

Is There a Viable Account of Well-Founded Belief?

Maria Lasonen-Aarnio

Received: 21 April 2009 / Accepted: 10 November 2009 / Published online: 20 December 2009
© Springer Science+Business Media B.V. 2009

Abstract My starting point is some widely accepted and intuitive ideas about justified, well-founded belief. By drawing on John Pollock's work, I sketch a formal framework for making these ideas precise. Central to this framework is the notion of an inference graph. An inference graph represents everything that is relevant about a subject for determining which of her beliefs are justified, such as what the subject believes based on what. The strengths of the nodes of the graph represent the degrees of justification of the corresponding beliefs. There are two ways in which degrees of justification can be computed within this framework. I argue that there is not any way of doing the calculations in a broadly probabilistic manner. The only alternative looks to be a thoroughly non-probabilistic way of thinking wedded to the thought that justification is closed under competent deduction. However, I argue that such a view is unable to capture the intuitive notion of justification, for it leads to an uncomfortable dilemma: either a widespread scepticism about justification, or drawing epistemically spurious distinctions between different types of lotteries. This should worry anyone interested in well-founded belief.

1 Well-Founded Belief and Structuralism

Here is a train of thought that has enjoyed considerable popularity in epistemology, especially among those who defend an internalist notion of justification:

It is one thing for a subject to have a justification to believe a proposition, and quite another for her to justifiably believe that proposition. Justified – and not just justifiable – beliefs must be *well-founded*: they must be based in the right sort of way on adequate reasons. Moreover, justification is defeasible. This is

M. Lasonen-Aarnio (✉)
Fitzjames Research Fellow in Philosophy, Merton College, Oxford OX1 4JD, UK
e-mail: maria.lasonen@philosophy.ox.ac.uk; maria.lasonen@merton.ox.ac.uk

to say that even if a certain set of reasons justifies a belief in a proposition p based on those reasons at one time, that belief, based on those reasons, might fail to be justified at a later time.

I will express this as the view that a justified belief in a proposition p is appropriately based on reasons that are adequate and undefeated as reasons to believe p .¹ This thought is to be contrasted, for instance, with a view on which a belief is justified as long as a subject's total set of reasons, or total evidence, sufficiently support that belief.² Such a view would allow for a belief to count as justified even if it was in fact based on poor or defeated reasons. By contrast, the present thought is that justification does not require merely *having* good reasons, but *basin* one's belief on such reasons. In effect, I suspect that most epistemologists—at least those concerned with an internalist notion of justification—would simply take something along the above lines to be a statement of what it is for a belief to be *doxastically* (and not just *propositionally*) justified.

I call views endorsing the idea just sketched *structuralist*, since on these views facts about justification for a subject depend on the structure of her reasoning, i.e., on just how, or based on what, she came to believe what she believes:

Structuralism

A subject s is justified in believing a proposition p just in case (i) s 's belief is appropriately based on reasons r_1, \dots, r_n that make up a (sufficiently strong) prima facie reason to believe p and (ii) r_1, \dots, r_n are not defeated as reasons to believe p .³

What is it for a belief to be *appropriately based* on sufficiently strong prima facie reasons? Assume for now that every time a subject bases a belief on something she applies some rule of inference or inference schema, and that what (i) requires is basing a belief on adequate reasons by using a sound, or acceptable, rule of inference. What is it for a reason to be defeated *as a reason to believe p* ? As an example, take a reason that consists of Suzy's testimony to the effect that it is raining. That reason might be defeated as a reason to believe that it is raining (for instance, by Fred's conflicting testimony), but need not be defeated as a reason to believe that Suzy said something about the weather.

¹ We might term a reason that is adequate for believing p a *prima facie reason* to believe p . Pollock (1974) discusses the idea of *prima facie* justification, and Pryor (2004), for instance, make use of the notion. However, note that, unlike Pollock, I am not intending to restrict *prima facie* reasons to non-deductive reasons.

² For instance, some of the things Earl Conee and Richard Feldman say could be interpreted as an endorsement of such a view. As an example, they claim that as long as believing a proposition fits the evidence a subject has, the subject's belief in that proposition is justified, despite failing to be founded or based on adequate evidence (see Feldman and Conee 2004, pp. 92–93).

³ Very many epistemologists flag ideas along these lines. For instance, in the contemporary debate the ideas about justification sketched by Jim Pryor (2004) come at least close. What complicates things a bit is that Pryor allows irrational, unjustified doubts to turn justified beliefs into unjustified ones (see 2004, p. 365). Perhaps such doubts could just be viewed as reasons the subject has (after all, Pryor restricts his claim to doubts that the subject does not recognise as unjustified), in which case the view is perfectly compatible with Structuralism.

The structuralist must say just how facts about what a subject is justified in believing are fixed by the structure of her reasoning, where one piece of reasoning can produce a defeater for another. Within the epistemology literature, the most serious attempt to answer this question is a series of views defended by John Pollock. Pollock's approach constitutes an *argument-based* theory of defeasible, non-monotonic epistemic reasoning within the area of artificial intelligence. His ultimate goal is to construct an artificial intellect capable of defeasible reasoning. However, he explicitly assimilates the notion of epistemic rationality that theories of non-monotonic reasoning try to capture with an internalist notion of justification. Though I cannot examine all theories of non-monotonic reasoning within AI, Pollock's theory strikes me as the most promising in capturing this notion of justification. However, the mere fact that it is the only detailed theory of just when a belief is well-founded within the area of epistemology should suffice to justify my interest in it.

In what follows I will sketch a broad formal framework for making the above ideas about justified, well-founded belief more precise. Though this framework is not committed to Pollock's views about justification, it draws heavily on his work. I then draw a contrast between two broad ways of calculating degrees of justification within this framework, *probabilistic* and *non-probabilistic* approaches. Pollock's own view is an instance of the latter approach. Such approaches are committed to *The Weakest Link Principle*, a principle entailing that justification is closed under competent deduction, whether from a single or from multiple premises. I say why a probabilistic approach fits together very awkwardly with the basic motivation behind structuralism, and then proceed to raise problems for the non-probabilistic approach, arguing that it fails to capture the intuitive notion of epistemic justification. These problems are brought forth by looking at how the view deals with the Lottery Paradox and the Paradox of the Preface.

In the next section I define an *inference graph*. An inference graph records everything that is relevant about a subject's epistemic state for determining which of her beliefs are justified. In particular, it records just what the subject believes based on what—i.e., what she has *inferred* from what—as well as relations of defeat among her reasons. To simplify the discussion, I will bracket the issue of conditional reasoning, such as reasoning by *reductio ad absurdum* or conditional proof.

2 Inference Graphs

The facts about a subject's epistemic state that are relevant for determining the justificatory status of her beliefs can be recorded by an *inference graph*. Define an inference graph as follows:

Definition

An inference graph G is an ordered septuple $\langle N, R, B, D, f, \rho, \sigma \rangle$ such that

1. N is a nonempty set. Call members of N *nodes* of G .

2. R is a set of ordered pairs the first member of which is a subset of N and the second of which is a member of N . Call R the *support relation* of G .
3. B is a relation such that $R \supseteq B$, and such that for any $i \in N$, there is at most one set J such that $\langle J, i \rangle \in B$. Call B the *inference relation* of G .
4. D is a set of ordered pairs the first member of which is a subset of N and the second of which is a member of N . Call D the *defeat relation* of G .
5. f is a function with N as its domain. Call f the *interpretation function* of G .
6. ρ is a function from members of R to real values in the interval $[0, 1]$.
7. σ is a function from members of D to real values in the interval $[0, 1]$.

Call the things that the interpretation function f assigns to nodes the *occupiers* of the nodes. These are the sorts of things that can serve as reasons by entering into support relations or reasons-for relations with each other. I shall assume that reasons are either mental states or propositions. If reasons are mental states, then the range of f will be a set of mental state types; if reasons are propositions, then the range of f will be a set of propositions.⁴

The relation R encodes relations of support among occupiers of nodes. If $\langle J, i \rangle \in R$, then I shall assume that occupiers of the nodes in J support the occupier of i to a sufficiently high degree. A consequence of this is that an inference graph only records inferences that are successful or *acceptable* in the sense that they have premises that sufficiently support their conclusions. In particular, I will take an acceptable inference to be one that can yield a justified belief in a conclusion if the beliefs in its premises are justified to a sufficiently high degree. Note that I intend to draw a contrast between being *a reason for something*, on the one hand, and being *a reason a subject has*, on the other. Roughly, the idea is that anything occupying a node in an inference graph can stand in the reason-for relation to items occupying other nodes, but—I shall assume—only items occupying sufficiently strong, undefeated nodes count as reasons the subject has. Reasons are the sorts of things on which further justified beliefs can be based.

At this point it might be helpful to have in mind an intuitive if crude picture of how one could think of an inference graph as representing a subject's epistemic state. On this picture, inference graphs record the history of a subject's reasoning, or at least a relevant part of that history extending beyond the present moment. The nodes of the graph correspond to beliefs or other mental states a subject is or has been in. R represents relations of prima facie support among these states (or their contents), ρ indicating how strong the support relations are. B represents the (acceptable) inferences the subject has performed. D represents relations of defeat among the subject's past and present mental states, σ indicating their strength.

Given that facts about (doxastic) justification for a subject are supposed to be fixed by the inference graph she instantiates, it is not at all clear whether the above

⁴ I will talk about beliefs corresponding to nodes in an inference graph, intending to remain neutral on this choice. If the occupiers of nodes are propositions, then a belief corresponds to a node just in case the occupier of the node is a proposition p , the belief has p as its content, and the belief has the right sort of inferential history, a history encoded by the relevant part of the inference graph. If the occupiers of nodes are mental state types, then a belief corresponds to a node just in case it falls under the type assigned to that node, and the belief has the right sort of inferential history.

view of instantiation could be reconciled with a *time-slice internalist* theory according to which facts about justification for a subject at a time supervene only on the subject's internal states at that time. A time-slice internalist would either have to maintain a view on which subjects rebase all of their beliefs at every instant, or else resort not to an actual inferential history, but a history that is somehow reconstructed from the subject's time-slice states.⁵ Also, note that even those who do not endorse time-slice internalism would plausibly have to restrict the inferential history recorded by a graph in some well-motivated manner—one might not, for instance, want long forgotten beliefs to act as defeaters for current beliefs.

Let me turn back to the above definition of an inference graph. Some elements in the definition might be redundant for the task of computing which nodes are defeated, and someone might even argue that the definition does not build in enough elements. Here is an example of a possible redundancy. For now I will assume a criterion for what makes a node a defeater for another node that those with structuralist inclinations tend to adopt. My main points will not essentially depend on this choice, but it might affect whether or not one needs to build the defeat relation *D* into the definition of an inference graph, and it will definitely affect whether *D* can be defined by means of *R* and *B*. In informal terms, this criterion is the following. Assume that a subject believes a proposition *p* on the basis of *r*. Then, *d* is a defeater for the subject's belief in *p* just in case *d* and *r* together do not make up a (prima facie) reason to believe *p*.⁶ This rough idea provides the basis for a very economical treatment of defeat, since it allows fixing the defeat relation by fixing the support and inference relations *R* and *B*. Alternative criteria for being a defeater do not allow for such economy.⁷

Before saying just how fixing *R* and *B* suffices to fix the defeat relation *D* given the criterion for being a defeater that I am here adopting, let me introduce some more terminology. It will be helpful to talk about support-links, inference-links, and defeat-links between sets of nodes and individual nodes in an inference graph. These are just members of the support, inference, and defeat relations *R*, *B*, and *D*.

Definition

For any node *i*, set of nodes *J*, and ordered pair *x*, *x* is a *support-link* from *J* to *i* just in case $x = \langle J, i \rangle$ and $x \in R$. Call *J* the *root* and *i* the *target* of the support-link.

⁵ Pollock (personal correspondence) draws a contrast between the actual historical basis of a belief and the basis recorded by the reasoning system.

⁶ See Pollock and Cruz (1999, p. 195). However, the basic idea goes back much further. For instance, Chisholm (1964) defends a similar thought. Instead of prima facie reasons, Chisholm talks about a proposition completely justifying a belief in another proposition. Assume that a proposition *p* completely justifies a belief in a proposition *q*. The idea is that *d* defeats this justification if and only if *d* is a true proposition and the conjunction *d* & *q* does not completely justify the subject in believing *q*. In effect, this criterion for being a defeater is assumed pretty much by anyone defending the sort of defeasibilist analysis of knowledge popular in the post-Gettier debate. See, for instance, Lehrer and Paxson (1978) and Annis (1978).

⁷ There is also another reason for adding *D* to the definition of an inference graph, which has to do with the function σ . In particular, one might worry that the strengths of the members of the defeat-relation cannot be read from the rest of the inference graph, but must be specified separately through σ .

For any node i , set of nodes J , and ordered pair x , x is an *inference-link* from J to i just in case $x = \langle J, i \rangle$ and $x \in B$. Call J the *root* and i the *target* of the inference-link.

Given the criterion for being a defeater adopted, a defeat-link can be defined by means of the relations R and B . Given that defeat-links are members of the defeat relation D , the following definition shows how fixing R and B fixes D :

Definition

For any node i , set of nodes J , and ordered pair x , x is a *defeat-link* from J to i just in case $x = \langle J, i \rangle$, for some set of nodes K , $\langle K, i \rangle \in B$ but $\langle K \cup J, i \rangle \notin R$, and for every proper subset J^* of J , $\langle K \cup J^*, i \rangle \in R$. Call J the *root* and i the *target* of the defeat-link.

The last condition assures that the root J of a defeat-link is minimal in the sense that no node in J is redundant for J to have its defeating force.⁸

Talk of defeat-links translates into talk of defeaters as follows: whenever a node is the target of a defeat-link, that node *has a defeater*. The root of a defeat-link is always a defeater for the target of the link. But having a defeater does not entail that a node *is defeated*, since the defeater of a node might itself have a defeater. Moreover, sometimes a defeater will not be strong enough to defeat its target.

Given the above terminology, we can say that ρ assigns strengths to support- and inference-links, and that σ assigns strengths to defeat-links. For instance, if ρ assigns some value δ to an ordered pair $\langle J, i \rangle \in R$, then I will say that the strength of the support-link between J and i is δ , or that the occupiers of the set of nodes J support the occupier of node i to degree ρ . Now, it is possible to base a belief on a given set of reasons, or to infer a proposition from those reasons, by employing numerous different rules of inference, some of which may be acceptable or sound while others are not. These rules themselves can be thought of as having different strengths. The strengths of non-deductive rules could be determined in a variety of ways, but deductively valid inference rules have maximum strengths. Very roughly, what I mean by the strength of an inference rule is how likely it is that inferences employing that rule avoid giving false conclusions from true premises. Talk of the degree to which the premises of an inference support its conclusion, or of the strength of a support-link, should always be understood in terms of the strength of the strongest acceptable inference rule for getting from those premises to that conclusion. Unless I state otherwise, I will assume that a subject always infers by the strongest acceptable rule.

Finally, talk of inference parents, inference ancestors, and ancestral inferences will be helpful:

Definition

For any set of nodes J and any nodes i and k , i is a *support-parent* of k just in case $\langle J, k \rangle \in R$ and $i \in J$.

⁸ This definition allows the roots of defeat-links to be sets consisting of more than one node. Someone who thought that a defeater was always the occupier of a single node (such as Pollock 1995), could require that J always contains only a single node.

For any nodes j and k , k is a *support-ancestor* of j just in case k occurs before j in a finite sequence of nodes each of which is a support-parent of the next.⁹

For any set of nodes J and any nodes i and k , i is an *inference-parent* of k just in case $\langle J, k \rangle \in B$ and $i \in J$.

For any nodes j and k , k is an *inference-ancestor* of j just in case k occurs before j in a finite sequence of nodes each of which is an inference-parent of the next.

For any node j , the *ancestral inference-links* of j are all of the inference-links that have inference-ancestors of j as their targets.

This concludes my definition of an inference graph and of related notions.

Now, the basic idea of structuralism is compatible with various views of the types of chains of reasoning that yield justified beliefs. On *foundationalist* views, any chain terminates with a foundational reason, a reason that is not supported by any other reason. On *coherentist* versions of structuralism, chains of reasoning can form loops. *Infinitist* views allow for chains of reasoning along which there is no reason that is not supported by another reason, despite the chains not being circular.

Though I will not argue for this here, and my arguments below will not essentially rely on this assumption, I believe the most plausible structuralist views to be foundationalist. The core foundationalist thought is that every chain of justification ends after finitely many steps in basic or foundational reasons. The foundationalist will need to add an additional constraint on inference graphs. The constraint is the following. For any inference graph G , let P be the inference parent relation of G : for any nodes i and j of G , $\langle i, j \rangle \in P$ just in case i is an inference parent of j . To rule out inference-loops, the constraint must entail that no node is an inference-ancestor of itself. In addition, it must be possible to get from any node to a foundational node in finitely many steps by tracing inference-links backwards. The constraint needed is that P , the relation of being an inference parent, is *well-founded*. This amounts to the following:

Foundationalism

For any inference graph G and any non-empty subset J of nodes of G , J contains a *P-minimal element*, that is, an element i such that J doesn't contain any inference-parent of i .¹⁰

Well-foundedness rules out inference-chains that form loops, since in such chains, no element is P -minimal. It also rules out infinite inference-chains of the sort allowed by infinitists, since such chains cannot have a first member.¹¹

Recall that facts about which beliefs of a subject are justified at a time t are fixed by the inference graph instantiated by the subject at t . Once the question of which

⁹ Note that I am not here defining the reflexive ancestral relation: a node k is not automatically a support-ancestor of itself.

¹⁰ See, for instance, Enderton (1977) for the definition of a well-founded relation.

¹¹ One might think that in addition to the conditions given, every node should either be basic or the target of a support-link. However, it's not clear what work ruling out isolated nodes would do, as long as beliefs corresponding to such nodes would not be justified, and as long as such nodes could not be roots of defeat-links.

inference graph represents a subject's epistemic state at a time is settled, the structuralist faces the difficult task of stating rules for computing the strength of each node of the graph based on the information contained in the graph, in particular, the relations B and D, and the assignments of strengths to the inference- and defeat-links. The idea will then be that the (doxastically) justified beliefs of a subject correspond to the sufficiently strong nodes of the inference graph she instantiates.¹² It is very difficult to come up with a set of rules determining the strengths of the nodes of any inference graph in a satisfactory way.¹³ I will not even attempt to formulate a precise set of rules (a task that is for the structuralist to take up) but will, instead, look at guiding principles behind two different ways of computing strengths. In particular, I now turn to a distinction between *probabilistic* and *non-probabilistic* approaches.

3 Two Approaches to Computing Degrees of Justification

The difference between probabilistic and non-probabilistic approaches is clearest if one sets aside the issue of how to take defeaters into account. So for now, consider nodes that have no defeaters. A node might have a single inference parent or numerous inference parents, and its inference parents might support the node to a maximal degree (which is the case when an inference is deductive¹⁴) or to some non-maximal degree. This gives rise to four types of cases:

- (1) single inference parent, deductive inference;
- (2) single inference parent, non-deductive inference;
- (3) numerous inference parents, deductive inference;
- (4) numerous inference parents, non-deductive inference.

Any views I consider agree on cases of type (1). When a node is maximally supported by a single inference parent and has no defeaters, the strength of the node just equals the strength of its inference parent. Potential divergences can arise in cases of types (2)–(4).

I will adopt an assumption entailing that how one treats cases falling under any of (2)–(4) will simply follow once a policy has been settled for treating cases of type (2) and specific cases of type (3) that involve inferring a conjunction from individual conjuncts. The assumption is the following. Take a multi-premise inference by which a proposition q is inferred from a set of premises p_1, \dots, p_n . Let the belief in q

¹² Hence, though I talk about degrees of justification, the picture being sketched is compatible with the thought that being justified in believing a proposition is an all-or-nothing affair. As an analogy, one could talk about different degrees of tallness, while holding that only people above, say, 180 cm are tall.

¹³ Pollock (2001) attempted to formulate set of rules, or what is referred to within AI as a “semantics”, that takes into account degrees of justification, instead of just assigning one of two defeat-statuses to each node, but retracted this semantics. The semantics he proposed in (1995) and (1999) involving just two defeat-statuses have also turned out to have counterexamples. He subsequently worked on a new semantics.

¹⁴ I am here simply ignoring possible cases in which there is a maximal degree of support without an entailment.

correspond to a node j , and the beliefs in p_1, \dots, p_n correspond to nodes i_1, \dots, i_n . The assumption is that the strength of j equals the strength of a node arrived at from i_1, \dots, i_n through two inferential steps. In the first, one infers the conjunction $p_1 \& \dots \& p_n$ from the premises p_1, \dots, p_n . This is an inference of type (3). In the second, one infers the proposition q from the conjunction $p_1 \& \dots \& p_n$.¹⁵ I am assuming that the degree to which the conjunction supports q just equals the degree to which the individual premises p_1, \dots, p_n taken together support q . This second inference is of type (2). It follows from the assumption made that if the premises p_1, \dots, p_n entail q , then the degree to which a belief in q is justified will equal the degree to which a belief in the conjunction $p_1 \& \dots \& p_n$ based on the premises p_1, \dots, p_n is justified. Hence, the two types of cases left to deal with are cases of type (2), non-deductive single-premise inferences, and cases of type (3) that involve deducing a conjunction from individual conjuncts.

So let me now state the difference between probabilistic and non-probabilistic approaches to calculating strengths of nodes. The crucial dividing question is the following (I am here assuming that the threshold level for justification is set at some non-maximal level):

Assume that j is a node with no defeaters, and that its inference parents are nodes i_1, \dots, i_n (where $n \geq 1$). Assume that the strength of each of i_1, \dots, i_n is at least as great as some non-maximal value δ . And assume that the degree to which j is supported by its inference-parents is at least δ , that is, $\rho(\langle\{i_1, \dots, i_n\}, j\rangle) \geq \delta$. In this situation, can the strength of j be lower than δ ?

Non-probabilistic approaches to calculating strengths of nodes entail a negative answer to this question, and probabilistic approaches entail a positive answer. Of course, there is more than one possible view that answers the dividing question positively, and more than one possible view that answers it negatively. However, I think that the only motivation for answering the question in a positive manner is a broadly probabilistic framework, and that an uncontroversial principle narrows down ways of answering the question negatively to views committed to the Weakest Link Principle (WLP) discussed below.

Here is the uncontroversial principle. Assume that i_1, \dots, i_n are the inference parents of j . The principle states that the strength of j cannot be greater than the weakest of the strengths of i_1, \dots, i_n and the value that ρ assigns to $\langle\{i_1, \dots, i_n\}, j\rangle$ (that is, the degree to which i_1, \dots, i_n support j). For instance, a belief in the conclusion of a deductive inference can never be justified to a higher degree than the weakest of the individual premises it is based on. This principle, together with a negative answer to the dividing question formulated above, entails the *Weakest Link Principle*:

¹⁵ To be more precise, the strength of j in inference graph G in which q is inferred from p_1, \dots, p_n in one step equals the strength of a node j^* in a different inference graph G^* that differs from G only in that q is inferred through two steps, first one in which a subject infers the conjunction $p_1 \& \dots \& p_n$ from p_1, \dots, p_n and second a step in which she infers q from $p_1 \& \dots \& p_n$.

The Weakest Link Principle (WLP)

If j is a node with no defeaters and its inference parents are the nodes i_1, \dots, i_n , then the strength of j is the weakest of the strengths of i_1, \dots, i_n and the value assigned by ρ to $\langle \{i_1, \dots, i_n\}, j \rangle$.

WLP entails that if a subject comes to believe a proposition p by a chain of inferences none of which have defeaters, then the degree to which she is justified in believing p is the minimum of the degrees to which she is justified in believing the ancestral premises of p (that is, the strengths of the nodes that are inference ancestors of the node that the belief in p corresponds to) and the degrees to which the premises of the component inferences of the chain support their conclusions.

Together with the sort of criterion for being a defeater assumed so far, WLP entails that deductively valid inferences from nodes each of which is stronger than or at least as strong as some value δ always produce nodes at least as strong as δ . Indeed, Pollock, who defends the Weakest Link Principle, takes this to be one of the most decisive considerations in its favour.¹⁶ The principle yields the following closure principle:

Justification Closure

For any subject s and propositions p_1, \dots, p_n, q , if s is justified in believing each of p_1, \dots, p_n , and s comes to believe q solely based on competent deduction from p_1, \dots, p_n , and her beliefs in p_1, \dots, p_n remain justified throughout, then s 's belief in q is justified.

If the subject's beliefs in p_1, \dots, p_n are justified, then the nodes that these beliefs correspond to must be sufficiently strong. Assume that the weakest of the strengths of these nodes is at least of an arbitrary strength δ . By the criterion for being a defeater that was assumed above, there can be no node that acts as a defeater for this node. A defeater would have to be some d such that $\{d, p_1, \dots, p_n\}$ is not a reason to believe q . But because of the monotonicity of logical entailment, for any d , $\{d, p_1, \dots, p_n\}$ will entail q , thereby supporting it to a maximal degree. Then, by WLP, the node corresponding to the subject's belief in q that is based on inference from p_1, \dots, p_n is likewise at least of strength δ .

Closure ensures that a subject can never be justified in believing each proposition in a set of propositions entailing a contradiction. For instance, a subject cannot both be justified in believing of each ticket in a lottery that it will lose, and in believing that at least one of the tickets will win. I discuss lotteries more below. I argue that in effect, WLP commits the structuralist to a sceptical treatment of the Lottery Paradox, a treatment on which a subject is not justified in believing of *any* ticket that it will lose. But let me first say why it is difficult to see how the probabilistic approach could be reconciled with structuralism.

In particular, I take one of the basic tenets of a structuralist way of thinking to be that if a belief in a proposition q is based on further beliefs in propositions p_1, \dots, p_n , then the degree to which the subject is justified in believing q ought to depend on nothing but the following factors: the degrees to which she is justified in believing the premises p_1, \dots, p_n , the degree to which these premises support q , and the degree

¹⁶ No doubt numerous philosophers would disagree, arguing that WLP is false for the reason that risks can pile up over multi-premise deductions in a way that it does not allow for.

to which she is justified in believing any propositions defeating the inference from p_1, \dots, p_n to q . As will become clear, the problem of the probabilistic approach is that determining the strength of a node even in a simple case in which it has no defeaters threatens to bring in more global considerations that go beyond the strengths of the inference ancestors and ancestral inferences of that node.

4 Problems with Probabilistic Approaches

I will assume that probabilistic approaches subscribe to the following rough idea: the degree to which a belief in a proposition based on certain premises is justified equals the probability of all of the premises being true—i.e., the probability of the conjunction of the premises—times the degree to which the premises support the conclusion, or *how likely* the premises make the conclusion.¹⁷

The most obvious way of thinking about this degree of support is as the conditional probability of the conclusion on the premises. For instance, if the degree to which I am justified in believing that the object looks red (say, based on Suzy's testimony) is 0.99, and the probability of the object being red conditional on its looking red is similarly 0.99, then the degree to which I am justified in believing that it is red based on its looking red is $0.99 \times 0.99 = 0.9801$. It follows that even in the absence of defeaters, deductively valid inferences from multiple justified premises can fail to result in justified beliefs. Deductions exploiting rules as mundane as *modus ponens* will sometimes fail to extend a subject's justification.

However, this suggestion runs into an immediate problem. In particular, take a case in which a subject infers a proposition q from a set of multiple premises p_1, \dots, p_n (and there are no defeaters). The thought is that the degree to which the subject is justified in believing q is the probability of the conjunction $p_1 \& \dots \& p_n$ times the conditional probability of q on this conjunction. But according to standard probability theory, the probabilities of the individual conjuncts p_1, \dots, p_n are not enough to determine the probability of their conjunction, since the probability of the conjunction depends on relations of probabilistic dependence that can be expressed in terms of conditional probabilities. It follows from the definition of conditional probabilities¹⁸ that

$$\Pr(p_1 \& \dots \& p_n) = \Pr(p_1) * \Pr(p_2|p_1) * \dots * \Pr(p_n|p_1 \& \dots \& p_{n-1}).$$

¹⁷ When I talk about the degree to which a set of premises supports a conclusion, it is worth underlining that this notion of support is different from the notion of degree of evidential support that arises in connection with different measures of confirmation. The latter is a comparative notion, having to do with what impact acquiring certain evidence has on the probability of a proposition. For instance, a well-known idea is one on which the degree of confirmation of a proposition h by a body of evidence E is the probability of h conditional on E minus the unconditional probability of h . If we translate this talk of probabilities to talk of degrees of justification, the question is *how much more* justified acquiring a piece of evidence makes one in believing a given proposition than one previously was.

¹⁸ The definition is the following: $\Pr(p|q) = \Pr(p \& q)/\Pr(q)$. Some prefer to treat this as a probability axiom instead of a definition.

But now the problem is that there is not any obvious way of reading the required conditional probabilities off an inference graph that just has nodes corresponding to the individual conjuncts!

One possible answer to how the required conditional probabilities are determined is that there is a prior probability function Pr_{PRIOR} fixing their values. Informally, we could think of Pr_{PRIOR} as giving the objective degree to which any proposition lends support to another proposition, thereby also fixing the function ρ assigning strengths to support-links.¹⁹ Then, the suggestion is that the conditional probabilities needed for fixing probabilities of conjunctions are always prior conditional probabilities. However, this view cannot be right. For instance, take a case in which a subject believes a proposition p_1 based on a basic reason ε_1 (an experience, perhaps), a proposition p_2 based on a basic reason ε_2 , and then comes to believe p_1 & p_2 based on p_1, p_2 . Let i_1 and i_2 be the nodes to which the beliefs in p_1 and p_2 correspond, and $\text{Str}(i_1)$ and $\text{Str}(i_2)$ the strengths of these nodes. According to the proposed view, the degree to which the subject's belief in p_1 & p_2 is justified equals $\text{Str}(i_1) * \text{Pr}_{\text{PRIOR}}(p_2|p_1)$. But first, note that then the degree to which the subject's belief in p_1 & p_2 is justified does not depend in any way on $\text{Str}(i_2)$. Whatever support experience ε_2 gave to p_2 has been erased by the time the subject has inferred p_1 & p_2 from p_1, p_2 . And second, there is no guarantee that $\text{Str}(i_1) * \text{Pr}_{\text{PRIOR}}(p_2|p_1) = \text{Str}(i_2) * \text{Pr}_{\text{PRIOR}}(p_1|p_2)$. This would be the case if we were dealing with a single probability distribution, but as it is, there is no guarantee that $\text{Str}(i_1) = \text{Pr}_{\text{PRIOR}}(p_1)$ and $\text{Str}(i_2) = \text{Pr}_{\text{PRIOR}}(p_2)$. Any asymmetry in how the calculation should be done is very difficult to motivate.

One might be led to the thought that what is needed for determining the probability of the conjunction is a probability function that results from the prior function by some update. One might suggest, for instance, that when calculating the strength of a node i , the prior function is updated on all of the basic inputs that are inference ancestors of i . So, for instance, in the above case the degree to which the subject comes to be justified in believing p_1 & p_2 equals $\text{Str}(i_1) * \text{Pr}_y(p_2|p_1)$ where Pr_y is the probability function that results from conditionalising Pr_{PRIOR} on both ε_1 and ε_2 . But this account still faces the problem that the degree to which the subject comes to be justified in believing p_1 & p_2 will depend on which way the calculation is done: there is no guarantee that $\text{Str}(i_1) * \text{Pr}_y(p_2|p_1)$ will equal $\text{Str}(i_2) * \text{Pr}_y(p_1|p_2)$, for the simple reason that $\text{Str}(i_1)$ and $\text{Str}(i_2)$ need not correspond to the same probability function.

In light of such problems, it looks as though the probabilistically inclined structuralist must give up calculating the probabilities of conjunctions in the standard way. What seems to be needed is an alternative that exploits only the degrees of justification assigned to the individual conjuncts. I take the main suggestion to be one on which the probability of a conjunction is calculated on the default assumption that all of the individual conjuncts are probabilistically independent, by multiplying their degrees of justification.²⁰ But one could reasonably object that this suggestion gives the wrong verdicts in some cases, for

¹⁹ Of course, this idea would have to be qualified in light of non-propositional items of evidence.

²⁰ I am assuming here that none of the conjuncts logically entail others.

sometimes there are non-logical dependences that intuitively ought to be taken into account.²¹

It is worth pointing to a further, general worry about reconciling structuralism with ideas about degrees of belief that at least most philosophers who are probabilistically inclined in the first place will want to subscribe to. In particular, many philosophers swayed by probabilistic considerations will think that the probability axioms impose constraints of rationality on a subject's degrees of belief. But another line of thought that is at least *prima facie* plausible is that a subject's degree of belief in a proposition ought to track the degree to which she is justified in believing that proposition.²² The problem is that within a structuralist framework, there is no way of reconciling these two ideas.

The problem is created by the way in which, on the structuralist way of thinking, the strength of a node depends just on those branches of the inference graph leading up to that node, together with any defeaters for the nodes along those branches. For instance, the probability axioms entail that $\Pr(p \vee q) = \Pr(p) + \Pr(q) - \Pr(p \& q)$. However, even when a subject's inference graph has nodes corresponding to beliefs in the relevant propositions, there is no guarantee that the strengths of these nodes will obey the above equation. For instance, assume that a subject's inference graph has a node corresponding to a belief in p and a node corresponding to a belief in q , that both nodes are of strength 0.95, and that there are no defeaters. Assume that she believes the disjunction $p \vee q$ on the basis of p , so that the strength of the node corresponding to a belief in the disjunction is likewise 0.95 (since it is deduced from a single premise). Now, she also believes the conjunction $p \& q$ based on the individual conjuncts. The only strength that could be assigned to the node corresponding to this belief that would make her degrees of beliefs obey the above equation is 0.95. But assuming that neither p nor q entail the other, this would violate any of the vaguely probabilistic candidate rules for calculating the degree to which a subject is justified in believing a conjunction based on the individual conjuncts. Or, take a case in which the degree to which a subject is justified in believing a proposition p is 0.9. By a single premise deduction, she infers from p the proposition $p \vee \sim p$. By any probabilistic rule, the degree to which she is justified in believing the disjunction equals the degree to which she is justified in believing p (times 1, since the inference is deductive). Hence, the degree to which she is justified in believing $p \vee \sim p$ is 0.9, and not 1, but if degrees of belief obey the probability axioms, then the subject's degree of belief in $p \vee \sim p$ had better be 1!

²¹ For instance, assume that p is the proposition that it will snow, and q is the proposition that the temperature will be freezing. These propositions are positively probabilistically relevant, and intuitively, the degree to which one is justified in believing their conjunction ought to be higher than the degree of justification that the rule mentioned would yield.

²² The idea that one should proportion one's belief to the evidence has been championed by Hume, Quine, and many others. In so far as our concern is a finer-grained notion of belief, this leads very naturally to the thought that one's degree of belief in a proposition should mirror how likely it is on one's evidence. The present framework replaces talk of evidence by talk of reasons. Then, proportioning one's beliefs to one's reasons will require believing a proposition to a degree that reflects the strength of one's reasons for that proposition—or, more precisely, to a degree that reflects the strength of the node corresponding to that belief, this strength just being the degree to which one is justified in believing that proposition.

Other problems also arise, such as a subject's degrees of belief over propositions making up a logical partition not summing up to 1.

In light of the above problems for the probabilistic approach, the most workable version of structuralism seems to be a view committed to the Weakest Link Principle. I now turn to such a view.

5 Some Consequences of the Weakest Link Principle for the Lottery Paradox

The WLP is the basic tenet of a non-probabilistic approach to calculating the strengths of the nodes in an inference graph. In effect, it is alone sufficient to fix the strengths of the nodes of a simple inference graph with no defeat-links. However, the interesting, challenging cases are graphs containing defeat-links, and it is such graphs that I will be interested in below.

For now, the discussion will be simplified by restricting attention to inference graphs each node of which can be assigned one of two strengths. For instance, assuming a foundationalist framework and WLP, graphs in which every basic node is of the same strength δ , and in which each inference-link is at least of strength δ , are plausibly like this. If a graph satisfies these conditions and has no defeat-links, it is clear that, assuming WLP, any non-basic node in the graph will have a strength of δ . I will also be assuming that within a graph satisfying the above conditions, a node is defeated just in case it is assigned the minimal value, and that a defeater cannot weaken a node without defeating it.²³ These assumptions allow assigning to each node just one of two values, either δ or a minimal value. Assuming that δ is a value that is above the threshold-level for justification, nodes of strength δ can just be assigned "undefeated", and nodes of the minimal value "defeated".

One of the most interesting and potentially worrying consequences of WLP has to do with the Lottery and Preface Paradoxes. I will argue, first, that the structuralist committed to WLP is forced to give a *sceptical* treatment of the Lottery Paradox on which a subject is not justified in believing of any ticket in the lottery that it will lose. Instead of providing a precise set of rules for computing the strengths of the nodes in an inference graph, I will argue that assuming WLP, a couple of basic principles that provide constraints on *any* adequate set of rules suffice to yield a treatment of the paradox along these lines. I then look at the problem created by the fact that this treatment of the Lottery Paradox is in danger of generalising to the Paradox of the Preface, thus leading to a sceptical result. I point to a difference between standard Lottery cases and Preface cases that John Pollock exploits to prevent a sceptical treatment of the Lottery Paradox from generalising to the Paradox of the Preface. However, I show how this commits Pollock to treating different types of lotteries differently. I argue that this yields unmotivated divisions between lotteries in which a subject is justified in believing that her ticket will lose and ones in which she is not.

²³ The idea is that defeat will be an all-or-nothing affair, and that defeated nodes will have minimal strengths.

At this point it is important to say something about what the discussion below is supposed to show about what actual subjects are or are not justified in believing. Whatever instantiating a graph involves, I have assumed that it involves in some sense having performed the inferences corresponding to the inference-links of the graph. But the inference graphs representing Lottery and Preface Paradox situations discussed below will have many inference-links corresponding to inferences that, one might think, actual subjects very rarely perform. Moreover, these links will be crucial for obtaining the desired results. I do not know exactly what structuralists such as Pollock would say about this, but I will just restrict the discussion to a class of subject satisfying certain constraints. I will take these constraints to include performing somewhat obvious inferences, or seeing somewhat obvious logical entailments. The idea will then be that if such a subject is not justified in believing that her lottery ticket will lose, neither can an ordinary subject with less logical acumen be so justified. Let me now finally turn to the Lottery Paradox.

All along, I have assumed that an inference graph only records *acceptable* inferences performed by a subject. Within a framework that assumes WLP, these inferences can be characterised as follows: an inference from premises p_1, \dots, p_n to a conclusion q is acceptable just in case an inference of that type yields a justified belief in q whenever one's beliefs in p_1, \dots, p_n are justified, and whenever there are no undefeated defeaters for q . Given WLP, all deductive inferences preserve justification in the sense just described. In particular, this applies to an inference from any number of propositions each of which states of an individual ticket in a lottery that it will lose to the conclusion that all of the tickets in question will lose.

For now, I will have in mind lotteries with some number n of tickets, each of which has the same chance of winning, and exactly *one of which is bound to win*.²⁴ For simplicity, I will consider a lottery with just three tickets. As long as the threshold for justification is set at a low enough level, one ought to end up with a case that has exactly the same structure as a more realistic case exemplifying the Lottery Paradox. Assuming that degrees of justification can be represented by numbers in the interval $[0, 1]$, the minimal threshold for a justified belief (or for a node being assigned "undefeated" by the computation of defeat statuses) would have to be as low as $2/3$. I will assume that it is exactly $2/3$.

Let r be a proposition stating what the structure of the lottery is: there are exactly 3 tickets in the lottery, each ticket has an equal chance of winning, and exactly one ticket will win. Let $\sim t_1$, $\sim t_2$, and $\sim t_3$ be propositions stating of each individual ticket in the lottery that it will lose. Assume that a subject believes r on the basis of an inference from a node that has no defeaters, and that the premises of the inference support r to degree $2/3$. But r supports each of $\sim t_1$, $\sim t_2$ and $\sim t_3$ to degree $2/3$, which means that inferring any of these propositions from r is acceptable. All of the inferences below are deductive and hence, acceptable:

$r, \sim t_1, \sim t_2$; Therefore, t_3 .

$r, \sim t_2, \sim t_3$; Therefore, t_1 .


$r, \sim t_1, \sim t_3$; Therefore, t_2 .

²⁴ Ultimately, the structuralist wedded to WLP will have to give the same type of sceptical treatment of lotteries in which *at least* one ticket will win.

So are the following inferences:

$\sim t_1, \sim t_2, \sim t_3$; Therefore, $\sim (t_1 \vee t_2 \vee t_3)$.
 t_1, t_2, t_3 ; Therefore, $\sim r$.

The latter inference is acceptable, since t_1, t_2 , and t_3 entail the negation of one of the conjuncts of r , namely, the conjunct stating that exactly one ticket will win. Then, t_1 is a defeater for $\sim t_1$, t_2 is a defeater for $\sim t_2$, and t_3 is a defeater for $\sim t_3$. $\sim (t_1 \vee t_2 \vee t_3)$ is a defeater for r , since r entails $t_1 \vee t_2 \vee t_3$. Moreover, $\sim r$ is a defeater for r . The situation can be graphically represented as follows:



5.1 Graph G_L

The lines with no dashes represent inference-links, and the dashed lines represent defeat-links. The graph contains nodes $k, h, i_1, i_2, i_3, j_1, j_2, j_3$, and l . What ' $i_1: \sim t_1$ ', for instance, indicates is that corresponding to node i_1 is a belief in proposition $\sim t_1$. In other words, the interpretation-function f of G_L assigns to node i_1 either just the proposition $\sim t_1$, or a mental state type involving the content $\sim t_1$. The defeat-links in G_L are $\langle j_1, i_1 \rangle, \langle j_2, i_2 \rangle, \langle j_3, i_3 \rangle, \langle l, k \rangle, \langle h, k \rangle$.²⁵ Now, there are numerous deductive, and hence acceptable, inferences that a subject could perform in addition to the ones that are recorded by G_L . For instance, from r and $\sim r$ she could infer the contradiction \perp , and from \perp she could infer anything at all. I will ignore such

²⁵ The defeat-links $\langle j_1, i_1 \rangle, \langle j_2, i_2 \rangle$, and $\langle j_3, i_3 \rangle$ do not go both ways: for instance, node j_1 is considered to be a defeater for node i_1 , but not vice versa. First, for present purposes I am assuming a criterion for being a defeater on which deductive inferences are not defeasible, and the inferences to t_1, t_2 , and t_3 are deductive. Second, even if an alternative criterion for being a defeater was adopted on which deductive inferences were defeasible, I do not think a case could be made for thinking that i_1 , for instance, was a defeater for j_1 . We do not seem to be dealing with a typical case of defeat of a deductive inference: the subject does not, for instance, have any reason to think that she has made a mistake in inferring any of the propositions t_1, t_2 , and t_3 from the premises that logically entail them. Finally, even if the defeat-links did go both ways, I think the right verdict on the graph would be the one I reach below.

additional inferences, since I assume them to be irrelevant for fixing the defeat statuses of the nodes in graph G_L .

I will now argue that assuming WLP, any acceptable assignment of defeat-statuses to the individual nodes in G_L will have to assign “defeated” to all of $i_1, i_2, i_3, j_1, j_2, j_3$, and l . The consequence is that a subject whose epistemic state is represented by G_L is not justified in believing of any of the tickets in the lottery that it will lose.²⁶ Such a treatment of the Lottery Paradox follows from principles that any set of rules for assigning defeat statuses to the nodes of an inference graph committed to WLP ought to subscribe to, together with a symmetry assumption. The symmetry assumption is that nodes i_1, i_2, i_3 are symmetric so that either none or all of them are defeated, and similarly for nodes j_1, j_2, j_3 . I take this to follow from the assumption that inference graphs record all of the factors relevant for determining whether or not a belief a subject holds is justified. To admit that two nodes that are perfectly symmetrical differ in their defeat statuses would, it seems, be to admit that the graph does not suffice for determining facts about justification after all.²⁷ The general principles are the following:

Inheritance

If a node j has a (non-redundant²⁸) inference ancestor i that is defeated, then j is defeated.

No Defeaters

If a node j has no defeaters, and all of the inference ancestors of j are undefeated, then j is undefeated.

Undefeated Defeaters

If a node j has a defeater i that is undefeated, then j is defeated.

Inheritance says roughly that one cannot base a justified belief on unjustified beliefs or other unjustified mental states. That is, a node always inherits the status “defeated” from one of its non-redundant ancestors. I take this principle to be immensely plausible in itself, but note also that it is entailed by WLP. To see why *No Defeaters* is plausible, recall that an inference graph only records acceptable inferences. Within a framework that accepts WLP, these are inferences that produce justified beliefs when the premises are justified and there are no defeaters for the conclusion. *Undefeated Defeaters* follows from the basic thought that given WLP, in the sorts of inference graphs that I am looking at (each node of which can be assigned one of two defeat statuses) there are only two ways for a node to be undefeated: either by having no defeaters, or by all of its defeaters being defeated (in particular, we do not have to worry about a node i having a defeater that is strong enough to be undefeated, but not strong enough to defeat i). These principles do not themselves suffice to fix the defeat statuses of every node of every inference graph.

²⁶ I think it is also extremely plausible that given these assignments, k will be undefeated, but do not really need this assumption, and will not complicate things by including more principles.

²⁷ I am assuming that subjects who instantiate inference graph G_L hold the same attitude towards each ticket in the lottery, and perform the same inferences regarding each ticket.

²⁸ One could argue that in cases in which a node j is based on a defeated node i that is completely redundant for the inference, j is not defeated.

But I will assume that any acceptable semantics wedded to WLP ought to entail them. And they have interesting consequences for graph G_L representing the structure of the Lottery Paradox.

First, note that node h has to be defeated. For assume that it was not. Then, k would have an undefeated defeater, and by *Undefeated Defeaters*, would have to be defeated. But k is an inference ancestor of h . By *Inheritance*, h would have to be defeated after all. Now, I think there are plausible rules showing that if the inference ancestors of k are undefeated, and if its only defeaters are l and h , then it is undefeated. But I do not really need to make this assumption. For if k was defeated, then by *Inheritance*, all of i_1, i_2, i_3, j_1, j_2 , and j_3 (and l) would be defeated, since k is an inference ancestor of all of them. What remains to be shown is that these nodes are defeated even on the assumption that k is undefeated. So let us assume that k is undefeated.

Showing that i_1, i_2 , and i_3 are all defeated in G_L relies on the symmetry assumption. After all, it is perfectly consistent with the rules stated above that, for instance, i_1 is undefeated, i_2 is undefeated, i_3 is defeated, j_1 is defeated, j_2 is defeated, and j_3 is undefeated. But by the symmetry assumption, either all or none of i_1, i_2 , and i_3 are defeated. Assume that they are all undefeated. Consider node j_1 . By *No Defeaters*, j_1 is undefeated, since all of its inference ancestors, k, i_2 , and i_3 , are undefeated. This would make j_1 an undefeated defeater for i_2 . By *Undefeated Defeater*, i_2 would have to be defeated after all, which contradicts the assumption made. Hence, i_1, i_2 , and i_3 must all be defeated. Note that by *Inheritance*, it also follows that node l is defeated. This gives the required result: a structuralist view that accepts WLP is committed to a sceptical treatment of lottery cases that have the structure represented by G_L , namely, a treatment on which a subject is not justified in believing of any of the tickets that it will lose.

I hope that the judgment that given WLP, all of nodes $i_1, i_2, i_3, j_1, j_2, j_3$, and l in G_L are defeated was plausible without further defence. But we saw how any set of rules for computing defeat statuses based on WLP would have to subscribe to principles that yield this result. Below I look at how this treatment of the Lottery Paradox threatens to generalise to the Paradox of the Preface. However, before doing so, it is worth raising an immediate, and to my mind very serious overgeneralisation problem that this treatment of the Lottery Paradox has. The problem is that it looks as though the same structure of collective defeat can be generated for any proposition whatsoever.

Let q be a proposition that a subject ought to have a justification to believe. Assume that the degree to which the subject is justified in believing q is δ , where δ is close to 1. Let a *minimally inconsistent* set of propositions be a set of propositions that is inconsistent but has no inconsistent proper subset. For instance, the following set of propositions is minimally inconsistent: $\{t_1 \vee t_2 \vee t_3, \sim t_1, \sim t_2, \sim t_3\}$. Assume that $\{p_1, \dots, p_n\}$ is such a minimally inconsistent set of propositions, and that the subject has a reason to believe each proposition in the set (as in the Lottery Paradox). We might assume, for instance, that the subject knows that each of these propositions is very probable. Then, each of the disjunctions $\sim q \vee p_1, \sim q \vee p_2, \dots, \sim q \vee p_n$ is likewise at least as probable and hence, the subject has at least as strong a reason to believe each of these disjunctions. But the set of propositions

$\{\sim q \vee p_1, \sim q \vee p_2, \dots, \sim q \vee p_n, q\}$ consisting of q and these disjunctions is itself minimally inconsistent: for any subset that contains all but one member of the set, that subset entails the negation of that member. In particular, the propositions in set $\{\sim q \vee p_1, \sim q \vee p_2, \dots, \sim q \vee p_n\}$ entail $\sim q$. We have generated a defeater for q . The situation is formally analogous to the Lottery Paradox. The problem is that q can be any proposition whatsoever.²⁹

This problem strikes me as very serious. Pollock attempts to circumvent it by arguing that the sorts of disjunctions employed in the above argument are not projectable in the suitable sense: even if a subject has a probabilistic *prima facie* reason to believe each of p_1, \dots, p_n , and the probability of any disjunction $\sim q \vee p_i$ is at least as high as that of p_i , it does not follow that the subject has a probabilistic reason to believe the disjunction.³⁰ I am not convinced by this response, but below I discuss a separate worry to the effect that the above sceptical treatment of the Lottery Paradox threatens to generalise to the Paradox of the Preface.

6 The Paradox of the Preface

It is not at all uncommon for a subject to be in a situation in which she has ample reasons to believe each proposition in a set of propositions $\{p_1, \dots, p_n\}$, while also having what seems like a very strong reason for believing that at least one of these propositions is false. For instance, assume that I write a book titled “1000 Hard Facts About Global Warming”. I have done years of research and consulted top experts. I have ample, strong evidence for each claim p_i expressing one of the (putative) facts, none of them being particularly controversial. It might nevertheless seem rational for me to believe on inductive or statistical grounds that at least one of the claims being made is false, and to prefix my book with a Preface asserting this.³¹ Now, the problem is that such a situation looks to be structurally similar to the Lottery Paradox. But if the above treatment of the Lottery Paradox was applied to Preface Paradox type cases, the result would be a far-reaching scepticism: whenever a subject has a reason to believe that not every belief in a large set of beliefs is true, none of the beliefs in that set are justified.

Given a large set of propositions $\{p_1, \dots, p_n\}$ each of which is made very likely by a subject’s evidence, exactly what sort of reason does a subject have for believing that at least one proposition in that set is false? Following Pollock, I shall assume that it is typically a statistical reason of the following sort: given any set X of propositions or beliefs that has some feature **F**, there is a certain statistical probability that not every proposition in X is true. Such a reason can be expressed as follows³²:

²⁹ This is a problem that Pollock (1995, pp. 64–65) raises for his own account.

³⁰ See Pollock and Cruz (1999, pp. 231–233), Pollock (1995, pp. 64–67).

³¹ See Makinson (1965).

³² Throughout the discussion of the statistical syllogism and its relevance for the Paradox of the Preface, I am using notation that is essentially the same that Pollock uses. See, for instance, Pollock (1995, p. 125, § 9).

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X)) \geq \delta.$$

For instance, \mathbf{F} might be the feature that a set of propositions X has just in case X is large enough, the propositions in X are contingent, they are at least logically independent, and the evidence supporting them is of a certain kind. By combining the above with a premise stating that a particular set of propositions has feature \mathbf{F} , a subject can reason as follows:

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X)) \geq \delta \ \& \ \mathbf{F}(\{p_1, \dots, p_n\})$$

Therefore,

$$\exists p_i(p_i \in \{p_1, \dots, p_n\} \ \& \ \sim p_i).$$

The conclusion is, of course, equivalent to $\sim(p_1 \ \& \ \dots \ \& \ p_n)$. This is an instance of a kind of inference that Pollock calls the *statistical syllogism*.³³ With Pollock, I will assume that the strength of such an inference (and hence, of the appropriate inference link within an inference graph) is δ . Given WLP, it follows that as long as the degree to which the subject is justified in believing the premise is also at least as high as δ , she is justified in believing the conclusion to degree δ .


Again, to simplify things, consider a case involving just three propositions p_1 , p_2 and p_3 , and let the threshold level for justification be $2/3$.³⁴ Assume that a subject's reasons for believing p_1 , p_2 , and p_3 , respectively are e_1 , e_2 , and e_3 . Assume that the degree to which she is justified in believing each of e_1 , e_2 , and e_3 is above threshold, that none of them have any defeaters, and that the degree to which each e_i supports the corresponding p_i is $2/3$. The subject's reason for believing the proposition $\sim(p_1 \ \& \ p_2 \ \& \ p_3)$ is a statistical reason of the sort described above. Just as in the Lottery Paradox case described, all of the following inferences are acceptable, since they are all deductive:

$$\begin{aligned} &\sim(p_1 \ \& \ p_2 \ \& \ p_3), p_1, p_2; \text{ Therefore, } \sim p_3. \\ &\sim(p_1 \ \& \ p_2 \ \& \ p_3), p_2, p_3; \text{ Therefore, } \sim p_1. \\ &\sim(p_1 \ \& \ p_2 \ \& \ p_3), p_1, p_3; \text{ Therefore, } \sim p_2. \\ &p_1, p_2, p_3; \text{ Therefore, } p_1 \ \& \ p_2 \ \& \ p_3. \end{aligned}$$

This generates a situation that is structurally very similar to the simple Lottery Paradox type case examined above. Let G_p be the inference graph so far described. It can be represented as follows:

³³ See, for instance, Pollock and Cruz (1999, pp. 229–234).

³⁴ In the lottery case discussed, r , a description of the structure of the lottery, provides a subject with both a reason to believe each of $\sim t_1$, $\sim t_2$, and $\sim t_3$ stating of individual tickets that they will lose, and a deductive reason to believe that not all of $\sim t_1$, $\sim t_2$, and $\sim t_3$ are true. The kind of Preface Paradox situation I will describe is one in which a subject has independent reasons for believing each of p_1, \dots, p_n , and a yet separate reason of the kind just discussed for believing that not all of p_1, \dots, p_n are true, simply because such situations are more typical. However, this is not essential to Preface Paradox cases.



6.1 Graph G_P

Again, lines without dashes represent inference links, and dashed lines represent defeat links. Corresponding to node m is a belief in the following proposition: $\Pr(\exists p_i(p_i \in X \& \sim p_i) \mid \mathbf{F}(X)) \geq 2/3 \& \mathbf{F}(\{p_1, p_2, p_3\})$.

But now the obvious worry is that the same sort of reasoning that was used to show that i_1, i_2, i_3, j_1, j_2 , and j_3 are all defeated in G_L can be used to show that $i_1^*, i_2^*, i_3^*, j_1^*, j_2^*$, and j_3^* are all defeated in G_P . If G_P represented the structure of a Preface Paradox situation, this would be a rather devastating result, leading to a far-reaching scepticism. In the next section I discuss a disanalogy between the Lottery and Preface Paradoxes (in fact, what I see as the only possible structural disanalogy). In particular, Pollock argues that graph G_P does not represent the full structure of a Preface Paradox situation: the graph does not record all of the relevant defeaters, defeaters for node k^* , which is an inference ancestor of all of j_1^*, j_2^* , and j_3^* . I argue that this attempt to block a sceptical treatment of the Paradox of the Preface has a very unhappy consequence, namely, treating different types of lotteries differently. In some lotteries, a subject will be justified in believing that her ticket will lose purely based on probabilistic considerations, and in others she will not, even if the chance of her ticket losing is at least as high. But let me first discuss Pollock’s solution to the Paradox of the Preface.

7 A Proposed Solution to the Paradox of the Preface

Why might one think that G_P does not represent the full structure of the Preface Paradox type situation described? Recall that in a Preface case, a subject typically reasons to the conclusion that at least one proposition in a set of propositions $\{p_1, \dots, p_n\}$ is false by employing what was termed the statistical syllogism. In such an inference, a subject infers that some thing a has a feature \mathbf{G} based on a having some feature \mathbf{F} and the probability of an arbitrary thing having \mathbf{G} conditional on its having \mathbf{F} being above some level δ . Such an inference is defeated by learning that a also has some further feature \mathbf{H} , and that conditional on something having both \mathbf{F} and \mathbf{H} ,

it is less likely to have **G**.³⁵ For instance, if all I know about Peggy is that she is of a species of animal that cannot fly, and I know that the probability of not being a bird conditional on not being able to fly is high, then I might be justified in believing that Peggy is not a bird. However, if I then learn that Peggy lives in Antarctica, I am no longer justified in believing that Peggy is not a bird.

Recall that in the Preface Paradox case described, the subject reasons as follows:

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X)) \geq 2/3 \ \& \ \mathbf{F}(\{p_1, p_2, p_3\})$$

Therefore,

$$\sim(p_1 \ \& \ p_2 \ \& \ p_3).$$

Something of the following form would be a defeater for this inference:

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X) \ \& \ \mathbf{H}(X)) < 2/3 \ \& \ \mathbf{F}(\{p_1, p_2, p_3\}) \ \& \ \mathbf{H}(\{p_1, p_2, p_3\}).$$

What might feature **H** be?

The Lottery Paradox case described above involved a simple lottery in which one of the tickets is more or less guaranteed to be drawn. It follows that propositions $\sim t_1$, $\sim t_2$, and $\sim t_3$ are negatively probabilistically relevant, since the probability of one ticket losing conditional on any number of the others losing is lower than its unconditional probability of losing.³⁶ For instance, the unconditional probability of ticket #1 losing is 2/3, but the probability of ticket #1 losing conditional on ticket #2 losing (and the lottery having the structure assumed) is just 1/2. However, one could argue that this is typically not true of Preface cases. Take, for instance, a random proposition I assert in my book “1000 Hard Facts About Global Warming”. Rather than making the rest *less* likely to be true, would not a more typical situation be one in which it makes the rest of the claims in my book *more* likely to be true? Pollock argues that this provides the key to just the sort of feature **H** we are looking for to defeat the inference to $\sim(p_1 \ \& \ p_2 \ \& \ p_3)$: though the probability of a set containing a false member conditional on its having feature **F** is sufficiently high, the probability of the set containing a false member conditional on its having feature **F** and (at least) $n - 1$ of its members being true is not.

Let **H** be a feature that a set of propositions X has just in case at least $n - 1$ of its n members are true. Now take set $\{p_1, p_2, p_3\}$. A reason to believe any of $p_1 \ \& \ p_2$, $p_2 \ \& \ p_3$, or $p_1 \ \& \ p_3$ is a reason to believe that set $\{p_1, p_2, p_3\}$ has feature **H**. But inferences to each of these propositions from their individual conjuncts are acceptable; hence, three different inferences to the conclusion that set $\{p_1, p_2, p_3\}$ has feature **H** are acceptable. We can also add to a graph representing the structure of the Preface Paradox situation described a node corresponding to the following belief—let this be node g ³⁷:

³⁵ Generally, $[r \ \& \ \Pr(p|q \ \& \ r) \neq \Pr(\Pr(p|q))]$ is an undercutting defeater for the reason for believing p provided by $[q \ \& \ \Pr(p|q) \geq r]$. See Pollock (1995, pp. 66–67).

³⁶ Note that this is also true of lotteries in which there is guaranteed to be *at least* (but maybe not exactly) one winning ticket.

³⁷ I am bracketing the issue of what the inference parents of this node are, just as I bracketed the issue of the inferential history of the node supporting the belief in the following proposition: $\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X)) \geq 2/3$.

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X) \ \& \ \mathbf{H}(X)) < 2/3.$$

Then, an inference (or three different inferences) to the following proposition are acceptable:

$$\Pr(\exists p_i(p_i \in X \ \& \ \sim p_i) \mid \mathbf{F}(X) \ \& \ \mathbf{H}(X)) < 2/3 \ \& \ \mathbf{F}(\{p_1, p_2, p_3\}) \ \& \ \mathbf{H}(\{p_1, p_2, p_3\}).$$

Let the node corresponding to this belief be o . But by what was said above about defeaters for inferences employing the statistical syllogism, o is a defeater for node k^* , which corresponds to the belief $\sim(p_1 \ \& \ p_2 \ \& \ p_3)$.

The thought is that the defeater for node i_1^* , node j_1^* , has i_2^* , i_3^* , and k^* as its inference parents, but two of its inference parents, i_2^* and i_3^* , can be used to construct a defeater for the third, k^* . Similarly, j_2^* is the defeater for i_2^* , and has as its inference parents nodes i_1^* , i_3^* , and k^* . But i_1^* and i_3^* can be used to construct a defeater for k^* . Finally, j_3^* is the defeater for i_3^* , and two of the inference parents of j_3^* , i_1^* , i_2 , can be used to construct a defeater for its third inference parent k^* . In a graph representing the full structure of the Preface Paradox case discussed, k^* is defeated. Then, j_1^* , j_2^* , and j_3^* all have a defeated inference ancestor and by *Inheritance*, are defeated. And then, they ought not to have the power to defeat i_1^* , i_2^* , and i_3^* . I will not attempt to draw the resulting graph.

Hence, a case can be made for thinking that in an inference graph representing the full structure of the Preface Paradox case described, nodes i_1^* , i_2^* and i_3^* are defeated. But for my purposes, I do not need to show that this follows from rules that anyone committed to WLP will have to accept. For whether or not this is right, the non-probabilistic structuralist wedded to WLP is in trouble. If this was not the correct assignment of defeat-statuses to i_1^* , i_2^* and i_3^* , then the version of structuralism committed to WLP would have a very serious sceptical consequence that I would take to be a *reductio* of the view. I will now argue that if it is correct that in an inference graph representing the structure of the Preface Paradox case described nodes i_1^* , i_2^* and i_3^* are undefeated, there is trouble nevertheless, since some lotteries have a relevantly similar structure to such cases.

8 Some Odd Results

Not all lotteries are like the ones discussed in connection with the Lottery Paradox. In the Lottery Paradox, a subject's reason for believing of each of the n tickets in the lottery that it will lose consisted of the lottery having the following structure: there are n tickets, exactly one ticket will be drawn, and each ticket has an equal chance of being drawn. Call such lotteries *standard lotteries*. Here is a structure that is an example of what I will term a *non-standard* lottery. Take n distinct and independent standard lotteries occurring in different parts of the world. Take ticket #1 from the first standard lottery, ticket #2 from the second standard lottery, and so forth. These will make up the tickets of the non-standard lottery. For any ticket # i , ticket # i wins the non-standard lottery just in case it wins the i th standard lottery; otherwise it loses the non-standard lottery. Now, provided that n is large enough, there will be a high

chance that at least one of tickets #1 – # n wins. Moreover, in such a lottery, the losing of any number of tickets is not negatively relevant to whether other tickets lose. It is simplest to just assume complete probabilistic independence.³⁸


I will now say why non-standard lotteries exhibit the structure of Preface-cases. However, in doing so I will relax the constraints on the sorts of inference graphs that I am interested in. This is to avoid having to construct a non-standard lottery in which the probability of each ticket losing equals the probability that at least one of the tickets will win.³⁹ It will suffice to give the structure of a non-standard lottery, a structure that is relevantly similar to that of the Preface case. For instance, we might look at cases in which a subject has a stronger reason to believe that at least one of the tickets in the lottery will win than she has to believe of any of the tickets that it will lose. An analogous situation is not at all uncommon in Preface cases: given enough propositions, the probability that at least one of them is false might be higher than the probability of any of the individual propositions. A sceptical treatment of such situations would be very worrying, since stringing together enough suitably independent propositions each of which is individually highly probable will result in a conjunction with an even higher probability of containing at least one false member. But a non-sceptical treatment would mean that a subject is justified in believing of each ticket in a non-standard lottery that it will lose, even if the probability that at least one will win is higher than in a standard lottery with an equal number of tickets!

Let $\sim t_1$, $\sim t_2$, and $\sim t_3$ be propositions stating of individual tickets in a non-standard lottery that they will lose. Let p be a proposition giving a description of the lottery. I will assume that the lottery is of the kind described above: each ticket of the non-standard lottery is part of a standard lottery with 3 tickets, and it is a winner of the non-standard lottery just in case it wins the standard lottery that it participates in.⁴⁰ Assume that the strengths of all of the inference ancestors and ancestral inference links of p are above threshold, and that p supports each of $\sim t_1$, $\sim t_2$, and $\sim t_3$ to a degree above the threshold. Here is an inference graph representing the structure of this case:

³⁸ There are numerous other possible forms of lotteries where there is negative relevance, but of a degree lesser than in standard lotteries.

³⁹ However, there are such lotteries, since the equation $1 - r^n = r$ has real solutions.

⁴⁰ To be more precise, we might assume p to state the following facts: There are three standard lotteries, and each of these has exactly three tickets. t_1 is true just in case ticket #1 wins the first standard lottery, t_2 is true just in case ticket #2 wins the second standard lottery, and t_3 is true just in case ticket #3 wins the third standard lottery. The outcomes of the standard lotteries are independent. Then, for each $\sim t_i$, p makes $\sim t_i$ probable to degree $2/3$ (≈ 0.67). This is the strength of the relevant inference link. The probability that at least one of the tickets will win is given by $1 - (2/3)^3$ (≈ 0.70). This would be a case in which p supports the proposition that at least one of the tickets will win to a higher degree than it supports any of the propositions stating of the individual tickets that they will lose. Again, set the minimal threshold level for justification at $2/3$.



8.1 Graph G_N

Note that though n is an inference parent of both k^* and of each of i_1^* , i_2^* , and i_3^* , k^* is *not* an inference parent of any of i_1^* , i_2^* , and i_3^* , and neither is the inference-link between n and k^* deductive. We can assume the inference from n to k^* to be the sort of statistical, probabilistic inference discussed in connection with the Preface case. In this non-standard lottery situation, defeaters can be generated for node k^* in the same way as in the Preface case.⁴¹ For, conditional on the description of the lottery, the probability that at least one of the propositions in set $\{\sim t_1, \sim t_2, \sim t_3\}$ is false is above threshold, but conditional on the description of the lottery *and* two of the members of $\{\sim t_1, \sim t_2, \sim t_3\}$ being true, this probability is below threshold.

The simple situation I have described is one in which the strengths of the inference links $\langle n, i_1^* \rangle$, $\langle n, i_2^* \rangle$, and $\langle n, i_3^* \rangle$ are weaker than the strength of the inference link $\langle n, k^* \rangle$. I have said nothing about how a set of rules for assigning strengths to the nodes of an inference graph should treat such cases. However, the point is just that the structure of graph G_N is analogous to the structure of graph G_P , and perfectly identical to the structure of a graph representing some Preface situations. For any non-standard lottery case of the sort described above, it is possible to describe a Preface case in which the strengths of the analogous inference-links are just the same, and treating such non-standard lottery cases differently from the corresponding Preface cases would be wholly unmotivated.

⁴¹ This inference graph is identical to a graph representing a Preface situation in which the subject's reason for believing each of three propositions p_1, p_2 , and p_3 and her reason for believing $\sim(p_1 \& p_2 \& p_3)$ is the same. Assume that this reason consists of believing that the set of propositions $\{p_1, \dots, p_n\}$ has the following feature **F**: each proposition is supported to degree $2/3$ by the evidence and the propositions are probabilistically independent. Then, the set having feature **F** is a reason to believe each of p_1, p_2 , and p_3 , and it is a reason to believe $\sim(p_1 \& p_2 \& p_3)$.

The upshot of all this is the following dilemma. One horn is a sceptical treatment of at least some Preface cases. Because Preface cases are only too easy to construct, such scepticism should be a last resort. The other horn is treating different types of lotteries differently. If a subject holds a ticket in a lottery she is justified in believing to be a standard one, she is not justified in believing of any of the participants in the lottery that they will lose, but if she holds a ticket in a lottery she is justified in believing to be a non-standard one she *is* justified in believing of each of the participants that they will lose. This is so even if she has a lower chance of losing in the non-standard lottery than she does in the standard lottery. Moreover, given WLP, in independent lotteries, the subject is justified in believing that *none* of the participants will win, despite the fact that there is a high chance that at least one will win, and the subject knows this. This all looks very awkward.

In effect, Pollock himself uses the following argument against a Kyburg-style treatment of the Lottery Paradox on which a subject is justified in believing of each ticket that it will lose: If the subject was justified in believing of each ticket that it will lose, then she could reason to the conclusion that no matter what the price, she should not buy a lottery ticket.⁴² But it clearly *is* rational to buy a lottery ticket. The problem is that Pollock's own account has the same consequence for non-standard lotteries. In such lotteries, a subject is justified in believing of each ticket that it will lose. Then, if she is offered the possibility of buying a ticket, she can reason to the conclusion that she should not buy the ticket no matter what the price. But surely it is just as rational to buy a ticket for a standard lottery for a certain price as it is to buy a ticket for an independent lottery for that price, provided that the chances of both tickets winning are the same! Moreover, the fact that the subject is justified in believing that *no* ticket will win is in itself extremely implausible, given that there is a very high chance that at least one ticket will win. In particular, situations in which a subject knows that a proposition *p* is extremely likely to be false tend to elicit defeat-intuitions: should not the subject's knowledge that it is extremely likely that at least one ticket will win defeat her justification for believing that no ticket will win? Hence, the additional worry is that the resulting picture cannot take into account intuitions about defeat after all.

The peculiarity of distinguishing between standard and independent lotteries is made most pressing when we consider the fact that one and the same ticket can be part of both a standard and an independent lottery. In effect, this was true of the non-standard lottery described. Each ticket #*i* was part of a standard lottery, winning the non-standard lottery just in case it won the standard lottery. What should the structuralist wedded to WLP say about such cases? Presumably, when considering ticket #1, for instance, as part of the standard lottery, a subject is not justified in believing that it will lose. But by the above account, is not she justified in believing that it will lose when she considers it as part of the non-standard lottery, even if the ticket loses the non-standard lottery just in case it loses the standard lottery? This seems absurd.

⁴² See Pollock (1995, pp. 61–62).

9 Conclusions

I began with some widely accepted and at first sight intuitive ideas about well-founded or doxastically justified belief. The basic thought was that for a belief to be justified, it must be based on adequate, undefeated reasons. I sketched a formal framework within which this basic structuralist insight could be made more precise. I looked at two approaches to computing degrees of justification within this framework. I argued that a probabilistic approach is difficult to reconcile with the basic tenets of a structuralist way of thinking. I then looked at the alternative, non-probabilistic approach committed to the WLP and justification closure. I argued that a view based on WLP must give a sceptical treatment of the Lottery Paradox, a treatment on which a subject is not justified in believing of any of the individual tickets in the lottery that it will lose. The problem is that either this treatment of the Lottery Paradox overgeneralises into a widespread scepticism about justification, or else unmotivated, epistemically spurious distinctions have to be drawn between different types of lotteries. All this casts severe doubt on the ability of the non-probabilistic structuralist framework sketched to provide a theory of doxastic justification. Given that both probabilistic and non-probabilistic varieties of structuralism appear to be untenable, this gives at least some reason to question the coherence of the very idea of well-founded belief, and of the role that the resulting notion of doxastic justification ought to play in epistemology. To say the very least, the burden is on those interested in the idea of well-founded belief to develop a viable theory.

References

- Annis, D. (1978). Knowledge and defeasibility. In G. S. Pappas & M. Swain (Eds.), *Essays on knowledge and justification* (pp. 155–159). London: Cornell University Press.
- Chisholm, R. (1964). The ethics of requirement. *American Philosophical Quarterly*, 1, 147–153.
- Enderton, H. B. (1977). *Elements of set theory*. London: Academic Press.
- Feldman, R., & Conee, E. (2004). *Evidentialism*. Oxford: Clarendon Press. (Reprinted, with Afterword, in E. Conee & R. Feldman, *Evidentialism* (pp. 83–107).
- Lehrer, K., & Paxson, T. D., Jr. (1978). Knowledge: Undefeated justified true belief. In G. S. Pappas & M. Swain (Eds.), *Essays on knowledge and justification* (pp. 146–154). London: Cornell University Press.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25, 205–207.
- Pollock, J. L. (1974). *Knowledge and justification*. Princeton: Princeton University Press.
- Pollock, J. L. (1995). *Cognitive carpentry: A blueprint for how to build a person*. Cambridge, MA: MIT Press.
- Pollock, J. L. (2001). Defeasible reasoning with variable degrees of justification. *Artificial Intelligence*, 133, 233–282.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Lanham, MD: Rowman & Littlefield.
- Pryor, J. (2004). What's wrong with Moore's argument? *Philosophical Issues*, 14, 349–378.