



Defeasible reasoning with variable degrees of justification

John L. Pollock

Department of Philosophy, University of Arizona, PO Box 210027, Tucson, AZ 85721, USA

Received 3 February 2001

Abstract

The question addressed in this paper is how the degree of justification of a belief is determined. A conclusion may be supported by several different arguments, the arguments typically being defeasible, and there may also be arguments of varying strengths for defeaters for some of the supporting arguments. What is sought is a way of computing the “on sum” degree of justification of a conclusion in terms of the degrees of justification of all relevant premises and the strengths of all relevant reasons.

I have in the past defended various principles pertaining to this problem. In this paper I reaffirm some of those principles but propose a significantly different final analysis. Specifically, I endorse the weakest link principle for the computation of argument strengths. According to this principle the degree of justification an argument confers on its conclusion in the absence of other relevant arguments is the minimum of the degrees of justification of its premises and the strengths of the reasons employed in the argument. I reaffirm my earlier rejection of the accrual of reasons, according to which two arguments for a conclusion can result in a higher degree of justification than either argument by itself. This paper diverges from my earlier theory mainly in its treatment of defeaters. First, it argues that defeaters that are too weak to defeat an inference outright may still diminish the strength of the conclusion. Second, in the past I have also denied that multiple defeaters can result in the defeat of an argument that is not defeated by any of the defeaters individually. In this paper I urge that there are compelling examples that support a limited version of this “collaborative” defeat.

The need to accommodate diminishers and collaborative defeat has important consequences for the computation of degrees of justification. The paper proposes a characterization of degrees of justification that captures the various principles endorsed and constructs an algorithm for computing them. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Defeasible; Defeat; Justification; Reasoning; Nonmonotonic

E-mail address: pollock@arizona.edu (J.L. Pollock).

1. Introduction

I have argued at length elsewhere that a rational agent operating in a complex environment must reason about its environment defeasibly.¹ For example, perceptual input is not always accurate, so an agent forming beliefs about its environment on the basis of its sensors must be prepared to withdraw such beliefs in the face of further information. A sophisticated agent should be able to discover generalizations about its environment by reasoning inductively, but inductive reasoning is defeasible—new evidence may overturn earlier generalizations. Because perception only enables an agent to monitor small parts of its environment at any one time, in order to build a coherent world model the agent must combine conclusions drawn on the basis of different perceptual experiences occurring at different times. But that requires a defeasible assumption that the world does not change too rapidly, so that what was perceived a moment ago is still true. The ability to maneuver through a rich environment requires an agent to be able to reason about the causal consequences of both its own actions and other events that it observes. That requires a solution to the frame problem, and it is generally agreed that any such solution will require defeasible reasoning.² I have also argued that planning with incomplete information requires a defeasible assumption that different plans do not destructively interfere with each other.³ Although most of these “theoretical” observations are fairly obvious, they have not had much impact on the actual practice of AI, because for the most part people have not tried to build autonomous agents of sufficient sophistication to encounter these problems. However, that is changing and we are getting close to the point where we can construct practical agents that will not work satisfactorily unless their designers take these observations to heart.

The OSCAR architecture for rational agents is based upon a general theory of defeasible reasoning. OSCAR implements the system of defeasible reasoning described in Pollock [28].⁴ That system in turn derives from thirty years of theorizing in philosophical epistemology. The basic idea is that the agent constructs arguments using both deductive and defeasible reason-schemes (inference-schemes). The conclusions of some of these arguments may constitute defeaters for steps of some of the other arguments. Given the set of interacting arguments that represent the agent’s epistemological state at a given time, an algorithm is run to compute degrees of justification, determining which arguments are undefeated and the level of support they provide their conclusions. What the agent should believe at any particular time are the conclusions of the undefeated arguments. The hard part of a theory of defeasible reasoning is to give an account of which arguments are undefeated. This is a topic I have addressed before [25,27,28], but some of the considerations adduced later in this paper have convinced me of the need to revisit it.

¹ The argument spans three decades. My most recent papers in this vein are [29] and [30], but see also [22,24,26,28,31].

² I have proposed and implemented a solution to the frame problem in [29].

³ A system of planning based upon this observation was described in [30], and has been implemented in OSCAR.

⁴ For a presentation of OSCAR and its current capabilities, see the OSCAR web page at <http://oscarhome.soc-sci.edu/ftp/OSCAR-web-page/OSCAR.htm>.

My analysis will turn on the taxonomy of defeaters that I introduced in [21] and [22] and that has been endorsed by most subsequent work on defeasibly reasoning (see Prakken and Vreeswijk [33] and Chesñevar, Maguitman, and Loui [4]). According to this taxonomy, there are two importantly different kinds of defeaters. Where P is a defeasible reason for Q , R is a *rebutting defeater* iff R is a reason for denying Q . All work on nonmonotonic logic and defeasible reasoning has recognized the existence of rebutting defeaters, but there are other defeaters as well. For instance, suppose x looks red to me, but I know that x is illuminated by red lights and red lights can make objects look red when they are not. Knowing this defeats the defeasible reason, but it is not a reason for thinking that x is *not* red. After all, red objects look red in red light too. This is an *undercutting defeater*. Undercutting defeaters attack the *connection* between the reason and the conclusion rather than attacking the conclusion directly. For example, an undercutting defeater for the inference from x 's looking red to x 's being red attacks the connection between “ x looks red to me” and “ x is red”, giving us a reason for doubting that x wouldn't look red unless it were red. I will symbolize the negation of “ P wouldn't be true unless Q were true” as “ $P \otimes Q$ ”. A shorthand reading is “ P does not guarantee Q ”. If Γ (a set of propositions) is a defeasible reason for P , then where $\Pi\Gamma$ is the conjunction of the members of Γ , any reason for believing “ $\Pi\Gamma \otimes P$ ” is a defeater. Thus I propose to characterize undercutting defeaters as follows:

If Γ is a defeasible reason for P , an *undercutting defeater* for Γ as a defeasible reason for P is any reason for believing “ $(\Pi\Gamma \otimes P)$ ”.

Are there any defeaters other than rebutting and undercutting defeaters? A number of authors have advocated *specificity defeaters* (e.g., Touretzky [39], Poole [32], Simari and Loui [38]). These have been formulated differently by different authors, but the general idea is that if two arguments lead to conflicting conclusions but one argument is based upon more information than the other then the “more informed” argument defeats the “less informed” one.

A phenomenon like this is common in several different kinds of probabilistic reasoning. To illustrate, consider the statistical syllogism. The statistical syllogism can be formulated as follows (see [26]):

(SS) If $r > 0.5$, then “ $Fc \ \& \ \text{prob}(G/F) \geq r$ ” is a defeasible reason for believing “ Gc ”, the strength of the reason depending upon the value of r .

When reasoning in accordance with (SS), there is a kind of “total evidence requirement” according to which we should make our inference on the basis of the most comprehensive facts regarding which we know the requisite probabilities. This can be accommodated by endorsing the following undercutting defeater for (SS):

“ $Hc \ \& \ \text{prob}(G/F \ \& \ H) \neq \text{prob}(G/F)$ ” is an undercutting defeater for (SS).

I refer to these as subproperty defeaters.⁵

⁵ I first pointed out the need for subproperty defeaters in [23]. Touretzky [40] subsequently introduced similar defeaters for use in defeasible inheritance hierarchies.

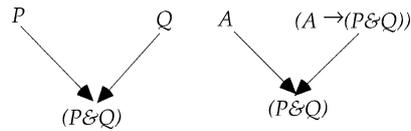


Fig. 1. An inference-graph.

Early work in AI on defeasible reasoning tended to concentrate on examples that were instances of the statistical syllogism (e.g., the venerable “Tweety” arguments), and this led people to suppose that something like subproperty defeat was operative throughout defeasible reasoning. However, I do not see any reason to believe that. There are several specific kinds of defeasible reasoning that are subject to total-evidence requirements. The statistical syllogism is one. Direct inference (discussed in section eight below) is another. Don Nute has pointed out to me that various kinds of legal reasoning and deontic reasoning are subject to a similar requirement. These can all be accommodated by acknowledging similar subproperty defeaters (which are undercutting defeaters) for the defeasible inference-schemes involved in the reasoning. To the best of my knowledge, there has never been an intuitive example of specificity defeat presented anywhere in the literature that is not an example of the operation of the total-evidence requirement in one of these special varieties of defeasible inference, and the latter are all instances of undercutting defeat. Accordingly, I will assume that undercutting defeaters and rebutting defeaters are the only possible kinds of defeaters.

I have defended the preceding remarks at length in numerous publications over the past 30 years, so for the purposes of this paper I will regard them as ancient history and take them for granted without further discussion. They are not the topic of this paper. This paper is about how to compute defeat statuses. The literature on defeasible and nonmonotonic reasoning contains numerous proposals for how to do this.⁶ The current version of OSCAR computes defeat statuses in the manner described in [27] and [28].⁷ If we ignore the fact that some arguments provide stronger support for their conclusions than other arguments, we can describe OSCAR’s defeat status computation as follows. We collect arguments into an *inference-graph*, where the nodes represent the conclusions of arguments, *support-links* tie nodes to the nodes from which they are inferred, and *defeat-links* indicate defeat relations between nodes. These links relate their *roots* to their *targets*. The root of a defeat-link is a single node, and the root of a support-link is a set of nodes. The analysis is somewhat simpler if we construct the inference-graph in such a way that when the same conclusion is supported by two or more arguments, it is represented by a separate node for each argument. For example, consider the inference-graph diagrammed in figure one, which represents two different arguments for $(P \& Q)$ given the premises, P , Q , A , and $(Q \rightarrow (P \& Q))$. The nodes of such an inference-graph represent arguments rather than just representing their conclusions. In such an inference-graph, a node has at most one support-link. When it is unambiguous to do so, I will refer to the nodes in terms of the conclusions they encode.

⁶ Two good surveys are Prakken and Vreeswijk [33] and Chesñevar, Maguitman, and Loui [4].

⁷ For comparison with default logic and circumscription, see [28], chapter three. For comparison with more recent systems of defeasible argumentation, see Prakken and Vreeswijk [33].

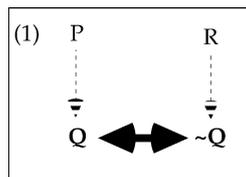
The *node-basis* of a node is the root of its support-link (if it has one), i.e., the set of nodes from which the node is inferred in a single step. If a node has no support-link (i.e., it is a premise) then the node-basis is empty. The *node-defeaters* are the roots of the defeat-links having the node as their target. We define:

A node of the inference-graph is *initial* iff its node-basis and list of node-defeaters is empty.

The *defeat status* of a node is either “defeated” or “undefeated”. It is initially tempting to try to characterize defeat statuses recursively using the following three rules:

- (1) Initial nodes are undefeated.
- (2) A node is undefeated if all the members of its node-basis are undefeated and all node-defeaters are defeated.
- (3) A node is defeated if either some member of its node-basis is defeated or some node-defeater is undefeated.

However, this recursion turns out to be ungrounded because we can have nodes of an inference-graph that defeat each other, as in inference-graph (1), where dashed arrows indicate defeasible inferences and heavy arrows indicate defeat-links. In computing defeat statuses in inference-graph (1), we cannot proceed recursively using rules (1)–(3), because that would require us to know the defeat status of Q before computing that of $\sim Q$, and also to know the defeat status of $\sim Q$ before computing that of Q . The problem is more generally that a node P can have an inference/defeat-descendant that is a defeater of P , where an inference/defeat-descendant of a node is any node that can be reached from the first node by following support-links and defeat-links. I will say that a node is P -dependent iff it is an inference/defeat-descendant of a node P . So the recursion is blocked in inference-graph (1) by there being Q -dependent defeaters of Q and $\sim Q$ -dependent defeaters of $\sim Q$.



Inference-graph (1) is a case of “collective defeat”. For example, let P be “Jones says that it is raining”, R be “Smith says that it is not raining”, and Q be “It is raining”. Given P and Q , and supposing you regard Smith and Jones as equally reliable, what should you believe about the weather? It seems clear that you should withhold belief, accepting neither. In other words, both Q and $\sim Q$ should be defeated. This constitutes a counter-example to rule (2). So not only do rules (1)–(3) not provide a recursive characterization of defeat statuses—they are not even true. The failure of these rules to provide a recursive characterization of defeat statuses suggests that no such characterization is possible, and that in turn suggested to me (in [27] and [28]) that rules (1)–(3) might be used to characterize defeat statuses in another way. Reiter’s [35] default logic proceeded in terms of multiple “extensions”, and “skeptical default logic” characterizes a conclusion as following

nonmonotonically from a set of premises and defeasible inference-schemes iff it is true in every extension. The currently popular stable model semantics [7] is derived from this approach. There are simple examples showing that these semantics are inadequate for the general defeasible reasoning of epistemic agents (see section two), but the idea of having multiple extensions suggested to me that rules (1)–(3) might be used to characterize multiple “status assignments”. On this approach, a status assignment is an assignment of defeat statuses to the nodes of the inference-graph in accordance with three simple rules:

An assignment σ of “defeated” and “undefeated” to a subset of the nodes of an inference-graph is a *partial status assignment* iff:

1. σ assigns “undefeated” to any initial node;
2. σ assigns “undefeated” to a node α iff σ assigns “undefeated” to all the members of the node-basis of α and all node-defeaters of α are assigned “defeated”; and
3. σ assigns “defeated” to a node α iff either some member of the node-basis of α is assigned “defeated”, or some node-defeater of α is assigned “undefeated”.

σ is a *status assignment* iff σ is a partial status assignment and σ is not properly contained in any other partial status assignment.

My proposal was then:

A node is undefeated iff every status assignment assigns “undefeated” to it; otherwise it is defeated.

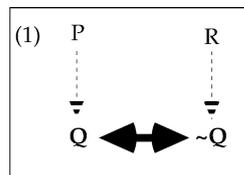
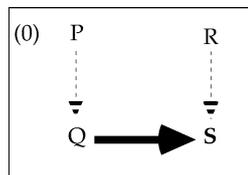
Belief in P is justified for an agent iff P is encoded by an undefeated node of the inference-graph representing the agent’s current epistemological state.

2. Examples

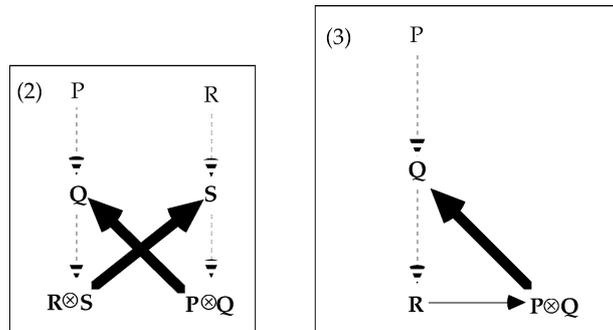
The ultimate test of a semantics is that it validates the intuitively right reasoning. It is human intuitions about correct reasoning that we want to capture. (See chapter six of [31] for further discussion of this.) With this in mind, it will be useful to have a number of examples of inference-graphs to test the analysis I will propound below. I will assume throughout this section that all initial nodes (premises) have the same degree of justification, and all reason-schemes have the same reason-strength.

The simplest case is inference-graph (0). Presumably, any semantics for defeasible reasoning will yield the result that S is defeated, and P , Q , and R are undefeated.

Inference-graph (1) illustrates “collective defeat”. It was discussed above, and again all semantics for defeasible reasoning yield the result that Q and $\sim Q$ are defeated, but P and R are undefeated.



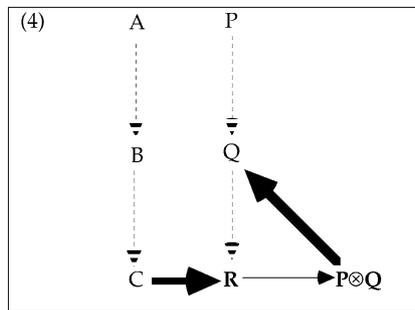
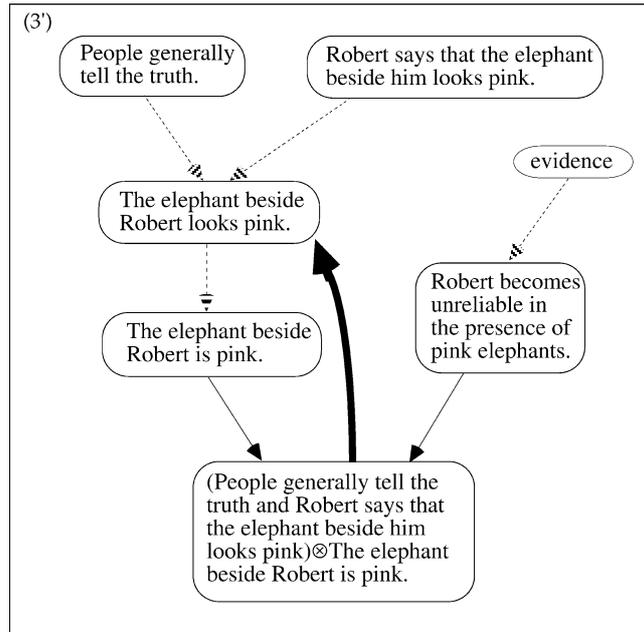
Inference-graph (2) is another instance of collective defeat, differing from inference-graph (1) in that the defeaters are undercutting defeaters rather than rebutting defeaters. The result should be that Q , S , $R \otimes S$, and $P \otimes Q$ are defeated, and P and R are undefeated. For an example, let P = “Jones says Smith is untrustworthy”, R = “Smith says Jones is untrustworthy”, Q = “Smith is untrustworthy”, S = “Jones is untrustworthy”. The semantics of section one produces two status assignments, one in which Q and $R \otimes S$ are assigned “defeated” and all other nodes are assigned “undefeated”, and one in which S and $P \otimes Q$ are assigned “defeated” and all other nodes are assigned “undefeated”.



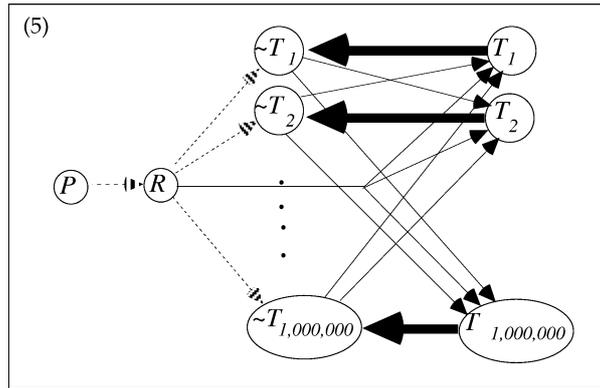
Inference-graph (3) is a simple example of a “self-defeating” argument. A partial status assignment must assign “undefeated” to P . If it assigned “undefeated” to Q then it would assign “undefeated” to R and $P \otimes Q$, in which case it would have to assign “defeated” to Q . So it cannot assign “undefeated” to Q . If it assigned “defeated” to Q it would have to assign “defeated” to R and $P \otimes Q$, in which case it would have to assign “undefeated” to Q . So that is not possible either. Thus a partial status assignment cannot assign anything to Q , R , and $P \otimes Q$. Hence there is only one status assignment (i.e., maximal partial status assignment). Accordingly, P is undefeated and the other nodes are defeated. An intuitive example having approximately the same form is shown in inference-graph (3’). Inference-graphs (3) and (3’) constitute intuitive counterexamples to default logic [35] and the stable model semantics [7] because there are no extensions. Hence on those semantics, P has the same status as Q , R , and $P \otimes Q$. It is perhaps more obvious that this is a problem for those semantics if we imagine this self-defeating argument being embedded in a larger inference-graph containing a number of otherwise perfectly ordinary arguments. On these semantics, all of the nodes in all of the arguments would have to have the same status, because there would still be no extensions. But surely the presence of the self-defeating argument should not have the effect of defeating all other (unrelated) arguments.

Inference-graph (4), illustrates that self-defeat can be “cancelled” by external defeat. Here R is defeated by C , so $P \otimes Q$ is defeated and Q is undefeated. Accordingly, there is just one status assignment, assigning “undefeated” to A , B , C , P , and Q , and “defeated” to R and $P \otimes Q$.

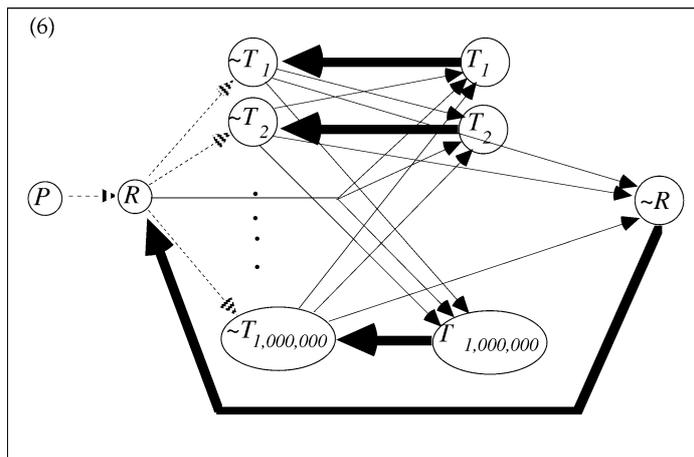
Inference-graph (5) illustrates the so-called “lottery paradox” (Kyburg [13]). Here P reports a description (e.g., a newspaper report) of a fair lottery with one million tickets. P constitutes a defeasible reason for R , which is the description. In such a lottery, each ticket



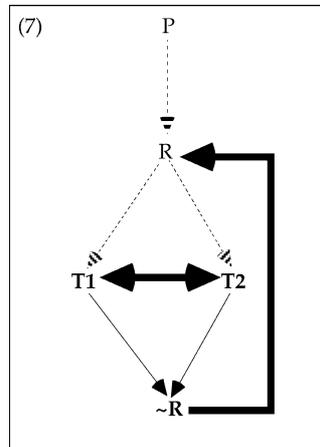
has a probability of one in a million of being drawn, so for each i , the statistical syllogism gives us a reason for believing $\sim T_i$ ("ticket i will not be drawn"). The supposed paradox is that although we thusly have a reason for believing of each ticket that it will not be drawn, we can also infer on the basis of R that some ticket will be drawn. Of course, this is not really a paradox, because the inferences are defeasible and this is a case of collective defeat. This results from the fact that for each i , we can infer T_i from the description R (which entails that some ticket will be drawn) and the conclusions that none of the other tickets will be drawn. This gives us a defeating argument for the defeasible argument to the conclusion that $\sim T_i$, as diagrammed in inference-graph (5). The result is that for each i , there is a status assignment on which $\sim T_i$ is assigned "defeated" and the other $\sim T_j$'s are all assigned "undefeated", and hence none of them are assigned "undefeated" in every status assignment.



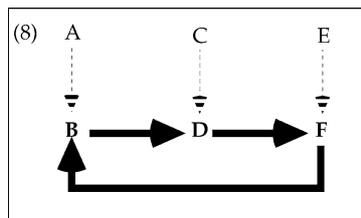
I believe that all semantics for defeasible reasoning get the lottery paradox right. A more interesting example is the “lottery paradox *paradox*”, diagrammed in inference-graph (6). This results from the observation that because R entails that some ticket will be drawn, from the collection of conclusions of the form $\sim T_i$ we can infer $\sim R$, and that is a defeater for the defeasible inference from P to R . This is another kind of self-defeating argument. Clearly, the inferences in the lottery paradox should not lead us to disbelieve the newspaper’s description of the lottery, so R should be undefeated. Circumscription (McCarthy [17]), in its simple non-prioritized form, gets this example wrong, because one way of minimizing abnormalities would be to block the inference from P to R . My own early analysis [25] also gets this wrong. This was the example that led me to the analysis of section one. That analysis gets this right. We still have the same status assignments as in inference-graph (5), and $\sim R$ is defeated in all of them because it is inferred from the entire set of $\sim T_i$ ’s, and one of those is defeated in every status assignment.



It will be convenient to have a simpler example of an inference-graph with the same general structure as the lottery paradox paradox. For that purpose we can use inference-graph (7). Here P and R should be undefeated, but T_1 , T_2 , and $\sim R$ should be defeated.



A final example that creates interesting problems involves “odd-length defeat cycles”, as in inference-graph (8). For example, let A = “Jones says that Smith is unreliable”, B = “Smith is unreliable”, C = “Smith says that Robinson is unreliable”, D = “Robinson is unreliable”, E = “Robinson says that Jones is unreliable”, F = “Jones is unreliable”. Intuitively, this should be another case of collective defeat, with A , C , and E being undefeated and B , D , and F being defeated. The semantics of section one does yield this result, but it does it in a peculiar way. A , C , and E must be assigned “undefeated”, but there is no consistent way to assign defeat statuses of B , D , and F . Accordingly, there is only one status assignment (maximal partial status assignment), and it leaves B , D , and F unassigned. We get the right answer, but it seems puzzling that we get it in a different way than we do for even-length defeat cycles like that in inference-graph (1). This difference has always bothered me.



3. Varying degrees of justification

The preceding account of defeat status assumes that all arguments support their conclusions equally strongly. However, this assumption is unrealistic. For example,

increasing the degrees of justification of the premises of an argument may increase the degree of justification of the conclusion, and increasing the strengths of the reasons employed in the argument may increase the degree of justification of the conclusion. This phenomenon has been ignored in most AI work on defeasible and nonmonotonic reasoning, but it is of considerable importance in applications of such reasoning. For example, in Pollock [29] I discussed temporal projection, wherein it is assumed defeasibly that if something is true at one time it is still true at a later time. This is, in effect, a defeasible assumption that fluents are stable, tending not to change truth values as time passes. The stability of a fluent is measured by the probability ρ that if it is true at time t then it is still true at time $t + 1$. More generally, if it is true at time t , then the probability of its being true at $t + \Delta t$ is $\rho^{\Delta t}$. So the strength of the defeasible expectation supported by temporal projection is a monotonic decreasing function of the time interval. This can be captured in a system of defeasible reasoning by employing a reasoning scheme of the following sort:

“*P-at-t*” is a defeasible reason for believing “*P-at-(t + Δt)*”, the strength of the reason being a monotonic decreasing function of Δt .⁸

The decreasing strength is important in understanding perceptual updating, wherein on the basis of new perceptual experience the agent overrides temporal projection and concludes that the fluent has changed truth value. Perception is not infallible, so perception should provide only a defeasible reason for believing that the environment is as represented by the percept.⁹ Suppose an object looks red at one time t_1 and blue at a later time t_2 . The agent should assume defeasibly that the object is initially red, but should also conclude defeasibly that it changes color later and is blue at t_2 . The object’s being red at t_1 provides a defeasible reason for expecting it to be red at t_2 , and its looking blue at t_2 provides a defeasible reason for thinking it blue and hence not red at t_2 . If these reasons were of the same strength, there would be no basis for preferring one conclusion to the other and the agent would be unable to draw a justified conclusion about the color of the object at t_2 . The situation is resolved by noting that the reason for thinking the object is still red at t_2 is weaker than the reason for thinking it was red at t_1 , and hence weaker than the reason for thinking the object is blue (and so not red) at t_2 . Because the agent has a stronger reason for thinking the object is blue at t_2 than for thinking it is red at t_2 , that becomes the justified conclusion and the agent is able to conclude that the object has changed color.

The preceding example illustrates the importance of incorporating an account of degrees of justification into a system of defeasible reasoning. There are many other examples illustrating the same point. For instance, in the statistical syllogism the strength of the reason is a function of the probability employed. Autonomous agents capable of engaging in sophisticated defeasible reasoning must accommodate varying degrees of justification. The question addressed in this paper is how the degree of justification of a conclusion should be determined. A conclusion may be supported by several different arguments. The arguments are typically defeasible, and there may also be arguments of varying strengths for defeaters for some of the supporting arguments. What is sought is a way of computing

⁸ The reason-schema proposed in [29] involves some additional qualifications, but they are not relevant to the present discussion.

⁹ This is discussed in detail in Pollock [29].

the “on sum” degree of justification of a conclusion. It is clear that three variables, at least, are involved in determining degrees of justification. The reason-strengths of the reason-schemes employed in the argument are relevant. The degrees of justification of the premises are relevant. And the degrees of justification of any defeaters for defeasible steps of the argument are relevant. Other variables might be relevant as well. I am going to assume that reason-strengths and degrees of justification are measurable as non-negative extended real numbers (i.e., the non-negative reals together with ∞). The justification for this assumption will be provided in section nine.

4. Argument-strengths

For the sake of completeness, this section and the next repeat arguments given in [28]. Let us begin by looking at arguments for which we have no arguments supporting defeaters. Let the *strength of an argument* be the degree of justification it would confer on its conclusion under those circumstances. A common and seductive view would have it that argument strength can be modeled by the probability calculus. On this view, the strength a conclusion gains from the premises can be computed in accordance with the probability calculus from the strength of the reason (a conditional probability) and the probabilities of the premises. I, and many other authors, have argued against this view at length, but it has a remarkable ability to keep attracting new converts.

There are a number of familiar arguments against the probabilistic model. The simplest argument proceeds by observing that the probabilistic model would make it impossible to be justified in believing a conclusion on the basis of a deductive argument from numerous uncertain premises. This is because as you conjoin premises, if degrees of support work like probabilities, the degree of support decreases. Suppose you have 100 independent premises, each highly probable, having, say, probability 0.99. According to the probability calculus, the probability of the conjunction will be only 0.37, so we could never be justified in using these 100 premises conjointly in drawing a conclusion. But this flies in the face of human practice. For example, an engineer building a bridge will not hesitate to make use of one hundred independent measurements to compute (deduce) the correct size for a girder. I have discussed this issue at length elsewhere [24,31], so in this paper I am just going to assume that deductive arguments provide one way of arriving at new justified conclusions on the basis of earlier ones. A corollary is that the probabilistic model is wrong.¹⁰

If deductive reasoning automatically carries us from justified conclusions to justified conclusions, then the degree of support a deductive argument confers on its conclusion cannot decrease as the number of premises increases. The degree of justification for the conclusion must be no less than that of the most weakly justified premise. This is the *Weakest Link Principle for Deductive Arguments*, according to which a deductive argument is as good as its weakest link. More precisely:

The argument strength of a deductive argument is the minimum of the degrees of justification of its premises.

¹⁰ The same objection can be leveled against the Dempster–Shafer theory [6,37].

This formulation of the weakest link principle applies only to deductive arguments, but we can use it to obtain an analogous principle for defeasible arguments. If P is a defeasible reason for Q , then we can use conditionalization to construct a simple defeasible argument for the conclusion ($P \rightarrow Q$), and this argument turns upon no premises:

$$\frac{\text{Suppose } P}{\text{Then (defeasibly) } Q.}$$

Therefore, ($P \rightarrow Q$).

As this argument has no premises, the degree of support of its conclusion should be a function of nothing but the strength of the defeasible reason. The next thing to notice is that any defeasible argument can be reformulated so that defeasible reasons are only used in subarguments of this form, and then all subsequent steps of reasoning are deductive. The conclusion of the defeasible argument is thus a deductive consequence of the premises together with a number of conditionals justified in this way. By the weakest link principle for deductive arguments, the degree of support of the conclusion should then be the minimum of (1) the degrees of justification of the premises used in the argument and (2) the strengths of the defeasible reasons:

The argument strength of a defeasible argument is the minimum of the strengths of the defeasible reasons employed in it and the degrees of justification of its premises.

This is the general *Weakest Link Principle*. The problem of computing argument strengths is thus computationally simple.

5. The accrual of reasons

If we have two independent reasons for a conclusion, does that make the conclusion more justified than if we had just one? It is natural to suppose that it does,¹¹ but upon closer inspection that becomes unclear. Cases that seem initially to illustrate such accrual of justification appear upon reflection to be better construed as cases of having a single reason that subsumes the two separate reasons. For instance, if Brown tells me that the president of Fredonia has been assassinated, that gives me a reason for believing it; and if Smith tells me that the president of Fredonia has been assassinated, that also gives me a reason for believing it. Surely, if they both tell me the same thing, that gives me a better reason for believing it. However, there are considerations indicating that my reason in the latter case is not simply the conjunction of the two reasons I have in the former cases. Reasoning based upon testimony is a straightforward instance of the statistical syllogism. We know that people tend to tell the truth, and so when someone tells us something, that gives us a defeasible reason for believing it. This turns upon the following probability being reasonably high:

$$(1) \quad \text{prob}(P \text{ is true} / S \text{ asserts } P).$$

¹¹ See, for example, Verheij [41].

Given that this probability is high, I have a defeasible reason for believing that the president of Fredonia has been assassinated if Brown tells me that the president of Fredonia has been assassinated.

In the discussion of the weakest link principle, I urged that argument strengths do not conform to the probability calculus. However, that must be clearly distinguished from the question of whether probabilities license defeasible inferences. In fact, I think that a large proportion of our defeasible inferences are based upon probabilities. Such inferences proceed in terms of the statistical syllogism, described in section one and formulated as principle (SS). When we have the concurring testimony of two people, our degree of justification is not somehow computed by applying a predetermined function to the latter probability. Instead, it is based upon the quite distinct probability

$$(2) \quad \text{prob}(P \text{ is true}/S_1 \text{ asserts } P \text{ and } S_2 \text{ asserts } P \text{ and } S_1 \neq S_2).$$

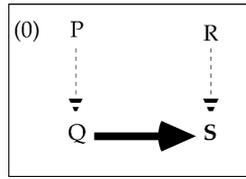
The relationship between (1) and (2) depends upon contingent facts about the linguistic community. We might have one community in which speakers tend to make assertions completely independently of one another, in which case $(2) > (1)$; and we might have another community in which speakers tend to confirm each other's statements only when they are fabrications, in which case $(2) < (1)$. Clearly our degree of justification for believing P will be different in the two linguistic communities. It will depend upon the value of (2), rather than being some function of (1).

It is important to distinguish *epistemic reasoning*—reasoning about what to believe—from *practical reasoning*—reasoning about what actions to perform. These two kinds of reasoning have quite different logical properties, as is illustrated at length in [31]. Something like the accrual of reasons seems to hold for practical reasoning. Normally, two independent reasons for choosing a particular action provide a stronger justification for choosing it than either reason by itself. But we cannot conclude from this that the same thing is true of epistemic reasoning.

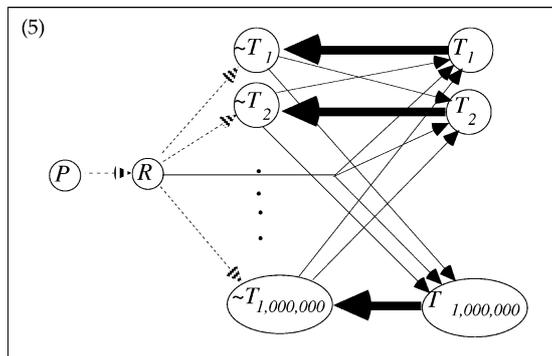
All examples I have considered that seem initially to illustrate the accrual of reasons in epistemic reasoning turn out in the end to have this same form. They are all cases in which we can estimate probabilities analogous to (2) and make our inferences on the basis of the statistical syllogism rather than on the basis of the original reasons. Accordingly, I doubt that epistemic reasons do accrue. If we have two separate undefeated arguments for a conclusion, the degree of justification for the conclusion is the maximum of the strengths of the two arguments. This will be my assumption.

6. The influence of defeaters

Thus far I have considered how reason-strengths and the degrees of justification of premises affect the degree of justification of a conclusion. The third variable we must consider is the presence of arguments supporting defeaters. Suppose we have only two arguments to consider, and the conclusion of one of them is a defeater for the final step of the other, as diagrammed in inference-graph (0). How should this affect the degree of justification of S ?



It seems clear that if the argument strength of the left argument (for Q) is as great as that of the right argument (for S), then the degree of justification of S should be 0. But what if the argument strength of the left argument is less than that of the right argument? In [28] I maintained that defeat was an all-or-nothing matter, and hence weaker defeaters leave arguments unaffected. In the scenario just described, this has the consequence that the degree of justification of S is the same as the argument strength of the right argument. However, there are some examples that now convince me that this is incorrect. The simplest examples have to do with biased lotteries. To see how these examples work, recall the earlier analysis of reasoning about fair lotteries. Consider a fair lottery consisting of 1 million tickets, and suppose it is known that one and only one ticket will win. Observing that the probability is only 0.000001 of any particular ticket being drawn given that it is a ticket in the lottery, it seems initially reasonable to employ the statistical syllogism and accept the conclusion regarding any particular ticket that it will not be drawn. This reasoning is completely general and applies to each ticket. However, these conclusions conflict jointly with something else we are justified in believing, viz., that some ticket will be drawn. We cannot be justified in believing each member of an explicitly contradictory set of propositions, and we have no way to choose between them, so it follows intuitively that we are not justified in believing of any ticket that Jones did not hold that ticket. The formal reconstruction of this reasoning proceeds by observing that this is a case of collective defeat. For each n , the statistical syllogism provides a defeasible reason for believing “ $\sim T_n$ ”. But for each k , we have an equally strong defeasible reason for believing each “ $\sim T_k$ ”. We know that some ticket will be drawn. Thus we can construct the counterargument diagrammed in inference-graph (5) for the conclusion that “ T_n ” is true. Our reason for believing each “ $\sim T_k$ ” is as good as our reason for believing “ $\sim T_n$ ”, so we



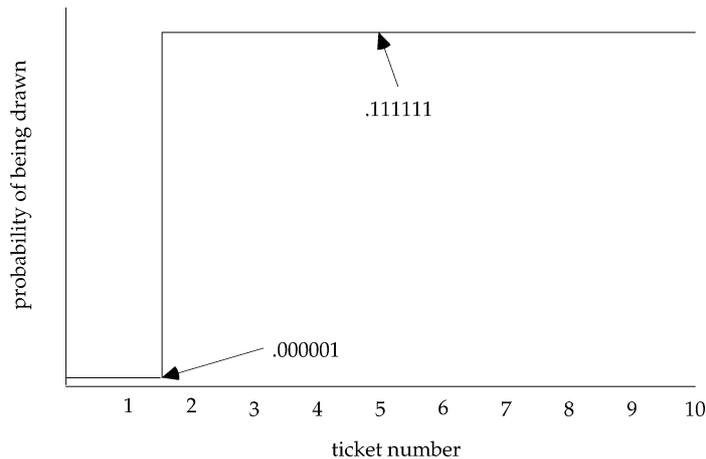


Fig. 2. Lottery 2.

have as strong a reason for “ T_n ” as for “ $\sim T_n$ ”. Hence our defeasible reason for the latter is defeated and we are not justified in believing “ $\sim T_n$ ”.

Next, consider lottery 2, which is a biased lottery consisting of just ten tickets. The probability of ticket 1 being drawn is 0.000001, and the probability of any other ticket being drawn is 0.111111. It is useful to diagram these probabilities as in Fig. 2. In *this* lottery, it seems reasonable to infer that ticket 1 will not be drawn, because the probability of any other ticket being the drawn is more than 100,000 times greater. This can be justified as follows. As before, we have a defeasible reason for believing “ $\sim T_n$ ”, for each n . But these reasons are no longer of equal strength. Because ticket 1 is much less likely to be drawn than any other ticket, we have a much stronger reason for believing that ticket 1 will not be drawn. As before, for $n > 1$, we have the counterargument diagrammed in inference-graph (5) for “ T_n ”, and that provides as good a reason for believing “ T_n ” as we have for believing “ $\sim T_n$ ”. Thus the defeasible reason for “ $\sim T_n$ ” is defeated. But we do not have as good a reason for believing “ T_1 ” as we do for believing “ $\sim T_1$ ”. An argument is only as good as its weakest link, and the counterargument for “ T_1 ” employs the defeasible reasons for “ $\sim T_n$ ” for $n > 1$. These reasons are based upon lower probabilities (of value 0.888889) and hence are not as strong as the defeasible reason for “ $\sim T_1$ ” (based upon a probability of value 0.999999). Thus, although we have a reason for believing “ T_1 ”, we have a better reason for believing “ $\sim T_1$ ”, and so on sum we are justified in believing the latter.

Now contrast lottery 2 with lottery 3, which consists of 10,000 tickets. In lottery 3, the probability of ticket 1 being drawn is still 0.000001, but the probability of any other ticket being drawn is 0.000011. This is diagrammed as in Fig. 3. It may still be reasonable to infer that ticket 1 will not be drawn, but, and this is the crucial observation, the justification for this conclusion does not seem to be nearly so strong. This is because although we have the same defeasible argument for “ $\sim T_1$ ”, the reasons involved in the counterargument for “ T_1 ” are now much better, being based upon a probability of 0.999989. They are still not strong enough to defeat the argument for “ $\sim T_1$ ” outright, but they seem to weaken the justification. Thus the degree of justification for “ $\sim T_1$ ” is lower in lottery 3 than it is in

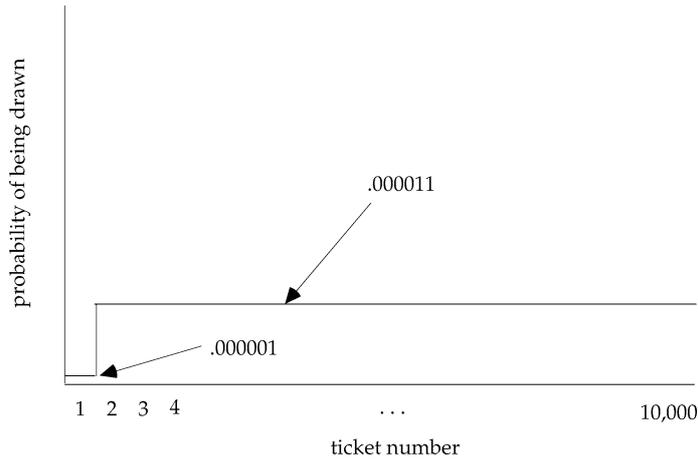


Fig. 3. Lottery 3.

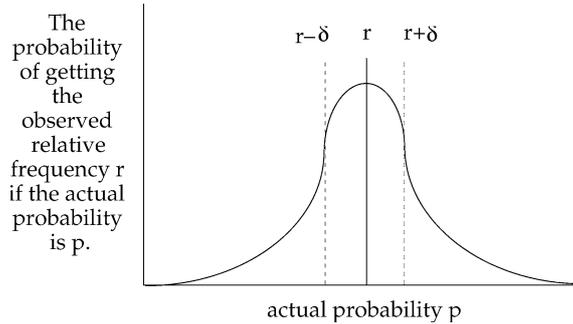


Fig. 4. Statistical induction.

lottery 2. The difference between lottery 2 and lottery 3 seems to illustrate that defeaters that are too weak to defeat a conclusion outright may still lower the degree of justification. In other words, they act as *diminishers*.

In [26] I argued that reasoning similar to this treatment of biased lotteries is what underlies statistical induction. In statistical induction, having observed a sample of *A*'s and found that the proportion of members of the sample that are *B*'s is r , we infer defeasibly that the probability of an arbitrary *A* being a *B* is approximately r , i.e., lies in an interval $[r - \delta, r + \delta]$ around the observed relative frequency r . The logic of the reasoning that allows us to conclude this is the same as that involved in reasoning about biased lotteries. We know that the actual probability p is in the interval $[0, 1]$. This is like knowing that some ticket will be drawn. If the sample is large, then for each choice of p the probability of getting the observed relative frequency is low given that p is the actual probability, but for some choices of p (those further from r) it is lower than for others. So we can reason as in the biased lottery that the actual probability does not lie on the tails of the bell curve.

Biased lotteries and statistical induction illustrate that given an undefeated argument for P and an otherwise undefeated weaker argument for $\sim P$, the degree of justification for P should be the argument strength of the first argument decremented by an amount determined by the argument strength of the second argument. That is, there should be a function \mathcal{J} such that given two arguments that rebut one another, if their strengths are x and y , the degree of justification for the conclusion of the former is $\mathcal{J}(x, y)$. \mathcal{J} must satisfy the following conditions:

- (3) $\mathcal{J}(x, y) \leq x$.
- (4) $\mathcal{J}(x, 0) = x$.
- (5) If $y \geq x$ then $\mathcal{J}(x, y) = 0$.
- (6) If $x \geq z$ and $w \geq y$ then $\mathcal{J}(x, y) \geq \mathcal{J}(z, w)$.

This leaves undetermined how $\mathcal{J}(x, y)$ and $\mathcal{J}(z, w)$ compare in cases in which $x \geq z$ and $w < y$. To resolve these cases, let us replace (3) and (6) by a stronger assumption:

- (7) If $\varepsilon \geq 0$, $\mathcal{J}(x + \varepsilon, y + \varepsilon) = \mathcal{J}(x, y)$.

This is a “linearity assumption”. It tells us that increasing the argument strength and the defeat strength by the same amount ε leaves the resulting degree of justification unchanged. With this further assumption, it becomes determinate how $\mathcal{J}(x, y)$ and $\mathcal{J}(z, w)$ compare, for any choice of x, y, z, w . Let us define:

$$x \sim y = \begin{cases} x - y & \text{if } y < x, \\ 0 & \text{otherwise.} \end{cases}$$

Then we have in general:

Theorem 1. *If (4), (5), and (7) hold, $\mathcal{J}(x, y) = x \sim y$.*

Proof. If $x \leq y$, then by (5) $\mathcal{J}(x, y) = 0 = x \sim y$. Suppose instead that $x \geq y$. Then by (7) and (4), $\mathcal{J}(x, y) = \mathcal{J}(x - y, 0) = x - y = x \sim y$. \square

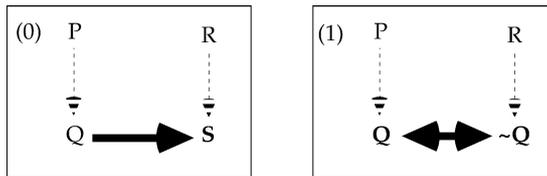
So I will assume:

- (8) Given an otherwise undefeated argument of strength x supporting P , and an otherwise undefeated argument of strength y supporting $\sim P$, and no other relevant arguments, the degree of justification of P is $x \sim y$.

The argument for diminishers is based on intuitive examples. In particular, I have appealed to biased lotteries. If one ticket is *much* more improbable than the others, it seems reasonable to conclude that it will not be drawn. But if it is only minutely less probable, that does not seem reasonable. In general, if the argument for a defeater is only minutely weaker than the argument for a conclusion, it does not seem reasonable to regard the conclusion as unscathed. These observations seem to me to be unassailable.

7. Computing defeat statuses

How should diminishers affect our defeat status computation? I assume that rather than merely assigning “defeated” or “undefeated”, a status assignment should now assign numerical degrees of justification $j(P)$ to nodes. How is $j(P)$ determined? In the following examples, unless I explicitly say otherwise, I assume throughout that the reason-strengths are at least as great as the degrees of justification of the initial nodes so that they can be ignored in computing degrees of justification. If we consider a very simple inference-graph like (0), it seems clear that we should have $j(Q) = j(P)$, $j(S) = j(R) \sim j(Q)$. This is in accordance with principle (8). So if $j(P) \geq j(R)$, S is defeated, otherwise it is diminished.



Consider the marginally more complicated inference-graph (1). Here there seem to be two possibilities regarding how the degrees of justification are to be computed:

- (a) We could have $j(Q) = j(P) \sim j(\sim Q)$ and $j(\sim Q) = j(R) \sim j(Q)$.
- (b) We could have $j(Q) = j(P) \sim j(R)$ and $j(\sim Q) = j(R) \sim j(P)$.

These seem to be the only two possibilities for this simple inference-graph. However, (a) is not a genuine possibility. We have the following theorem:

Theorem 2. *If $j(P) > j(R)$, $j(Q) = j(P) \sim j(\sim Q)$ and $j(\sim Q) = j(R) \sim j(Q)$ then $j(Q) = j(P)$ and $j(\sim Q) = 0$.*

Proof. Suppose $j(Q) \neq j(P)$. As $j(Q) = j(P) \sim j(\sim Q)$, $j(Q) < j(P)$, so $j(\sim Q) \neq 0$. Then as $j(\sim Q) = j(R) \sim j(Q)$, $j(\sim Q) = j(R) - j(Q) \leq R$. By assumption, $j(P) > j(R)$, so $j(P) > j(\sim Q)$, and hence

$$j(Q) = j(P) \sim j(\sim Q) = j(P) - j(\sim Q) = j(P) - j(R) + j(Q).$$

Thus $j(R) = j(P)$, which is impossible given the assumption that $j(P) > j(R)$. So by reductio, $j(Q) = j(P)$. As $j(P) > j(R)$, $j(P) > 0$. But $j(Q) = j(P) \sim j(\sim Q)$, so $j(\sim Q) = 0$. □

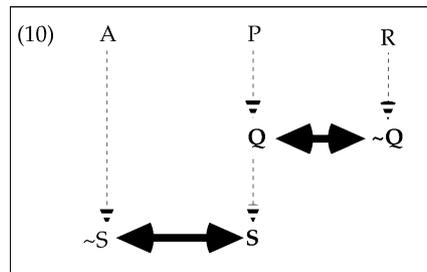
Thus (a) is incompatible with diminishing. That leaves only (b) as a possible computation of defeat statuses in inference-graph (1). That is,

$$j(Q) = j(P) \sim j(R) \quad \text{and} \quad j(\sim Q) = j(R) \sim j(P).$$

This means that in computing $j(Q)$, we take the strength of the argument supporting it, in this case determined by $j(P)$, and then subtract the strength of the argument supporting

the defeater, i.e., $j(R)$. We do not subtract the strength of the defeater itself, i.e., we do not subtract $j(\sim Q)$.

If we apply (b) to the case in which $j(P) = j(R)$, we get a single status assignment in which $j(Q) = j(\sim Q) = 0$. This has a surprising consequence. The semantics for defeasible reasoning described in section one, as well as default logic, the stable model semantics, circumscription, and almost all standard semantics for defeasible reasoning and nonmonotonic logic, support what I have called [25] “presumptive defeat”.¹² For example, consider inference-graph (10). A defeated conclusion like Q that is assigned “defeated” in some status assignment and “undefeated” in another retains the ability to defeat. In the case of inference-graph (10) this has the consequence that S is assigned “defeated” in those status-assignments in which Q is assigned “defeated”, but S is assigned “undefeated” and $\sim S$ is assigned “defeated” in those status-assignments in which Q is assigned “undefeated”. Touretzky, Horty, and Thomason [40] called this “ambiguity propagation”, and Makinson and Schlechta ([16] called such arguments “Zombie arguments” (they are dead, but they can still get you). However, computation (b) precludes presumptive defeat. It entails that there is a single status assignment in which $j(S) = j(Q) = j(\sim Q) = 0$, and $j(\sim S) = j(A)$. So Q , $\sim Q$, and S are all defeated, and $\sim S$ is undefeated. Is this the right answer? Consider an example: Jones and Smith hear one weather forecast, but disagree about whether rain was predicted (Q). Q gives one reason to believe it will rain (S). You didn’t hear the first forecast, but you hear another forecast (A), which says it will not rain. Should you believe $\sim S$? I have some inclination to think you should, but I don’t find the answer obvious, nor have I ever found another example of presumptive defeat where the answer is obvious. At a recent workshop on defeasible reasoning held at the University of Georgia, I found there to be no consensus among the participants as to the intuitive status of presumptive defeat.

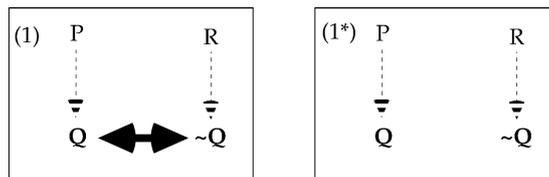


In the absence of clear intuitions, how can we decide whether presumptive defeat is a genuine phenomenon of defeasible reasoning, or an undesirable artifact of existing semantics? The preceding considerations constitute a rather strong argument for the conclusion that presumptive defeat is incompatible with diminishing. If it is granted that a correct semantics must accommodate diminishing, this means in turn that most standard semantics for defeasible reasoning produce incorrect assessments of inference-graph (10)

¹² The only semantics I know about that does not support presumptive defeat are certain versions of Nute’s [19] defeasible logic. See also Covington, Nute, and Vellino [5], and Nute [20].

even when all premises have the same degree of justification and all reason-strengths are the same. It is generally assumed that by ignoring variations in degrees of justification and reason-strengths we are simplifying the problem of constructing a semantics for defeasible reasons, but the present considerations indicate that by doing so we may also obscure phenomena that have important implications for the semantics even in this simplified case.

The fact that we get a single assignment of defeat statuses in inference-graph (1) suggests that we are not really computing status assignments. Rather, we are computing degrees of justification directly. The appeal to status assignments was motivated by the thought that we could not compute degrees of justification recursively because a node P of an inference-graph can have P -dependent defeaters (i.e., defeaters that are inference/defeat-descendants of P). But what the present approach may yield is a way of doing a different kind of recursive computation of degrees of justification. In inference-graph (1), we cannot compute $j(Q)$ without having a value for $\sim Q$, and we cannot compute $j(\sim Q)$ without having a value for Q . So we cannot do a straightforward recursive computation of degrees of justification. However, the values required for Q and $\sim Q$ need not be their degrees of justification. Indeed, they cannot be because those would be zero for each. In computing $j(Q)$, what we want to subtract from $j(P)$ is not $j(\sim Q)$ but rather a measure of the strength of the argument for $\sim Q$. The value of the latter should be $j(R)$. This measure can be obtained by removing the mutual defeat between the two arguments, as in inference-graph (1*), and computing $j(Q)$ and $j(\sim Q)$ in the new inference-graph. Call those values $j^*(Q)$ and $j^*(\sim Q)$. Then returning to inference-graph (1), the values we subtract when we compute $j(Q)$ and $j(\sim Q)$ are $j^*(\sim Q)$ and $j^*(Q)$. That is, $j(Q) = j(P) \sim j^*(\sim Q) = j(P) \sim j(R)$, and $j(\sim Q) = j^*(\sim Q) \sim j^*(Q) = j(R) \sim j(P)$. In constructing inference-graph (1*), we remove two defeat-links. Each link has the characteristic that removing it results in Q no longer having a Q -dependent defeater, and also in $\sim Q$ no longer having a $\sim Q$ -dependent defeater.

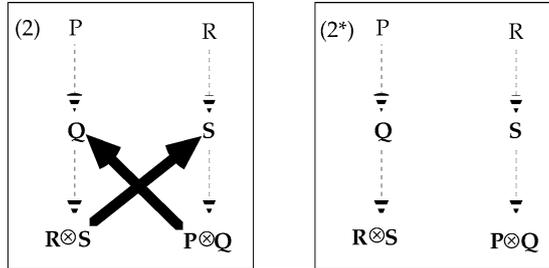


Now consider inference-graph (2). Again, there seem to be two ways the computation of degrees of justification might go. Presumably $j(R \otimes S) = j(Q)$ and $j(P \otimes Q) = j(S)$. Then we might have either:

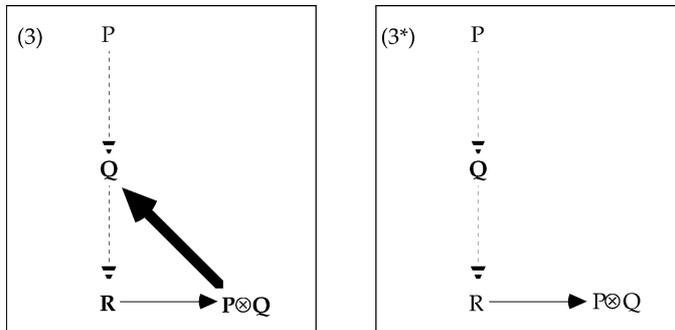
- (a) $j(S) = j(R) \sim j(R \otimes S) = j(R) \sim j(Q)$ and $j(Q) = j(P) \sim j(P \otimes Q) = j(P) \sim j(S)$; or
- (b) $j(S) = j(R) \sim j(P)$ and $j(Q) = j(P) \sim j(R)$.

As in Theorem 1, (a) is incompatible with diminishing, so (b) seems to be the only possibility. This means that in computing $j(S)$ we begin with the strength of the argument supporting S , i.e., $j(R)$, and then subtract the strength of the argument supporting the defeater $R \otimes S$ would have in the absence of the defeater $P \otimes Q$ that is obtained from S . We compute $j(Q)$ analogously. This is the same thing as computing $j(R \otimes S)$ and

$j(P \otimes S)$ in inference-graph (2*) (call the resulting values $j^*(R \otimes S)$ and $j^*(P \otimes Q)$) and then setting $j(Q) = j^*(Q) \sim j^*(P \otimes Q)$ and $j(S) = j^*(S) \sim j^*(R \otimes S)$. Again, the defeat-links we remove in constructing inference-graph (2*) are those such that if we remove either, Q no longer has a Q -dependent defeater, and similarly for S .

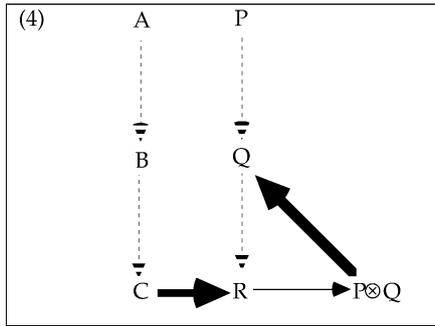


Consider inference-graph (3). Presumably $j(P \otimes Q) = j(R) = j(Q) = 0$. We can get that by ruling that $j(Q) = j(P) \sim j(P)$. So here we begin with the strength of the argument supporting Q , i.e., $j(P)$, and then subtract the strength the argument supporting $P \otimes Q$ would have in the absence of the defeater (itself) that is obtained from Q . This is the value assigned to $(P \otimes Q)$ in inference-graph (3*). Again, the defeat-link we remove in constructing (3*) is the only defeat-link whose removal results in Q no longer having a Q -dependent defeater.

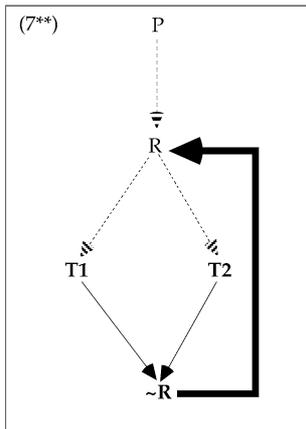
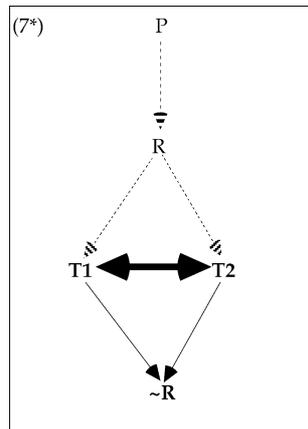
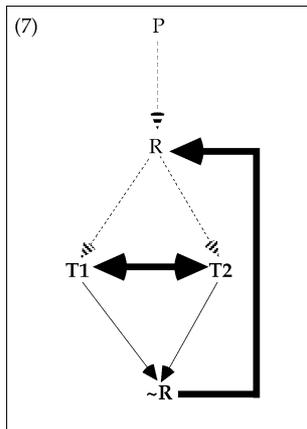


Inference-graph (4) is analogous. In (4) we compute $j^*(P \otimes Q)$ by removing the defeat-link between $P \otimes Q$ and Q . So $j^*(P \otimes Q) = j(P) \sim j(A)$. That has the result that if $j(A) = j(P)$ then $j^*(P \otimes Q) = j^*(R) = 0$ and so $j(Q) = j^*(Q) \sim j^*(P \otimes Q) = j(P)$. I will continue to refer to $j^*(P \otimes Q)$ as the argument-strength of $(P \otimes Q)$, but note that this is a somewhat more complex notion than the argument-strength discussed in section four. In the present sense, argument-strengths take account not only of the strengths of the premises and inference-schemes but also of the effect of those defeaters that are independent of the node being evaluated.

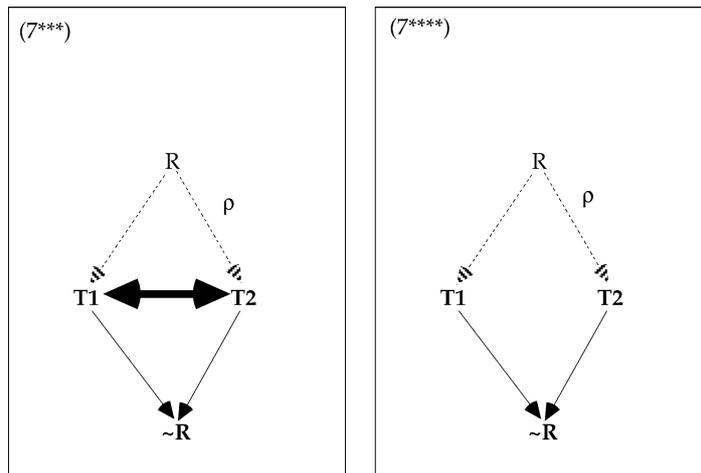
Finally, consider the lottery paradox paradox, in the guise of inference-graph (7). We construct (7*) by removing the only defeat-link whose removal results in R no



longer having an R -dependent defeater. In (7^*) , the triangle consisting of R , $T1$ and $T2$ is analogous to inference-graph (1), and so $j^*(T1) = j^*(R) \sim j^*(R) = 0$, and $j^*(T2) = j^*(R) \sim j^*(R) = 0$. $j^*(\sim R) = \min\{j^*(T1), j^*(T2)\} = 0$. Then $j(R) = j^*(R) \sim j^*(\sim R) = j(P)$.



When we turn to the computation of $j(T1)$ and $j(T2)$ in (7), we encounter an additional complication. Removing the defeat-links whose removal results in $T1$ no longer having a $T1$ -dependent defeater produces inference-graph (7**). On the strength of the preceding examples we might expect that $j(T1) = j^{**}(T1) \sim j^{**}(T2)$. However, inference-graph (7**) is analogous to inference-graph (3). Accordingly, $j^{**}(T1) = j^{***}(T2) = j^{**}(R) = j^{**}(\sim R) = 0$. This produces the intuitively correct answer that $j(T1) = 0$, but it seems to do so in the wrong way. To see this, let us modify the example slightly by taking it to represent a biased lottery. More precisely, let us suppose that the reason-strength of the reason supporting the inference of $T1$ from R is at least as great as the degree of justification of P (and hence of R), but suppose the reason-strength ρ of the reason supporting the inference of $T2$ from R is less than the degree of justification of P . In this case $T2$ should be defeated, and hence $\sim R$ should be defeated, but $T1$ should only be diminished. However, this change does not affect the computation of degrees of justification in inference-graph (7**). We still get that $j^{**}(T1) = j^{**}(T2) = j(R) = j^{**}(\sim R) = 0$.



The source of the difficulty is that the computation of degrees of justification is not being done recursively. We should first compute the degree of justification of R , as above. Then holding that degree of justification fixed we should go on to compute the degrees of justification of its inference-children, in this case $T1$ and $T2$. This amounts to regarding R as an initial node whose degree of justification is that computed in inference-graph (7). This yields inference-graph (7**). The degrees of justification for $T1$ and $T2$ are then those computed in inference-graph (7**). That computation proceeds by constructing inference-graph (7****) and computing $j^{****}(T1) = j(R)$, $j^{****}(T2) = r$, and so $j(T1) = j^{**}(T1) = j^{****}(T1) \sim j^{****}(T2) = j(R) \sim \rho$. Analogously, $j(T2) = j^{**}(T2) \sim j^{****}(T1) = 0$. Continuing recursively, $j(\sim R) = \min\{j(T1), j(T2)\} = 0$.

The defeat status computation that emerges from these examples proceeds in accordance with two rules. We begin with an inference-graph \mathcal{G} . Where φ is a node of \mathcal{G} , let $j(\varphi, \mathcal{G})$

be the degree of justification of φ in \mathcal{G} . The first rule governs the case in which P has no P -dependent defeaters. In that case, the computation proceeds in accordance with the originally proposed rules (1)–(3). This has the effect of computing degrees of justification in terms of the degrees of justification of the basis of P and the defeaters of P . The second rule governs the case in which P has P -dependent defeaters. In that case the computation of the degree of justification of P proceeds instead in terms of the argument-strength for P and maximal argument-strength of the P -dependent defeaters. Those argument-strengths are computed by constructing a new inference-graph \mathcal{G}_P by removing each defeat-link of \mathcal{G} whose removal results in P no longer having a P -dependent defeater. More generally, there can be parallel routes from one node to another, with the result that we must delete multiple defeat-links to ensure that P no longer has a P -dependent defeater. So let us define:

A defeat-link is *P-critical* in an inference-graph \mathcal{G} iff it is a member of some minimal set of defeat-links such that the removal of all the links in the set results in P no longer having a P -dependent defeater.

Let \mathcal{G}_P be the inference-graph that results from removing all P -critical defeat-links from \mathcal{G} and making all P -independent nodes φ initial with $j(\varphi, \mathcal{G}_P) = j(\varphi, \mathcal{G})$. The argument strengths in \mathcal{G} are then the degrees of justification in \mathcal{G}_P .

Putting this altogether, the two rules are as follows (where $\max(\emptyset) = 0$):

(DJ1) If P is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the defeaters for P , and no D_i is P -dependent, then

$$j(P, \mathcal{G}) = \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\ \sim \max\{j(D_1, \mathcal{G}), \dots, j(D_k, \mathcal{G})\}.$$

(DJ2) If P has P -dependent defeaters D_1, \dots, D_k in \mathcal{G} and \mathcal{G}_P results from deleting all P -critical defeat-links from \mathcal{G} and making all P -independent nodes φ initial with $j(\varphi, \mathcal{G}_P) = j(\varphi, \mathcal{G})$, then

$$j(P, \mathcal{G}) = j(P, \mathcal{G}_P) \sim \max\{j(D_1, \mathcal{G}_P), \dots, j(D_k, \mathcal{G}_P)\}.$$

The general idea is that the computation of degrees of justification is made recursive by appealing to argument strengths rather than degrees of justification in ungrounded cases. Argument strengths are computed by computing degrees of justification in simpler inference-graphs from which the source of ungroundedness has been removed. The result is a recursive computation of degrees of justification.

8. Collaborative defeat

The rules (DJ1) and (DJ2) have the effect of dividing the defeaters of P into two sets—those that are P -dependent and those that are not. These two sets of defeaters then affect $j(P, \mathcal{G})$ separately. Those that are P -independent will diminish the value of P in \mathcal{G}_P , reducing $j(P, \mathcal{G}_P)$, and then those that are P -dependent will further diminish the value of P in \mathcal{G} . So if D_1 is the most strongly supported defeater that is P -independent, and D_2 is the most strongly supported defeater that is P -dependent, the result will be that

$$\begin{aligned}
j(P, \mathcal{G}) &= j(P, \mathcal{G}_P) \sim j(D_2, \mathcal{G}_P) \\
&= (\min\{\rho, j(B_1, \mathcal{G}_P), \dots, j(B_n, \mathcal{G}_P)\} \sim j(D_1, \mathcal{G}_P)) \sim j(D_2, \mathcal{G}_P) \\
&= (\min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \sim j(D_1, \mathcal{G})) \sim j(D_2, \mathcal{G}_P) \\
&= \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \sim (j(D_1, \mathcal{G}) + j(D_2, \mathcal{G})).
\end{aligned}$$

In this way, we can replace the sequential application of (DJ2) and (DJ1) by the application of a single principle:

(DJ) If P is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the P -independent defeaters for P , D_{k+1}, \dots, D_m are the P -dependent defeaters of P , and \mathcal{G}_P results from deleting all P -critical defeat-links from \mathcal{G} and making all P -independent nodes φ initial with $j(\varphi, \mathcal{G}_P) = j(\varphi, \mathcal{G})$, then

$$\begin{aligned}
j(P, \mathcal{G}) &= \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\
&\sim [\max\{j(D_1, \mathcal{G}), \dots, j(D_k, \mathcal{G})\} \\
&\quad + \max\{j(D_{k+1}, \mathcal{G}_P), \dots, j(D_m, \mathcal{G}_P)\}].
\end{aligned}$$

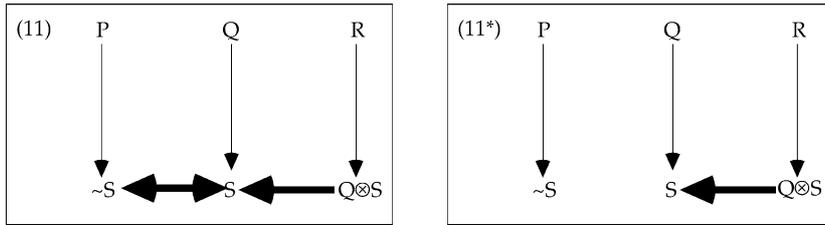
This produces a kind of double counting of defeaters. It has the consequence that although no single defeater may be sufficient to defeat the inference to P , two defeaters can accomplish that by acting in unison. I will call this *collaborative defeat*.¹³

Collaborative defeat might seem suspect. However, I will now argue that there are examples which illustrate that it actually occurs in defeasible reasoning. As an autobiographical note, I first became convinced of the need for collaborative defeat on the basis of examples like those I will present below. This was long before I had a theory of diminishers that entailed collaborative defeat. My expectation was that I would have to construct an initial theory to accommodate diminishers, and then embellish it in some rather ad hoc way to make it compatible with collaborative defeat. Thus I am delighted to find that collaborative defeat simply falls out of my theory of diminishers without any ad hoc treatment.

To set the stage, note that you cannot have two rebutting defeaters such that one is P -dependent and the other is not. Rebutting defeaters of P are always P -dependent. This is because rebutting defeat is symmetrical, so if you can follow defeat-links in one direction, you can follow them back again. Thus you can have one defeater that is P -dependent and another that is not only if one of them is an undercutting defeater. The general form is that of inference-graph (11). To compute $j(S, \mathcal{G}_{11})$, we construct inference-graph (11*). $j(S, \mathcal{G}_{11}) = j(Q, \mathcal{G}_{11}) \sim j(R, \mathcal{G}_{11})$, and so $j(S, \mathcal{G}_{11}) = (j(Q, \mathcal{G}_{11}) \sim j(S, \mathcal{G}_{11})) \sim j(P, \mathcal{G}_{11}) = j(Q, \mathcal{G}_{11}) \sim (j(P, \mathcal{G}_{11}) + j(R, \mathcal{G}_{11}))$.

To confirm that the computation in inference-graph (11) is correct, we must consider examples that mix rebutting defeat and undercutting defeat. One place in which this is common is in inferences involving a total-evidence requirement. Suppose the total

¹³ Normally, the collaborating defeaters will be reasons for different conclusions, viz., $\sim P$ and $(P \otimes Q)$. But in some cases they can both be reasons for $(P \otimes Q)$, in which case this has a similar flavor to the accrual of reasons. It is not anything that simple, however, because each defeater is the strongest defeater (not a sum) from a whole class of defeaters—the dependent and independent ones.



evidence consists of P_1, \dots, P_n , and this supports the conclusion Q . Typically, some of our evidence will be somewhat better justified than the rest. Suppose the maximally justified subset of our evidence consists of P_1, \dots, P_k . This may support a conflicting conclusion R . If the rest of the evidence is only slightly less justified, we will want to draw the conclusion Q rather than R . The logic of this is that we will have three arguments. The first argument infers Q from P_1, \dots, P_n . The second argument infers R from P_1, \dots, P_k . The third argument appeals to the fact that we have the evidence $\{P_1, \dots, P_n\}$ and it is a larger set of evidence than $\{P_1, \dots, P_k\}$, and it infers an undercutting defeater for the second argument. The first two arguments rebut one another, and the third argument undercuts the second argument. However, the first and third arguments are weaker than the second argument, because they depend upon all of P_1, \dots, P_n and hence depend upon a weaker link than the second argument. Without collaborative defeat, the second argument would defeat the first argument, and the first argument and third argument would each do nothing but diminish the justification of the conclusion R of the second argument. Thus we would be led draw the conclusion R on the basis of a subset of the evidence rather than drawing the conclusion Q that is supported by the total evidence. That is the intuitively wrong answer. Collaborative defeat, however, can produce the right answer. First, argument three supports the R -independent undercutting defeater, which diminishes the degree of justification of R . Then it is that weaker conclusion that squares off against the support of Q by the first argument (i.e., the R -dependent rebutting defeater), and hence Q can be justified and R unjustified.

To illustrate this general phenomenon, consider a case of direct inference. In direct inference we reason from general probabilities (symbolized using “prob”) to single case probabilities (symbolized using “PROB”). The basic idea behind direct inference was first articulated by Hans Reichenbach [34]: in determining the probability that an individual c has a property F , we find the narrowest reference class X for which we have reliable statistics and then infer that $\text{PROB}(Fc) = \text{prob}(Fx/x \in X)$. For example, insurance rates are calculated in this way. Although there is general agreement that direct inference is based upon some such principle as this, there is little agreement about the precise form the theory should take.¹⁴ In [26] I proposed reconstructing direct inference as defeasible reasoning that proceeds primarily in terms of the following two principles:

- (DI) “ $\text{prob}(F/G) = r \ \& \ \mathbf{J}(Gc) \ \& \ \mathbf{J}(P \leftrightarrow Fc)$ ” is a defeasible reason for “ $\text{PROB}(P) = r$ ”.
- (SD) “ $\text{prob}(F/H) \neq \text{prob}(F/G) \ \& \ \mathbf{J}(Hc) \ \& \ \square \forall (H \rightarrow G)$ ” is an undercutting defeater for (DI).

¹⁴ For instance, see Kyburg [14] and Levi [15].

Here “ \mathbf{JP} ” means “ P is justified”, $\forall P$ is the universal closure of P , and \Box is logical necessity.¹⁵ The defeaters provided by (SD) are *subproperty defeaters* (reminiscent of the subproperty defeaters for the statistical syllogism). To illustrate, suppose you are an insurance actuary computing auto insurance rates for Helen. You know that Helen is female, and the probability of a female driver having an accident within one year is 0.03. This gives you a defeasible reason for concluding that the probability of Helen having an accident is 0.03. But you also know that Helen is a reckless driver, and the probability of a female reckless driver having an accident in a year is 0.1. This gives you a defeasible reason for the conflicting conclusion that the probability of Helen having an accident is 0.1. Because the latter inference is based upon more information, you will accept it and reject the former inference. Formally, this is because the latter inference is based upon information that provides a subproperty defeater for the former inference, so the former inference is defeated and the latter is left undefeated. Let us examine this formal reconstruction carefully. There are three arguments involved:

- *Argument 1*—for the conclusion that the probability of Helen having an accident is 0.03, based upon the fact that Helen is female.
- *Argument 2*—for the conclusion that the probability of Helen having an accident is 0.1, based upon the fact that Helen is female and a reckless driver.
- *Argument 3*—for the conclusion that Helen is a reckless driver, where the probability of a female driver having an accident is different from the probability of a female reckless driver having an accident.

Arguments 1 and 2 provide rebutting defeaters for each other. If they were of equal strength and there were no other relevant arguments, then both would be defeated and you would be unable to draw any conclusion about the probability of Helen having an accident. However, the conclusion of argument 3 supports an undercutting defeater for argument 1. So if all the arguments are of equal strength, argument 1 will be defeated, leaving argument 2 undefeated. The inference-graph representing all three arguments can be diagrammed as in Fig. 5.

The preceding analysis of the problem assumes that we are equally justified in believing that Helen is female and that Helen is a reckless driver, but typically we will be better justified in believing the former (although adequately justified in believing the latter as well). Arguments 2 and 3 both depend upon knowing that Helen is a reckless driver. Argument 1 depends only upon knowing that Helen is female, which is more strongly justified. This implies that the argument strengths for arguments 2 and 3 are both less than the argument strength for argument 1. The result is that neither the rebutting defeater provided by argument 2 nor the undercutting defeater provided by argument 3 should be sufficient by itself to defeat argument 1. On the other hand, argument 2 would be defeated outright by the more strongly supported argument 1. This is the intuitively wrong result.

It appears that the only way to get argument 1 defeated and argument 2 undefeated is to allow the rebutting defeater provided by argument 2 and the undercutting defeater provided by argument 3 to work in unison rather than separately. That is precisely what (DJ) accomplishes. Because the undercutting defeater is undefeated, it lowers the degree

¹⁵ These two principles do not by themselves provide a complete account, but see [26] for more details.

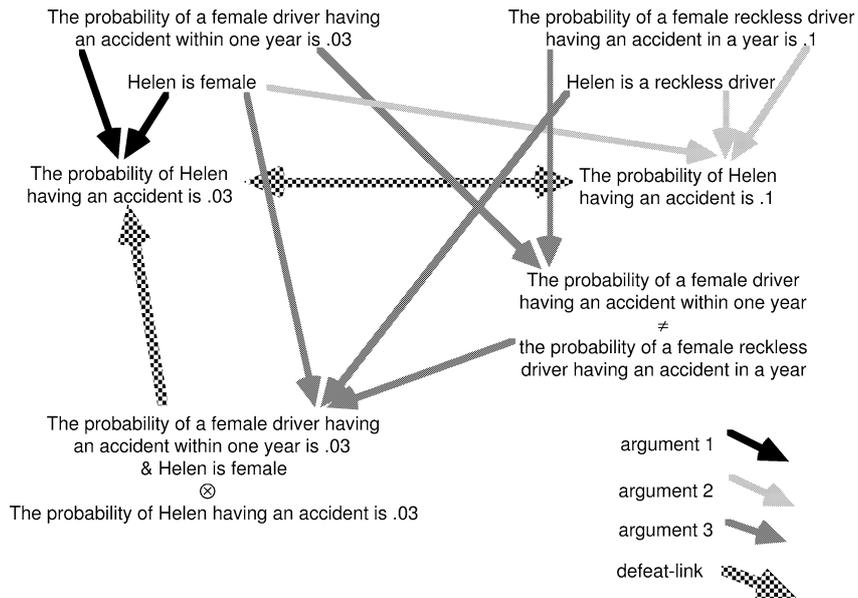


Fig. 5. Direct inference.

of support for the conclusion of argument 1. It is that lowered support that should then be compared with the support for the rebutting defeater provided by argument 2, and as the latter is greater, argument 1 is defeated. So the intuitively correct answer is reached by collaborative defeat.

There are numerous other cases in which the reasoning has an analogous structure and variations in the degrees of justification of the premises cause similar prima facie difficulties. In each case, collaborative defeat resolves the prima facie difficulties. For example, the Yale Shooting Problem has played an important role in discussions of the Frame Problem.¹⁶ In the Yale Shooting Problem, we are given that a gun is initially loaded. Then it is pointed at Jones and the trigger is pulled. We suppose we know (simplistically) that if a loaded gun is pointed at someone and the trigger pulled, that person will shortly become dead. The conclusion we are supposed to draw in this case is that Jones will die. The Yale Shooting Problem is the problem of showing how this conclusion can be justified. There are two problems. First, our reason for thinking that the gun is loaded when the trigger is pulled is that it was loaded moments earlier when we checked it. So we are employing temporal projection. Temporal projection provides a defeasible reason for thinking that if something is true at one time then it will still be true later. Some form of temporal projection has been endorsed by most recent authors discussing the frame problem.¹⁷ Using temporal projection, we can construct:

¹⁶ Hanks and McDermott [11].

¹⁷ This was proposed by Sandewall [36] and subsequently endorsed by McDermott [18], McCarthy [17] and virtually all subsequent authors.

Argument 1—for the conclusion that Jones will be dead after the shooting, based upon causal knowledge and the temporal projection that the gun will still be loaded when the trigger is pulled.

However, as Hanks and McDermott [11,12] were the first to note, with the help of temporal projection we can also construct a second argument supporting a conflicting conclusion:

Argument 2—for the conclusion that Jones will still be alive after the shooting, based upon temporal projection from the fact that Jones was initially alive.

In the absence of further arguments, these arguments will defeat each other, leaving us with no justified conclusion to draw about the state of Jones' health. To get the intuitively correct answer, we need a third argument:

Argument 3—supporting an undercutting defeater for argument 1.

The Yale Shooting Problem is resolved by explaining the details of argument 3. I have proposed such a solution in [29]. For present purposes, the details are not important. Suffice it to say that the undercutter gives inferences based upon causal knowledge priority over those based upon temporal projection and turns on the premise that the gun is still loaded when the trigger is pulled.

A problem derives from the fact that the strength of an inference by temporal projection decreases as the time interval increases. That is, temporal projection gives us a reason for thinking that things won't change, but the greater the time interval the weaker the reason. This is the *temporal decay* of temporal projection discussed at the beginning of section three (see [29]). If we ignore the temporal decay of temporal projection, the foregoing constitutes a solution to the Yale Shooting Problem. But now suppose, as will typically be the case, that we observe Jones to be alive at a later time (normally, right up to the time we pull the trigger) than we observe the gun to be loaded. I cannot observe the latter at the time I pull the trigger without shooting myself in the face. In this case the strengths of arguments 1 and 3, depending as they do on inferring that the gun is still loaded when the trigger is pulled, may both be less than the strength of argument 2.

The temporal profile we should get is the following. If we observe the gun to be loaded long before observing Jones to be alive (e.g., years ago), our justification for believing Jones to remain alive might be somewhat weakened but not defeated. On the other hand, if we observe the gun to be loaded just shortly before observing Jones to be alive, we should be able to conclude that Jones will die. The intuitive rationale for this profile seems to be as follows. First, we have the undefeated argument 3 for the undercutting defeater for argument 2, but it is weaker than argument 2. Instead of defeating argument 2 outright, it weakens it seriously. This leaves us with only a weak reason for thinking that Jones remains alive. Then argument 1 provides a reason for thinking Jones is dead. If argument 1 turns upon a temporal projection from the distant past, it will not be strong enough to defeat even the weakened argument 2, but if the temporal projection is from a recent observation then argument 1 will be strong enough to defeat the weakened argument 2. This is exactly the computation that results from collaborative defeat.

The upshot is that what seemed initially like a suspicious consequence of (DJ1) and (DJ2) turns out to be a very important logical phenomenon that is often crucial for computing defeat statuses correctly.

9. Measuring strengths

It has been assumed throughout that reason-strengths and degrees of justification can be measured using the extended reals, although nothing has been said about how that is done. If we are to take strength seriously, we must have some way of measuring it. One way is to compare reasons with a set of standard equally good reasons that have numerical values associated with them in some determinant way. I propose to do that by taking the set of standard reasons to consist of instances of the statistical syllogism (SS). For any proposition p , we can construct a standardized argument for $\sim p$ on the basis of the pair of suppositions “ $\text{prob}(F/G) \geq r \ \& \ Gc$ ” and “ $(p \leftrightarrow \sim Fc)$ ”:

1. Suppose $\text{prob}(F/G) \geq r \ \& \ Gc$;
2. Suppose $(p \leftrightarrow \sim Fc)$;
3. Fc from 1;
4. $\sim p$ from 2,3,

where the strength of the argument is a function of r . We can measure the strength of a defeasible reason for p in terms of that value of r such that the conflicting argument from the suppositions “ $\text{prob}(F/G) \geq r \ \& \ Gc$ ” and “ $(p \leftrightarrow \sim Fc)$ ” exactly counteracts it. The value r determines the reason-strength in the sense that the reason-strength is some function $dj(r)$ of r . It is tempting to identify $dj(r)$ with r , but that will not work. The difficulty is that by identifying \mathcal{J} with \sim , we have required reason-strength to be a cardinal measure that can be meaningfully added and subtracted. Adding and subtracting reason-strengths may not be the same thing as adding and subtracting the corresponding probabilities. To illustrate the general point, should it turn out (it won't) that the reason-strength corresponding to a probability r is given by $\log(r)$, then adding reason-strengths would be equivalent to multiplying probabilities rather than adding them.

I do not have an a priori argument to offer regarding what function $dj(r)$ produces the reason-strength corresponding to r . The only way to determine this is to look for proposals that work plausibly in concrete examples. Perhaps the most illuminating example is that of the biased lotteries diagrammed in Figs. 2 and 3. Suppose reason-strengths could be identified with the corresponding probabilities. In lottery 2, the probability of ticket 1 being drawn is 0.000001, and the probability of any other ticket being drawn is 0.111111. We wanted to conclude in this case that we are justified in believing that ticket 1 will not be drawn. The probability corresponding to the argument-strength for this conclusion is 0.999999, however the probability corresponding to the counter-argument for the conclusion that ticket 1 will be drawn (because no other ticket will) is 0.888889. The difference between these probabilities is 0.111111, which is a very low probability. If probabilities and degrees of justification could be identified, i.e., $dj(r) = r$, this would produce too low a degree of justification for it to be reasonable to believe that ticket 1 will not be drawn. So apparently we cannot compute degrees of justification by adding and subtracting probabilities.

There is statistical lore suggesting that in probabilistic reasoning degrees of justification can be compared in terms of likelihood ratios.¹⁸ When (as in the biased lotteries) we have

¹⁸ This is known as the likelihood principle. It is due to R.A. Fisher [9], and versions of it have been endorsed by a variety of authors, including G.A. Barnard [1,2], Alan Birnbaum [3], A.W.F. Edwards [8] and Ian Hacking [10].

an argument for P based on a probability r , and a weaker argument for $\sim P$ based on a probability r^* , the likelihood ratio is $(1 - r)/(1 - r^*)$. The suggestion is that the degree of justification for $\sim P$ is determined by the likelihood ratio. For example, in lottery 2 the likelihood ratio is 0.000009, while in lottery 3 it is 0.09. Note that likelihood ratios are defined so that higher likelihood ratios correspond to lower degrees of justification. An equivalent but more intuitive way of measuring degrees of justification is by using the inverse of the likelihood ratios.

In [26] I argued that likelihood ratios seem to yield the intuitively correct answers in many cases of statistical and inductive reasoning, and on that basis I am prepared to tentatively endorse their use in measuring degrees of justification. If we take the degree of justification $dj(r)$ resulting from an application of the statistical syllogism with probability r to be $\log(0.5) - \log(1 - r)$, then the result of subtracting degrees of justification is the same as taking the logarithm of the inverse of the likelihood ratio, i.e., $dj(r) - dj(r^*) = \log(1 - r^*) - \log(1 - r) = \log((1 - r^*)/(1 - r))$. Thus this choice of $dj(r)$ yields congenial results. This will be my proposal:

If X is a defeasible reason for p , the strength of this reason is $\log(0.5) - \log(1 - r)$ where r is that real number such that an argument for $\sim p$ based upon the suppositions “ $\text{prob}(F/G) \geq r \ \& \ Gc$ ” and “ $(p \leftrightarrow \sim Fc)$ ” and employing the statistical syllogism exactly counteracts the argument for p based upon the supposition X .

The reason for the “ $\log(0.5)$ ” term is that this produces the result that $dj(0.5) = 0$. An application of the statistical syllogism based on a probability of 0.5 should produce no justification for the conclusion, and this is captured by setting the degree of justification to be 0. On the other hand, note that $dj(1) = \infty$. That is, the strongest reasons have infinite reason-strength. This could create problems if we ever wanted to subtract the strengths of such arguments from each other, because $\infty - \infty$ is undefined, but in fact we will never have occasion to do that.

10. Simplifying the computation

Principles (DJ1) and (DJ2) provide a recursive characterization of degrees of justification relative to an inference-graph. However, this characterization does not lend itself well to implementation because it requires the construction of modified inference-graphs. The objective of this section is to produce an equivalent recursive characterization that appeals only to the given inference-graph.

We defined a defeat-link to be φ -critical in an inference-graph \mathcal{G} iff it is a member of some minimal set of defeat-links such that the removal of all the links in the set results in φ no longer having a φ -dependent defeater in \mathcal{G} . Recall that a defeat-link or support-link extends from its *root* to its *target*. The root of a defeat-link is a single node, and the root of a support-link is a set of nodes. A node θ is an *inference/defeat-descendant* of a node ψ iff θ can be reached by traveling along support-links and defeat-links from ψ . Then a necessary condition for a defeat-link L to be φ -critical is that L 's root is equal to or is an inference/defeat-descendant of the node φ and φ is equal to or is an inference/defeat-

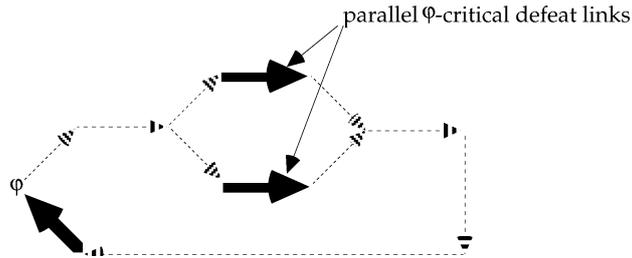


Fig. 6. Parallel φ -critical defeat-links.

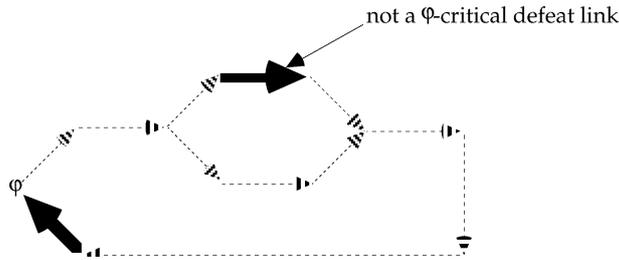


Fig. 7. Defeat link that is not φ -critical.

descendant of L 's target. But that is not a sufficient condition. For example, in inference-graph (7) the defeat-links between $T1$ and $T2$ are not R -critical, but they satisfy this condition. The reason these defeat-links fail to be R -critical is that there is another route from R to $\sim R$ that does not pass along the defeat-links. This route consists exclusively of support-links.

An *inference-defeat-path* from a node φ to a node θ is a sequence of support-links and defeat-links such that (1) if the first link is a defeat-link, its root is φ ; and if it is a support-link then φ is a member of its root; (2) the target of the last link in the path is θ ; (3) the last link in the path is a defeat-link; (4) the root of each defeat-link after the first member of the path is the target of the preceding link; (5) some member of the basis of each support-link after the first member of the path is the target of the preceding link; and (6) the path does not contain an internal loop, i.e., no two links in the path have the same root. A φ -critical defeat-link must lie on an inference-defeat-path from φ to φ . There can be diverging and reconverging paths with several “parallel” defeat-links, as in Fig. 6. The set of parallel defeat-links may then be a minimal set of defeat-links such that the removal of all the links in the set results in φ no longer having a φ -dependent defeater in \mathcal{G} , in which case the defeat-links are all φ -critical. The only way in which a defeat-link on an inference-defeat-path can fail to be φ -critical is when there is a path around it consisting entirely of support-links, as diagrammed in Fig. 7. This is what happens in inference-graph (7), and it is crucial to the computation of degrees of justification that such defeat-links not be regarded as φ -critical.

Let us say that a node α *precedes* a node β on an inference-defeat-path iff α and β both lie on the path and either $\alpha = \beta$ or the path contains a subpath originating on α

and terminating on β . *Inference-descendants* of a node are nodes that can be reached by following support-links. We can characterize the φ -critical defeat-links as follows:

A defeat-link L is φ -critical in \mathcal{G} iff (1) L lies on an inference-defeat-path σ in \mathcal{G} from φ to φ , and (2) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in \mathcal{G} .

This characterization makes it straightforward to search for φ -critical defeat-links in an inference-graph.

Recall that (DJ) was formulated as follows:

(DJ) If P is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the P -independent defeaters for P , D_{k+1}, \dots, D_m are the P -dependent defeaters of P , and \mathcal{G}_P results from deleting all P -critical defeat-links from G and making all P -independent nodes φ initial with $j(\varphi, \mathcal{G}_P) = j(\varphi, \mathcal{G})$, then

$$\begin{aligned} j(P, \mathcal{G}) &= \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\ &\sim [\max\{j(D_1, \mathcal{G}), \dots, j(D_k, \mathcal{G})\} \\ &\quad + \max\{j(D_{k+1}, \mathcal{G}_P), \dots, j(D_m, \mathcal{G}_P)\}]. \end{aligned}$$

In most of the cases we have considered, \mathcal{G}_φ is an inference-graph in which no node θ has a θ -dependent defeater. Thus $j(\theta, \mathcal{G}_\varphi)$ can be computed using just (DJ1) in \mathcal{G}_φ . This is equivalent to applying (DJ1) in \mathcal{G} but ignoring φ -critical defeat-links:

If ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in the inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the defeaters for ψ in \mathcal{G} that are not φ -critical, and no D_i is ψ -dependent, then

$$\begin{aligned} j(\psi, \mathcal{G}_\varphi) &= \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\ &\sim \max\{j(D_1, \mathcal{G}), \dots, j(D_k, \mathcal{G})\}. \end{aligned}$$

So for this computation, it is unnecessary to modify the inference-graph.

In those cases in which there is a node θ having a θ -dependent defeater in \mathcal{G}_φ , we must instead apply (DJ2):

If ψ has ψ -dependent defeaters D_1, \dots, D_k in \mathcal{G}_φ and $\mathcal{G}_{\psi, \varphi}$ results from deleting all ψ -critical defeat-links from \mathcal{G}_φ and making all ψ -independent nodes θ initial with $j(\theta, \mathcal{G}_{\psi, \varphi}) = j(\theta, \mathcal{G}_\varphi)$, then

$$j(\psi, \mathcal{G}_\varphi) = j(\psi, \mathcal{G}_{\psi, \varphi}) \sim \max\{j(D_1, \mathcal{G}_{\psi, \varphi}), \dots, j(D_k, \mathcal{G}_{\psi, \varphi})\}.$$

Applying this computation recursively leads to a sequence of inference-graphs of the form $\mathcal{G}_{\varphi_1, \dots, \varphi_n}$. The inference-graph \mathcal{G} is finite because it represents the actual state of an agent's reasoning. It follows that the computation eventually terminates because the inference-graph \mathcal{G} has only finitely many defeat-links, and each subsequent $\mathcal{G}_{\varphi_1, \dots, \varphi_n}$ has fewer defeat-links than its predecessor. The sequence of inference-graphs constructed in this way can be characterized as follows. Where $\varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} , define recursively:

\mathcal{G}_φ results from deleting all φ -critical defeat-links from \mathcal{G} and making all nodes θ that are φ -independent in \mathcal{G} initial with $j(\theta, \mathcal{G}_\varphi) = j(\theta, \mathcal{G})$;

$\mathcal{G}_{\varphi_1, \dots, \varphi_n}$ results from deleting all φ_1 -critical defeat-links from $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ and making all nodes θ that are φ_1 -independent in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ initial with $j(\theta, \mathcal{G}_{\varphi_1, \dots, \varphi_n}) = j(\theta, \mathcal{G}_{\varphi_2, \dots, \varphi_n})$.

To reformulate the recursion so as to avoid constructing modified inference-graphs, we define some new concepts. Where $\varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} , define recursively:

A defeat-link L of \mathcal{G} is $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical in \mathcal{G} iff either (1) L is $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical in \mathcal{G} or (2a) L lies on an inference-defeat-path σ in \mathcal{G} from φ to φ containing no $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical defeat-links and (2b) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in \mathcal{G} .

Where $\alpha, \varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} , define:

α is φ -dependent in \mathcal{G} iff there is an inference-defeat-path σ in \mathcal{G} from φ to α .

α is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} iff there is an inference-defeat-path σ in \mathcal{G} from φ_1 to α containing no $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical defeat-links.

Theorem 3. L is φ_1 -critical in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ iff L is $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical in \mathcal{G} but not $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical in \mathcal{G} .

Proof (By induction on n). For $n = 1$, this is true by definition. Suppose it is true for $n - 1$. Then:

L is $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical in \mathcal{G} but not $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical in \mathcal{G}

iff (1) L lies on an inference-defeat-path σ in \mathcal{G} from φ_1 to φ_1 containing no $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical defeat-links and (2) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in \mathcal{G}

iff (1) L lies on an inference-defeat-path σ in \mathcal{G} from φ_1 to φ_1 containing no $\langle \varphi_i, \dots, \varphi_n \rangle$ -critical defeat-links for $i \leq 2$ and (2) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in \mathcal{G}

iff (1) L lies on an inference-defeat-path σ in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ from φ_1 to φ_1 (2) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in \mathcal{G}

iff (1) L lies on an inference-defeat-path σ in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ from φ_1 to φ_1 , and (2) there is no node α preceding the root of L on σ and node β preceded by the target of L on σ such that β is an inference-descendant of α in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$

iff L is φ_1 -critical in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$. \square

Theorem 4. α is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} iff α is φ_1 -dependent in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$.

Proof (By induction on n). For $n = 1$, this is true by definition. Suppose it is true for $n - 1$. Then:

α is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G}

iff there is an inference-defeat-path σ in \mathcal{G} from φ_1 to a containing no $\langle \varphi_2, \dots, \varphi_n \rangle$ -critical defeat-links

iff there is an inference-defeat-path σ in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$ from φ_1 to α

iff φ_1 -dependent in $\mathcal{G}_{\varphi_2, \dots, \varphi_n}$. \square

The next three theorems follow immediately from (DJ1), (DJ2), and Theorems 3 and 4.

Theorem 5. *If ψ is initial in \mathcal{G} then $j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n}) = j(\psi, \mathcal{G})$.*

Theorem 6. *If ψ is $\langle \varphi_1, \dots, \varphi_n \rangle$ -independent in \mathcal{G} then $j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n}) = j(\psi, \mathcal{G}_{\varphi_2, \dots, \varphi_n})$.*

Theorem 7. *If $\psi \neq \varphi_1$ and ψ is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} and ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the $\langle \psi, \varphi_1, \dots, \varphi_n \rangle$ -independent defeaters for ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, and D_{k+1}, \dots, D_m are the $\langle \psi, \varphi_1, \dots, \varphi_n \rangle$ -dependent defeaters of ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, then*

$$\begin{aligned} j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n}) &= \min\{\rho, j(B_1, \mathcal{G}_{\varphi_1, \dots, \varphi_n}), \dots, j(B_n, \mathcal{G}_{\varphi_1, \dots, \varphi_n})\} \\ &\sim [\max\{j(D_1, \mathcal{G}_{\varphi_1, \dots, \varphi_n}), \dots, j(D_k, \mathcal{G}_{\varphi_1, \dots, \varphi_n})\} \\ &\quad + \max\{j(D_{k+1}, \mathcal{G}_{\psi, \varphi_1, \dots, \varphi_n}), \dots, j(D_m, \mathcal{G}_{\psi, \varphi_1, \dots, \varphi_n})\}]. \end{aligned}$$

Theorem 8. *If $\psi = \varphi_1$ and ψ is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} and ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the $\langle \varphi_1, \dots, \varphi_n \rangle$ -independent defeaters for ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, then*

$$\begin{aligned} j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n}) &= \min\{\rho, j(B_1, \mathcal{G}_{\varphi_1, \dots, \varphi_n}), \dots, j(B_n, \mathcal{G}_{\varphi_1, \dots, \varphi_n})\} \\ &\sim \max\{j(D_1, \mathcal{G}_{\varphi_1, \dots, \varphi_n}), \dots, j(D_k, \mathcal{G}_{\varphi_1, \dots, \varphi_n})\}. \end{aligned}$$

In computing $j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n})$, we can mirror the application of (DJ1) and (DJ2) by applications of Theorems 5–8. However, in order to apply Theorems 5–8, we never have to construct the modified inference-graphs $\mathcal{G}_{\varphi_1, \dots, \varphi_n}$. This becomes obvious if we rewrite $j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n})$ as $j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G})$. In light of Theorems 5–8, this function can be defined recursively as follows, where $\psi, \varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} and $j_0(\theta, \mathcal{G})$ is the assignment of values to initial nodes in \mathcal{G} :

Theorem 9.

- If ψ is initial in \mathcal{G} then $j(\psi, \mathcal{G}) = j_0(\psi, \mathcal{G})$.*
- If ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the ψ -independent defeaters for ψ , and D_{k+1}, \dots, D_m are the ψ -dependent defeaters of ψ , then*

$$\begin{aligned} j(\varphi, \mathcal{G}) &= \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\ &\sim [\max\{j(D_1, \mathcal{G}), \dots, j(D_k, \mathcal{G})\} \\ &\quad + \max\{j_{\varphi}(D_{k+1}, \mathcal{G}), \dots, j_{\psi}(D_m, \mathcal{G})\}]. \end{aligned}$$

- (c) If ψ is $\langle \varphi_1, \dots, \varphi_n \rangle$ -independent in \mathcal{G} then $j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G}) = j_{\varphi_2, \dots, \varphi_n}(\psi, \mathcal{G})$.
 (d) If $\psi \neq \varphi_1$ and is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} and ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the $\langle \psi, \varphi_1, \dots, \varphi_n \rangle$ -independent defeaters for ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, and D_{k+1}, \dots, D_m are the $\langle \psi, \varphi_1, \dots, \varphi_n \rangle$ -dependent defeaters of ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, then

$$j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G}) = \min\{\rho, j_{\varphi_1, \dots, \varphi_n}(B_1, \mathcal{G}), \dots, j_{\varphi_1, \dots, \varphi_n}(B_n, \mathcal{G})\} \\ \sim [\max\{j_{\varphi_1, \dots, \varphi_n}(D_1, \mathcal{G}), \dots, j_{\varphi_1, \dots, \varphi_n}(D_k, \mathcal{G})\} \\ + \max\{j_{\psi, \varphi_1, \dots, \varphi_n}(D_{k+1}, \mathcal{G}), \dots, j_{\psi, \varphi_1, \dots, \varphi_n}(D_m, \mathcal{G})\}].$$

- (e) If $\psi = \varphi_1$ and ψ is $\langle \varphi_1, \dots, \varphi_n \rangle$ -dependent in \mathcal{G} and ψ is inferred from the basis $\{B_1, \dots, B_n\}$ in an inference-graph \mathcal{G} in accordance with a reason-scheme of strength ρ , D_1, \dots, D_k are the $\langle \varphi_1, \dots, \varphi_n \rangle$ -independent defeaters for ψ that are not $\langle \varphi_1, \dots, \varphi_n \rangle$ -critical, then

$$j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G}) = \min\{\rho, j_{\varphi_1, \dots, \varphi_n}(B_1, \mathcal{G}), \dots, j_{\varphi_1, \dots, \varphi_n}(B_n, \mathcal{G})\} \\ \sim \max\{j_{\varphi_1, \dots, \varphi_n}(D_1, \mathcal{G}), \dots, j_{\varphi_1, \dots, \varphi_n}(D_k, \mathcal{G})\}.$$

Alternatively, if we reconstrue Theorem 9 as the definition of $j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G})$ then we have:

Theorem 10. Where $\psi, \varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} , $j_{\varphi_1, \dots, \varphi_n}(\psi, \mathcal{G}) = j(\psi, \mathcal{G}_{\varphi_1, \dots, \varphi_n})$.

Theorem 10 constitutes a recursive characterization of $j(\varphi, \mathcal{G})$. To illustrate the computation, consider inference-graph (13). Here we have two separate arguments for C , from which $A \otimes B$ is inferred with reason-strength ρ . Nodes supporting the same conclusions are distinguished by subscripting the conclusions. To compute $j(B, \mathcal{G}_{13})$ using (DJ1) and (DJ2), we remove the B -critical link to get (13*).

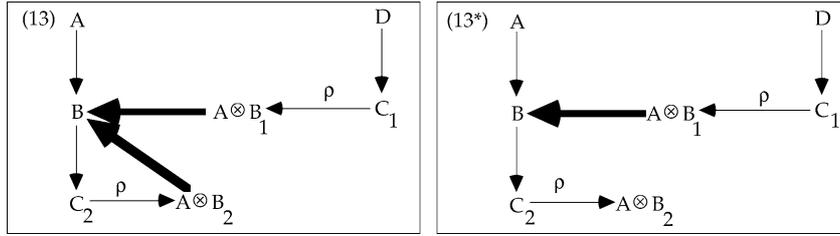
$$j(B, \mathcal{G}_{13}) = j_0(A, \mathcal{G}_{13}) \sim [j(A \otimes B_1, \mathcal{G}_{13}) + j(A \otimes B_2, \mathcal{G}_{13}^*)], \\ j(A \otimes B_1, \mathcal{G}_{13}) = \min\{\rho, j(C_1, \mathcal{G}_{13})\} = \min\{\rho, j_0(D, \mathcal{G}_{13})\}, \\ j(A \otimes B_2, \mathcal{G}_{13}^*) = \min\{\rho, j(C_2, \mathcal{G}_{13}^*)\} = \min\{\rho, j(B, \mathcal{G}_{13}^*)\} \\ = \min\{\rho, (j_0(A, \mathcal{G}_{13}) \sim \min\{\rho, j_0(D, \mathcal{G}_{13})\})\}.$$

Thus

$$j(B, \mathcal{G}_{13}) = j_0(A, \mathcal{G}_{13}) \sim [\min\{\rho, j_0(D, \mathcal{G}_{13})\} \\ + \min\{\rho, (j_0(A, \mathcal{G}_{13}) \sim \min\{\rho, j_0(D, \mathcal{G}_{13})\})\}].$$

For example, if $j_0(A, \mathcal{G}_{13}) = j_0(D, \mathcal{G}_{13})$, and $\rho < 0.5 \cdot j_0(A, \mathcal{G}_{13})$. Then

$$j(B) = j_0(A, \mathcal{G}_{13}) \sim 2\rho.$$



To repeat the computation using Theorem 9, by (b)

$$j(B, \mathcal{G}_{13}) = j_0(A, \mathcal{G}_{13}) \sim [j(A \otimes B_1, \mathcal{G}_{13}) + j_B(A \otimes B_2, \mathcal{G}_{13})],$$

$$j(A \otimes B_1, \mathcal{G}_{13}) = \min\{\rho, j(C_1, \mathcal{G}_{13})\} = \min\{\rho, j_0(D, \mathcal{G}_{13})\}.$$

By (d),

$$j_B(A \otimes B_2, \mathcal{G}_{13}) = \min\{\rho, j_B(C_2, \mathcal{G}_{13})\} = \min\{\rho, j_B(B, \mathcal{G}_{13})\}.$$

$A \otimes B_1$ is the only B -independent defeater of B , and it is not B -critical. Thus by (e),

$$\begin{aligned} j_B(B, \mathcal{G}_{13}) &= j_B(A, \mathcal{G}_{13}) \sim j_B(A \otimes B_1, \mathcal{G}_{13}) \\ &= j_0(A, \mathcal{G}_{13}) \sim j(A \otimes B_1, \mathcal{G}_{13}) \\ &= j_0(A, \mathcal{G}_{13}) \sim \min\{\rho, j_0(D, \mathcal{G}_{13})\}. \end{aligned}$$

Therefore,

$$j_B(A \otimes B_2, \mathcal{G}_{13}) = \min\{\rho, j_0(A, \mathcal{G}_{13}) \sim \min\{\rho, j_0(D, \mathcal{G}_{13})\}\}.$$

Hence once again,

$$\begin{aligned} j(B, \mathcal{G}_{13}) &= j_0(A, \mathcal{G}_{13}) \sim [\min\{\rho, j_0(D, \mathcal{G}_{13})\} \\ &\quad + \min\{\rho, j_0(A, \mathcal{G}_{13}) \sim \min\{\rho, j_0(D, \mathcal{G}_{13})\}\}]. \end{aligned}$$

Theorem 10 constitutes the desired recursive characterization of $j(\varphi, \mathcal{G})$. This could be implemented straightforwardly. However, as the next section shows, further improvements are possible.

11. And/or inference-graphs

11.1. Two kinds of inference-graphs

In the interest of theoretical clarity, inference-graphs were defined in such a way that different arguments for the same conclusion are represented by different nodes. This made it clearer how the algorithm for computing defeat status works. However, for the purpose of implementing defeasible reasoning, using different nodes to represent different arguments for the same conclusion is an inefficient representation, because it leads to

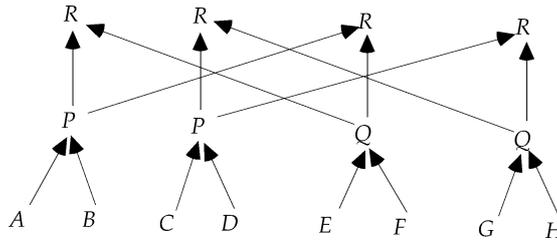


Fig. 8. Inference-graph with multiple arguments for a single conclusion.

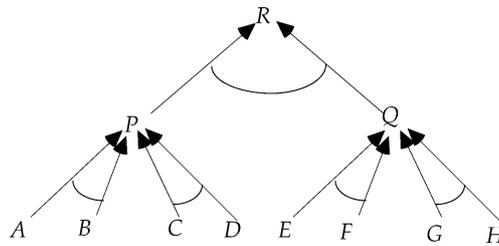


Fig. 9. An and/or inference-graph.

needless duplication. If we have two arguments supporting a single conclusion, then any further reasoning from that conclusion will generate two different nodes. If we have two arguments for each of two conclusions, and another inference proceeds from those two conclusions, the latter will have to be represented by four different nodes in the inference-graph, and so on. This is illustrated in Fig. 8, where P and Q are each inferred in two separate ways, and then R is inferred from P and Q .

A more efficient representation of reasoning would take the inference-graph to be an and/or graph rather than a standard graph.¹⁹ In an and/or graph, nodes are linked to *sets* of nodes rather than individual nodes. This is represented diagrammatically by connecting the links with arcs. In an and/or inference-graph, when we have multiple arguments for a conclusion, the single node representing that conclusion will be tied to different bases by separate groups of links. (Note that this has nothing to do with accommodating conjunction and disjunction in our language—there are no syntactical restrictions on the formulas expressing the conclusions encoded in either simple inference-graphs or and/or graphs.) This is illustrated in Fig. 9 by an and/or inference-graph encoding the same reasoning as the standard inference-graph in Fig. 8. In an and/or inference-graph, a support-link will be represented by a *set* of supporting-arrows connected with an arc.

Although and/or graphs provide an efficient representation of reasoning, they complicate the computation of degrees of justification. Using simple inference-graphs, we can use

¹⁹ This observation derives from [27]. This section adapts the approach taken there to the new definition of “status assignment”.

principles (DJ1) and (DJ2) to compute degrees of justification. If we are to use and/or graphs in the implementation, we must find a computation for and/or graphs that is equivalent to that for simple inference-graphs. A simple inference-graph can be rewritten as an and/or graph in which each node of the and/or graph corresponds to a set of nodes of the simple graph. A node of the simple inference-graph corresponds to an *argument* in the and/or graph. An argument is a kind of connected sub-tree of the graph. More precisely:

An *argument* in an and/or inference-graph \mathcal{G} for a node φ is a minimal subset \mathbb{A} of the nodes and support-links of the graph such that (1) if a node in \mathbb{A} has any support-links in \mathcal{G} , exactly one of them is in \mathbb{A} , (2) if a support-link is in \mathbb{A} then the nodes in its support-link-basis are also in \mathbb{A} , and (3) φ is in \mathbb{A} .

Nodes in the simple inference-graph correspond one-one to arguments in the and/or inference-graph.

11.2. Computing degrees of justification

It was argued in section five that the degree of justification of a conclusion should be the maximum of the degrees of justification supplied by the different arguments supporting that conclusion. Thus the result we want for and/or inference-graphs is the following *Correspondence Principle*:

The degree of justification of a node of the and/or inference-graph is equal to the maximum of the degrees of justification of the corresponding nodes of the simple inference-graph.

To accomplish this, it helps to assign degrees of justification to support-links as well as nodes. Some nodes will be initial, in which case they are simply assigned a degree of justification $j_0(P, \mathcal{G})$ in an inference-graph \mathcal{G} . I assume that such nodes have no support-links. If a node is not initial, its degree of justification is the maximum of the degrees of justification of its support-links. If a node is not initial and it has no support-links, we stipulate that its degree of justification is zero. This cannot happen naturally, but certain constructions that occur in the computation of degrees of justification can produce inference-graphs having such nodes. Having characterized the degree of justification of a node of the inference-graph in terms of the degrees of justification of its support-links, the remaining problem is how to compute the degrees of justification of support-links.

In an and/or inference-graph, undercutting defeat-links attach to support-links rather than nodes. A rebutting defeat-link could be regarded as attaching to either a node or to all defeasible support-links for the node. It will be more convenient to adopt the latter convention so that both undercutting and rebutting defeaters can be treated uniformly. Where λ is a support-link, let $\otimes\lambda$ be its undercutting defeater. That is, if the root of λ is $\{B_1, \dots, B_n\}$ and the target is φ , $\otimes\lambda$ is $[(B_1 \& \dots \& B_n) \otimes \varphi]$.

In and/or inference-graphs, we can define inference-defeat-paths pretty much as we did for simple inference-graphs. That is:

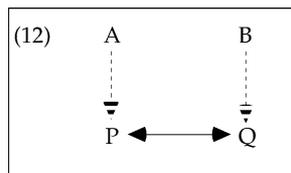
An *inference-defeat-path* from a node φ to a node θ is a sequence of support-links and defeat-links such that (1) if the first link is a defeat-link, its root is φ ; and if it is a support-link then φ is a member of its root; (2) the target of the last link in the path is a support-link for θ ; (3) the last link in the path is a defeat-link; (4) the root of each defeat-link after the first member of the path is the target of the preceding link; (5) some member of the basis of each support-link after the first member of the path is either the target of the preceding link or has a support-link that is the target of the preceding link; and (6) the path does not contain an internal loop, i.e., no two links in the path have the same root.

Let us say that a simple graph *corresponds* to an and/or graph iff the corresponding nodes support the same conclusions and the support-links and defeat-links are analogous. Representing links by ordered pairs, we can define precisely:

A simple inference-graph \mathcal{G} *corresponds* to an and/or inference-graph $\&\mathcal{G}$ iff there is a function μ mapping the nodes of \mathcal{G} onto the nodes of $\&\mathcal{G}$ such that:

- (a) if α is a node of \mathcal{G} and β is a node of $\&\mathcal{G}$, $\mu(\alpha) = \mu(\beta)$ iff α and β support the same proposition;
- (b) if α is a node of \mathcal{G} and $\Gamma = \{\beta \mid \langle \beta, \alpha \rangle \text{ is a support-link for } \alpha \text{ in } \mathcal{G}\}$ then $\langle \{\mu(\gamma) \mid \gamma \in \Gamma\}, \mu(\alpha) \rangle$ is a support-link in $\&\mathcal{G}$;
- (c) if $\langle \Gamma, \alpha \rangle$ is a support-link in $\&\mathcal{G}$ then for all nodes β, γ in \mathcal{G} , if $\mu(\beta) \in \Gamma$ and $\mu(\gamma) = \alpha$ then $\langle \beta, \gamma \rangle$ is a support-link in \mathcal{G} ;
- (d) $\langle \beta, \alpha \rangle$ is a defeat-link in \mathcal{G} iff, if $\langle \gamma_1, \alpha \rangle, \dots, \langle \gamma_n, \alpha \rangle$ are the support-links for α in \mathcal{G} then $\langle \mu(\beta), \langle \{\mu(\gamma_1), \dots, \mu(\gamma_n)\}, \mu(\alpha) \rangle \rangle$ is a defeat-link in $\&\mathcal{G}$.

Then given an inference-defeat-path in the simple graph, the corresponding sequence of links in the and/or graph is a corresponding inference-defeat-path in the and/or graph. Furthermore, given an inference-defeat-path in the and/or graph, every corresponding path in the simple graph is an inference-defeat-path.



Inference-paths must be defined somewhat differently in and/or inference-graphs than in simple inference-graphs. The difficulty is that if we simply define inference-paths to be sets of linked support-links, they can be circular. E.g., suppose P and Q are logically equivalent. Then if we have independent reasons for P and for Q , we can derive each from the other. This makes perfectly good sense. If the independent argument for P is subsequently defeated, the argument that derives P from Q will still support P as long as the argument supporting Q is not defeated, and vice versa. Thus we can have

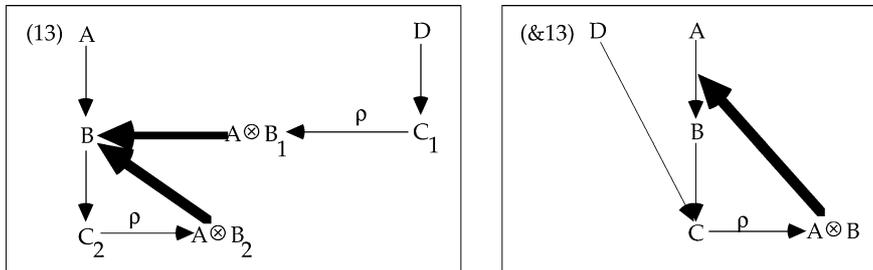
an inference-graph like (12). In this inference-graph, we do not want to count paths like $A \rightarrow P \rightarrow Q \rightarrow P$ as inference-paths. We can handle this by defining:

An *inference-path* from a node φ to a node θ is a sequence of support-links such that (1) φ is a member of the root of the first link, (2) the target of the last is θ , (3) some member of each support-link after the first is the target of the preceding link, and (4) the path is non-circular, i.e., $\varphi \neq \theta$ and no two links in the path have the same root.

Recall that the target of a defeat-link is a support-link, so the target of the target of a defeat-link is a node. Then we can define, for and/or graphs:

A defeat-link L is φ -critical in \mathcal{G} iff (1) L lies on an inference-defeat-path σ in \mathcal{G} from φ to φ , and (2) there is no node α preceding the root of L on σ and node β preceded by the target of the target of L on σ such that there is an inference-path σ^* from α to β such that the path resulting from splicing σ^* into σ in place of L (attaching it to the target of the target of L) is still an inference-defeat-path.

The somewhat simpler definition (in terms of inference-descendants) employed in simple inference-graphs no longer works, because it will not guarantee that the path resulting from splicing σ^* into σ does not contain an internal loop.



Relating and/or inference-graphs to simple inference-graphs is complicated by the fact that a defeat-link in an and/or inference-graph can correspond to several different defeat-links in the simple inference-graph. Consider the simple inference-graph (13) again and the corresponding and/or inference-graph (&13). Here ρ is the reason-strength for the inference from C to $A \otimes B$. The defeat-link $\langle A \otimes B, B \rangle$ in (&13) corresponds to the two defeat-links $\langle A \otimes B_1, B \rangle$ and $\langle A \otimes B_2, B \rangle$ in (13). If a defeat-link L is φ -critical in the and/or graph, then *some* corresponding defeat-link is φ -critical in the simple graph. More precisely, if L is φ -critical in the and/or graph by virtue of lying on the inference-defeat-path σ , then for every corresponding inference-defeat-path Σ in the simple graph, the corresponding defeat-link in Σ is φ -critical in the simple graph. However, there can also be corresponding defeat-links in the simple graph that do not lie on inference-defeat-paths corresponding to σ . For example, $\langle A \otimes B, B \rangle$ is B -critical in (&13), but only $\langle A \otimes B_2, B \rangle$, not $\langle A \otimes B_1, B \rangle$, is B -critical in (13).

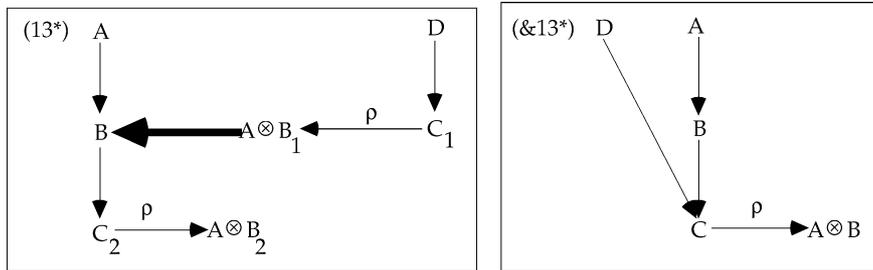
Our objective is to find a way of computing degrees of justification in the and/or inference-graph that agrees with the computation in the simple inference-graph. To compute $j(B, \mathcal{G}_{13})$ we removed the B -critical link to get (13*), and then computed that

$$\begin{aligned} j(B, \mathcal{G}_{13}) &= j(B, \mathcal{G}_{13}^*) \sim j(A \otimes B_2, \mathcal{G}_{13}^*) \\ &= j(A, \mathcal{G}_{13}) \sim [\min\{\rho, j(D, \mathcal{G}_{13})\} \\ &\quad + \min\{\rho, (j(A, \mathcal{G}_{13}) \sim \min\{\rho, j(D, \mathcal{G}_{13})\})\}]. \end{aligned}$$

For example, if $j(A, \mathcal{G}_{13}) = j(D, \mathcal{G}_{13})$, and $\rho < 0.5 \cdot j(A, \mathcal{G}_{13})$, then $j(B, \mathcal{G}_{13}) = j(A, \mathcal{G}_{13}) \sim 2\rho$. It might be supposed that we can compute degrees of justification analogously in the and/or inference-graph by deleting B -critical defeat-links. However, this does not work. Removing B -critical defeat-links would produce inference-graph (&13*). The computation analogous to that employed in (13) would yield

$$\begin{aligned} j(B, \mathcal{G}_{\&13}) &= j(B, \mathcal{G}_{\&13}^*) \sim j(A \otimes B, \mathcal{G}_{\&13}^*) = j(A, \mathcal{G}_{\&13}) \sim \min\{\rho, j(C, \mathcal{G}_{\&13}^*)\} \\ &= j(A, \mathcal{G}_{\&13}) \sim \min\{\rho, \max\{j(A, \mathcal{G}_{\&13}), j(D, \mathcal{G}_{\&13})\}\}. \end{aligned}$$

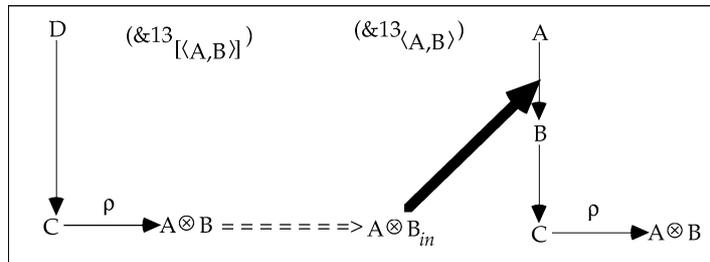
If $\rho < j(A, \mathcal{G}_{\&13}) = j(D, \mathcal{G}_{\&13})$ then $j(B, \mathcal{G}_{\&13}) = j(A, \mathcal{G}_{\&13}) \sim \rho$. So this does not validate the Correspondence Principle.



This illustrates a general point. In an and/or inference-graph $\&\mathcal{G}$, if a defeat-link for a support-link with target P has multiple arguments supporting it, some P -dependent and others not, then it represents both P -dependent and P -independent defeaters in the corresponding simple inference-graph \mathcal{G} , and so by (DJ) it will be counted twice in computing degrees of justification in \mathcal{G} . The computation for $\&\mathcal{G}$ must work similarly. We must compute both a P -dependent and a P -independent degree of justification for the defeat-link, and then subtract their sum from the argument-strength for P . The P -independent degree of justification for a defeater D in $\&\mathcal{G}$ should be the maximum degree of justification of P -independent defeaters corresponding to D in \mathcal{G} , and the P -dependent degree of justification for a defeater D in $\&\mathcal{G}$ should be the maximum degree of justification in \mathcal{G}_P of P -dependent defeaters corresponding to D in \mathcal{G} . The Correspondence Principle will then follow.

Using simple inference-graphs we were able to compute the P -independent and P -dependent values for nodes by modifying the inference-graph. The P -independent values were computed in the original inference-graph \mathcal{G} , and then the P -dependent values were

computed in the inference-graph \mathcal{G}_P . However, the analogous strategy does not work for and/or inference-graphs. For example, there is no way to represent the simple inference-graph (13*) as an and/or inference-graph. The best we can do is construct (&13), but that corresponds to (13), not (13*). This is because in an and/or inference-graph, there is no way to remove just one of the two defeat-links between $A \otimes B$ and B . To get around this difficulty, we must construct “hybrid” inference-graphs which are and/or inference-graphs except that we allow undercutting defeaters to be represented by two separate nodes, one having the P -independent value and the other having the P -dependent value. Let me explain.



We can compute P -independent values for nodes by deleting arguments containing support-links having P -dependent defeaters. Thus in inference-graph (&13), the B -independent strength of $A \otimes B$ can be computed by removing from the inference-graph all support-links having B -dependent defeaters, producing (&13 $_{\langle A, B \rangle}$). To compute the B -dependent strength of $A \otimes B$, we use the B -independent strength to attenuate the strength of the support-link $\langle A, B \rangle$, as in (&13 $_{\langle A, B \rangle}$). The latter inference-graph is constructed from (&13) by (1) removing B -independent arguments for $A \otimes B$, (2) removing B -critical defeat-links, and (3) adding a dummy node $A \otimes B$ in that is treated as an initial node whose value is inherited from inference-graph (&13 $_{\langle A, B \rangle}$). By way of explanation, note that separate B -independent arguments for $A \otimes B$ in (&13) correspond to separate B -independent defeaters for B in the simple inference-graph (13), and the strength assigned to $A \otimes B_{in}$ is then the maximum strength of those B -independent defeaters for $\langle A, B \rangle$. The strength computed for $A \otimes B$ in (&13 $_{\langle A, B \rangle}$) will then be the maximum strength of the B -dependent defeaters for $\langle A, B \rangle$. It then follows by (DJ) that $j(B, \&13) = j(A, \&13) \sim [j(A \otimes B, \&13_{\langle A, B \rangle}) + j(A \otimes B, \&13_{\langle A, B \rangle})]$. This should then be the value computed for $j(\langle A, B \rangle, \&13)$.

In constructing (&13 $_{\langle A, B \rangle}$) we remove B -independent arguments for $A \otimes B$. This is done by removing support-links that occur in B -independent arguments but not also in B -dependent arguments. Let us say that a support-link is a B -link iff it occurs in some B -dependent argument. Precisely:

A support-path in an and/or inference-graph \mathcal{G} is an inference-path in \mathcal{G} from an initial node to some other node.

A support-link is a B -link in an and/or inference-graph \mathcal{G} iff it occurs in a support-path in \mathcal{G} some member of which has a B -dependent defeater.

We can define analogously:

A support-link is a $[B]$ -link in an and/or inference-graph \mathcal{G} iff it occurs in a support-path in \mathcal{G} no member of which has a B -dependent defeater.

The B -links are the support-links included in $(\&13_{\langle A, B \rangle})$ and the $[B]$ -links are the support-links included in $(\&13_{[\langle A, B \rangle]})$. So, for example $\langle C, A \otimes B \rangle$ is both a $[B]$ -link and a B -link, but $\langle D, C \rangle$ is only a $[B]$ -link and $\langle A, B \rangle$ and $\langle B, C \rangle$ are only B -links.

Where \mathcal{G} is an and/or inference-graph, λ is a support-link in \mathcal{G} , and φ is the target of λ , we can then define:

$\mathcal{G}_{[\lambda]}$ results from removing all support-links that are not φ -links from \mathcal{G} and removing all non-initial nodes having no support-link that is a φ -link.

\mathcal{G}_λ results from (1) removing all support-links that are not $[\varphi]$ -links from \mathcal{G} and removing all non-initial nodes having no support-link that is a $[\varphi]$ -link, (2) removing φ -critical defeat-links, (3) adding an initial node $\otimes\lambda_{in}$ with $j(\otimes\lambda_{in}, \mathcal{G}_\lambda) = j(\otimes\lambda, \mathcal{G}_{[\lambda]})$ and a defeat-link to λ , and (4) making all nodes θ that are φ -independent in \mathcal{G} initial with $j(\theta, \mathcal{G}_\lambda) = j(\theta, \mathcal{G})$.

We can then replicate the computation of degrees of justification in simple inference-graphs by adopting the following principle for and/or inference-graphs (possibly containing extra “dummy” defeaters):

(DJ&) If λ is a support-link with basis $\{B_1, \dots, B_n\}$, target φ , and reason-strength ρ , \mathcal{U} and R are the undercutting and rebutting defeaters for λ , and D is the “dummy” defeater for λ (if there is one in \mathcal{G}) then

$$j(\lambda, \mathcal{G}) = \min\{\rho, j(B_1, \mathcal{G}), \dots, j(B_n, \mathcal{G})\} \\ \sim [\max\{j(\mathcal{U}, \mathcal{G}_{[\lambda]}), j(D, \mathcal{G}_{[\lambda]})\} + \max\{j(R, \mathcal{G}_\lambda), j(\mathcal{U}, \mathcal{G}_\lambda)\}].$$

(DJ&) is to be understood so that if any of \mathcal{U} , D , or R is not present in \mathcal{G} then their contributions are omitted from the summation. Assuming that the initial nodes in the inference-graph representing an agent’s cognitive state cannot be defeaters, the dummy defeaters in the inference-graphs constructed by this computation can be detected easily by the fact that they are initial nodes in the constructed inference-graphs.

To illustrate, consider the inference-graph $(\&13)$ again.

$$\begin{aligned} j(B, \&13) &= j(\langle A, B \rangle, \&13) \\ &= j_0(A, \&13) \sim [j(A \otimes B, \&13_{[\langle A, B \rangle]}) + j(A \otimes B, \&13_{\langle A, B \rangle})], \\ j(A \otimes B, \&13_{[\langle A, B \rangle]}) &= j(\langle C, A \otimes B \rangle, \&13_{[\langle A, B \rangle]}) \\ &= \min\{\rho, j(C, \&13_{[\langle A, B \rangle]})\} = \min\{\rho, j(\langle D, C \rangle, \&13_{[\langle A, B \rangle]})\} \\ &= \min\{\rho, j(D, \&13_{[\langle A, B \rangle]})\} = \min\{\rho, j(D, \&13)\} \\ &= \min\{\rho, j_0(D, \&13)\}, \\ j(A \otimes B, \&13_{\langle A, B \rangle}) &= j(\langle C, A \otimes B \rangle, \&13_{\langle A, B \rangle}) \\ &= \min\{\rho, j(C, \&13_{\langle A, B \rangle})\} = \min\{\rho, j(\langle B, C \rangle, \&13_{\langle A, B \rangle})\} \\ &= \min\{\rho, j(B, \&13_{\langle A, B \rangle})\} = \min\{\rho, j(\langle A, B \rangle, \&13_{\langle A, B \rangle})\}, \end{aligned}$$

$$\begin{aligned}
j(\langle A, B \rangle, \&13_{\langle A, B \rangle}) &= j(A, \&13_{\langle A, B \rangle}) \sim j(A \otimes B_{in}, \&13_{\langle A, B \rangle}) \\
&= j(A, \&13_{\langle A, B \rangle}) \sim j(A \otimes B, \&13_{[\langle A, B \rangle]}) \\
&= j_0(A, \&13_{\langle A, B \rangle}) \sim \min\{\rho, j_0(D, \&13)\}.
\end{aligned}$$

Thus

$$\begin{aligned}
j(B, \&13) = j_0(A, \&13) &\sim [\min\{\rho, j_0(D, \&13)\} \\
&+ \min\{\rho, [j_0(A, \&13_{\langle A, B \rangle}) \sim \min\{\rho, j_0(D, \&13)\}]\}].
\end{aligned}$$

11.3. Refining the computation

To refine the recursive computation and make it more readily implementable we can more or less repeat the analysis given earlier for simple inference-graphs. Where σ is a sequence of support-links λ and bracketed support-links $[\lambda]$, let σ_1 be the first member of σ and let σ^* be the rest of σ (its cdr). Where $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$, let $\lambda^\wedge \sigma = \langle \lambda, \sigma_1, \dots, \sigma_n \rangle$. Applying (DJ&) repeatedly generates inference-graphs \mathcal{G}_σ that can be characterized recursively. Where \mathcal{G} is an and/or inference-graph, λ is a support-link in \mathcal{G} , σ is a finite sequence of support-links and/or bracketed support-links, and φ is the target of λ :

$$\mathcal{G}_\emptyset = \mathcal{G};$$

$\mathcal{G}_{[\lambda]^\wedge \sigma}$ results from removing all support-links that are not φ -links from \mathcal{G}_σ and removing all non-initial nodes having no support-link that is a φ -link;

$\mathcal{G}_{\lambda^\wedge \sigma}$ results from (1) removing all support-links that are not $[\varphi]$ -links from \mathcal{G}_σ and removing all non-initial nodes having no support-link that is a $[\varphi]$ -link, (2) removing defeat-links that are φ -critical in \mathcal{G}_σ , (3) adding an initial node $\otimes \lambda_{in}$ with $j(\otimes \lambda_{in}, \mathcal{G}_{\lambda^\wedge \sigma}) = j(\otimes \lambda, \mathcal{G}_{[\lambda]^\wedge \sigma})$ and a defeat-link to λ , and (4) making all nodes θ that are φ -independent in \mathcal{G}_σ initial with $j(\theta, \mathcal{G}_{\lambda^\wedge \sigma}) = j(\theta, \mathcal{G}_\sigma)$.

To reformulate the recursion so as to avoid constructing modified inference-graphs, we define some new concepts. Where $\varphi_1, \dots, \varphi_n$ are nodes of an inference-graph \mathcal{G} , define recursively:

A defeat-link L of \mathcal{G} is σ -critical in \mathcal{G} iff $\sigma \neq \emptyset$ and either (1) L is σ^* -critical in \mathcal{G} or (2a) σ_1 is a support-link λ (not a bracketed support-link) and where φ is the target of λ , L lies on an inference-defeat-path σ in \mathcal{G} from φ to φ containing no σ^* -critical defeat-links and (2b) there is no node α preceding the root of L on σ and node β preceded by the target of the target of L on σ such that there is an inference-path σ^* from α to β such that the path resulting from splicing σ^* into σ in place of L is still an inference-defeat-path.

Where σ_1 is a support-link λ and φ is the target of λ , α is σ -independent in \mathcal{G} iff there is no inference-defeat-path in \mathcal{G} from φ to α containing σ -critical defeat-links.

A support-path μ in \mathcal{G} is a σ -path iff (1) if $\sigma^* \neq \emptyset$ then μ is a σ^* -path and (2a) if $\sigma_1 = \lambda$ then some member of μ has a σ -dependent defeater, and (2b) and if $\sigma_1 = [\lambda]$ then no member of μ has a σ -dependent defeater.

A support-link is a σ -link iff it occurs in a σ -path.

The import of these definitions is given by the following theorems (whose proofs are either trivial or analogous to the proofs of the corresponding theorems for simple inference-graphs):

Theorem 11. *L is λ -critical in \mathcal{G}_σ iff L is $\lambda^\wedge\sigma$ -critical in \mathcal{G} but not σ -critical in \mathcal{G} .*

Theorem 12. *α is σ -independent in \mathcal{G} iff α is σ_1 -independent in \mathcal{G}_{σ^*} .*

Theorem 13. *If ψ is initial in \mathcal{G} then $j(\psi, \mathcal{G}_\sigma) = j_0(\psi, \mathcal{G})$.*

Theorem 14. *If ψ is σ -independent in \mathcal{G} then $j(\psi, \mathcal{G}_\sigma) = j(\psi, \mathcal{G}_{\sigma^*})$.*

Theorem 15. *A support-path in \mathcal{G} is a σ -path iff it occurs in \mathcal{G}_{σ^*} .*

Theorem 16. *If ψ is not σ -independent in \mathcal{G} then $j(\psi, \mathcal{G}_\sigma) = \max\{j(\lambda, \mathcal{G}_\sigma) \mid \lambda \text{ is a } \sigma\text{-link for } \psi\}$.*

Theorem 17. *If $\lambda = \sigma_1$ and λ is a support-link with basis $\{B_1, \dots, B_n\}$ and reason-strength ρ , then*

$$j(\lambda, \mathcal{G}_\sigma) = \min\{\rho, j(B_1, \mathcal{G}_{\sigma^*}), \dots, j(B_n, \mathcal{G}_{\sigma^*})\} \sim j(\otimes\lambda, \mathcal{G}_{[\lambda]^\wedge\sigma^*}).$$

Theorem 18. *If $\lambda \neq \sigma_1$ and λ is a support-link with basis $\{B_1, \dots, B_n\}$, target φ , and reason-strength ρ , and \mathcal{U} and R are the undercutting and rebutting defeaters for λ , then*

$$j(\lambda, \mathcal{G}_\sigma) = \min\{\rho, j(B_1, \mathcal{G}_\sigma), \dots, j(B_n, \mathcal{G}_\sigma)\} \\ \sim [j(\mathcal{U}, \mathcal{G}_{[\lambda]^\wedge\sigma}) + \max\{j(R, \mathcal{G}_{\lambda^\wedge\sigma}), j(\mathcal{U}, \mathcal{G}_{\lambda^\wedge\sigma})\}].$$

By virtue of Theorem 17, rather than talking about dummy defeaters in \mathcal{G}_σ we can talk about real nodes in $\mathcal{G}_{[\lambda]^\wedge\sigma}$.

Now let us define:

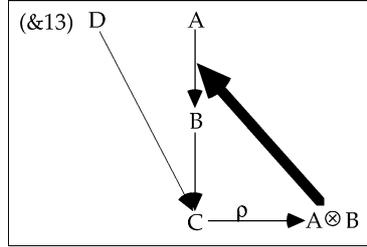
$$j_\sigma(\varphi, \mathcal{G}) = j(\varphi, \mathcal{G}_\sigma), \quad j_\sigma(\lambda, \mathcal{G}) = j(\lambda, \mathcal{G}_\sigma).$$

Then we have:

Theorem 19.

- (a) *If ψ is initial in \mathcal{G} then $j_\sigma(\psi, \mathcal{G}) = j_0(\psi, \mathcal{G})$.*
- (b) *If ψ is not initial in \mathcal{G} then $j_\sigma(\psi, \mathcal{G}) = \max\{j_\sigma(\lambda, \mathcal{G}) \mid \lambda \text{ is a support-link for } \psi\}$.*
- (c) *If ψ is σ -independent in \mathcal{G} then $j_\sigma(\psi, \mathcal{G}) = j_{\sigma^*}(\psi, \mathcal{G})$.*
- (d) *If ψ is not σ -independent in \mathcal{G} then $j_\sigma(\psi, \mathcal{G}) = \max\{j_\sigma(\lambda, \mathcal{G}) \mid \lambda \text{ is a } \sigma\text{-link for } \psi\}$.*
- (e) *If λ is a support-link with basis $\{B_1, \dots, B_n\}$, target φ , and reason-strength ρ , and \mathcal{U} and R are the undercutting and rebutting defeaters for λ , then*
 - (i) *if $\lambda = \sigma_1$ then $j_\sigma(\lambda, \mathcal{G}) = \min\{\rho, j_{\sigma^*}(B_1, \mathcal{G}), \dots, j_{\sigma^*}(B_n, \mathcal{G})\} \sim j_{[\lambda]^\wedge\sigma^*}(\otimes\lambda, \mathcal{G})$;*
 - (ii) *if $\lambda \neq \sigma_1$ then $j_\sigma(\lambda, \mathcal{G}) = \min\{\rho, j_\sigma(B_1, \mathcal{G}), \dots, j_\sigma(B_n, \mathcal{G})\} \sim [j_{[\lambda]^\wedge\sigma}(\mathcal{U}, \mathcal{G}) + \max\{j_{\lambda^\wedge\sigma}(R, \mathcal{G}), j_{\lambda^\wedge\sigma}(\mathcal{U}, \mathcal{G})\}]$.*

Theorem 19 constitutes a recursive definition of $j_\sigma(\psi, \mathcal{G})$. Because $j(\psi, \mathcal{G}) = j_\emptyset(\psi, \mathcal{G})$, this also constitutes a recursive definition of $j(\psi, \mathcal{G})$.



To illustrate the recursive computation, consider the and/or inference-graph (&13) again. Let \mathcal{G} be this inference-graph. By Theorem 19(a), $j(D, \mathcal{G}) = j_0(D, \mathcal{G})$ and $j(A, \mathcal{G}) = j_0(A, \mathcal{G})$. Writing support-links and defeat-links as ordered pairs of nodes, by Theorem 19(e.ii), $j(\langle D, C \rangle, \mathcal{G}) = j(D, \mathcal{G}) = j_0(D, \mathcal{G})$. By Theorem 19(d), $j(B, \mathcal{G}) = j(\langle A, B \rangle, \mathcal{G})$. By Theorem 19(e.ii),

$$\begin{aligned} j(\langle A, B \rangle, \mathcal{G}) &= j(A, \mathcal{G}) \sim [j_{[\langle A, B \rangle]}(A \otimes B, \mathcal{G}) + j_{\langle A, B \rangle}(A \otimes B, \mathcal{G})] \\ &= j_0(A, \mathcal{G}) \sim [j_{[\langle A, B \rangle]}(A \otimes B, \mathcal{G}) + j_{\langle A, B \rangle}(A \otimes B, \mathcal{G})]. \end{aligned}$$

$\langle A \otimes B, B \rangle$ is not the root of an $\langle [\langle A, B \rangle] \rangle$ -critical defeat-link, so $A \otimes B$ is $\langle [\langle A, B \rangle] \rangle$ -independent. Thus by Theorem 19(c),

$$j_{[\langle A, B \rangle]}(A \otimes B, \mathcal{G}) = j(A \otimes B, \mathcal{G}) = \max\{j(\lambda, \mathcal{G}) \mid \lambda \text{ is a support-link for } A \otimes B\}.$$

The only support-link for $A \otimes B$ is $\langle C, A \otimes B \rangle$, so $j_{[\langle A, B \rangle]}(A \otimes B, \mathcal{G}) = j(\langle C, A \otimes B \rangle, \mathcal{G})$. $\langle C, A \otimes B \rangle$ has no defeaters, so by Theorem 19(e.ii), $j(\langle C, A \otimes B \rangle, \mathcal{G}) = \min\{\rho, j(C, \mathcal{G})\}$. By Theorem 19(e.ii), $j(C, \mathcal{G}) = \max\{j(\lambda, \mathcal{G}) \mid \lambda \text{ is a support-link for } C\}$. The only support-link for C is $\langle D, C \rangle$, so $j(C, \mathcal{G}) = j(\langle D, C \rangle, \mathcal{G}) = j_0(D, \mathcal{G})$ (by Theorems 19(e.ii) and 19(a)). Thus $j_{[\langle A, B \rangle]}(A \otimes B, \mathcal{G}) = \min\{\rho, j_0(D, \mathcal{G})\}$.

$\langle A \otimes B, B \rangle$ is a $\langle \langle A, B \rangle \rangle$ -critical defeat-link, so $A \otimes B$ is not $\langle \langle A, B \rangle \rangle$ -independent. Thus by Theorem 19(d),

$$j_{\langle A, B \rangle}(A \otimes B, \mathcal{G}) = \max\{j_{\langle A, B \rangle}(\lambda, \mathcal{G}) \mid \lambda \text{ is an } \langle \langle A, B \rangle \rangle\text{-link for } A \otimes B\}.$$

The only support-link for $A \otimes B$ is $\langle C, A \otimes B \rangle$ and it is a $\langle \langle A, B \rangle \rangle$ -link, so $j_{\langle A, B \rangle}(A \otimes B, \mathcal{G}) = j_{\langle A, B \rangle}(\langle C, A \otimes B \rangle, \mathcal{G})$. $\langle C, A \otimes B \rangle$ has no defeaters, so by Theorem 19(e.ii), $j_{\langle A, B \rangle}(\langle C, A \otimes B \rangle, \mathcal{G}) = \min\{\rho, j_{\langle A, B \rangle}(C)\}$. By Theorem 19(d),

$$j_{\langle A, B \rangle}(C, \mathcal{G}) = \max\{j_{\langle A, B \rangle}(\lambda, \mathcal{G}) \mid \lambda \text{ is a } \langle \langle A, B \rangle \rangle\text{-link for } C\}.$$

The only $\langle \langle A, B \rangle \rangle$ -link for C is $\langle B, C \rangle$, so

$$j_{\langle A, B \rangle}(C, \mathcal{G}) = j_{\langle A, B \rangle}(\langle B, C \rangle, \mathcal{G}) = j_{\langle A, B \rangle}(B, \mathcal{G}).$$

By Theorem 19(d),

$$j_{\langle A, B \rangle}(B, \mathcal{G}) = \max\{j_{\langle A, B \rangle}(\lambda, \mathcal{G}) \mid \lambda \text{ is an } \langle \langle A, B \rangle \rangle\text{-link for } B\} = j_{\langle A, B \rangle}(\langle A, B \rangle, \mathcal{G}).$$

By Theorem 19(e.i),

$$j_{\langle A, B \rangle}(\langle A, B \rangle, \mathcal{G}) = j_0(A, \mathcal{G}) \sim j_{\langle A, B \rangle}(A \otimes B, \mathcal{G}) = j_0(A, \mathcal{G}) \sim \min\{\rho, j_0(D, \mathcal{G})\}.$$

It follows that

$$j(B, \mathcal{G}) = j_0(A, \mathcal{G}) \sim [\min\{\rho, [j_0(A, \mathcal{G}) \sim \min\{\rho, j_0(D, \mathcal{G})]\} \\ + \min\{\rho, j_0(D, \mathcal{G})\}].$$

12. Conclusions

The topic of degrees of justification resulting from defeasible reasoning is virtually unexplored in AI. There is a massive literature on degrees of probability (and the related Dempster–Shafer theory), but this paper partly argues and partly assumes (referring to arguments given elsewhere) that degrees of justification are not probabilities, in the sense that they do not conform to the probability calculus.

The starting point of the present theory is the observation that defeaters that are too weak to defeat an inference may nevertheless diminish the degree of justification of its conclusion. The search for a semantics for defeasible reasoning that is compatible with diminishing leads to a new way of computing degrees of justification recursively. A consequence of this analysis is the principle of collaborative defeat, wherein a pair of defeaters can defeat an inference when they are individually too weak to do that. Work is currently underway to implement this new computation of degrees of justification in OSCAR.

Acknowledgements

This work was supported by NSF grants nos. IRI-9634106 and IRI-IIS-0080888. I owe a special debt of gratitude to Don Nute for organizing a workshop on defeasible reasoning at the University of Georgia in 2000 that inspired me to develop the present theory.

References

- [1] G.A. Barnard, Statistical inference, *J. Roy. Statist. Soc. B* II (1949) 115–149.
- [2] G.A. Barnard, The use of the likelihood function in statistical practice, in: *Proc. V Berkeley Symposium on Mathematical Statistics and Probability I*, 1966, pp. 27–40.
- [3] A. Birnbaum, On the foundations of statistical inference, *J. Amer. Statist. Assoc.* 57 (1962) 269–326.
- [4] C. Chesñevar, A.B. Maguitman, R. Loui, Logical models of argument, *ACM Comput. Surv.* 32 (2000) 337–383.
- [5] M. Covington, D. Nute, A. Vellino, *Prolog Programming in Depth*, 2nd edn., Prentice-Hall, Englewood Cliffs, NJ, 1997.
- [6] A.P. Dempster, A generalization of Bayesian inference, *J. Roy. Statist. Soc. Ser. B* 30 (1968) 205–247.
- [7] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n -person games, *Artificial Intelligence* 77 (1995) 321–357.
- [8] A.W.F. Edwards, *Likelihood*, Cambridge University Press, Cambridge, 1972.

- [9] R.A. Fisher, On the mathematical foundations of theoretical statistics, *Philos. Trans. Roy. Soc. A* 222 (1922) 309–368.
- [10] I. Hacking, *Logic of Statistical Inference*, Cambridge University Press, Cambridge, 1965.
- [11] S. Hanks, D. McDermott, Default reasoning, nonmonotonic logics, and the frame problem, in: *Proc. AAAI-86*, Philadelphia, PA, 1986.
- [12] S. Hanks, D. McDermott, Nonmonotonic logic and temporal projection, *Artificial Intelligence* 33 (1987) 379–412.
- [13] H. Kyburg, *Probability and the Logic of Rational Belief*, Wesleyan University Press, Middletown, CT, 1961.
- [14] H. Kyburg, *The Logical Foundations of Statistical Inference*, Reidel, Dordrecht, 1974.
- [15] I. Levi, Direct inference, *J. Philos.* 74 (1977) 5–29.
- [16] D. Makinson, K. Schlechta, Floating conclusions and zombie paths: Two deep difficulties in the ‘directly skeptical’ approach to inheritance nets, *Artificial Intelligence* 48 (1991) 199–209.
- [17] J. McCarthy, Applications of circumscription to formalizing common sense knowledge, *Artificial Intelligence* 26 (1986) 89–116.
- [18] D. McDermott, A temporal logic for reasoning about processes and plans, *Cognitive Sci.* 6 (1982) 101–155.
- [19] D. Nute, Basic defeasible logic, in: L. Fariás del Cerro, M. Penttonen (Eds.), *Intensional Logics for Programming*, Oxford University Press, Oxford, 1992, pp. 125–154.
- [20] D. Nute, Norms, priorities, and defeasibility, in: P. McNamara, H. Prakken (Eds.), *Norms, Logics and Information Systems*, IOS Press, Amsterdam, 1999, pp. 201–218.
- [21] J. Pollock, The structure of epistemic justification, *Amer. Philos. Quart.* 4 (1970) 62–78.
- [22] J. Pollock, *Knowledge and Justification*, Princeton University Press, Princeton, NJ, 1974.
- [23] J. Pollock, Epistemology and probability, *Synthese* 55 (1983) 231–252.
- [24] J. Pollock, *Contemporary Theories of Knowledge*, Rowman and Littlefield, 1986.
- [25] J. Pollock, Defeasible reasoning, *Cognitive Sci.* 11 (1987) 481–518.
- [26] J. Pollock, *Nomic Probability and the Foundations of Induction*, Oxford University Press, Oxford, 1990.
- [27] J. Pollock, Justification and defeat, *Artificial Intelligence* 67 (1994) 377–408.
- [28] J. Pollock, *Cognitive Carpentry*, MIT Press, Cambridge, MA, 1995.
- [29] J. Pollock, Perceiving and reasoning about a changing world, *Computational Intelligence* 14 (1998) 498–562.
- [30] J. Pollock, The logical foundations of goal-regression planning in autonomous agents, *Artificial Intelligence* 106 (1998) 267–335.
- [31] J. Pollock, J. Cruz, *Contemporary Theories of Knowledge*, 2nd edn., Rowman and Littlefield, 1999.
- [32] D. Poole, A logical framework for default reasoning, *Artificial Intelligence* 36 (1988) 27–47.
- [33] H. Prakken, G.A.W. Vreeswijk, Logics for defeasible argumentation, in: D. Gabbay, F. Guenther (Eds.), *Handbook of Philosophical Logic*, Vol. 5, 2nd edn., Kluwer, Dordrecht, to appear.
- [34] H. Reichenbach, *A Theory of Probability*, University of California Press, Berkeley, 1949 (Original German edition 1935).
- [35] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13 (1980) 81–132.
- [36] E. Sandewall, An approach to the frame problem and its implementation, in: B. Metzger, D. Michie (Eds.), *Machine Intelligence* 7, Edinburgh University Press, Edinburgh, 1972.
- [37] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ, 1976.
- [38] G.R. Simari, R.P. Loui, A mathematical treatment of defeasible reasoning and its implementation, *Artificial Intelligence* 53 (1992) 125–158.
- [39] D. Touretzky, Implicit orderings of defaults in inheritance systems, in: *Proc. AAAI-84*, Austin, TX, 1984.
- [40] D. Touretzky, J. Horty, R. Thomason, A clash of intuitions: The current state of nonmonotonic multiple inheritance systems, in: *Proc. IJCAI-87*, Milan, Italy, 1987, pp. 476–482.
- [41] B. Verheij, *Rules, reasons, arguments*, Ph.D. Dissertation, University of Maastricht, The Netherlands, 1996.