

JON WILLIAMSON

AN OBJECTIVE BAYESIAN ACCOUNT OF CONFIRMATION

ABSTRACT

This paper revisits Carnap's theory of degree of confirmation, identifies certain shortcomings, and argues that a new approach based on objective Bayesian epistemology can overcome these shortcomings.

Rudolf Carnap can be thought of as one of the progenitors of Bayesian confirmation theory (§1). Bayesian confirmation theory is construed in §2 as a four-step process, the third step of which results in the identification of the degree to which e confirms h , $c(h, e)$, with the probability of h conditional on e in the total absence of further evidence, $P_{\emptyset}(h|e)$. The fourth step of this process involves isolating an appropriate candidate for P_{\emptyset} ; Carnap rejected the most natural construal of P_{\emptyset} on the grounds that it leads to a confirmation function c^{\dagger} that fails to adequately capture the phenomenon of learning from experience (§3). This led him, and subsequent confirmation theorists, to more elaborate interpretations of P_{\emptyset} , resulting in certain continua of confirmation functions (§§4, 5). I argue in §§5, 6 that this was a wrong move: the original construal of P_{\emptyset} is in fact required in order that degree of confirmation can capture the phenomenon of partial entailment. There remains the problem of learning from experience. I argue that this problem is best solved by revisiting the third—rather than the fourth—step of the four-step Bayesian scheme (§7) and that objective Bayesianism, which is outlined in §8, offers the crucial insight as to how this step can be rectified. This leads to an objective Bayesian confirmation theory that can capture both partial entailment and learning from experience (§9).

§1 CARNAPIAN CONFIRMATION

Our current understanding of confirmation owes much to Rudolf Carnap's pioneering work of the 1940s and beyond. Carnap (1950, §8) distinguishes three concepts of confirmation: a classificatory concept which applies when evidence e qualitatively confirms a hypothesis h , a comparative concept which applies when h is confirmed by e at least as highly as h' by e' , and a quantitative concept according to which h is confirmed by e to degree q , written $c(h, e) = q$. Carnap also distinguishes two principal notions

of probability: probability₁, or degree of confirmation, and probability₂, or relative frequency (Carnap, 1950, §9). Carnap was chiefly concerned with the quantitative concept of confirmation (i.e., probability₁) and we will likewise restrict our attention here to this quantitative notion.

For Carnap, as for Keynes (1921) before him, this notion of probability is fundamentally a logical relation between a body of evidence and a proposition. It is clear that Carnap viewed this relation as objective, not as an expression of subjective degree of belief. On the other hand, Keynes was emphatic that the logical concept of probability underwrites *rational* degrees of belief; Carnap went along with this view but was more ambivalent:

Many logicians prefer formulations which may be regarded as a kind of *qualified psychologism*. They admit that logic is not concerned with the actual processes of believing, thinking, inferring, because then it would become a part of psychology. But, still clinging to the belief that there must somehow be a close relations between logic and thinking, they say that logic is concerned with correct or rational thinking. Thus they might explain the relation of logical consequence as meaning: ‘if somebody has sufficient reasons to believe in the premise *i*, then the same reasons justify likewise his belief in *j*.’ It seems to me that psychologism thus diluted has virtually lost its content; the word ‘thinking’ or ‘believing’ is still there, but its use seems gratuitous. . . . The characterization of logic in terms of correct or rational or justified belief is just as right but not more enlightening than to say that mineralogy tells us how to think correctly about minerals. The reference to thinking may just as well be dropped in both cases. (Carnap, 1950, pp. 41–42)

Some years later, however, Carnap came to be less ambivalent and took the rational degree of belief approach more seriously (see, e.g., Carnap, 1971). It is fair to say, then, that while Carnap cannot be considered an advocate of what is now called the *Bayesian interpretation of probability*, which takes probability to be *fundamentally* interpretable in terms of rational degree of belief, he can be considered to be a pioneer of what is now called *Bayesian confirmation theory*, which typically admits an identity (whether fundamental or not) between degree of confirmation and rational degree of belief and which proceeds along the following lines.

§2 THE BAYESIAN APPROACH TO CONFIRMATION

The Bayesian approach to confirmation might broadly be characterised in terms of the following four steps.

Step 1. Consider probability functions defined over a language \mathcal{L} .

Step 2. Identify $c(h, e) = P_{\{e\}}(h)$ for some suitable probability function P on \mathcal{L} , where $P_{\{e\}}(h)$ is the probability of h on evidence e .

Step 3. Identify $P_{\{e\}}(h) = P_{\emptyset}(h|e)$.

Step 4. Find an appropriate P_{\emptyset} that represents confirmation in the absence of evidence.

Let us examine these four steps in turn.

Step 1. Consider probability functions defined over a language \mathcal{L} .

Degree of confirmation is taken to be a relation between evidence and a hypothesis and these are naturally construed as propositions (or sometimes, in the case of evidence, sets of propositions). Hence the functions we need to consider—confirmation functions and probability functions—should be defined on propositions. But probability functions are normally defined on events construed as sets of possible outcomes (Kolmogorov, 1933). One of Carnap’s important contributions to this area was his work on defining probability functions on logical languages as opposed to fields of sets. One might, for example, consider a propositional language $\mathcal{L} = \{A_1, \dots, A_n\}$ on elementary propositions A_1, \dots, A_n , with compound propositions formed by the usual connectives $\neg, \vee, \wedge, \rightarrow, \leftrightarrow$. The set of *atomic states* of \mathcal{L} is defined as $\Omega_n = \{\pm A_1 \wedge \dots \wedge \pm A_n\}$, where $+A_i$ is just A_i and $-A_i$ is $\neg A_i$. A probability function on \mathcal{L} is then a function P , from propositions of \mathcal{L} to real numbers, that satisfies the properties:

P1. $P(\omega) \geq 0$ for each $\omega \in \Omega_n$,

P2. $P(\tau) = 1$ for some tautology τ , and

P3. $P(\theta) = \sum_{\omega \models \theta} P(\omega)$ for each proposition θ .

Alternatively one might consider a predicate language rather than a propositional language. There are various ways of proceeding here, but perhaps the simplest goes as follows (see Williamson, 2010b, Chapter 5). Construe a predicate language as $\mathcal{L} = \{A_1, A_2, \dots\}$ where the A_i enumerate the atomic propositions of the form Ut for some predicate U and tuple t of constant symbols. (There is assumed to be a constant symbol for each domain individual.) A finite sublanguage $\mathcal{L}_n = \{A_1, \dots, A_l\}$ uses only constant symbols t_1, \dots, t_n . The set of atomic states of \mathcal{L}_n is $\Omega_n = \{\pm A_1 \wedge \dots \wedge \pm A_l\}$. A probability function on a predicate language \mathcal{L} is then a function from propositions of \mathcal{L} to real numbers that satisfies the properties:

PP1. $P(\omega) \geq 0$ for each $\omega \in \Omega_n$ and each n ,

PP2. $P(\tau) = 1$ for some tautology τ ,

PP3. $P(\theta) = \sum_{\omega \in \Omega_n, \omega \models \theta} P(\omega)$ for each quantifier-free proposition θ , where n is large enough that \mathcal{L}_n contains all the atomic propositions occurring in θ , and

PP4. $P(\exists x\theta(x)) = \sup_m P(\bigvee_{i=1}^m \theta(t_i))$.

Note in particular that a probability function P on predicate language \mathcal{L} is determined by its values on the ω for $n = 1, 2, \dots$ (see, e.g., Paris, 1994, Theorem 11.2). PP4 is known as *Gaifman's condition*, and PP1–4 imply that $P(\exists x\theta(x)) = \lim_{m \rightarrow \infty} P(\bigvee_{i=1}^m \theta(t_i))$ and $P(\forall x\theta(x)) = \lim_{m \rightarrow \infty} P(\bigwedge_{i=1}^m \theta(t_i))$.

Step 2. Identify $c(h, e) = P_{\{e\}}(h)$ for some suitable probability function P on \mathcal{L} , where $P_{\{e\}}(h)$ is the probability of h on evidence e .

For Carnap, this step is just his explication of the quantitative concept of confirmation in terms of probability₁. One thing that makes Bayesian confirmation theory *Bayesian* is that the probability of h on evidence e is, in turn, interpretable as the degree to which one should believe h if one were to grant just e . (Bayes (1764) wrote of ‘expectation’ for belief or credence.) It should be reiterated that the proponent of a logical interpretation of probability, such as Keynes or Carnap, would want to say that this Bayesian construal of the probabilities is derivative rather than fundamental: that $P_{\{e\}}(h) = x$ means that there is a logical probability-relation between $\{e\}$ and h of degree x , and it is this fact that makes it rational to believe h to degree x if one were to grant just e . Proponents of a Bayesian interpretation, on the other hand, would take the rational degree of belief interpretation as fundamental. According to subjective Bayesianism, x largely depends on the whim of the agent in question, while according to objective Bayesianism, the agent’s evidence plays the leading role in determining x .

Step 3. Identify $P_{\{e\}}(h) = P_\emptyset(h|e)$.

It is usual for the Bayesian to identify a conditional belief with a conditional probability: the degree to which one should believe h if one were to grant just e is identified with the probability of h conditional on e (granting nothing at all). As with other rules of Bayesian probability, the justification for such a move normally proceeds via the betting interpretation of degrees of belief. In this case, the degree to which one should believe h if one were to grant just e is first interpreted in terms of a certain conditional bet and then it is shown that under this interpretation the identity posited in Step 3 must hold.

The argument proceeds as follows. Interpret $P_{\{e\}}(h) = q$ as saying that one is prepared to offer a betting quotient q for h (i.e., one is prepared to bet qS for a return of S if h is true), with the bet called off if e is false. (The stake S depends on the betting quotient and may be positive or negative.) The loss one incurs on such a bet is $I_e(q - I_h)S$, where I_θ is the indicator function for proposition θ , which takes the value 1 if θ is true and 0 if θ is false. If one also offers betting quotient $P(h \wedge e) = q'$ on $h \wedge e$ and $P(e) = q''$ on e then one’s total loss is

$$I_e(q - I_h)S + (q' - I_e I_h)S' + (q'' - I_e)S''.$$

If $q' < qq''$ then the stake-maker can choose $S' = -S = 1$ and $S'' = q$ to ensure certain loss $qq'' - q'$. Similarly if $q' > qq''$ the stake-maker can choose $S' = -S = -1$ and $S'' = -q$ to ensure certain loss $q' - qq''$. Hence unless $P_{\{e\}}(h)P(e) = P(h \wedge e)$ one can be made to lose money whatever happens. But if $P_{\{e\}}(h)P(e) = P(h \wedge e)$ then one avoids the possibility of sure loss, for the following reason. The expected loss is

$$qq'S' + q''S'' + q'(-S - S') + q''(qS - S'') = (q''q - q')S$$

but this is zero if $q''q - q' = 0$; if the expected loss is zero then the loss cannot be positive in every eventuality. So one avoids the possibility of sure loss if and only if $q''q - q' = 0$. Granting that avoiding the possibility of sure loss is a requirement of rationality, the identity $P_{\{e\}}(h) = P(h \wedge e)/P(e) = P(h|e)$ must hold for rational degrees of belief (as long as $P(e) \neq 0$). Assuming finally that e exhausts the available evidence, $P(h \wedge e) = P_{\emptyset}(h \wedge e)$ and $P(e) = P_{\emptyset}(e)$, and Step 3 follows.

Step 3 proposes the use of conditional probabilities in the explication of confirmation, and this yields another sense in which the approach can be described as Bayesian. In fact it is often easier to determine the probability of the evidence conditional on the hypothesis than the probability of the hypothesis conditional on the evidence, so Step 3 provides an avenue for Bayes' theorem to enter the picture:

$$c(h, e) = P_{\emptyset}(h|e) = \frac{P_{\emptyset}(e|h)P_{\emptyset}(h)}{P_{\emptyset}(e)}.$$

Although Step 3 proposes the use of conditional probabilities, it should not be confused with the principle of *Bayesian conditionalisation*, which relates degrees of belief at different points in time, and which says: if you adopt belief function P now and you come to learn just e , you should then change your belief function to $P(\cdot|e)$. While someone who endorses Step 3 might well endorse Bayesian conditionalisation and vice versa, they are in fact rather different principles, one dealing with conditional belief and the other with changes of belief. Bayesian conditionalisation is advocated by many proponents of a Bayesian interpretation of probability, but will not be relevant in our context of Bayesian confirmation theory.

Note that Steps 2 and 3 are sometimes conflated. Carnap himself ran the two steps together by making assumptions about c that directly ensure that $c(h, e) = P_{\emptyset}(h \wedge e)/P_{\emptyset}(e)$ (Carnap, 1950, §§53,54B). This is perhaps a mistake; as we shall see below, the key steps must be teased apart if we are to make progress with confirmation theory.

Step 4. Find an appropriate P_{\emptyset} that represents confirmation in the absence of evidence.

This step seems straightforward, although, as we shall see, Carnap had reservations about the following proposal. The natural choice for P_{\emptyset} is

the *equivocator*, $P_{=}$, on \mathcal{L} , i.e., the probability function that equivocates between the atomic states, giving each $\omega \in \Omega_n$ the same probability:

$$P_{\emptyset}(\omega) = P_{=}(\omega) \stackrel{\text{df}}{=} \frac{1}{|\Omega_n|}$$

for all $\omega \in \Omega_n$. (The equivocator can alternatively be defined in terms of models of \mathcal{L} rather than states of \mathcal{L} —see Kemeny (1953).)

Putting the four steps together we then have the recommendation that

$$c(h, e) = P_{=}(h|e).$$

Carnap used the notation c^{\dagger} or c_{∞} for this confirmation function.

Having characterised the two-place confirmation relation it is then usual to define a three-place *support* relation in terms of the confirmation relation (Jeffreys, 1936, p. 421; Good, 1960, pp. 146–147; Gillies, 1990, p. 144).¹ Degree of support $s(h, e, k)$ is supposed to capture the *added* confirmation that e offers to h , over and above the confirmation provided by background k . One possible measure of support is given by $s(h, e, k) = c(h, e \wedge k) - c(h, k)$, but there are many others and little consensus as to which is the most appropriate (see, e.g., Fitelson, 1999). Confusingly, the word ‘confirmation’ is often used to refer both to the two-place relation and to the three-place support relation. In this paper we restrict our attention to the two-place confirmation relation.

§3 LEARNING FROM EXPERIENCE

There is a difficulty with the approach to the problem of confirmation outlined in §2, as Carnap realised very early on in his research (see, e.g., Carnap, 1945, p. 81; Carnap, 1952, p. 38). This is the problem that the resulting choice of confirmation function, c^{\dagger} , renders learning from experience impossible. One can illustrate this general problem via the following example. Suppose that ravens r_1, \dots, r_{101} are being observed to see if they are black (B). Then

$$c^{\dagger}(Br_{101}, \emptyset) = P_{=}(Br_{101}) = \frac{1}{2},$$

where \emptyset represents an empty evidential statement—a tautology, say. This seems right—in the absence of any evidence it seems appropriate to say that Br_{101} and $\neg Br_{101}$ are equally confirmed. However it is also the case that

$$c^{\dagger}(Br_{101}, Br_1 \wedge \dots \wedge Br_{100}) = P_{=}(Br_{101} \mid Br_1 \wedge \dots \wedge Br_{100}) = \frac{1/2^{101}}{1/2^{100}} = \frac{1}{2}.$$

1 Carnap introduces the distinction between confirmation and support in §B.II (p. xvi) of the *Preface to the Second Edition* of Carnap (1950).

Hence, on evidence of the first 100 ravens being black, the degree of confirmation of Br_{101} remains stuck at $\frac{1}{2}$. This inability of evidence to change degree of confirmation is quite unacceptable and the confirmation function c^\dagger should be rejected, Carnap argued.

This problem was in fact recognised by George Boole, who considered drawing balls from an urn containing black and white balls:

It follows, therefore, that if the number of balls be infinite, and all constitutions of the system be equally probable, the probability of drawing m white balls in succession will be $\frac{1}{2^m}$, and the probability of drawing $m+1$ white balls in succession $\frac{1}{2^{m+1}}$; whence the probability that after m white balls have been drawn, the next drawing will furnish a white one, will be $\frac{1}{2}$. In other words, past experience does not in this case affect future expectation. (Boole, 1854, pp. 371–2)

§4 CARNAP'S RESOLUTION

Carnap's strategy for circumventing the problem of learning from experience was to tinker with Step 4 of the four-step scheme of §2: by isolating desiderata that P_\emptyset ought to satisfy, one can narrow down the functional form of P_\emptyset , without narrowing it down so much as to force the identity $P_\emptyset = P_+$ (Johnson, 1932; Carnap, 1952; Paris, 1994, pp. 189–197). Consider the following desiderata:

Constant Exchangeability. P_\emptyset should be invariant under permutations of the constant symbols t_i .

Johnson's Sufficientness Postulate. $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k)$ should depend only on k and the number r_k of positive observations.

It turns out that, for a predicate language with two or more predicates, all unary, there is a continuum of probability functions satisfying Constant Exchangeability and Johnson's Sufficientness Postulate, characterised by:

$$P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k) = \frac{r_k + \lambda/2^m}{k + \lambda},$$

where m is the number of predicates in the language and $\lambda \in [0, \infty]$ is an adjustable parameter, and where instances of different predicates are probabilistically independent. This is known as Carnap's *continuum of inductive methods*; given $\lambda \in [0, \infty]$, the corresponding confirmation function is denoted by c_λ .

Note that this characterisation is also supposed to apply to languages with a single unary predicate. In that case, if $\lambda = 0$ then $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k) = \frac{r_k}{k}$ and c_0 , sometimes called the *straight rule*, sets degrees of confirmation to observed frequencies. If $\lambda = 1$ then $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge$




Figure 1: Carnap's inductive methods for $\lambda \in [0, 50]$, $m = 1$ and $k = 10$.

$\pm Ut_k) = \frac{r_k+1/2}{k+1}$ and c_1 is called the *Jeffreys-Perks' rule of succession*. If $\lambda = 2$ then $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k) = \frac{r_k+1}{k+2}$ and c_2 is known as *Laplace's rule of succession*. If $\lambda = \infty$ then $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k) = 1/2$, and we have $c_\infty = c^\dagger$, the function that fails to admit learning from experience. $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k)$ is depicted in Fig. 1 for $k = 10$ and a range of λ and r_k .

§5 PROBLEMS WITH CARNAP'S RESOLUTION

There are several concerns one might have about Carnap's resolution to the problem of learning from experience; in this section we will consider three.

Determining λ . One question that immediately arises is, how should λ be determined? Carnap himself suggested that the choice of λ will depend on empirical performance, simplicity and formal elegance of the resulting inductive method (Carnap, 1952, §18), but he gave no clear indication as to how this balance should be achieved. One might suggest that λ should be treated as a meta-inductive parameter: one should attach a prior probability distribution over λ and update in the light of new evidence (see, e.g., Good, 1980). But then there is a danger of regress: if there is a continuum, with parameter λ' , of suitable prior distributions over λ , one needs

to formulate a prior over λ' , and so on (Howson and Urbach, 1989, §4.c.2). To get round this problem one might try taking an arbitrary initial value of λ , and changing that as evidence e is gathered in order to minimise the distance between the inductive probability function $P_{\{e\}}$ and the physical probability function P^* (Carnap, 1952, §§19–24; Kuipers, 1986). A choice has to be made concerning the most appropriate distance function—mean square error seems to be the usual choice in this context—and of course since the physical probability function is unknown, one must estimate these probabilities on the basis of available evidence. This leads to an iterative approximation method for updating λ that does not require a prior over λ and that consequently avoids the regress problem. The difficulty with this line of attack is that, since λ varies, the resulting sequence of inductive probabilities cannot be captured by a single member of the λ -continuum—the resulting inductive method is thus irrational according to the norms laid down by Carnap himself. Hence this avenue undermines the whole basis of Carnap's resolution to the problem of learning from experience.

The δ -continuum. A second worry about Carnap's resolution is that a very similar—and apparently equally justifiable—strategy leads to a totally different continuum of inductive methods, namely the Nix-Paris δ -continuum (Nix, 2005; Nix and Paris, 2006). This continuum takes parameter $\delta \in [0, 1)$ and is the only set of probability functions satisfying:

Regularity. $P_\theta(\theta) = 0$ iff $\models \neg\theta$.

Constant Exchangeability. P_θ should be invariant under permutations of the t_i .

Predicate Exchangeability. P_θ should be invariant under permutations of the predicate symbols U .

Strong Negation. P_θ should be invariant under negating each occurrence of some predicate.

Generalised Principle of Instantial Relevance. If $\theta \models \varphi$ and $\varphi(t_{i+1}) \wedge \psi$ is consistent then $P_\theta(\theta(t_{i+2})|\varphi(t_{i+1}) \wedge \psi) \geq P_\theta(\theta(t_{i+1})|\psi)$.

For a language with a single unary predicate we have that

$$P_\theta(\pm Ut_1 \wedge \cdots \wedge \pm Ut_k) = \frac{1}{2} \left(\frac{1-\delta}{2} \right)^k \left[\left(\frac{1+\delta}{1-\delta} \right)^{r_k} + \left(\frac{1+\delta}{1-\delta} \right)^{k-r_k} \right]$$

and

$$P_\theta(Ut_{k+1}|\pm Ut_1 \wedge \cdots \wedge \pm Ut_k) = \frac{1+\delta}{2} - \frac{\delta}{\left(\frac{1+\delta}{1-\delta} \right)^{r_k-s_k} + 1},$$




Figure 2: The Nix-Paris inductive methods for $\delta \in [0, 1)$, a single unary predicate and $k = 10$.

where $s_k = k - r_k$ is the number of observed negative instances of U . This last function is depicted in Fig. 2.

In general the δ -continuum only agrees with the λ -continuum at point $\delta = 0$, which corresponds to $\lambda = \infty$. This point is the equivocator function of §2—the function that gave rise to the problematic c^\dagger . (Nix and Paris stipulate that $\delta = 1$ should correspond to $\lambda = 0$, but this stipulation is rather counterintuitive when one compares the graph of the δ -continuum, given in Fig. 2, with that of the λ -continuum given in Fig. 1.)

If one takes the principles characterising the δ -continuum to be just as plausible as those characterising the λ -continuum, then Carnap's resolution to the problem of learning from experience faces an important problem: underdetermination.

The pre-eminence of the equivocator. The last but most important of the problems facing Carnap's resolution is that—setting aside, for the moment, the argument that it gives rise to c^\dagger and the problem of learning from experience—the equivocator function P_- (i.e., $\lambda = \infty, \delta = 0$) stands out by a long shot as the only viable candidate for P_θ .

As noted above, if the Johnson-Carnap justification of the λ -continuum is convincing at all, then so is the Nix-Paris justification. And, putting

all the desiderata together, we have that the point where the two continua coincide—i.e., the equivocator $P_=_$ —is the only function satisfying

Regularity. $P_\emptyset(\theta) = 0$ iff $\models \neg\theta$.

Constant Exchangeability. P_\emptyset should be invariant under permutations of the t_i .

Predicate Exchangeability. P_\emptyset should be invariant under permutations of the predicate symbols U .

Strong Negation. P_\emptyset should be invariant under negating each occurrence of some predicate.

Generalised Principle of Instantial Relevance. If $\theta \models \varphi$ and $\varphi(t_{i+1}) \wedge \psi$ is consistent then $P_\emptyset(\theta(t_{i+2})|\varphi(t_{i+1}) \wedge \psi) \geq P_\emptyset(\theta(t_{i+1})|\psi)$.

Johnson's Sufficientness Postulate. $P_\emptyset(Ut_{k+1}|\pm Ut_1 \wedge \dots \wedge \pm Ut_k)$ should depend only on k and the number r_k of positive observations.

Thus the equivocator stands out as the only viable candidate for P_\emptyset .

One might respond to this line of argument that several of the above desiderata are invariance conditions and can be thought of as applications of the principle of indifference, which says that if one is indifferent concerning which member of a partition will occur then all members of the partitions should receive the same probability, and which is notorious principally for the problems that arise when it is applied over different partitions. Perhaps, then, the line of argument should not be trusted.

This response can lead in two directions. If one thinks that the problems generated by multiple applications of the principle of indifference are reason enough to reject the principle straight off, then one will, indeed, reject the above line of argument. But one will also reject the applications of the principle of indifference that lead to the λ -continuum and δ -continuum respectively. Hence Carnap's resolution of §4 does not get off the ground and there is no serious alternative to the claim at Step 4 of §2 that $P_\emptyset = P_=_$.

But the response can go in another direction. The problems generated by multiple applications of the principle of indifference are more plausibly taken as reasons to restrict the principle of indifference rather than reject it straight off. After Keynes (1921, §4.21) it is usual to restrict the principle of indifference to the *finest* partition over which one is indifferent. In our case there is no evidence at all (we are considering P_\emptyset) and the finest partition over which there is indifference is the finest partition simpliciter—i.e., the partition Ω_n of the atomic states. This leads to the following desideratum:

State Exchangeability. P_\emptyset should be invariant under permutations of the states $\omega \in \Omega_n$.

But State Exchangeability clearly implies that $P_\emptyset(\omega) = P_=(\omega) = 1/|\Omega_n|$ for all $\omega \in \Omega_n$. And it is clear that taking different values of n will not lead to inconsistent applications of the principle of indifference. Hence we have, again, that $P_\emptyset = P_=_$.

In sum, in either direction in which one takes concerns about applying the principle of indifference, we are left with the equivocator as the only viable candidate for P_\emptyset .

§6 A ROCK AND A HARD PLACE

A theory of degree of confirmation needs to capture two concepts. On the one hand, it should capture the ampliative concept of degree of *inductive plausibility*, e.g., the degree to which an observed sample of ravens all being black renders plausible the conclusion that the next observed raven will be black. On the other hand, degree of confirmation should also capture the non-ampliative concept of degree of *partial entailment*, e.g., the degree to which $A \vee B$ entails the conclusion A .² We apply the concept of confirmation in both cases—a sample of ravens can *confirm* the conclusion that the next raven will be black; $A \vee B$ *confirms* A —so a theory of confirmation should be able to cope with both kinds of case.

Carnap was rightly concerned that employing the equivocator as a basis for confirmation—by using c^\dagger —would mean that confirmation theory would not be able to capture the concept of inductive plausibility. But by rescinding Step 4 of §2—i.e., by rejecting the identification of P_\emptyset with $P_=_$ and by developing his continuum of inductive methods—he threw the baby out with the bath water, because the equivocator is the only function able to capture partial entailment in the total absence of evidence. If there is no evidence to distinguish interpretations of a logical language then the degree to which premisses entail a conclusion can only viably be identified with the proportion of models of the premisses that also satisfy the conclusion—equivalently, with the proportion of those atomic states logically implying the premisses that also logically imply the conclusion (Wittgenstein, 1922, §5.1.5).

One way to argue for this claim is to appeal to the reasons given in §5 for the pre-eminence of the equivocator. In order to determine the degree to which $A \vee B$ entails A , consider the following.

$$P_\emptyset(A \mid A \vee B) = \frac{P_\emptyset(A \wedge (A \vee B))}{P_\emptyset(A \vee B)} = \frac{P_\emptyset(A \wedge B) + P_\emptyset(A \wedge \neg B)}{P_\emptyset(A \wedge B) + P_\emptyset(A \wedge \neg B) + P_\emptyset(\neg A \wedge B)}$$

² This concept is called *structural confirmation* by Kuipers (2001, pp. 208–9).

but, as argued in §5, these atomic states should all have the same probability in the absence of evidence, so

$$P_{\emptyset}(A \mid A \vee B) = \frac{1/2}{3/4} = \frac{2}{3}.$$

Thus $A \vee B$ partially entails A to degree $2/3$. ($A \vee B$ might be said to *support* A to degree $P_{\emptyset}(A \mid A \vee B) - P_{\emptyset}(A) = 2/3 - 1/2 = 1/6$.)

Kemeny and Oppenheim (1952, p. 314) provide a rather different argument for the claim that the equivocator is required to capture partial entailment (which they use to derive a measure of *degree of factual support*). First they point out that if A and B are logically independent atomic propositions then it must be the case that they are probabilistically independent, $P_{\emptyset}(\pm A \wedge \pm B) = P_{\emptyset}(\pm A)P_{\emptyset}(\pm B)$: ‘Two atomic statements which are logically independent cannot support each other factually since they express distinct facts’. Also, A and $A \leftrightarrow B$ must be probabilistically independent since $A \leftrightarrow B$ is just as favourable to A as to $\neg A$. But $A \wedge (A \leftrightarrow B)$ is logically equivalent to $A \wedge B$. Hence, $P_{\emptyset}(A)P_{\emptyset}(B) = P_{\emptyset}(A \wedge B) = P_{\emptyset}(A \wedge (A \leftrightarrow B)) = P_{\emptyset}(A)P_{\emptyset}(A \leftrightarrow B)$. In which case $P_{\emptyset}(B) = P_{\emptyset}(A \leftrightarrow B)$. Moreover $\neg A \wedge (A \leftrightarrow B)$ is logically equivalent to $\neg A \wedge \neg B$ so $P_{\emptyset}(\neg A)P_{\emptyset}(\neg B) = P_{\emptyset}(\neg A \wedge \neg B) = P_{\emptyset}(\neg A \wedge (A \leftrightarrow B)) = P_{\emptyset}(\neg A)P_{\emptyset}(A \leftrightarrow B)$ and $P_{\emptyset}(\neg B) = P_{\emptyset}(A \leftrightarrow B)$. Hence $P_{\emptyset}(B) = P_{\emptyset}(\neg B) = 1/2$. Similarly $P_{\emptyset}(A) = P_{\emptyset}(\neg A) = 1/2$ and, since A and B are probabilistically independent, $P_{\emptyset}(\pm A \wedge \pm B) = 1/4$. Similarly the other atomic propositions are all probabilistically independent and have probability $\frac{1}{2}$, so $P_{\emptyset}(\omega) = 1/|\Omega_n|$ for $\omega \in \Omega_n$. Hence $P_{\emptyset} = P_{=}$, the equivocator.

We are thus stuck between a rock and a hard place: on the one hand, the equivocator seems to preclude learning by experience, and so fails to capture the concept of inductive plausibility, while on the other, the equivocator seems to be required to capture the concept of partial entailment. Wesley Salmon recognised this dilemma very clearly. He pointed out that if q entails p then p partially entails q because it entails a part of q , and he argued:

if degree of confirmation is to be identified with partial entailment, then c^{\dagger} is the proper confirmation function after all, for it yields the result that p is probabilistically irrelevant to q whenever p and q are completely independent and there is no partial entailment between them. . . . (Salmon, 1967, p. 731)

But Salmon despaired of finding a way out of this dilemma:

. . . Unfortunately for induction, statements strictly about the future (unobserved) are completely independent of statements strictly about the past (observed). Not only are they deductively independent of each other, but also they fail to exhibit any partial entailment. The force of Hume’s insight that the future is logically independent of the past is very great indeed. It rules out both full entailment and partial entailment. If partial entailment were the fundamental concept of inductive

logic, then it would in fact be impossible to learn from experience. (Salmon, 1967, pp. 731–2)

While Carnap sacrificed partial entailment for inductive plausibility, Kemeny focussed on explicating partial entailment (Kemeny, 1953). For both Carnap and Kemeny, the quest for a theory of confirmation that adequately handles the two concepts at once is left empty-handed. Salmon thought that there is no way of satisfying these apparently conflicting demands. But we shall see that there is, by taking another look at the four-step Bayesian approach to confirmation.

§7 THE BAYESIAN APPROACH REVISITED

Let us revisit the scheme of §2 in the light of our discussion so far.

Step 1. Consider probability functions defined over a language \mathcal{L} .

Step 2. Identify $c(h, e) = P_{\{e\}}(h)$ for some suitable probability function P on \mathcal{L} , where $P_{\{e\}}(h)$ is the probability of h on evidence e .

Step 3. Identify $P_{\{e\}}(h) = P_{\emptyset}(h|e)$.

Step 4. Find an appropriate P_{\emptyset} that represents confirmation in the absence of evidence. Here $P_{\emptyset}(\omega) = P_{=}(\omega) \stackrel{\text{df}}{=} 1/|\Omega_n|$ for all $\omega \in \Omega_n$, the *equivocator* on \mathcal{L} .

We saw that together these steps have the unhappy consequence that $c(h, e) = c^{\dagger}(h, e) = P_{=}(h|e)$, which precludes learning from experience.

Although the focus of the last 60 years of work on confirmation theory and inductive logic has been on Step 4, we have seen that it is not Step 4 that is at fault: the equivocator does indeed stand out as the only viable confirmation function in the total absence of evidence. If Step 4 is not at fault then we must look elsewhere. Revising Steps 1 or 2 would take us away from Bayesian confirmation theory and the remit of this paper;³ instead we will focus on Step 3:

Step 3. Identify $P_{\{e\}}(h) = P_{\emptyset}(h|e)$.

This says that the degree to which you should believe h if you were to grant e is exactly the degree to which you should believe $h \wedge e$ were you to grant nothing, divided by the degree to which you should believe e were you to grant nothing.

This claim is far from obvious, and, given that one of the four steps must be revised if confirmation theory is to capture learning from experience, Step 3 could do with closer scrutiny.

³ Popper (1934, Appendix *ix), for one, argued against Step 2.

We saw in §2 that the standard justification of Step 3 is in terms of conditional bets: if we interpret the degree to which one should believe h , were one to grant only e , as a betting quotient for h where the bet is called off if e is false, then Step 3 must hold to avoid the possibility of sure loss. It looks at first sight like the case for Step 3 is as compelling as that for the other axioms of probability, which rely on very similar betting justifications. Accordingly, if one were to cast aspersions on this kind of betting justification then one would seem to undermine the whole Bayesian programme.

But this is too quick. It is clear that there is something special about Step 3, for it is clear that there are at least two cases in which one cannot explicate the relevant conditional degree of belief as a conditional probability. First, if e is not expressible in the relevant language \mathcal{L} then while $P_{\{e\}}(h)$ may be well-defined, $P_{\emptyset}(h|e)$ clearly is not. To take a trivial example, if \mathcal{L} is a propositional language with a single propositional variable A , and e says that A has probability 0.8 then $P_{\{e\}}(A)$ is arguably 0.8 although $P_{\emptyset}(h|e)$ is undefined because e is not a proposition of \mathcal{L} . Second, if e is expressible in \mathcal{L} but has probability 0 then $P_{\emptyset}(h|e)$ is undefined but $P_{\{e\}}(A)$ may be well-defined. For example, the probability that a dart will hit a particular point of a dartboard may be 0, but on evidence e that the dart hit that point, the hypothesis h that the resulting score increased by 20 has a well-defined probability ($P_{\{e\}}(h) = 0$ or 1); yet $P_{\emptyset}(h|e)$ is undefined, so it is not possible that $P_{\{e\}}(h) = P_{\emptyset}(h|e)$. In response to this second case, one might point out that, as an alternative to taking conditional probability to be undefined, one can construe the conditional probability as *unconstrained* when the condition has zero probability: $P_{\emptyset}(h|e)$ can be any value in the unit interval. But the main point goes through as before: $P_{\{e\}}(h)$ is well-defined and fully constrained by P_{\emptyset}, h and e , yet $P_{\emptyset}(h|e)$ is unconstrained, so the two quantities cannot be identified.

In sum, it is apparent that it is *not* always appropriate to explicate $P_{\{e\}}(h)$ in terms of a conditional probability. This conclusion leads naturally to two questions. First, under what conditions, exactly, is this explication (and hence Step 3) plausible? Second, if we articulate these conditions to reformulate Step 3, will the problem of learning from experience remain? In order to answer these questions we will need to invoke the machinery of objective Bayesian epistemology.

§8 OBJECTIVE BAYESIAN EPISTEMOLOGY

Bayesian epistemology addresses the following question: how strongly should an agent believe the various propositions expressible in her language? There are various kinds of Bayesian epistemology; in this section we will sketch *objective Bayesian epistemology*. The reader is referred to Williamson (2010b) for the details of this particular version of Bayesian epistemology.

According to objective Bayesian epistemology, an agent with evidence \mathcal{E} and language \mathcal{L} should apportion the strengths of her beliefs according to three norms:

Probability. Her belief function $P_{\mathcal{E}}$ should be a probability function on \mathcal{L} .

Calibration. Her belief function should be calibrated with her evidence. For example, her degrees of belief should be set to frequencies where known.

Equivocation. Her belief function should otherwise equivocate sufficiently between basic possibilities expressible in \mathcal{L} .

The Probability norm requires that rational degrees of belief satisfy the axioms of probability given in §2. The norm says that $P_{\mathcal{E}} \in \mathbb{P}$ where \mathbb{P} is the set of probability functions on \mathcal{L} . (We need not assume that \mathcal{E} itself is expressible as a set of sentences of \mathcal{L} .) The usual justification of this norm is in terms of betting behaviour: if degrees of belief are interpreted in terms of betting quotients, then, in order to avoid the possibility of certain loss, they must be probabilities. Note that this justification only needs to appeal to an interpretation of *unconditional* degrees of belief as betting quotients—conditional beliefs will be analysed separately below—and the problems facing the interpretation of conditional beliefs in terms of conditional bets, alluded to in §7, can be set aside for the moment.

The Calibration norm says that the agent's belief function should lie within some subset of probability functions that are calibrated with her evidence, $P_{\mathcal{E}} \in \mathbb{E} \subseteq \mathbb{P}$. This can be cashed out as follows. The agent's evidence, construed as everything she takes for granted in her current operating context, may contain information about physical chances that constrains her degree of belief, and it may contain information that constrains degrees of belief in a way that is not mediated by facts about chances. To handle the latter kind of constraint, we may suppose that \mathcal{E} imposes a set of structural, non-chance constraints which are satisfied by a subset \mathbb{S} of all probability functions, and we insist that $P_{\mathcal{E}} \in \mathbb{S}$; since this kind of constraint is not central to the points of this paper, there is no need to go into further detail here. To handle the former kind of constraint, we may suppose that the agent's evidence narrows down the chance function P^* on \mathcal{L} to a subset \mathbb{P}^* of \mathcal{L} . Now this information will typically be pertinent to the agent's degrees of belief, for if she neglects to bet according to the known chances a shrewd stake-maker can force her to lose money in the long run. But it is too simplistic to say that the agent's belief function should itself be in \mathbb{P}^* : she might, for instance, have evidence that θ refers to an event in the past, in which case its chance is 0 or 1 and $\mathbb{P}^* \subseteq \{P \in \mathbb{P} : P(\theta) = 0 \text{ or } 1\}$, but it would be absurd to insist that $P_{\mathcal{E}} \in \mathbb{P}^*$, i.e., to insist that she should either fully believe or fully disbelieve θ , because she might have no other evidence bearing on the truth of θ . For this reason $P_{\mathcal{E}}$ is only constrained to lie in the

convex hull $\langle \mathbb{P}^* \rangle$ of \mathbb{P}^* . (The whole convex hull is admitted because, while the agent can be made to lose money in the long run if she bets according to degrees of belief outside the hull, as long as she stays within the hull then she avoids this possibility of loss.) In sum, the Calibration norm says that $P_{\mathcal{E}} \in \mathbb{E} = \langle \mathbb{P}^* \rangle \cap \mathbb{S}$.

The Equivocation norm says that the agent's belief function should equivocate sufficiently between the basic possibilities expressible in \mathcal{L} . The basic possibilities expressible in \mathcal{L} are just the atomic states ω ; the probability function that is maximally equivocal is the equivocator $P_{=}$, so the Equivocation norm can be read as saying that the agent's belief function should be a function in \mathbb{E} that is sufficiently close to $P_{=}$. If we write $\Downarrow \mathbb{E}$ for the subset of functions in \mathbb{E} that are sufficiently close to the equivocator, then the Equivocation norm says that $P_{\mathcal{E}} \in \Downarrow \mathbb{E}$. It is usual to measure distance between probability functions by what has come to be known as the *Kullback-Leibler divergence*, $d_n(P, Q) = \sum_{\omega \in \Omega_n} P(\omega) \log(P(\omega)/Q(\omega))$. (For a predicate language, one can deem P to be closer to R than Q if there is some N such that for all $n \geq N$ the divergence $d_n(P, R)$ is strictly less than the divergence $d_n(Q, R)$.) Why should a belief function be equivocal? Because the equivocal belief functions turn out to be those that, under the betting interpretation, minimise worst-case expected loss, for a natural default loss function (Williamson, 2010a). Why should the belief function be *sufficiently* equivocal rather than *maximally* equivocal? Because in certain cases there may not be a maximally equivocal belief function in \mathbb{E} ; in such cases contextual considerations (such as the required numerical accuracy of predictions) can be used to determine what is to count as close enough to the equivocator. In general, if $\Downarrow \mathbb{E}$ is the set of maximally equivocal probability functions in \mathbb{E} then $\Downarrow \mathbb{E} \subseteq \Downarrow \mathbb{E} \subseteq \mathbb{E}$. If there *are* maximally equivocal functions and if $\Downarrow \mathbb{E} = \Downarrow \mathbb{E}$ then one can derive the *maximum entropy principle* of Jaynes (1957): $P_{\mathcal{E}} \in \Downarrow \mathbb{E} = \{P \in \mathbb{E} : \text{entropy } H(P) = -\sum_{\omega} P(\omega) \log P(\omega) \text{ is maximised}\}$. We shall suppose, in this paper, that if $\Downarrow \mathbb{E}$ is non-empty then $\Downarrow \mathbb{E} = \Downarrow \mathbb{E}$, so that the maximum entropy principle is applicable in this case.

There are two important consequences of this framework that set objective Bayesianism apart from other versions of Bayesian epistemology. First, no further rule of updating is required. If evidence \mathcal{E} changes to \mathcal{E}' then $P_{\mathcal{E}}$ changes to $P_{\mathcal{E}'}$ accordingly, where the latter function is determined afresh by the requirement that $P_{\mathcal{E}'} \in \Downarrow \mathbb{E}'$. Thus belief change is said to be *foundational*, with beliefs constantly tracking their evidential grounds, rather than *conservative* (independent rules for updating such as Bayesian conditionalisation tend to conserve prior belief, keeping new beliefs as close as possible to old beliefs). Having said all that, there are many natural circumstances under which the objective Bayesian update will match an update generated by Bayesian conditionalisation, and the cases in which there is disagreement between the two forms of updating can be thought of as pathological cases—cases in which it would be inappropriate to condi-

tionalise (Williamson, 2009). So under objective Bayesianism one can often think in terms of conditionalisation if one wishes, as long as one is aware of the pathological cases.

The second important consequence concerns the treatment of conditional belief. Conditional degrees of belief are already determined by the above scheme: the degree to which one should believe h were one to grant only e , $P_{\{e\}}(h)$, is determined by the objective Bayesian protocol $P_{\{e\}}(h) = P_{\mathcal{E}}(h)$ where $P_{\mathcal{E}} \in \Downarrow\mathbb{E}$ and $\mathcal{E} = \{e\}$. There is thus no need to resort to conditional probabilities or conditional bets in order to handle conditional beliefs. Under the objective Bayesian scheme, then, conditional probabilities are much less central than under other versions of Bayesian epistemology—they simply abbreviate quotients of unconditional probabilities, $P(\theta|\varphi) \stackrel{\text{df}}{=} P(\theta \wedge \varphi)/P(\varphi)$, and are not to be interpreted in terms of special, conditional betting quotients. Having said all that, there are natural circumstances under which the objective Bayesian view of conditional beliefs will match the conditional bet view. Since these circumstances are important from the point of view of the present paper, we shall dwell on them.

We have supposed that evidence \mathcal{E} imposes a set of constraints that ought to be satisfied by an agent with that evidence. (There may be more than one way to formulate this set of constraints, but this will not matter for our purposes.) We will use $\chi_{\mathcal{E}}$ to denote this set of constraints; hence $\mathbb{E} = \{P \in \mathbb{P} : P \text{ satisfies the constraints in } \chi_{\mathcal{E}}\}$. Should evidence be inconsistent, i.e., should it determine a set $\chi_{\mathcal{E}}^0$ of *prima facie constraints* that is unsatisfiable, one cannot identify $\mathbb{E} = \{P \in \mathbb{P} : P \text{ satisfies the constraints in } \chi_{\mathcal{E}}^0\} = \emptyset$ because in such a situation one can hardly preclude an agent from holding any beliefs at all. Rather, some consistency maintenance procedure needs to be invoked, to generate a set $\chi_{\mathcal{E}}$ of constraints that are jointly satisfiable. One might take $\chi_{\mathcal{E}}$ to be a disjunction of maximal consistent subsets of $\chi_{\mathcal{E}}^0$, for example, or one might use a consistency maintenance procedure that retains the more entrenched evidence and revokes the less entrenched evidence; we need not decide this question here.

Consider two sets of evidence, \mathcal{E} and $\mathcal{E}' = \mathcal{E} \cup \{e\}$, where e is some sentence of \mathcal{L} . We shall call e *simple* with respect to \mathcal{E} iff $\chi_{\mathcal{E}'}$ is equivalent to (isolates the same set of probability functions as) $\chi_{\mathcal{E}} \cup \{P(e) = 1\}$, i.e., iff the only constraint that e imposes in the context of \mathcal{E} is $P(e) = 1$. Call e *consistent* with respect to \mathcal{E} iff $\chi_{\mathcal{E}} \cup \chi_{\{e\}}^0$ is satisfiable by some probability function (so that $\chi_{\mathcal{E}'}$ is equivalent to $\chi_{\mathcal{E}} \cup \chi_{\{e\}}^0$). We then have the following useful result (Seidenfeld, 1986, Result 1; Williamson, 2009):

Theorem 8.1 *If*

1. e is expressible in \mathcal{L} ,
2. e is simple with respect to \mathcal{E} ,

3. e is consistent with respect to \mathcal{E} , and

4. $P_{\mathcal{E}}(\cdot|e)$ satisfies $\chi_{\mathcal{E}}$,

then $P_{\mathcal{E}'}(h) = P_{\mathcal{E}}(h|e)$.

We see, then, that if the above four conditions are satisfied, a conditional degree of belief will match a corresponding conditional probability.

§9 OBJECTIVE BAYESIAN CONFIRMATION THEORY

Having taken a detour into objective Bayesian epistemology, we are now in a position to return to the central concern of the paper—developing an account of confirmation that can capture both inductive plausibility (in particular, learning from experience) and partial entailment (in particular, the fact that the equivocator function captures confirmation in the total absence of evidence). In §7 we suggested that it is Step 3 of the Bayesian scheme—rather than Step 4—that needs reformulating. Here we apply objective Bayesian epistemology to see how Step 3 should be revised.⁴

The original Step 3 was,

Step 3. Identify $P_{\{e\}}(h) = P_{\emptyset}(h|e)$.

We have seen that objective Bayesianism has a rather different conception of conditional beliefs. Conditional beliefs are to be determined by the norms of objective Bayesianism, rather than via an interpretation in terms of conditional bets. This motivates a new version of Step 3:

Step 3'. Determine $P_{\{e\}}(h)$ using $P_{\{e\}} \in \Downarrow\mathbb{E}$, where $\Downarrow\mathbb{E}$ is the set of sufficiently equivocal probability functions satisfying constraints imposed by e .

According to this conception, the Bayesian scheme becomes:

Step 1. Consider probability functions defined over a language \mathcal{L} .

Step 2. Identify $c(h, e) = P_{\{e\}}(h)$ for some suitable probability function P on \mathcal{L} , where $P_{\{e\}}(h)$ is the probability of h on evidence e .

Step 3'. Determine $P_{\{e\}}(h)$ using $P_{\{e\}} \in \Downarrow\mathbb{E}$, where $\Downarrow\mathbb{E}$ is the set of sufficiently equivocal probability functions satisfying constraints imposed by e .

Step 4. Find an appropriate P_{\emptyset} that represents confirmation in the absence of evidence. Here $P_{\emptyset}(\omega) = P_{=}(\omega) \stackrel{\text{df}}{=} 1/|\Omega_n|$ for all $\omega \in \Omega_n$, the equivocator on \mathcal{L} .

4 The approach of this section is a development of that taken in Williamson (2010b) and supersedes that of Williamson (2007, 2008).

Partial entailment and inductive plausibility

It is not hard to see that this revised scheme does what we need of confirmation.

For one thing, partial entailment is captured because confirmation in the total absence of evidence is implemented using the equivocator. In fact Step 4 is a consequence of Step 3'. According to Step 3', $P_\emptyset(h)$ is determined by the function in \mathbb{P} that is closest to the equivocator. But this is just the equivocator itself (since there is no evidence here). Hence Step 4 follows. Indeed we can calculate that $c(A, A \vee B) = P_{\{A \vee B\}}(A) = P_\emptyset(A|A \vee B) = P_=(A|A \vee B) = 2/3$, just as suggested in §6. Here the identity $P_{\{A \vee B\}}(A) = P_\emptyset(A|A \vee B)$ follows by Theorem 8.1.

For another thing, inductive plausibility can also be captured by this theory of confirmation: learning from experience is no longer impossible. Suppose that an agent grants that a hundred ravens were sampled and all found to be black and that all outcomes are independent and identically distributed (iid) with respect to physical probability. This yields an evidence base \mathcal{E} and tells her something about the physical probabilities: there is high probability that the probability of a raven being black is close to the sample mean, i.e., to 1. Statistical theory can be used to quantify this probability and to derive conclusions of the form $P^*(P^*(Br_{101}) \geq 1 - \delta) = 1 - \epsilon$.⁵ Now fix $1 - \epsilon_0$ to be the minimum degree of belief to which the agent

5 Note that frequentist statistical theory only yields claims about repeatably instantiatable events—not about single cases such as Br_{101} . Thus frequentist statistics yields statements of the form $freq_S(|\bar{X} - freq_R(B)| < \delta) = 1 - \epsilon$, where here the reference class R of the innermost frequency statement is that of all ravens, the reference class S of the outermost frequency statement is that of all samples of a hundred ravens, and \bar{X} is the sample mean, i.e., the proportion of sampled ravens that are black (1 in the case of the agent's particular sample). Such statements are read: if one were to repeatedly sample a hundred ravens then the proportion of samples which have sample mean within δ of the proportion of ravens that are black, is $1 - \epsilon$. While the normal approximation to the binomial distribution might be applied to yield δ or ϵ in many such cases, in the case of extreme sample frequencies, such as the frequency 1 in our example, interval estimation is rather subtle—see, e.g., Brown et al. (2001). The frequencies in such statements are normally understood as counterfactual rather than actual frequencies—i.e., the reference classes include possible ravens and possible samples other than those that are actually appear (Venn, 1866, p. 18; Kolmogorov, 1933, §2).

Such a frequency statement must then be specialised to the single case before the Calibration norm can be used to constrain the single-case belief function $P_{\mathcal{E}}$ by appealing to the single-case chance function P^* . The specialisation to the single case is itself a subtle question, not least because frequencies involving different reference classes can yield conflicting information about single-case probabilities (the so-called *reference-class problem*). The machinery of *evidential probability* was developed for the task of specialising fre-

would need to believe $P^*(Br_{101}) \geq x$ for her to *grant* it (i.e., for her to add that proposition to her evidence base). Then apply statistical theory to determine a δ_0 such that $P^*(P^*(Br_{101}) \geq 1 - \delta_0) = 1 - \epsilon_0$. By the Calibration norm of §8, the agent's rational degrees of belief should be calibrated to this physical probability and so she should strongly believe that the chance is close to 1, $P_{\mathcal{E}}(P^*(Br_{101}) \geq 1 - \delta_0) = 1 - \epsilon_0$. Accordingly the agent grants that the chance is close to 1, thereby increasing her evidence base from \mathcal{E} to $\mathcal{E}' = \mathcal{E} \cup \{P^*(Br_{101}) \geq 1 - \delta_0\}$. Applying the Calibration norm again, the agent should strongly believe that the raven in question will be black, $P_{\mathcal{E}'}(Br_{101}) \geq 1 - \delta_0$. The Equivocation norm will then incline the agent to a sufficiently equivocal point in the interval $[1 - \delta_0, 1]$, e.g., $P_{\mathcal{E}'}(Br_{101}) = 1 - \delta_0$. We then have that $c(Br_{101}, \mathcal{E}') = P_{\mathcal{E}'}(Br_{101}) = 1 - \delta_0$. Thus gaining evidence \mathcal{E}' does raise the degree of confirmation of the next raven being black and we do have learning from experience.

Note that in this account of inductive plausibility, quite a lot is packed into \mathcal{E} and \mathcal{E}' . In particular, the evidence base needs to include not only facts about the observed sample but also facts about the sampling process in order to derive useful consequences about the chances. However, as pointed out in §8, we do not need to presume that \mathcal{E} or \mathcal{E}' is expressible as a proposition e of \mathcal{L} . This is a decided advantage of the objective Bayesian approach over other versions of Bayesian confirmation theory: while, when we are deciding how strongly to believe a proposition h , it is important to be able to express that proposition, the task of expressing everything we take for granted is a hopeless, if not in principle impossible, task.⁶

Note too that statistical theory plays a leading role in implementing the Calibration norm. Hence it is statistical theory that accounts for the inductive plausibility component of confirmation. This contrasts with Carnap's view that inductive plausibility is a question of logic rather than of mathematical statistics. But it is surely partial entailment, rather than inductive plausibility, that is the logical notion: partial entailment deals with the extent to which premisses entail a conclusion—and entailment is clearly a logical notion—while inductive plausibility deals with the extent to which a hypothesis which goes well beyond the evidence (i.e., which may have little or no deductive support from the evidence) is nevertheless warranted by that evidence—and this goes beyond logic.

quentist statements to the single case (Kyburg Jr and Teng, 2001)—this kind of machinery can integrate into the objective Bayesian framework to permit calibration (Wheeler and Williamson, 2009).

6 On the other hand, in the above example \mathcal{L} is taken to be rich enough to express claims, such as $P^*(Br_{101})$, about physical probabilities. It is often possible to draw useful consequences about chance on less expressive languages, but one should not expect conclusions drawn on a more impoverished language to agree with those drawn on a richer language (Williamson, 2010b, §9.2).

Broadly speaking, then, the Equivocation norm of objective Bayesian epistemology captures partial entailment and the Calibration norm captures inductive plausibility.

Step 3 and Step 3'

To what extent does Step 3' differ from Step 3? I.e., when will $P_{\{e\}}(h) = P_{\emptyset}(h|e)$ under an objective Bayesian construal? And should Step 3' or Step 3 be preferred where they disagree? Theorem 8.1 can help us answer these questions. Applying Theorem 8.1 in the context of Step 3 and Step 3', $\mathcal{E} = \emptyset$, $\mathcal{E}' = \{e\}$, and the four conditions of Theorem 8.1 are the conditions under which $P_{\{e\}}(h) = P_{\emptyset}(h|e)$. Should Step 3 and Step 3' disagree, $P_{\{e\}}(h) \neq P_{\emptyset}(h|e)$, and one or more of these four conditions must fail. Let us examine such failures to see whether Step 3 or Step 3' is to be preferred in each case.

Condition 1. Suppose e is not expressible in \mathcal{L} . Then, as noted at the end of §7, $P_{\emptyset}(h|e)$ is undefined. Hence $P_{\{e\}}(h) \neq P_{\emptyset}(h|e)$. Of course in this case Step 3' is more plausible than Step 3, because Step 3 cannot be implemented.

Condition 2. Suppose then that e is expressible in \mathcal{L} but that e is not simple with respect to $\mathcal{E} = \emptyset$: i.e., e does not merely impose the constraint $P(e) = 1$. To take a rather trivial example, suppose e says that $P^*(h) = 0.9$. This e clearly imposes at least two constraints: $P(e) = 1$ (i.e., $P(P^*(h) = 0.9) = 1$) and, via the Calibration norm, $P(h) = 0.9$. Hence Step 3' sets $P_{\{e\}}(h) = 0.9$. Where there is disagreement between Step 3 and Step 3', $P_{\emptyset}(h|e) \neq 0.9$. Clearly it is more appropriate to use Step 3', which forces $c(h, e) = 0.9$, rather than Step 3, which forces $c(h, e) \neq 0.9$: the conditional probability simply gets it wrong.

The same point can be made in favour of Step 3' even if the details of the account of calibration of §8 are not adopted. Suppose e says that $P(h) = 0.9$ (so e talks of rational belief rather than chance). Again, e clearly imposes at least two constraints: $P(e) = 1$ (i.e., $P(P(h) = 0.9) = 1$) and $P(h) = 0.9$. Now there are two cases. If $P_{\emptyset}(h|e) = 0.9$ then Step 3' will agree with Step 3 and the question of which is to be preferred does not arise. Otherwise $P_{\emptyset}(h|e) \neq 0.9$, and Step 3' is clearly more appropriate because Step 3 will break one of the constraints imposed by e : Step 3' forces $c(h, e) = 0.9$ but Step 3 forces $c(h, e) \neq 0.9$. Again, the conditional probability simply gets it wrong.

Condition 3. Suppose e is inconsistent with respect to \mathcal{E} . Since $\mathcal{E} = \emptyset$ here, this means that e imposes a set $\chi_{\{e\}}^0$ of prima facie constraints that is not satisfiable by any probability function on \mathcal{L} . As mentioned in §8, the

objective Bayesian strategy is to invoke some consistency maintenance procedure to generate a consistent set $\chi_{\mathcal{E}'}$ of constraints, and to set $\mathbb{E}' = \{P : P \text{ satisfies } \chi_{\mathcal{E}'}\}$. Step 3' then selects some $P_{\mathcal{E}'} \in \mathbb{E}'$ that is sufficiently equivocal. How does Step 3 proceed? There are two cases here. First, e may be a logical contradiction. If so, e must have probability 0 and the conditional probability $P_{\emptyset}(h|e)$ must be undefined (or, just as bad, unconstrained). In this case Step 3' is more plausible than Step 3, because either Step 3 cannot be implemented or it offers no constraint—i.e., e confirms h to no degree at all, or e confirms h to any degree (admitting conclusions as bizarre as $c(e, e) = 0$).

The second possibility is that e is not a logical contradiction, but nevertheless it imposes unsatisfiable constraints. For instance, e may say $h \wedge P(h) = 0.9$, i.e., that h is true but you ought to believe it only to degree 0.9. While e is not a logical contradiction there is nevertheless something fishy about it, in the sense of Moore's paradox, because it imposes a set of prima facie constraints $\chi_{\{e\}}^0 = \{P(h) = 1, P(h) = 0.9\}$ that is unsatisfiable. While there might be some question as to which consistency maintenance procedure to adopt in this situation—one might identify $\chi_{\{e\}}$ with $\{P(h) = 1 \vee P(h) = 0.9\}$ or $\{P(h) \in [0.9, 1]\}$ or \emptyset , for example—it is clearly the right strategy to maintain consistency somehow, since an agent must be entitled to some belief function or other in such a situation. So Step 3' seems the right approach to take. Now if $P_{\emptyset}(e) = 0$ then, as before, $P_{\emptyset}(h|e)$ is undefined or unconstrained and Step 3' is clearly to be preferred over Step 3. But if $P_{\emptyset}(e) > 0$ and $P_{\{e\}}(h) \neq P_{\emptyset}(h|e)$, then intuitively one should go with $P_{\{e\}}(h)$ rather than $P_{\emptyset}(h|e)$ since only the former results from the appropriate consistency maintenance procedure. Indeed, in our example if $P_{\emptyset}(e) > 0$ then $P_{\emptyset}(h|e) = 1$ since $h \wedge e$ is logically equivalent to e , but it is clearly unacceptable to insist that $c(h, e) = 1$ when e is unsatisfiable, so Step 3' is to be preferred over Step 3.⁷

Condition 4. Suppose $P_{\mathcal{E}}(\cdot|e)$ does not satisfy $\chi_{\mathcal{E}}$. Since in the current context $\mathcal{E} = \emptyset$, $\chi_{\mathcal{E}}$ must also be empty. So the only way in which $P_{\mathcal{E}}(\cdot|e)$ can fail to satisfy $\chi_{\mathcal{E}}$ is if $P_{\mathcal{E}}(\cdot|e)$ is not a well-defined probability function. This occurs if $P_{\mathcal{E}}(e) = 0$ and conditional probability is taken as undefined

7 If $P_{\emptyset}(e)$ is understood as an objective Bayesian probability, this last situation perhaps does not arise. Arguably it cannot be that $P_{\emptyset}(e) > 0$ because the norms of objective Bayesianism should ensure that $P_{\emptyset}(e) = 0$ when e imposes unsatisfiable constraints. The idea here is that any set of evidence \mathcal{E} imposes the constraint $P(\theta) = 0$ for each θ inconsistent with respect to \mathcal{E} . Such a constraint is called a *structural constraint* (§8). If this policy is accepted then indeed the aforementioned situation does not arise under an objective Bayesian construal of $P_{\emptyset}(e)$. Note too, though, that if this policy is accepted then P_{\emptyset} will not agree with the equivocator function $P_{=}$ on those unsatisfiable propositions that are not logical contradictions.

when the condition has probability zero. As we just saw in the discussion of Condition 3, in this case Step 3 is not implementable and Step 3' is to be preferred.

At the end of §7 we encountered two situations in which Step 3 is inappropriate: the case in which e is not expressible in \mathcal{L} and the case in which e has probability 0. These cases correspond to infringements of Conditions 1 and 4 of Theorem 8.1 and one kind of infringement of Condition 3. We asked in §7 whether there are any other restrictions that need to be made to Step 3. We now have our answer: Conditions 2 and 3 spell out the only other restrictions that need to be made. Where these four conditions are satisfied the objective Bayesian account will agree with the original Bayesian scheme of §2. On the other hand, in each case in which these conditions fail, the objective Bayesian account, which replaces Step 3 by Step 3', is to be preferred.

Note that with the problem of learning from experience, it is Condition 2 that is pertinent: new evidence e tends not to be simple with respect to background \mathcal{E} . If e says that a hundred ravens were observed and all found to be black, and that the pertinent chances are iid, then, according to the above account, e does not merely impose the constraint $P(e) = 1$ but also constraints that imply $P(P^*(Br_{101}) \geq 1 - \delta) = 1 - \epsilon$. Consequently e is not simple and Step 3', rather than Step 3, must be applied.

Since Step 3 is abandoned in favour of Step 3', the question arises as to whether the resulting account is prone to the Dutch book argument of §2. Surely an agent who does not set $P_{\{e\}}(h) = P_{\emptyset}(h|e)$ opens herself up to the possibility of sure loss?

The natural response to this worry is just to point out that in the objective Bayesian framework conditional beliefs are not interpreted in terms of conditional bets, so infringing Step 3 does not expose an agent to sure loss. To put it another way, one would be advised not to place a conditional bet, conditional on evidence that is not simple with respect to current evidence, with a betting quotient matching one's rational degree of belief (as determined by Step 3'), for fear of sure loss as per the argument of §2. The interpretation of conditional beliefs in terms of conditional bets is therefore inappropriate in general.

Under the approach advocated here, conditional beliefs are explicated by considering unconditional probabilities relative to an evidence base that is expanded to include the conditioning proposition, rather than by considering conditional bets and conditional probabilities. Levi (2010, §4) also favours an approach based on expanding evidence rather than conditional bets. However, Levi imposes a principle—*Confirmational Conditionalisation*—that forces consistency between conditional beliefs and conditional probabilities. This principle is arguably too strong: according to the argu-

ment of this section, while one should expect considerable agreement between conditional degrees of belief and conditional probabilities, agreement should not be universal. In particular, if the conditioning evidence is not simple with respect to the rest of the evidence base then a conditional degree of belief may well disagree with the corresponding conditional probability, and for good reason. Hence Levi's principle of Confirmational Condition-alisation is arguably just as inappropriate as the Carnapian tradition of confirmation theory upon which he is trying to improve.⁸

§10 CONCLUSION

Let us recap the main line of argument. Of the four-step Bayesian scheme of §2, Step 4 has been the main locus of the debate concerning Bayesian confirmation theory, largely because it is commonly thought that Step 4 must be revised if confirmation theory is to adequately capture the problem of learning from experience. But revising Step 4 leads to another problem, namely a failure of confirmation theory to capture the phenomenon of partial entailment. In fact, learning from experience can be accounted for in a different way: by reformulating Step 3 in accordance with the prescriptions of objective Bayesian epistemology. This leads to an objective Bayesian confirmation theory and a new four-step scheme that is broadly preferably to the original scheme of §2.

During the course of this argument we have had to appeal to some subtle distinctions—the distinction between Bayesian confirmation theory, the Bayesian interpretation of probability and Bayesian epistemology, for

8 Proponents of an interpretation of conditional beliefs in terms of conditional bets might wonder whether one can force consistency between conditional beliefs and conditional probabilities on the objective Bayesian account. If successful, such a move might salvage Step 3. Perhaps the most promising suggestion in this regard is simply to impose a structural constraint of the form $P_\emptyset(h|e) = P_{\{e\}}(h)$ for each pair of sentences e and h of \mathcal{L} . If the resulting set of constraints is satisfiable then it would appear that conditional beliefs can be thought of as conditional probabilities after all.

However, it is doubtful that such a set of constraints is satisfiable. Note that $P_{\{A_1 \vee \neg A_1\}}(\omega) = 1/|\Omega_n|$ for any $\omega \in \Omega_n$. This is because a tautology fails to provide substantive information about chances, so $\mathbb{E} = \mathbb{P}$ and $\Downarrow \mathbb{E} = \{P_\equiv\}$. But according to the above suggestion we have a structural constraint of the form $P_{\{A_1 \vee \neg A_1\}}(\omega) = P_\emptyset(\omega|A_1 \vee \neg A_1)$. Now $P_\emptyset(\omega|A_1 \vee \neg A_1) = P_\emptyset(\omega \wedge (A_1 \vee \neg A_1))/P_\emptyset(A_1 \vee \neg A_1) = P_\emptyset(\omega)$. Therefore $P_\emptyset(\omega) = 1/|\Omega_n|$ for all $\omega \in \Omega_n$, i.e., $P_\emptyset = P_\equiv$. But then the problem of learning from experience reappears: under the proposed structural constraints, $P_{\{e\}}(h) = P_\emptyset(h|e) = P_\equiv(h|e) = P_\equiv(h) = P_\emptyset(h)$ if h and e are logically independent. This contradicts the observation above that it is possible to learn from experience on the objective Bayesian account, i.e., that $P_{\{e\}}(h) > P_\emptyset(h)$ for some logically independent e and h .

instance, as well as the distinction between conditional probabilities, conditional bets and conditional beliefs, and the distinction between inductive plausibility and partial entailment. But by teasing these concepts apart we create the conceptual space for a new and promising theory of confirmation.

Further work needs to be done to flesh out the theory, of course. The problem of reconciling learning from experience with partial entailment is but one problem for confirmation theory—others include the question of whether universal hypotheses can have positive confirmation and the question of whether language relativity infects confirmation theory. It would be interesting to see how an objective Bayesian confirmation theory might answer these questions.

Concerning the first question, it is well known that the equivocator awards zero probability to universally quantified statements. (Indeed, any function satisfying Johnson's Sufficientness Postulate, hence any function in Carnap's λ -continuum, awards zero probability to universally quantified statements—see, e.g., Hintikka and Niiniluoto (1980) and Paris (1994, Theorem 12.10).) But under the objective Bayesian account the equivocator only captures confirmation in very special cases (e.g., in the total absence of evidence, or in the case of tautological evidence)—cases in which it is by no means problematic that universal statements be given probability zero. There is clearly nothing in objective Bayesian theory that precludes awarding positive probability to universally quantified statements in the presence of less trivial evidence. Indeed it is clear that if evidence imposes constraints that force a universal hypothesis to have positive probability then it will have positive probability. Arguably an agent's evidence includes everything she takes for granted, including theory, assumptions and background knowledge as well as the results of observations (Williamson, 2010b, §1.4). If her theoretical evidence includes universal hypotheses, then those universal hypotheses will have positive probability, as will universal hypotheses that they partially entail or render inductively plausible. A detailed investigation of this phenomenon remains a topic for further research.

Concerning the second question, it appears that objective Bayesian probability does depend to some extent on the underlying language \mathcal{L} , and rightly so because an agent's language can, independently of any expressions formulated in that language, encapsulate factual information about the world in which the agent dwells. To take a simple example, if the agent's language has 20 different words for snow, that says something about her environment (Williamson, 2010b, §9.2). Under an objective Bayesian confirmation theory, this would imply that the degree to which e confirms h is relative to some extent on the perspective of the underlying language \mathcal{L} . As to whether this leads to any kind of problematic incommensurability of confirmation is another topic for further research.

Acknowledgements

I am very grateful to Donald Gillies, Theo Kuipers and an anonymous referee for helpful comments, and to the British Academy and the Leverhulme Trust for financial support.

BIBLIOGRAPHY

- Bayes, T. (1764). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418.
- Boole, G. (1854). *An investigation of the laws of thought on which are founded the mathematical theories of logic and probabilities*. Walton and Maberly, London.
- Brown, L. D., Cai, T. T., and DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statistical Science*, 16(2):101–117.
- Carnap, R. (1945). On inductive logic. *Philosophy of Science*, 12(2):72–97.
- Carnap, R. (1950). *Logical foundations of probability*. Routledge and Kegan Paul, London, second (1962) edition.
- Carnap, R. (1952). *The continuum of inductive methods*. University of Chicago Press, Chicago IL.
- Carnap, R. (1971). Inductive logic and rational decisions. In *Studies in inductive logic and probability*, volume 1, pages 5–31. University of California Press, Berkeley.
- Fitelson, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science*, 66:S362–S378.
- Gillies, D. (1990). The turing-good weight of evidence function and popper's measure of the severity of a test. *British Journal for the Philosophy of Science*, 41:143–146.
- Good, I. J. (1960). The paradox of confirmation. *British Journal for the Philosophy of Science*, 11:145–149.
- Good, I. J. (1980). Some history of the hierarchical Bayes methodology. *Trabajos de Estadística y de Investigación Operativa*, 31(1):489–519.
- Hintikka, J. and Niiniluoto, I. (1980). An axiomatic foundation for the logic of inductive generalisation. In Jeffrey, R. C., editor, *Studies in inductive logic and probability*, volume 2. University of California Press, Berkeley and Los Angeles.
- Howson, C. and Urbach, P. (1989). *Scientific reasoning: the Bayesian approach*. Open Court, Chicago IL, second (1993) edition.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *The Physical Review*, 106(4):620–630.
- Jeffreys, H. (1936). Further significance tests. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32:416–445.

- Johnson, W. E. (1932). Probability: the deductive and inductive problems. *Mind*, 41(164):409–423.
- Kemeny, J. G. (1953). A logical measure function. *Journal of Symbolic Logic*, 14(4):289–308.
- Kemeny, J. G. and Oppenheim, P. (1952). Degree of factual support. *Philosophy of Science*, 19(4):307–324.
- Keynes, J. M. (1921). *A treatise on probability*. Macmillan (1948), London.
- Kolmogorov, A. N. (1933). *The foundations of the theory of probability*. Chelsea Publishing Company (1950), New York.
- Kuipers, T. A. F. (1986). Some estimates of the optimum inductive method. *Erkenntnis*, 24:37–46.
- Kuipers, T. A. F. (2001). *Structures in science: heuristic patterns based on cognitive structures; an advanced textbook in neo-classical philosophy of science*. Kluwer, Dordrecht. Synthese Library volume 301.
- Kyburg Jr, H. E. and Teng, C. M. (2001). *Uncertain inference*. Cambridge University Press, Cambridge.
- Levi, I. (2010). Probability logic, logical probability, and inductive support. *Synthese*, 172:97–118.
- Nix, C. (2005). *Probabilistic induction in the predicate calculus*. PhD thesis, University of Manchester.
- Nix, C. J. and Paris, J. B. (2006). A continuum of inductive methods arising from a generalised principle of instantial relevance. *Journal of Philosophical Logic*, 35:83–115.
- Paris, J. B. (1994). *The uncertain reasoner's companion*. Cambridge University Press, Cambridge.
- Popper, K. R. (1934). *The Logic of Scientific Discovery*. Routledge (1999), London. With new appendices of 1959.
- Salmon, W. C. (1967). Carnap's inductive logic. *The Journal of Philosophy*, 64(21):725–739.
- Seidenfeld, T. (1986). Entropy and uncertainty. *Philosophy of Science*, 53(4):467–491.
- Venn, J. (1866). *Logic of chance: an essay on the foundations and province of the theory of probability*. Macmillan, London, second (1876) edition.
- Wheeler, G. and Williamson, J. (2009). Evidential probability and objective Bayesian epistemology. In Bandyopadhyay, P. S. and Forster, M., editors, *Handbook of the Philosophy of Statistics*. Elsevier.
- Williamson, J. (2007). Inductive influence. *British Journal for the Philosophy of Science*, 58(4):689–708.
- Williamson, J. (2008). Objective Bayesianism with predicate languages. *Synthese*, 163(3):341–356.
- Williamson, J. (2009). Objective Bayesianism, Bayesian conditionalisation and voluntarism. *Synthese*, 178(1):67–85.
- Williamson, J. (2010a). Bruno de Finetti: Philosophical lectures on probability. *Philosophia Mathematica*, 18(1):130–135.

Williamson, J. (2010b). *In defence of objective Bayesianism*. Oxford University Press, Oxford.

Wittgenstein, L. (1922). *Tractatus logico-philosophicus*. Routledge & Kegan Paul.

Philosophy
SECL, University of Kent
Canterbury, CT2 7NF
United Kingdom
j.williamson@kent.ac.uk