta}{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon)} \quad \text{Ax}$$

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\alpha, 0) \quad \text{Ref}$$

$$\frac{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_1) \qquad \Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_2)}{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta \wedge \gamma, \epsilon_1 + \epsilon_2)} \quad \text{And}$$

$$\frac{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_1) \qquad \Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_2)}{\Delta, \alpha \wedge \beta \mid\!\sim_{\mathrm{P}} (\gamma, \frac{\epsilon_2}{1-\epsilon_1})} \quad \text{CM}$$

$$\frac{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_1) \qquad \beta \vdash \gamma}{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1)} \quad \text{RW}$$

$$\frac{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1) \qquad \vdash \alpha \leftrightarrow \beta}{\Delta, \beta \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1)} \quad \text{LLE}$$

$$\frac{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1) \qquad \Delta, \beta \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_2)}{\Delta, \alpha \vee \beta \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1 + \epsilon_2)} \quad \text{Or}$$

$$\frac{\Delta, \alpha \wedge \beta \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1)}{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta \rightarrow \gamma, \epsilon_1)} \quad \text{S}$$

$$\frac{\Delta, \alpha \wedge \beta \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1) \qquad \Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_2)}{\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1 + \epsilon_2)} \quad \text{Cut}$$

Fig. 5. The consequence relation $\mid\!\sim_{\mathrm{P}}$.

Definition 6 *Given a set of defaults* Δ *and a set of conditional assertions* Δ'*,* Δ *is the* default dual *of* Δ' *if the default dual of every assertion in* Δ' *is in* Δ *and there are no additional defaults in* Δ*.*

Definition 7 *Given a set of defaults* Δ *and a set of conditional assertions* Δ'*,* Δ' *is the* assertion dual *of* Δ *if* Δ *is the default dual of* Δ'*.*

The reason for writing the defaults in this way is to distinguish between the conditional assertions themselves, and the consequence relation which defines what may be inferred from them—a distinction which is not always clear in work on System P. Assuming that we have a knowledge base Δ which consists of a set of *dwff*s, we can then define the valid set of conclusions which may be drawn from Δ as those sanctioned by the consequence relation $\mid\!\sim_{\mathrm{P}}$ defined in Figure 5. Note that this includes the two rules Cut and S which can be derived from the basic set of rules[5]. These rules are included as useful "macros" which are equivalent to applications of several other rules and help to shorten proofs as a result. We could equally well formulate $\mathcal{SP}$ in terms of the basic rules of System P, and exactly the same results would follow, though less concisely.

The proof rules that define $\mid\!\sim_{\mathrm{P}}$ may need a little explanation. The rule Ax is a form of "bootstrap" rule which says that if some default $\alpha \Rightarrow_\epsilon \beta$ is in Δ, then were α added to Δ, it would be possible to infer β with probability not less than $1 - \epsilon$. The rule And says that if adding α to Δ makes it possible to infer β with probability no less than $1 - \epsilon_1$ and γ with probability no less than $1 - \epsilon_2$, then adding α to Δ makes it possible to infer $\beta \wedge \gamma$ with probability no less than $1 - (\epsilon_1 + \epsilon_2)$. Thus the denotation of the consequence:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon)$$

is that on the basis of what is given in Δ, we can infer $\Pr(\beta \mid \alpha) \geq 1 - \epsilon$. Another way of viewing this is that if we add α to Δ, then we can infer β with a probability greater than $1 - \epsilon$.

The rules RW and LLE are a little unusual in that both have antecedents which involve $\vdash$, which stands for the consequence relation of standard propositional calculus. Thus RW says that you can replace any inference made by $\mid\!\sim_{\mathrm{P}}$ with any logical consequence, and LLE says that you can replace anything on the left-hand side of $\mid\!\sim_{\mathrm{P}}$ with something that is logically equivalent to it.

This proof system we will call $\mathcal{SP}$. As with any proof system we are interested in the soundness and completeness of the conclusions which may be drawn using $\mathcal{SP}$. We define:

Definition 8 *A* default base *is a set of default well-formed formulae.*

Definition 9 *A basic well-formed formula β is a* p-consequence *of a default base Δ,* conditional on *α, iff:*

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon)$$

By analogy with the strength of a default, the value ϵ associated with a p-consequence is known as the *strength* of the consequence. With these definitions, suitable soundness results are easy to obtain. The first relates what can be inferred using $\mid\!\sim_{\mathrm{P}}$ to System P:

Theorem 2 *For every p-consequence β, conditional on α, of a default base Δ, $\alpha \mid\!\sim \beta$ is p-entailed by the set of assertions Δ' which is the assertion dual of Δ.*

Proof: $\mathcal{SP}$ has a set of proof rules which mirror those of System P, and anything that may be proved using these rules is a p-consequence. Since Kraus *et al.* [5] have shown that anything proved using the rules of System P from a given set of conditional assertions Δ' is p-entailed by that set, it follows that any p-consequence of Δ, the default dual of Δ', is p-entailed by Δ'. $\Box$

Thus $\mathcal{SP}$ allows us to infer exactly the same things as System P. We also need to show the soundness of the mechanism for propagating the strength of the consequences. This is given by the following:

Theorem 3 *The strengths of the p-consequences of a default base are those justified by probability theory.*

Proof: The soundness of the propagation of ϵ-values with respect to probability theory follows from Theorem 1. □

Together these two results guarantee that $\mathcal{SP}$ is sound—it will generate conclusions sanctioned by System P with probabilistically correct strengths. Since Kraus *et al.* show that the rules of System P are sufficient to infer all the consequences of System P, the following completeness result is immediate:

Theorem 4 *For every $\alpha \mid\!\sim \beta$ which is p-entailed by a set of conditional assertions Δ, β is a p-consequence of the default dual of Δ conditional on α.*

What this theorem guarantees is any conditional assertion which is p-entailed by a given set of defaults will, when those defaults are translated into the language of $\mathcal{SP}$, be a consequence of the corresponding set of *dwff*s. However, this result gives no clue as to the kinds of conclusions we can draw from a given set of *dwff*s. It does not tell us if a particular p-consequence will be found, it just says that it will be found if its assertion dual is p-entailed.

What we would also like are results which say exactly what kind of conclusions we can infer from some initial set of defaults, and that is what we consider in the remainder of the paper.

5. Defining the scope of $\mathcal{SP}$

Our approach is to start by analysing what can be inferred from a set of the simplest kind of defaults, and then extend our scope to look at more complex defaults.

5.1. *Simple defaults*

We start by considering that we have a set of *simple defaults* of the form $\alpha \Rightarrow_{\epsilon_i} \gamma_i$ which all have the same antecedent. These form a simple default base:

Definition 10 *A* simple default base *for a language $\mathcal{L_S}$ is a default base:*

$$\Delta = \bigcup_{i=1,\ldots,n} \{\alpha \Rightarrow_{\epsilon_i} \gamma_i\}$$

where α and the γ_i are bwffs *in $\mathcal{L_S}$.*

We can think of the consequents of this set of defaults forming a set $\mathbf{G}$. In general, we have:

Definition 11 *The* consequent set *of a simple default base Δ is the set $\mathbf{G}$ such that:*

$$\mathbf{G} = \{\gamma_i | \{\alpha \Rightarrow_{\epsilon_i} \gamma_i\} \in \Delta\}$$

As we shall see, we need a way of referring to the conjunction of all the propositions in the consequent set, and we do this by means of the associated conjunction:

Definition 12 *The* associated conjunction Γ *of a set of propositions* $\mathbf{G}$ *is defined by*

$$\Gamma = \bigwedge_i \gamma_i \text{ for all } \gamma_{\mathrm{i}} \in \mathbf{G}$$

The set $\mathbf{G}$ *is called the* associated set *of* Γ.

Now, applying Ax and CM to $\alpha \Rightarrow_{\epsilon_1} \gamma_1$ and $\alpha \Rightarrow_{\epsilon_2} \gamma_2$, we obtain:

$$\Delta, \alpha \wedge \gamma_1 \mid\!\sim_{\mathrm{P}} \left(\gamma_2, \frac{\epsilon_2}{1-\epsilon_1}\right) \tag{9}$$

Using the same rules on $\alpha \Rightarrow_{\epsilon_1} \gamma_1$ and $\alpha \Rightarrow_{\epsilon_3} \gamma_3$ gives:

$$\Delta, \alpha \wedge \gamma_1 \mid\!\sim_{\mathrm{P}} \left(\gamma_3, \frac{\epsilon_3}{1-\epsilon_1}\right)$$

and combining the latter with (9) using CM will give:

$$\Delta, \alpha \wedge \gamma_1 \wedge \gamma_2 \mid\!\sim_{\mathrm{P}} \left(\gamma_3, \frac{\epsilon_3}{1-\epsilon_1-\epsilon_2}\right)$$

If we imagine repeating this process it is clear that given Δ we can recursively apply CM to obtain:

$$\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} (\gamma_{\mathrm{i}}, \epsilon)$$

for any $\gamma_i \in \mathbf{G}$, and for any B' which is the associated conjunction of a set $\mathbf{B}'$ such that $\mathbf{B}' \subseteq \mathbf{G}$ where ϵ is a function of the ϵ-values of the defaults to which CM has been applied. In fact we have:

Lemma 1 *Given a simple default base* Δ *with antecedent* α *and consequent set* $\mathbf{G}$ *the consequence relation* $\mid\!\sim_{\mathrm{P}}$ *will generate all consequences:*

$$\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} \left(\gamma_{\mathrm{i}}, \frac{\epsilon_{\mathrm{i}}}{1-\sum_{\mathrm{j}} \epsilon_{\mathrm{j}}}\right)$$

where $\gamma_i \in \mathbf{G}$, B' *has an associated set* $\mathbf{B}'$, *for every* γ_j *in* $\mathbf{B}'$ *there is a default* $\alpha \Rightarrow_{\epsilon_j} \gamma_j$ *with strength* ϵ_j *in* Δ, *and* $\mathbf{B}' \subseteq \mathbf{G}$.

Proof: This follows more or less directly from the previous discussion. Since it is possible to use CM to get $\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} (\gamma_{\mathrm{i}}, \epsilon)$ for any γ_i in the consequent set and any B' which is a conjunction of propositions from the consequent set, then it is possible to use it to obtain all such p-consequences. The relevant value of ϵ follows by simple arithmetic on the strengths of the relevant defaults. □

In other words, CM allows us to obtain as a p-consequence any proposition in the consequent set of Δ conditional on α conjoined with any other propositions in the consequent set. The ϵ-value that results is that of the p-consequence conditional on α alone, divided by 1 minus the sum of the ϵ-values of each of the propositions in the conjunction conditional on α alone.

Since the rule And makes it possible to build up conjunctions on the consequent side, similar reasoning makes it obvious that recursively applying the rule to the same initial set of defaults will give:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\Gamma', \epsilon)$$

for some ϵ, where Γ' is the associated conjunction of a set $\mathbf{G}'$ and $\mathbf{G}' \subseteq \mathbf{G}$. This time, we have:

Lemma 2 *Given a simple default base Δ with antecedent α and consequent set $\mathbf{G}$ the consequence relation $\mid\!\sim_{\mathrm{P}}$ will generate all consequences:*

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} \left(\Gamma', \sum_k \epsilon_k\right)$$

where Γ' has an associated set $\mathbf{G}'$, for every γ_k in $\mathbf{G}'$ there is a default $\alpha \Rightarrow_{\epsilon_k} \gamma_k$ with strength ϵ_k in Δ, and $\mathbf{G}' \subseteq \mathbf{G}$.

Proof: As with the discussion of CM, consider applying the rules Ax and And to $\alpha \Rightarrow_{\epsilon_1} \gamma_1$ and $\alpha \Rightarrow_{\epsilon_2} \gamma_2$. This gives:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma_1 \wedge \gamma_2, \epsilon_1 + \epsilon_2) \tag{10}$$

Using the same rules on $\alpha \Rightarrow_{\epsilon_3} \gamma_3$ and $\alpha \Rightarrow_{\epsilon_4} \gamma_4$ gives:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma_3 \wedge \gamma_4, \epsilon_3 + \epsilon_4)$$

and combining the latter with (10) using the rule And will give:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} \left(\bigwedge_{i=1,\ldots,4} \gamma_i, \sum_{i=1,\ldots,4} \epsilon_i\right)$$

Now, it is clearly possible to use And in this way to get $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\Gamma_i, \epsilon)$ for any conjunction Γ_i whose constituent propositions γ_i are in the consequent set. Thus it is possible to use it to obtain all such p-consequences and the result follows. $\Box$

Thus the And rule makes it possible to obtain as a p-consequence any conjunction of propositions from the consequence set of Δ, conditional on α. The ϵ-value which results is the sum of the ϵ-values of those propositions alone conditional on α. Clearly, then, if we use both rules together, we can derive conclusions of the form:

$$\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} (\Gamma', \epsilon)$$

where B' and Γ' have associated sets $\mathbf{B}'$ and $\mathbf{G}'$ such that $\mathbf{B}' \subseteq \mathbf{G}$ and $\mathbf{G}' \subseteq \mathbf{G}$. Note it is possible that $\mathbf{B}' \cap \mathbf{G}' \neq \emptyset$. To prove this formally, we first need to extend the notion of duality between different representations of defaults introduced above. There we had the notion of the assertion $\alpha \mid\!\sim \beta$ being the dual of a default $\alpha \Rightarrow_{\epsilon_i} \beta$. We extend this by noting that any such default, after the application of the proof rule Ax generates a p-consequence $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_i)$. We thus define:

Definition 13 *The default $\alpha \Rightarrow_{\epsilon_i} \beta$ which is part of Δ, is the* default dual *of the p-consequence $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_i)$ and the p-consequence $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\beta, \epsilon_i)$ is the* consequence dual *of $\alpha \Rightarrow_{\epsilon_i} \beta$.*

This overloads the term "default dual", but its meaning will always be clear from the context. As before we extend this definition to sets of defaults and p-consequences:

Definition 14 *Given a set of defaults Δ and a set of p-consequences Δ', Δ is the* default dual *of Δ' if the default dual of every p-consequence in Δ' is in Δ and there are no additional defaults in Δ.*

Definition 15 *Given a set of defaults Δ and a set of p-consequences Δ', Δ' is the* consequence dual *of Δ if Δ is the default dual of Δ'.*

With these definitions we can combine Lemmas 1 and 2 to obtain the following:

Theorem 5 *Given a simple default base with antecedent α and consequent set $\mathbf{G}$, the consequence relation $\mid\!\sim_{\mathrm{P}}$ will generate all consequences:*

$$\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} \left(\Gamma', \frac{\sum_{\mathrm{k}} \epsilon_{\mathrm{k}}}{1 - \sum_{\mathrm{j}} \epsilon_{\mathrm{j}}}\right)$$

where B' and Γ' have associated sets $\mathbf{B}'$ and $\mathbf{G}'$, for every γ_j in $\mathbf{B}'$ there is a default $\alpha \Rightarrow_{\epsilon_j} \gamma_j$ with strength ϵ_j in Δ, for every γ_k in $\mathbf{G}'$ there is a default $\alpha \Rightarrow_{\epsilon_k} \gamma_k$ with strength ϵ_k in Δ, and $\mathbf{B}', \mathbf{G}' \subseteq \mathbf{G}$.

Proof: First apply Lemma 1 to obtain a series of p-consequences:

$$\Delta, \alpha \wedge \mathrm{B}' \mid\!\sim_{\mathrm{P}} \left(\gamma_{\mathrm{k}}, \frac{\epsilon_{\mathrm{k}}}{1 - \sum_{\mathrm{j}} \epsilon_{\mathrm{j}}}\right)$$

for each γ_k which is one of the conjuncts in Γ'. Then apply Lemma 2 to the default dual of this set. The result follows. □

We refer to the set of consequences defined by Theorem 5 as the *simple consequences* of Δ.

5.2. *More complex defaults*

The results in the previous section characterise the kind of consequences we can prove using the rules And and CM on a set of simple defaults. It is possible to generalise these results to wider sets of defaults. Consider that instead of a set of simple defaults, we have, instead, a general set of *conjunctive consequent defaults* of the form $\alpha \Rightarrow_{\epsilon_i} \Gamma_i$ where α, as before, is a single proposition and Γ_i is conjunction of propositions, known as the *conjunctive consequent*. This set of defaults is a conjunctive consequent default base:

Definition 16 *A* conjunctive consequent default base *for a language $\mathcal{L}_\mathcal{S}$ is a default base:*

$$\Delta = \bigcup_{i=1,\ldots,n} \{\alpha \Rightarrow_{\epsilon_i} \Gamma_i\}$$

where α is a bwff *in $\mathcal{L}_\mathcal{S}$, and Γ_i is a conjunction of such* bwffs*.*

For such defaults we expand the notion of the consequent set to include all propositions which occur in a conjunctive consequent:

Definition 17 *The* consequent set *of a conjunctive consequent default base Δ is the set* $\mathbf{G}$ *such that:*

$$\mathbf{G} = \{\gamma_i | \{\alpha \Rightarrow_{\epsilon_i} \gamma_1 \wedge \ldots \wedge \gamma_i \wedge \ldots \wedge \gamma_n\} \in \Delta\}$$

Since a conjunctive consequent default base Δ can contain simple defaults, it is helpful to distinguish the *simple subset*, which is the set of all simple default rules in Δ. We denote this by Δ_s. Now, applying the rule Ax to any conjunctive consequent default in Δ will give:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\Gamma_i, \epsilon_i)$$

RW allows us to replace any p-consequence with any of its logical consequences. This makes it possible to obtain:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma_j, \epsilon_i)$$

for any γ_j which is one of the conjuncts in Γ_i. This immediately gives us:

Lemma 3 *Given a conjunctive consequent default base Δ, with consequent set* $\mathbf{G}$*, then $\mid\!\sim_{\mathrm{P}}$ will generate all consequences:*

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma', \epsilon)$$

where Δ contains a default $\alpha \Rightarrow_{\epsilon_i} \Gamma_i$, $\mathbf{G}_i$ is the associated set of Γ_i and $\gamma' \in \mathbf{G}_i$

Since this set of p-consequences is the consequence dual of the set of simple defaults $\alpha \Rightarrow_{\epsilon_i} \gamma'$, Lemma 3 suggests that any conjunctive consequent default base has a corresponding simple default base such that both default bases have a common set of p-consequences—the consequence dual of the simple default base.[‡] We call this

[‡]It should be noted that while the p-consequences of these two default bases are the same, the ϵ-values of these p-consequences will, in general, differ, with the ϵ-values of the p-consequences derived from the simple default base being higher. As an example, consider the default base $\{\alpha \Rightarrow_\epsilon \beta \wedge \gamma\}$, which has simple equivalent $\{\alpha \Rightarrow_\epsilon \beta, \alpha \Rightarrow_\epsilon \gamma\}$. The formula $\beta \wedge \gamma$ is a p-consequence of both default bases, but has strength ϵ when derived from the first and 2ϵ when derived from the second.

simple default base the *simple equivalent* of the conjunctive consequence default base.

Definition 18 *Given a conjunctive consequent default base Δ, with consequent set* $\mathbf{G}$*, then its* simple equivalent *is the set of defaults:*

$$\{\alpha \Rightarrow_{\epsilon_i} \gamma_j \mid \{\alpha \Rightarrow_{\epsilon_i} \gamma_1 \wedge \ldots \wedge \gamma_j \wedge \ldots \wedge \gamma_n\} \in \Delta\}$$

Thus to transform a conjunctive consequent default base into its simple equivalent we replace every conjunctive consequent default with a set of simple defaults, each with the same strength as the original default and a consequent which is one of the conjuncts in the consequent of the original default. Given Definition 18 we have:

Theorem 6 *Given a conjunctive consequent default base Δ, the consequence relation $\mid\!\sim_{\mathrm{P}}$ will generate all the simple consequences of the simple equivalent of Δ.*

Proof: Call the consequence dual of the simple subset Δ_s of Δ by the name P_1. Take the set $\Delta - \Delta_s$, and apply Lemma 3 to it to obtain a set of p-consequences of the form:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma_k, \epsilon_j)$$

which includes one such p-consequence for each γ_k which appears in the consequent set of $\Delta - \Delta_s$. Call this set of p-consequences P_2. The set $P_1 \cup P_2$ is then exactly the consequence dual of the simple equivalent of Δ. Thus anything which can be derived from the simple equivalent of Δ can also be derived from Δ itself. □

The reason that this result is important is that it allows us to apply Theorem 5 to conjunctive consequent default bases, by first turning the default base into its simple equivalent. This in turn means that we can immediately write down a subset of the p-consequences of any conjunctive consequent default base Δ—the simple consequences of its simple equivalent. We call these the *simple equivalent consequences* of Δ.

Now let's consider generalising the set of defaults Δ to what we will call a set of *general conjunctive defaults* of the form $\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}$ where α, as before, is a single proposition and $\mathrm{B_i}$ and Γ_i are conjunctions of propositions. This set of defaults is a general conjunctive default base:

Definition 19 *A* general conjunctive default base *for a language $\mathcal{L}_\mathcal{S}$ is a default base:*

$$\Delta = \bigcup_{i=1,\ldots,n} \{\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}\}$$

where α is a bwff *in $\mathcal{L}_\mathcal{S}$, and the* $\mathrm{B_i}$ *and* Γ_i *are conjunctions of such* bwffs*.*

Thus a general conjunctive default base is just a set of conjunctive consequence defaults which have a *conjunctive antecedent* $\mathrm{B_i}$ conjoined to their *base antecedent* α. Such a default base has a simple subset Δ_s as before, and a *conjunctive consequent* subset Δ_c which includes defaults of the form $\alpha \Rightarrow_{\epsilon_i} \Gamma_i$. From such a default base we can clearly generate all p-consequences which are simple equivalent consequences of $\Delta_s \cup \Delta_c$. These are all the p-consequences of Δ which can be obtained by applying CM, And, and RW alone. However, there are further p-consequences of a set of general conjunctive defaults. Applying the rules Ax and S to a default $\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_i} \Gamma_\mathrm{i}$ in $\Delta - (\Delta_s \cup \Delta_c)$ will give:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i} \rightarrow \Gamma_\mathrm{i}, \epsilon_\mathrm{i})$$

Now, if we can obtain:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i}, \epsilon_\mathrm{j})$$

Applying And will give us

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i} \wedge (\mathrm{B_i} \rightarrow \Gamma_\mathrm{i}), \epsilon_\mathrm{i} + \epsilon_\mathrm{j})$$

and then RW will allow us to obtain:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i} \wedge \Gamma_\mathrm{i}, \epsilon_\mathrm{i} + \epsilon_\mathrm{j})$$

and hence:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\varphi_{i_j}, \epsilon_i + \epsilon_j)$$

for any $\varphi_{i_j} \in \mathbf{B}_i \cup \mathbf{G}_i$ where $\mathbf{B}_i$ and $\mathbf{G}_i$ are the associated sets of $\mathrm{B_i}$ and Γ_i. One way that:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i}, \epsilon_\mathrm{j})$$

can be obtained at the crucial point is if this is a simple equivalent consequent of $\Delta_s \cup \Delta_c$. This gives us:

Lemma 4 *Given a general conjunctive default base Δ with base antecedent α, whose simple subset is Δ_s and whose conjunctive consequent subset is Δ_c, then $\mid\!\sim_\mathrm{P}$ will generate all the consequences of the form:*

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\gamma', \epsilon_i + \epsilon_j)$$

where Δ contains a default $\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_i} \Gamma_\mathrm{i}$, $\mathbf{G}_i$ is the associated set of Γ_i, $\gamma' \in \mathbf{G}_i$, and

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\mathrm{B_i}, \epsilon_\mathrm{j})$$

is one of the simple equivalent consequences of $\Delta_s \cup \Delta_c$.

Proof: This follows immediately from the previous discussion. □

As with Lemma 3, this is a "reduction" result, which allows us to use a complex set of defaults to obtain a set of p-consequences which we could obtain from a much simpler set of defaults. Following this analogy, in the same way as we defined the simple equivalent of a set of conjunctive antecedent defaults, we can define a set of simple defaults which, when we apply Ax to them, give us the same set of consequences as applying Lemma 4 to a general conjunctive default base. This set is defined as follows:

Definition 20 *Given a general conjunctive default base Δ with simple subset Δ_s and conjunctive consequent subset Δ_c, its* reduced equivalent *is the union of Δ_s, the simple equivalent of Δ_c, and the set of defaults:*

$$\{\alpha \Rightarrow_{\epsilon_i} \gamma_j \mid \{\alpha \wedge \beta_1 \wedge \ldots \wedge \beta_m \Rightarrow_{\epsilon_i} \gamma_1 \wedge \ldots \wedge \gamma_j \wedge \ldots \wedge \gamma_n\} \in \Delta, \beta_1, \ldots, \beta_m \in \mathbf{G}\}$$

where $\mathbf{G}$ is the consequent set of $\Delta_s \cup \Delta_c$.

With this notion of a reduced equivalent set, it is easy to identify one set of consequences of a general conjunctive default base.

Theorem 7 *Given a general conjunctive default base Δ, $\mid\!\sim_{\mathrm{P}}$ will generate all simple consequences of the reduced equivalent of Δ.*

Proof: Immediate by applying Lemma 4. □

We call these the *conservative consequences* of Δ, so named because they are only a subset of the full set of consequences. However, in establishing this set of consequences, we have characterised a significant portion of the consequences of general conjunctive default bases in a way which makes it easy to determine if a particular consequence is one of the number. For a given formula γ and general conjunctive default base Δ, we can answer the question "is γ a conservative consequence of Δ?" by inspection.

We can view the results we have obtained as forming a sequence of consonant sets of consequences. Consider a general conjunctive default base Δ. Theorem 5 identifies all the consequences of the simple subset of Δ which can be obtained using the rules And and CM, the two rules which allow arbitrary conjunctions to be established on either side of the turnstile. These are the simple consequences of Δ. Theorem 6 makes it possible to draw conclusions from a larger subset of Δ, namely the union of the simple subset and the set of conjunctive consequent defaults, again using the rules And and CM, along with RW. This set of consequences, the simple equivalent consequences, includes all the simple consequences. Finally, Theorem 7 makes it possible to use every default in Δ, obtaining the set of conservative consequences, which includes all the simple equivalent consequences.

Theorems 5, 6 and 7 complement Theorem 4. The latter says that anything provable will eventually be proved. It therefore defines what is provable from above, placing a limit on the set of consequences which can be proved without giving an indication of what they are. The former are a first step towards defining what is provable from below. Given a default base these results tell us what can be proved. However, they do not identify every possible consequence, since other consequences can be determined by the application of other proof rules. For example we have:

Theorem 8 *Given a general conjunctive default base $\Delta = \bigcup_i \{\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}\}$ with base antecedent α, the consequence relation $\mid\!\sim_{\mathrm{P}}$ will generate all p-consequences:*

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\Phi, \epsilon)$$

where $\Phi = \bigwedge_j \varphi_j$,

$$\bigwedge_i (\neg \mathrm{B_i} \vee \Gamma_\mathrm{i}) \vdash \varphi_\mathrm{j}$$

and

$$\epsilon = \sum_i \epsilon_i$$

Proof: For every $\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}$ we can apply Ax, S and RW to get $\Delta, \alpha \mid\!\sim_\mathrm{P} (\neg \mathrm{B_i} \vee \Gamma_\mathrm{i}, \epsilon_\mathrm{i})$. Applying And to all of these gives

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\bigwedge_i (\neg \mathrm{B_i} \vee \Gamma_\mathrm{i}), \sum_\mathrm{i} \epsilon_\mathrm{i})$$

and since applying RW to this latter does not change the value of ϵ, the result follows. □

The value of ϵ here is clearly an upper bound, giving a lower bound on the value of the conditional probability of the resulting p-consequences. A tighter bound could be obtained by "And"ing only those defaults which are actually used in the derivation of the φ_j. This raises the issue of what we should conclude if we obtain several p-consequences of the form:

$$\Delta, \alpha \mid\!\sim_\mathrm{P} (\Phi, \epsilon_i)$$

with different strengths ϵ_i. What these mean, of course is that:

$$\Pr(\Phi \mid \alpha) \geq 1 - \epsilon_i$$

for various ϵ_i. These results are entirely consistent, and we are justified in picking whichever ϵ_i we wish. Typically we will choose the smallest, since this gives us the highest value of $\Pr(\Phi \mid \alpha)$.

It should also be noted that Theorem 8 identifies a much larger set of potential consequences than the conservative consequences. However, to establish these it is necessary to invoke a standard propositional theorem prover.

5.3. *More than one default base*

All the results presented so far hold for sets of defaults with a single proposition α on the antecedent side. Clearly we can replace α with an arbitrary conjunction of propositions, recovering analogous results after the necessary changes in definition of terms such as "simple default"—in the interests of space we will not consider this extension in detail. Instead, taking the idea of having a conjunction as an antecedent somewhat further, one might imagine that there are more consequences that might be inferred from defaults with antecedents with several propositions in common, for instance:

$$\alpha \wedge \beta \Rightarrow_{\epsilon_1} \gamma'$$

and

$$\alpha \wedge \delta \Rightarrow_{\epsilon_2} \gamma''$$

However, if α is a single proposition, we have already analysed the situation since these two defaults are part of a general conjunctive default base with base antecedent α. If α is a conjunction, then this case is part of the simple extension already discussed. Thus there are no particularly interesting results that may be obtained here.

Another case that seems worth investigating is when we have more than one default base. Such a situation arises when we can partition a set of defaults into two or more subsets where every default in each subset has at least one antecedent proposition in common (the base antecedent of that subset), but there are no common antecedent propositions between defaults in different subsets. An example of such a situation is when we have:

$$\Delta = \bigcup_{i=1,\ldots,n} \{\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}\}$$

and

$$\Delta' = \bigcup_{j=1,\ldots,n} \{\beta \wedge \mathrm{B_j} \Rightarrow_{\epsilon_\mathrm{j}} \Gamma_\mathrm{j}\}$$

where $\mathbf{B}_i$ and $\mathbf{B}_j$ are the associated sets of $\mathrm{B_i}$ and $\mathrm{B_j}$ respectively, and

$$\bigcup_i \mathbf{B}_i \cap \bigcup_j \mathbf{B}_j = \emptyset$$

However, this is a situation when the very conservative nature of System P works against us. The only rule of $\mathcal{SP}$ which makes it possible to combine two p-consequences with such antecedents is Or. Or only applies to two defaults which have the same formula on the right of the turnstile. In other words, it only applies directly if Γ_i and Γ_j are identical. However, thanks to RW, we can convert anything on the right hand side to any of its logical consequences, so we can apply Or indirectly provided that either $\Gamma_i \rightarrow \Gamma_j$ or $\Gamma_j \rightarrow \Gamma_i$. Thus, given:

$$\alpha \Rightarrow_{\epsilon_1} \beta$$
$$\gamma \Rightarrow_{\epsilon_2} \delta$$

provided that:

$$\beta \vdash \delta$$

we can apply Ax to both defaults and then RW to the first to obtain:

$$\Delta, \alpha \mathrel{|\!\sim_{\mathrm{P}}} (\delta, \epsilon_1)$$
$$\Delta, \gamma \mathrel{|\!\sim_{\mathrm{P}}} (\delta, \epsilon_2)$$

and then use Or to infer:

$$\Delta, \alpha \vee \gamma \mathrel{|\!\sim_{\mathrm{P}}} (\delta, \epsilon_1 + \epsilon_2)$$

However, Or on its own does not modify the right hand side of the turnstile, and so does not make it possible to establish any new p-consequences—in this example we could already obtain:

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\delta, \epsilon_1)$$

using just the first default. What Or does is to make it possible to alter the states that existing p-consequences are conditional on, and since this is somewhat outside our interest we will say no more about it here.

In fact, the only way to draw substantial conclusions from several default bases is to turn them into a single default base. This is possible using a combination of the rules S and LLE. Given any default

$$\alpha \Rightarrow_{\epsilon_1} \gamma$$

we can use Ax and LLE to obtain:

$$\Delta, \alpha \wedge \top \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1)$$

which is not much use on its own, but allows us to apply S to get:

$$\Delta, \top \mid\!\sim_{\mathrm{P}} (\alpha \rightarrow \gamma, \epsilon_1)$$

Applying this transformation to several defaults from different default bases gives us a new default base with base antecedent $\top$. Clearly we can then establish similar results to those obtained above for defaults with base antecedent α, but writing $\top$ in place of α on the left of the turnstile, and $\alpha \wedge \mathrm{B_i}$ in place of $\mathrm{B_i}$ on the right. We can summarise everything which can be inferred using this particular combination of proof rules in the same way as is possible in Theorem 8:

Theorem 9 *Given a general conjunctive default base* $\Delta = \bigcup_i \{\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}\}$ *with base antecedent* α*, the consequence relation* $\mid\!\sim_{\mathrm{P}}$ *will generate all p-consequences:*

$$\Delta, \top \mid\!\sim_{\mathrm{P}} (\Phi, \epsilon)$$

where $\Phi = \bigwedge_j \varphi_j$*,*

$$\bigwedge_i (\neg\alpha \vee \neg\mathrm{B_i} \vee \Gamma_\mathrm{i}) \vdash \varphi_\mathrm{j}$$

and

$$\epsilon = \sum_i \epsilon_i$$

Proof: For every $\alpha \wedge \mathrm{B_i} \Rightarrow_{\epsilon_\mathrm{i}} \Gamma_\mathrm{i}$ we can apply Ax, LLE, S and RW to get $\Delta, \top \mid\!\sim_{\mathrm{P}} (\neg(\alpha \wedge \mathrm{B_i}) \vee \Gamma_\mathrm{i}, \epsilon_\mathrm{i})$. Applying And to all of these gives

$$\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\bigwedge_i (\neg(\alpha \wedge \mathrm{B_i}) \vee \Gamma_\mathrm{i}), \sum_\mathrm{i} \epsilon_\mathrm{i})$$

(i)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (steve, 0.1)$	$Ax, 4$
(ii)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (great, 0.01)$	$Ax, 2$
(iii)	$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (great, 0.011)$	$CM, (i), (ii)$
(iv)	$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (\neg noisy, 0.05)$	$Ax, 5$
(v)	$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (great \wedge \neg noisy, 0.061)$	$And, (iii), (v)$

Fig. 6. The proof of a conservative consequence about Linda

and since applying RW to this latter does not change the ϵ-value, the result follows. □

All of the results obtained in this section can be considered to be completeness results for $\mathcal{SP}$ in the sense that they identify all the possible consequences which fall into the particular classes given. In combination with the soundness result of Theorem 2 and 3 they give the usual guarantees for a particular set of consequences.

5.4. *Future work*

One way of looking at the results presented in this section is as a set of partial classifications of the kinds of consequences which can be derived from a set of defaults. Another way of considering them is as a set of results for transforming a set of defaults in a way which does not change their p-consequences. In this sense, we can consider Definition 18 as a way of transforming a conjunctive consequent default base into its simple equivalent without changing the set of simple consequences which can be derived from it (though, as mentioned above, the ϵ-values will in general be different).

The fact that this kind of transformation is possible suggests that it may be possible to identify additional transformations which are distinguished in some way (the transformation to a simple equivalent being distinguished by the fact that it is made up entirely of simple defaults). Two such distinguished transformations spring to mind. One is that which guarantees the lowest ϵ-values for all the p-consequences—this would clearly be useful since it would guarantee the strongest consequences (equally clearly it won't be the transformation which gives the simple equivalent). The second is that which gives the shortest proofs (in the sense of requiring the fewest applications of the proof rules of $\mathcal{SP}$) for the set of p-consequences. One interesting direction for future work in this area is the identification of such distinguished transformations.

6. An example

We now illustrate the use of $\mathcal{SP}$ on the following, inspired by examples given by Kraus *et al.*[5].

> Brian and Linda are two happy-go-lucky people who are normally the life and soul of any party (so if either go to a party it will normally be great). Until recently Brian and Linda were married, but then Linda

> ran off with a mime artist, Steve. As a result, if both Brian and Linda go to the same party they will probably have a screaming row and ruin it (so it will not be great and it will be noisy).
>
> If Linda goes to a party she will probably take her new boyfriend Steve and get him to entertain the guests with his marvellous miming. Thus if Linda goes to a party, Steve will probably go to the same party and if Linda and Steve go to a party together it will normally not be noisy because everyone will be watching his miming. Normally parties that are great are noisy, and those that are not noisy are not great.

We represent this by the following default base Δ^{party}. It should be understood that we are trying to ascertain the likelihood of any given party having various attributes (*brian* is present, it is *noisy*, and so on).

1. $brian \Rightarrow_{0.01} great$
2. $linda \Rightarrow_{0.01} great$
3. $brian \wedge linda \Rightarrow_{0.15} \neg great \wedge noisy$
4. $linda \Rightarrow_{0.1} steve$
5. $linda \wedge steve \Rightarrow_{0.05} \neg noisy$
6. $great \Rightarrow_{0.1} noisy$
7. $\neg noisy \Rightarrow_{0.1} \neg great$

As an example of the generation of a conservative consequence, consider the proof of Figure 6. As this proof demonstrates, we can conclude that if both Linda and Steve go to the party, then the probability that it will be both great and not noisy is greater than 0.939 (1 minus the strength of the p-consequence $linda \wedge steve$).

If we combine defaults from the different conjunctive default bases in Δ, we can obtain additional conclusions. For example, consider Figure 7 which gives a proof for the p-consequence *linda* conditional on $\top$. This tells us that the probability of Linda going to any particular party is at most 0.26. This last example neatly illustrates two points.

The first is a property of System P. We have shown that the probability of Linda going to any particular party is quite low. It certainly isn't likely enough to be a default conclusion. However, if we know that Linda *does* go to a party—a fact which makes the party somewhat abnormal—then we can draw conclusions which are very likely for such abnormal parties (they are very likely to be great, for instance). The second point is to do with the form of the proof. As stated above, the proof of the p-consequence $\neg$*linda* involves the use of defaults from different conjunctive default bases (in particular that with base antecedent *linda* and the single default with base antecedent $\neg noisy$). This is possible through the use of LLE and S to obtain p-consequences conditional on $\top$ which may then be combined using And. As mentioned above, this is an important mechanism for combining defaults from different default bases.

(i)	$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (\neg noisy, 0.05)$	$Ax, 4$
(ii)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (steve, 0.1)$	$Ax, 5$
(iii)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (\neg noisy, 0.15)$	$Cut, (i), (ii)$
(iv)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (great, 0.01)$	$Ax, 2$
(v)	$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (great \wedge \neg noisy, 0.16)$	$And, (iii), (iv)$
(vi)	$\Delta^{party}, \top \wedge linda \mid\!\sim_{\mathrm{P}} (great \wedge \neg noisy, 0.16)$	$LLE, (v)$
(vii)	$\Delta^{party}, \top \mid\!\sim_{\mathrm{P}} (\neg linda \vee (great \wedge \neg noisy), 0.16)$	$S, (vi)$
$(viii)$	$\Delta^{party}, \neg noisy \mid\!\sim_{\mathrm{P}} (\neg great, 0.1)$	$Ax, 7$
(ix)	$\Delta^{party}, \top \wedge \neg noisy \mid\!\sim_{\mathrm{P}} (\neg great, 0.1)$	$LLE(viii)$
(x)	$\Delta^{party}, \top \mid\!\sim_{\mathrm{P}} (\neg great \vee noisy, 0.1)$	$S, (ix)$
(xi)	$\Delta^{party}, \top \mid\!\sim_{\mathrm{P}} ((\neg great \vee noisy)$ $\wedge (\neg linda \vee (great \wedge \neg noisy)), 0.26)$	$And, (vii), (x)$
(xii)	$\Delta^{party}, \top \mid\!\sim_{\mathrm{P}} (\neg linda, 0.26)$	$RW, (xi)$

Fig. 7. The proof of a non-conservative consequence concerning Linda

The above treatment of the example is an illustration of applying the proof rules of $\mathcal{SP}$ directly. We can also consider the example from the perspective of the sets of simple, simple equivalent, and reduced equivalent consequences of Δ^{party}. To do this, we first identify the fact that the database contains four separate simple default bases. These are:

$\Delta^{party}_{s_1}$:	$linda \Rightarrow_{0.01} great$ $linda \Rightarrow_{0.1} steve$	$\Delta^{party}_{s_2}$:	$brian \Rightarrow_{0.01} great$
$\Delta^{party}_{s_3}$:	$great \Rightarrow_{0.1} noisy$	$\Delta^{party}_{s_4}$:	$\neg noisy \Rightarrow_{0.1} \neg great$

Of these only $\Delta^{party}_{s_1}$ has any interesting p-consequences beyond its consequent dual. Building the simple consequences of this set of defaults, and ignoring the consequence duals of the original defaults, and those p-consequences which have the same proposition on both sides of the turnstile, we get, by Theorem 5:

$$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (steve \wedge great, 0.11)$$
$$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (great, 0.011)$$
$$\Delta^{party}, linda \wedge great \mid\!\sim_{\mathrm{P}} (steve, 0.101)$$

We can also explore the consequences of more complex defaults in Δ^{party}. There are no conjunctive consequence defaults, so we will look at the general conjunctive default bases in Δ^{party}. Of these, again the most interesting is that with base antecedent *linda*. This is:

$$linda \Rightarrow_{0.01} great$$
$$linda \Rightarrow_{0.1} steve$$
$$linda \wedge steve \Rightarrow_{0.05} \neg noisy$$

which has the reduced equivalent:

$$linda \Rightarrow_{0.01} great$$
$$linda \Rightarrow_{0.1} steve$$
$$linda \Rightarrow_{0.05} \neg noisy$$

Once again we can easily generate a set of simple consequences from this, part of the set of conservative consequences of Δ^{party}. Ignoring, once again, consequence duals and p-consequences in which the same proposition appears on both sides of the turnstile, along with the simple consequences of Δ_{s1}^{party}, we can obtain the following conservative consequences:

$$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (steve \wedge great \wedge \neg noisy, 0.16)$$
$$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (steve \wedge \neg noisy, 0.15)$$
$$\Delta^{party}, linda \mid\!\sim_{\mathrm{P}} (great \wedge \neg noisy, 0.16)$$
$$\Delta^{party}, linda \wedge steve \mid\!\sim_{\mathrm{P}} (great \wedge \neg noisy, 0.06)$$
$$\Delta^{party}, linda \wedge great \mid\!\sim_{\mathrm{P}} (steve \wedge \neg noisy, 0.152)$$

As with the simple consequences of Δ_{s1}^{party} we can simply write these down without the need to use $\mathcal{SP}$ directly.

7. Related work

There are four main areas of closely related work. The first is the large body of work on System P as a mechanism for default reasoning. Most of this work has involved extending System P in various ways. The problem that all this work addresses is the fact that System P is too weak. The consequences it sanctions are correct, and are widely accepted as the minimum that any interesting nonmonotonic reasoning system should generate, but they are too conservative since they are guaranteed not to be false in the light of any subsequent information. Thus System P doesn't really provide nonmonotonic reasoning—it doesn't draw conclusions which are later withdrawn. In probability terms, what System P does is to look at all the constraints, embodied in the default assertions, on the probabilities of all the propositions in the database it is invoked on, and then identify the family of probability distributions which satisfy the constraints. It then sanctions any inference which represents an additional constraint that holds in *every* distribution. The various approaches to extending System P have looked for ways to choose a preferred distribution—they then sanction any assertion which satisfies that distribution. System Z[9], and equivalently rational closure[10], do this in a way which corresponds to adding the following proof rule, the rule of rational monotonicity, to System P:

$$\frac{a \mid\!\sim c, \quad a \not\mid\!\sim \neg b}{a \wedge b \mid\!\sim c}$$

Another approach to choosing the preferred distribution is to use the principle of maximum entropy, as initially suggested by Goldszmidt[11] and later extended by

Bourne[12]. More discussion of this line of work may be found elsewhere[13], but it should be clear from the above that this work has a rather different emphasis from that of this paper, being concerned with extending System P rather than working within in, and not being greatly concerned with either the strength of the conclusions or their proof.

The second main piece of related work is Bacchus'[14] inheritance reasoner. Bacchus' system allows two kinds of relation between formulae, $\Rightarrow$ and $\rightarrow$ which distinguish between strict and statistical set inclusion. Thus $\alpha \Rightarrow \beta$ denotes the fact that all αs are βs while $\beta \rightarrow \gamma$ denotes the fact that "most" βs are γs. The latter is true provided that more than half of all βs are also γs, in other words if the set of all individuals with both property β and property γ is at least half as large as the set with property β. These two relations, along with negation, are sufficient to capture a range of attractive properties for reasoning about inheritance.

Comparing Bacchus' system to $\mathcal{SP}$, there are two obvious remarks. The first is that $\mathcal{SP}$ is more expressive since Bacchus' system does not include conjunction or disjunction. The second is that $\mathcal{SP}$ is directly concerned with the bounds on the assertions which are derived, while Bacchus is only concerned with deriving whether "most" of a class of individuals have some property. The aim of his work is thus closer to System P where the value of the bounds is not an issue (though it is arguably more realistic not to depend on infinitesimal values as System P does).

It is also possible to compare the systems in more detail—though a full exploration of the differences and similarities would probably require a further paper—looking at the various properties of Bacchus' system and identifying whether they hold in $\mathcal{SP}$. The main properties of interest are those relating to deduction (captured in Bacchus' Theorem 4.1), those relating to resolving clashes between conclusions (captured in "subset preference" and "certainty preference") and the fact that inheritance is only sanctioned over one $\rightarrow$ link. The deductive properties hold in $\mathcal{SP}$, as does the failure to chain over "most" links (assuming a translation between Bacchus' $\alpha \rightarrow \beta$ and $\alpha \mid\!\sim \beta$)[§] The natural way to resolve clashes in $\mathcal{SP}$ is to look at the ϵ-values, allowing: $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\gamma, \epsilon_1)$ to be preferred to: $\Delta, \alpha \mid\!\sim_{\mathrm{P}} (\neg\gamma, \epsilon_2)$ provided that $\epsilon_1 < \epsilon_2$ (since the relevant probability is one minus the ϵ-value). Doing this ensures that certainty preference holds, but subset preference does not—that is we can resolve clashes between properties which are inherited from a class and those inherited from a superclass, but not necessarily in a way that respects specificity. This, of course, is a well-known limitation of System P[13]. Overall, then, while the two systems have some properties in common, neither captures the other.

The third piece of closely related work is that of Gilio[15], who has followed the approach adopted in the first part of this paper but using de Finetti's approach[16,17] to compute the bounds on the derived assertions. This approach makes it possible to derive the probably tightest bounds on the ϵ-values, and doing this allows Gilio

[§] This chaining property is investigated by Kraus *et al.*[5] and found not to hold for System P for much the same reason as it fails to hold for Bacchus' system.

to improve on our results for Cut and Or. The resulting rules are, respectively:

$$\frac{\alpha \wedge \beta \mid\!\sim_{\epsilon_1} \gamma, \alpha \mid\!\sim_{\epsilon_2} \beta}{\alpha \mid\!\sim_{\epsilon_1+\epsilon_2-\epsilon_1\epsilon_2} \gamma}$$

and

$$\frac{\alpha \mid\!\sim_{\epsilon_1} \gamma, \beta \mid\!\sim_{\epsilon_2} \gamma}{\alpha \vee \beta \mid\!\sim_{\frac{\epsilon_1+\epsilon_2-2\epsilon_1\epsilon_2}{1-\epsilon_1\epsilon_2}} \gamma}$$

It should be noted that the difference between the ϵ-values obtained using the improved bounds and those obtained using our bounds are small. For example, Gilio obtains an ϵ-value of 0.145 for the assertion $Linda \mid\!\sim \neg noisy$ as opposed to the 0.15 we obtain.

Finally, Snow [18,19] and Benferhat *et al.* [20] have investigated probabilistic semantics for System P which do not rely upon infinitesimal values. This work is clearly related to both our approach and Gilio's. However, it is more in line with the work on extensions to System P discussed above because it is not concerned with the actual probabilities of the assertions or their consequences. All Snow and Benferhat *et al.* are interested in is the fact that it is possible to construct suitable non-infinitesimal probability distributions which satisfy System P, they aren't particularly interested in the actual probabilities.

8. Conclusion

This paper has three main results. The first of these is to have shown that given the assumption that conditional assertions may be treated as conditional probabilities with lower bounds, we can obtain lower bounds for the derived consequences. The second main result is to have given a proof mechanism for obtaining these consequences, and to have explored its properties. Thus if we know the lower bounds on the conditional probabilities of a set of input assertions, we can identify which consequences may be derived, and establish the lower bounds on the conditional probabilities of these consequences. Moreover the lower bounds are given by simple functions of the initial bounds calculated for each proof step in System P. The third main result is to have identified three sets of consequences of a set of defaults—the simple consequences, the simple equivalent consequences, and the conservative consequences—in a way which enables them to be easily enumerated. Further sets of consequences can be obtained at the cost of some conventional theorem proving.

The advantages of these results are as follows. First they allow us to use real rather than infinitesimal probabilities since by keeping track of the bounds we can tell which consequences are justified—clearly any consequence with a low lower bound might be considered suspect. The second advantage is that only a lower bound conditional probability is required for each default rule rather than a point probability, and this may mean that the numerical values necessary for this approach are easier to assess than those necessary for approaches which use point values. Clearly we still require these values to be high or the results obtained will be useless since derived conditionals will only be known to have an associated conditional probability that is greater than some small value. The third advantage is that a subset of the full set of consequences of a given set of conditional assertions is immediately identifiable, without the need for any theorem proving. This makes it possible to both enumerate all such consequences, and to quickly establish if a particular consequence is a member of this subset.¶

Of course there are disadvantages to the use of our approach, and perhaps the worst of these stems from System P itself and our use of its proof rules. As mentioned above, System P is accepted as being a sceptical reasoning mechanism, that is, only conservative (and completely sound) conclusions can be obtained. This is insufficient for most purposes since we will often want to draw more tenuous conclusions. The fact that we use the rules of System P directly prevent us extending our approach to cover some of the specialisations of System P that have been suggested since these specialisations do not have explicit proof rules.

Finally, it is worth noting that because the initial set of lower bounded conditional probabilities are propagated through the proof, the output is a set of probability statements similar to:

$$\Pr(\alpha \wedge \neg\beta \mid \gamma \wedge \delta) \geq 1 - \epsilon$$

If the propositions γ and δ are pieces of evidence (in other words things which are known to have occurred), this output information is sufficient to establish the probability of the state $\alpha \wedge \neg\beta$. Thus the output of $\mathcal{SP}$ can be used, along with information on the utility of $\alpha \wedge \neg\beta$ as the basis of some decision making process, and this is the direction that our research on the topic of this paper is taking us now. This connection to decision theory also explains our focus on conjunctions and the fact that we have not made much use of the proof rule Or—in decision making we are not usually interested in probability statements like:

$$\Pr(\alpha \wedge \neg\beta \mid \gamma \vee \delta) \geq 1 - \epsilon$$

Acknowledgments

This work was partly funded by the EPSRC under grant GR/L84117. The second author was supported by an EPSRC PhD studentship. We are grateful to an anonymous reviewer for many useful suggestions on how to improve the paper.

References

1. M. L. Ginsberg, editor. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, San Mateo, CA, 1987.
2. R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
3. J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

¶Of course, because these results only cover a subset of the full set of consequences of a given set of assertions, the fact that a given formula is not in this subset doesn't mean that it isn't a consequence of the set of assertions.

4. Y. Shoham. Nonmonotonic logics: meaning and utility. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 388–393, 1987.
5. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
6. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.
7. E. Adams. *The Logic of Conditionals.* Reidel, Dordrecht, Netherlands, 1975.
8. D. Gabbay. *Labelled Deductive Systems.* Oxford University Press, Oxford, UK, 1996.
9. J. Pearl. System Z: a natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 121–135, 1990.
10. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.
11. M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15:220–232, 1993.
12. R. A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 50–55, San Francisco, CA, 1999. Morgan Kaufmann.
13. R. A. Bourne. *Default reasoning using maximum entropy and variable strength defaults.* PhD thesis, Department of Electronic Engineering, Queen Mary and Westfield College, London, 1999.
14. F. Bacchus. A modest, but semantically well founded, inheritance reasoner. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1104–1109, San Mateo, CA, 1989. Morgan Kaufmann.
15. A. Gilio. Precise propagation of upper and lower probability bounds in System P. In *8th International Workshop on Nonmonotonic Reasoning, Special Session: Uncertainty Frameworks in Nonmonotonic Reasoning*, 2000.
16. B. de Finetti. *Theory of Probability: A Critical Introductory Treatment.* John Wiley & Sons, Chichester, UK, 1975.
17. A. Gilio. Probabilistic consistency of conditional probability bounds. In B Bouchon-Meunier, R. R. Yager, and L. A. Zadeh, editors, *Advances in Intelligent Computing—IPMU'94*, pages 200–209. Springer Verlag, Berlin, Germany, 1994.
18. P. Snow. The emergence of ordered belief from initial ignorance. In *Proceedings of the 12th National Conference on Artificial Intelligence*, pages 281–286, Menlo Park, CA, 1994. AAAI Press/MIT Press.
19. P. Snow. Warranting non-demonstrative reasoning upon the common graound of contending uncertainty approaches. In *Proceedings of the 6th International Conference on Information Processing and the Management of Uncertainty*, pages 1381–1386, 1996.
20. S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation*, 9:873–895, 1999.