⋆

Where do we stand on maximal entropy?

⋆

Jon Williamson

Draft of September 22, 2023

Abstract

Edwin Jaynes' ***principle of maximum entropy*** holds that one should use the probability distribution with maximum entropy, from all those that fit the evidence, to draw inferences, because that is the distribution that is maximally non-committal with respect to propositions that are underdetermined by the evidence. The principle was widely applied in the years following its introduction in 1957, and in 1978 Jaynes took stock, writing the paper 'Where do we stand on maximum entropy?' to present his view of the state of the art. Jaynes' principle needs to be generalised to a principle of *maximal* entropy if it is to be applied to first-order inductive logic, where there may be no unique maximum entropy function. The development of this objective Bayesian inductive logic has also been very fertile and it is the task of this chapter to take stock. The chapter provides an introduction to the logic and its motivation, explaining how it overcomes some problems with Carnap's approach to inductive logic and with the subjective Bayesian approach. It also describes a range of recent results that shed light on features of the logic, its robustness and its decidability, as well as methods for performing inference in the logic.

Contents

§1
Introduction

In his pioneering work on information theory, Claude Shannon (1948, §6) argued that the amount of information carried by a discrete probability distribution should be measured by its *entropy*,

$$H(P) = -\sum_{\omega\in\Omega} P(\omega)\log P(\omega),$$

where Ω is a finite set of basic outcomes. In 1957, Edwin Jaynes put forward his 'principle of maximum entropy':

> [I]n making inferences on the basis of partial information we must use that probability distribution which has maximum entropy subject to whatever is known. This is the only unbiased assignment we can make; to use any other would amount to arbitrary assumption of information which by hypothesis we do not have. ... The maximum-entropy distribution may be asserted for the positive reason that it is uniquely determined as the one which is maximally noncommittal with regard to missing information. (Jaynes, 1957, p. 623.)

Jaynes' principle of maximum entropy, or 'maxent' for short, was quickly taken up in areas such as physics, engineering and statistics and theoretical developments were rapid. So much so that at a conference on the Maximum Entropy Formalism, held at MIT in 1978, Jaynes decided that the time was right to take stock and present the state of the art in a paper entitled 'Where do we stand on maximum entropy?' (Jaynes, 1979). In that paper, Jaynes presented the historical background to maxent and its key features, speculated about its future and discussed its application to irreversible statistical mechanics. Maxent has continued to be fruitfully applied to the sciences, and 2023 saw the 42nd International Conference on Bayesian and Maximum Entropy methods in Science and Engineering.

Jaynes argued that maxent can underpin a version of objective Bayesianism, which he viewed as providing an inductive 'logic of science' (Jaynes, 2003).[1] In this chapter, we will see that maxent and objective Bayesianism can indeed yield a viable inductive logic. This logic is called *objective Bayesian inductive logic*, or OBIL for short. The application of maxent to inductive logic has been an active area of collaborative research since the main ideas were set out in Williamson (2008) and Barnett and Paris (2008), and the time is now ripe to take stock and present the state of the art.

§2 presents an introduction to objective Bayesianism (§2.1) and its connection to inductive logic (§2.2). §3 highlights a number of special cases that are particularly well understood. In §4, we turn to the question of how to motivate OBIL, showing that it can be justified on the grounds that if one were to use an inductive logic to decide how to bet, OBIL would be needed to avoid incurring avoidable losses. Next, we compare OBIL to some alternative approaches to inductive logic: Carnap's programme (§5) and two subjective Bayesian approaches (§6). We then consider the question of how to perform inference in the logic, in §7. We see that there is a large class of entailment relationships in OBIL that are decidable (§7.1), and inference

[1] See Rosenkrantz (1977) for an early philosophical account of the role of maxent in an objective Bayesian approach to inductive inference.

can be performed by using augmented truth tables (§7.2) or probabilistic graphical models (§7.3), for example. We give a flavour of the logical properties of OBIL in §8 and discuss the question of language relativity in §9. In §10 we explore the connection between OBIL and a closely related approach, namely the 'entropy-limit' approach of Barnett and Paris (2008). Conclusions are drawn in §11.

This chapter will present key results without proof, so that the reader can quickly ascertain where we currently stand. While OBIL is now a mature theory, we shall see along the way that there remain many open questions, making it a potentially fruitful formalism for further research.

§2
Objective Bayesianism and inductive logic

§2.1. Objective Bayesianism

According to objective Bayesianism, the strengths of one's beliefs need to satisfy three kinds of norm to qualify as rational (Williamson, 2010):[2]

Structural. Degrees of belief should satisfy the laws of probability. For example, the degree to which one believes that the next card to be drawn from a stack of standard playing cards will be black should be the sum of the degree to which one believes that it will be a spade and the degree to which one believes it will be a club, given that spades and clubs are the black suits.

Evidential. Degrees of belief should satisfy constraints imposed by evidence. In particular, they should be calibrated to empirical probabilities, insofar as one has evidence of these empirical probabilities. For example, if one establishes that the stack of playing cards is a complete deck, one should believe to degree 0.25 that the next card drawn is a club.

Equivocation. Degrees of belief should otherwise equivocate as far as possible between outcomes. For example, if one knows only that an experiment has four possible (mutually exclusive) outcomes and that outcome 1 has empirical probability 0.1, then one should believe that the next outcome will be outcome 2 to degree 0.3.[3]

Maxent provides a natural way to explicate the Equivocation norm. Entropy can be interpreted as the measure of the extent to which a probability function equivocates between the basic possibilities. Hence, the Equivocation norm can be implemented by selecting the probability function that has maximum entropy, from all those that satisfy constraints imposed by the Evidential norm.

From a technical point of view, this works rather straightforwardly when the space Ω of basic possibilities is finite. From a philosophical point of view, however, it is open to the charge that the choice of Ω may be rather arbitrary and yet may

[2] Note that this view of objective Bayesianism is very different to the version developed by Jaynes (2003). In particular, Jaynes rejected the idea of empirical probabilities, to which this version of objective Bayesianism appeals.

[3] By the Evidential norm, one should believe that it will be outcome 1 to degree 0.1. By the Structural norm, one should believe that it will be one of the remaining outcomes to degree 0.9. By Equivocation, one should assign each of these remaining outcomes the same probability, 0.3, in the absence of any other relevant evidence.

affect the inferences that are drawn, undermining the purported objectivity of objective Bayesianism. To address this philosophical objection, the approach taken by Williamson (2010, p. 156) is to take Ω to be the set of basic possibilities expressible in one's language:[4] that inferences depend upon the underlying language in this way is relatively unproblematic, because languages evolve to represent and reason about the world efficiently, and thus can be thought of as providing implicit evidence about the world. The difficulty here is that languages tend to be very rich. If one's language could be modelled by a finite propositional language, then it would induce a finite set Ω of all states $\pm a_1 \wedge \cdots \wedge \pm a_n$ of the atomic propositions $a_1, \ldots, a_n$ of the language, and again it would be straightforward to maximise entropy (see, e.g., Paris, 1994). But more typically, one needs at least the richness of a first-order predicate language to express many of the propositions that feature in our inferences. A first-order predicate language does not induce a finite set of states, and it is no longer obvious how to maximise entropy. Hence there is a danger that this appeal to language mitigates a philosophical problem at the expense of introducing a technical problem.

Objective Bayesian inductive logic allows us to address this technical problem, however, as we shall now see.

§2.2. Objective Bayesian Inductive Logic

OBIL considers inductive entailment relationships of the form:

$$\varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \mathrel{|\!\!\!\approx^{\circ}} \psi^Y$$

Here, $\varphi_1, \ldots, \varphi_k, \psi$ are sentences of a first-order predicate language $\mathscr{L}$, and $X_1, \ldots,$ X_k, Y are sets of probabilities. The entailment relationship can be read: if $P(\varphi_1) \in X_1, \ldots,$ and $P(\varphi_k) \in X_k$ then $P(\psi) \in Y$, for any rational belief function P. The premisses on the left-hand side of the entailment relation are interpreted as all the constraints on rational degrees of belief imposed by the Evidential norm, while the conclusion on the right-hand side follows just when each maximally equivocal probability function P, from all those that satisfy the premisses, also satisfies the conclusion. The key task is to say what constitutes 'maximally equivocal'.

First, we need to specify the framework more precisely. Here $\mathscr{L}$ is a pure first-order predicate language: it has relation symbols $U_1, \ldots, U_l$, constant symbols $t_1, t_2, \ldots$ and variable symbols $x_1, x_2, \ldots$, but no function symbols or equality. Sentences θ, φ_i, ψ etc. are formed in the usual way using quantifiers $\forall, \exists$, and connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$. Atomic sentences $a_1, a_2, \ldots$ are ordered so that those involving constants $t_1, \ldots, t_n$ occur in the ordering before those involving t_{n+1}.[5] We will consider the sublanguages $\mathscr{L}_n$ that have all the syntactic apparatus of $\mathscr{L}$ but involve only the constants $t_1, \ldots, t_n$. The *n-states* $\omega \in \Omega_n$ of $\mathscr{L}$ are the states of $\mathscr{L}_n$, i.e., the sentences of the form $\pm a_1 \wedge \cdots \wedge \pm a_{r_n}$, where $a_1, \ldots, a_{r_n}$ are the atomic sentences of $\mathscr{L}_n$.

We will take $X_1, \ldots, X_k, Y$ to be intervals. This is because constraints imposed by the Evidential norm are convex: if the empirical probability of φ is known to be either x or y, where $y > x$, but it is not known which, then the Evidential norm deems any value in the interval $[x, y]$ to be an admissible degree of belief

[4] This line of response is motivated by the suggestion of Keynes (1921) that one should only equivocate between indivisible possibilities.

[5] See §9 for results that indicate that OBIL is invariant under the precise ordering.

(Williamson, 2010, §3.3). We will abbreviate the trivial interval $[x,x]$ by x. We will also abbreviate a premiss or conclusion statement of the form $\varphi^{[1,1]}$ by the categorical (i.e., unqualified) sentence φ.

A probability function P is a function defined on the sentences of $\mathscr{L}$ such that:

P1: If τ is a deductive tautology, i.e., $\models \tau$, then $P(\tau) = 1$.

P2: If θ and φ are mutually exclusive, i.e., $\models \neg(\theta \wedge \varphi)$, then $P(\theta \vee \varphi) = P(\theta) + P(\varphi)$.

P3: $P(\exists x \theta(x)) = \sup_m P\left(\bigvee_{i=1}^m \theta(t_i)\right)$.

Axiom P3, which is sometimes called *Gaifman's condition*, presupposes that each member of the domain of discourse is named by some constant symbol t_i. A probability function is uniquely determined by its values on the n-states (Williamson, 2017, Chapter 2). The set of all probability functions on $\mathscr{L}$ is denoted by $\mathbb{P}$. We will be particularly interested in the set of probability functions that satisfy the evidential constraints:

$$\mathbb{E} \stackrel{\text{df}}{=} [\varphi_1^{X_1}, \ldots, \varphi_k^{X_k}] \stackrel{\text{df}}{=} \{P \in \mathbb{P} : P(\varphi_1) \in X_1, \ldots, P(\varphi_k) \in X_k\}.$$

Now we are in a position to see what constitutes 'maximally equivocal'. We define the n-entropy:

$$H_n(P) \stackrel{\text{df}}{=} -\sum_{\omega \in \Omega_n} P(\omega) \log P(\omega).$$

We then say that P has greater entropy than Q iff the n-entropy of P eventually dominates that of Q, i.e., iff there is an $N \in \mathbb{N}$ such that for all $n \geq N$, $H_n(P) > H_n(Q)$. The greater-entropy relation yields a partial ordering of probability functions, which may contain maximal elements (undominated functions) but need not necessarily contain maximum elements (functions that dominate all others). We thus define the maximally equivocal functions in $\mathbb{E}$ to be those with maximal entropy:

$$\text{maxent}\,\mathbb{E} \stackrel{\text{df}}{=} \{P \in \mathbb{E} : \text{ there is no } Q \in \mathbb{E} \text{ that has greater entropy than } P\}.$$

This yields what we might call the 'principle of maximal entropy'—an extension of the principle of maximum entropy to the setting of an infinite predicate language:

Maximal Entropy Principle. In making inferences on the basis of partial information we must use the probability distributions which have maximal entropy, among all those that satisfy the evidential constraints.

We can then use the maximal entropy principle to provide semantics for OBIL:

$$\varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \mathrel{|\!\!\!\approx^{\circ}} \psi^Y \text{ iff } P(\psi) \in Y \text{ for all } P \in \text{maxent}\,\mathbb{E},$$

as long as $\text{maxent}\,\mathbb{E} \neq \emptyset$. There are two cases in which $\text{maxent}\,\mathbb{E} = \emptyset$, and we need some conventions to cover these cases.[6] The first is where the premisses are unsatisfiable, $\mathbb{E} = \emptyset$. In that case, it is desirable to avoid 'explosion', i.e., the phenomenon that any conclusion follows, because it is never rational to believe everything. We thus consider $\text{maxent}\,\mathbb{P}$ instead of $\text{maxent}\,\mathbb{E}$ when $\mathbb{E} = \emptyset$. $\mathbb{P}$ has a

[6]These cases will not be pursued further in this chapter, but see Williamson (2010, 2017) for further discussion and alternative conventions.

unique maximiser, namely the *equivocator function* $P_=$ defined by $P_=(\omega) = 1/|\Omega_n| = 1/2^{r_n}$ for all $\omega \in \Omega_n$ and $n \geq 1$. Hence an entailment relationship with inconsistent premisses holds iff $P_=(\psi) \in Y$. $P_=(\psi)$ is called the *measure* of ψ.[7] The second case occurs where $\mathbb{E} \neq \emptyset$ but for any probability function in $\mathbb{E}$ there is another function with greater entropy. In this case, we shall consider $\mathbb{E}$ instead of $\text{maxent}\,\mathbb{E}$, so the entailment relationship holds whenever $P(\psi) \in Y$ for all $P \in \mathbb{E}$.

If there are no premisses, the equivocator function is used for inference. If $\mathrel{\mid\!\approx^{\circ}} \psi$ then ψ is said to be an *inductive tautology*. Equivalently, $P_=(\psi) = 1$, i.e., ψ has measure 1. If $\mathrel{\mid\!\approx^{\circ}} \neg\psi$ (i.e., ψ has measure 0) then ψ is an *inductive contradiction*. If $\mathrel{\not\mid\!\approx^{\circ}} \neg\psi$ (i.e., ψ has positive measure) then ψ is *inductively consistent*. Similarly, if $\mathrel{\not\mid\!\approx^{\circ}} \neg(\psi \wedge \theta)$ (i.e., $P_=(\psi \wedge \theta) > 0$) then ψ is *inductively consistent* with θ. If $\mathrel{\mid\!\approx^{\circ}} \psi \leftrightarrow \theta$ then ψ and θ are *inductively equivalent*.

§3
Important cases

This section will survey some important special cases in which there is a unique maximal entropy function and where this function can be determined rather straightforwardly. As we shall see in subsequent sections, these are the cases that have been explored in most depth.

§3.1. $\mathbb{E}$ has finitely generated consequences

Recall that, given premisses $\varphi_1^{X_1}, \dots, \varphi_k^{X_k}$, $\mathbb{E} \stackrel{\text{df}}{=} [\varphi_1^{X_1}, \dots, \varphi_k^{X_k}] \stackrel{\text{df}}{=} \{P \in \mathbb{P} : P(\varphi_1) \in X_1, \dots, P(\varphi_k) \in X_k\}$. One special case occurs when the consequences of these premisses can be characterised as the consequences of some set of quantifier-free premisses:

Definition 1 (*Finitely generated*). $\mathbb{E}$ is *finitely generated* if there exist quantifier-free sentences $\theta_1, \dots, \theta_j$ and intervals $Z_1, \dots, Z_j$ such that $\mathbb{E} = [\theta_1^{Z_1}, \dots, \theta_j^{Z_j}]$. $\mathbb{E}$ has *finitely generated consequences* if there exist quantifier-free sentences $\theta_1, \dots, \theta_j$ and intervals $Z_1, \dots, Z_j$ such that $\text{maxent}\,\mathbb{E} = \text{maxent}[\theta_1^{Z_1}, \dots, \theta_j^{Z_j}]$. In this latter case, $\theta_1^{Z_1}, \dots, \theta_j^{Z_j}$ are called *generating statements* for (the consequences of) $\mathbb{E}$.

Clearly if $\mathbb{E}$ is finitely generated then it has finitely generated consequences, and if the premisses are themselves all quantifier-free then $\mathbb{E}$ is finitely generated. But even when the premisses are not quantifier-free, $\mathbb{E}$ often turns out to have finitely generated consequences. To see when this is so, we require a definition from Landes et al. (2024, §4):

Definition 2 (*Support*). Suppose $a_{i_1}, \dots, a_{i_m}$ include all the atomic propositions that appear in sentence φ of $\mathscr{L}$, and let $\Xi_\varphi \stackrel{\text{df}}{=} \{\pm a_{i_1} \wedge \dots \wedge \pm a_{i_m}\}$ be the set of states of these atomic propositions. If φ contains no atomic propositions, we take $\Xi_\varphi \stackrel{\text{df}}{=}$

[7] The equivocator function is an analogue of Lebesgue measure, via a bijective mapping between probability functions on $\mathscr{L}$ and probability measures on a σ-field of subsets of the unit interval (Williamson, 2017, §2.6). The restriction of $P_=$ to the sentences of any finite sublanguage $\mathscr{L}_n$ corresponds to the uniform distribution on the finite set of n-states Ω_n. There is no uniform distribution on $\mathscr{L}$ itself, because there are infinitely many n-states of $\mathscr{L}$, since $n = 1, 2, \dots$.

$\{a_1, \neg a_1\}$. The *support* $\check{\varphi}$ of φ is the disjunction of states in Ξ_φ that are inductively consistent with φ,

$$\check{\varphi} \stackrel{\text{df}}{=} \bigvee\{\xi \in \Xi_\varphi : P_=(\xi \wedge \varphi) > 0\}.$$

For example, if $\varphi = \exists x(Ut_1 \wedge Vt_1x)$ then $\check{\varphi} = Ut_1$.

Definition 3 (*Support-satisfiable*). Let $\check{\mathbb{E}} \stackrel{\text{df}}{=} [\check{\varphi}_1^{X_1}, \ldots, \check{\varphi}_k^{X_k}]$. We say that the premisses $\varphi_1^{X_1}, \ldots, \varphi_k^{X_k}$ have ***satisfiable support***, or that $\mathbb{E}$ is ***support-satisfiable***, if $\check{\mathbb{E}} \neq \emptyset$. An entailment relationship is ***support-satisfiable*** if $\mathbb{E}$ is support-satisfiable.

$\mathbb{E} = [\exists xVt_1x, (Ut_2 \vee \forall xRx)^{0.9}, Ut_1 \to Vt_1t_3, (Ut_1 \vee (\exists xVxt_3 \to Ut_2))^{[0.95,1]}]$, for example, is support-satisfiable, as we will see in §7. Support-satisfiability is violated in cases where the premisses force an inductive tautology to have probability greater than 0 or force an inductive contradiction to have probability less than 1 (Landes et al., 2024, §10.2). $\mathbb{E} = [\forall xUx^{0.7}]$, for instance, is not support-satisfiable, because the constraint $\forall xUx^{0.7}$ forces positive probability on the measure-zero sentence $\forall xUx$. Note that if $\mathbb{E}$ is non-empty and finitely generated then it is support-satisfiable. Landes et al. (2024, §5) show that:

Theorem 4. *If* $\mathbb{E}$ *is support-satisfiable then it has finitely generated consequences, with generating statements* $\check{\varphi}_1^{X_1}, \ldots, \check{\varphi}_k^{X_k}$.

We also have (see Landes et al., 2024):

Theorem 5. *If* $\mathbb{E}$ *is closed and has finitely generated consequences then it contains a unique maximal entropy function* $P^\dagger$*, i.e.,* $\operatorname{maxent}\mathbb{E} = \{P^\dagger\}$.

Moreover, $P^\dagger$ can be characterised as follows (Landes et al., 2024, §3.3). For any n, let $P_n^\dagger$ be the n-entropy maximiser, $P_n^\dagger \stackrel{\text{df}}{=} \arg\max_{P \in \mathbb{E}} H_n(P)$, if it exists. (Since $\mathbb{E}$ is convex and n-entropy is strictly concave, there can be no more than one n-entropy maximiser.) Now, if $\mathbb{E}$ is closed then $P_n^\dagger$ does indeed exist for each n. Consider any n large enough that all the quantifier-free generating statements for $\mathbb{E}$ are expressible in $\mathscr{L}_n$. Then the entropy maximiser $P^\dagger$ is the probability function that agrees with $P_n^\dagger$ on $\mathscr{L}_n$ but equivocates elsewhere: $P^\dagger$ is defined by $P^\dagger(\omega_m) = P_n^\dagger(\omega_n)P_=(\zeta)$ for each $m \geq n$ and m-state $\omega_m = \omega_n \wedge \zeta$ where $\omega_n \in \Omega_n$.

§3.2. $\mathbb{E}$ is closed in entropy

It turns out that we can relax both conditions of Theorem 5 by considering the convergence of the n-entropy maximisers.

Definition 6 (*Limit in Entropy*). $P \in \mathbb{P}$ is a ***limit in entropy*** of $\mathbb{E}$ if $|H_n(P_n^\dagger) - H_n(P)| \longrightarrow 0$ as $n \longrightarrow \infty$. $\mathbb{E}$ is ***closed in entropy*** if it contains some limit in entropy.

Landes et al. (2021, §5) provide several tests that can help to determine whether $\mathbb{E}$ is closed in entropy. For example, $\mathbb{E} = [\forall xUx^{0.7}]$ is closed in entropy but $\mathbb{E} = [Ut_1 \vee \exists x\forall yVxy]$ is not. The concept of closure in entropy is useful because (Landes et al., 2023, Theorem 16):

Theorem 7. *If* $\mathbb{E}$ *is closed in entropy then* $\operatorname{maxent}\mathbb{E} = \{P^\dagger\}$ *where* $P^\dagger$ *is the unique limit in entropy of* $\mathbb{E}$.

Thus $P^\dagger$ can be determined from the behaviour of $P^\dagger_n$ as n increases.

Let us consider an example (Landes et al., 2023, Example 17):

$$\forall x Ux^{0.7} \mathrel{|\!\!\!\approx^{\circ}} Ut_1^{0.85}$$

Here, $\mathbb{E}$ contains the following limit in entropy:

$$P(\omega) = \begin{cases} 0.7 + \frac{0.3}{2^n} & : \quad \omega = Ut_1 \wedge \ldots \wedge Ut_n \\ \frac{0.3}{2^n} & : \quad \omega \models \neg(Ut_1 \wedge \ldots \wedge Ut_n) \end{cases} .$$

Hence, $\text{maxent}\,\mathbb{E} = \{P\}$ and $P(Ut_1) = 0.7 + (0.3/2) = 0.85$. Thus the above entailment relationship does indeed hold.

Note that Theorem 7 generalises Theorem 5: if $\mathbb{E}$ is closed and has finitely generated consequences then it is closed in entropy. We thus have a sequence of increasingly general situations in which there is guaranteed to be a unique maximal entropy function:

$\mathbb{E}$ is closed, non-empty and finitely generated;

$\Rightarrow$ $\mathbb{E}$ is closed and support-satisfiable;

$\Rightarrow$ $\mathbb{E}$ is closed and has finitely generated consequences;

$\Rightarrow$ $\mathbb{E}$ is closed in entropy.

Important open questions remain. For example, is there a case in which $\text{maxent}\,\mathbb{E}$ is non-empty but no limit in entropy exists?

§4
Motivation

The norms of objective Bayesianism can be justified by appeal to the following considerations: the strengths of one's beliefs guide one's actions; in turn, one's actions can expose one to potential loss; one should not adopt beliefs that expose one to avoidable losses.

The Structural norm is the claim that rational degrees of belief are probabilities. This can be justified on the grounds that one should not adopt beliefs that expose one to *guaranteed* loss, however events turn out. In the finite case this is the well known Dutch Book argument of Ramsey (1926) and de Finetti (1937). This Dutch Book argument can be extended to the situation in which degrees of belief are defined over sentences of a first-order predicate language. Suppose that the loss incurred by believing sentence θ to degree x is:

$$L(\theta, x) = (x - I_\theta)S_\theta,$$

where $S_\theta \in \mathbb{R}$ is an unknown stake (positive or negative) and the indicator function I_θ takes the value 1 if θ is true and 0 if θ is false. A *Dutch book* on set Θ of sentences is a combination of stakes $S_\theta \in \mathbb{R}$ for each $\theta \in \Theta$ that guarantees some fixed loss, i.e., that ensures a positive finite loss $L(\Theta) \stackrel{\text{df}}{=} \sum_{\theta \in \Theta} L(\theta, x) \in (0, \infty)$ whichever sentences $\theta \in \Theta$ turn out to be true. Then (Williamson, 2017, Theorem 9.1):

Theorem 8 (*Dutch Book Theorem*). *Degrees of belief on sentences of $\mathcal{L}$ avoid the possibility of a Dutch book if and only if they satisfy the axioms of probability.*

The Evidential and Equivocation norms are explicated by means of the maximal entropy principle, and this principle can be justified on the grounds that one should not adopt beliefs that expose one to positive *expected* loss, with respect to a default loss function. A default loss function represents that losses one might reasonably anticipate, in the absence of any information about any particular decision scenario and hence the true losses to which one will be exposed. Let $L_\pi(\theta, P)$ signify the loss incurred by adopting belief function P when $\theta \in \pi$ turns out to be true, where π is a partition of sentences. For partitions π_1, π_2, π, write $\pi = \pi_1 \times \pi_2$ when for each $\theta_1 \in \pi_1, \theta_2 \in \pi_2$ there is some $\theta \in \pi$ such that $\theta \equiv \theta_1 \wedge \theta_2$. A probability function P renders π_1 and π_2 independent, written $\pi_1 \perp\!\!\!\perp_P \pi_2$, when $P(\theta_1 \wedge \theta_2) = P(\theta_1)P(\theta_2)$ for each $\theta_1 \in \pi_1, \theta_2 \in \pi_2$. Consider the following constraints on a default loss function:

L1: No loss is incurred when one fully believes the sentence that turns out to be true: $L_\pi(\theta, P) = 0$ if $P(\theta) = 1$.

L2: Loss $L_\pi(\theta, P)$ strictly increases as $P(\theta)$ decreases from 1 to 0.

L3: Loss $L_\pi(\theta, P)$ depends only on $P(\theta)$, not on $P(\varphi)$ for other partition members φ.

L4: Losses are additive over independent partitions: if $\pi = \pi_1 \times \pi_2$ where $\pi_1 \perp\!\!\!\perp_P \pi_2$, then for each $\theta \in \pi$, $L_\pi(\theta, P) = L_{\pi_1}(\theta_1, P) + L_{\pi_2}(\theta_2, P)$, where $\theta_1 \in \pi_1, \theta_2 \in \pi_2$ are such that $\theta \equiv \theta_1 \wedge \theta_2$.

L5: Loss $L_\pi(\theta, P)$ should not depend on the partition π in which θ occurs: there is some function L such that $L_\pi(\theta, P) = L(\theta, P)$ for all partitions π in which θ occurs.

These desiderata are enough to ensure that the default loss function is logarithmic (Williamson, 2017, Theorem 9.2):

Theorem 9. *L1–L5 imply that* $L(\theta, P) = -k \log P(\theta)$, *for some constant* $k > 0$.

Let P^* be the empirical probability function. Then we can consider the expected loss to which one is exposed by beliefs on $\mathscr{L}_n$, for $n \geq 1$, when one adopts $P \in \mathbb{P}$ as one's belief function:

$$\begin{aligned} S_n(P^*, P) &= \sum_{\omega \in \Omega_n} P^*(\omega) L(\omega, P) \\ &= -k \sum_{\omega \in \Omega_n} P^*(\omega) \log P(\omega). \end{aligned}$$

Now suppose that evidence establishes that the empirical probability function P^* lies in a non-empty, closed convex set $\mathbb{E}$ of probability functions. On $\mathscr{L}$ as a whole, $P \in \mathbb{P}$ has lower worst-case expected loss than $Q \in \mathbb{P}$ if there is some N such that for all $n \geq N$,

$$\sup_{P^* \in \mathbb{E}} S_n(P^*, P) < \sup_{P^* \in \mathbb{E}} S_n(P^*, Q).$$

Generalising a result of Topsøe (1979), it turns out that (Williamson, 2017, Theorem 9.3):

Theorem 10 (*Minimax*). *If* $\mathbb{E}$ *is finitely generated then there is a unique probability function that has minimal worst-case expected loss and this is the function* $P^\dagger \in \mathbb{E}$ *that has maximal entropy.*

This result thus provides a justification for the maximal entropy principle, at least in an important special case. The main open question here concerns how far one can relax the assumption that $\mathbb{E}$ is finitely generated.

It is worth noting that while this result can be viewed as providing a pragmatic justification of the maximal entropy principle, it can be recast as an epistemic justification by reinterpreting the function L_n as a measure of the epistemic inaccuracy of P, instead of as a default loss function (Williamson, 2018). Note too that the line of motivation presented in this section appeals to two arguments: a Dutch Book theorem to justify the claim that a rational belief function is a probability function and the Minimax theorem to justify the maximal entropy principle. Landes and Williamson (2013, 2015) explore the possibility of a unified argument that seeks to justify all the norms of OBIL at once.

§5
Relation to Carnap's programme

In this section, we briefly consider how OBIL compares to Rudolf Carnap's well known system of inductive logic (Carnap, 1952).

Carnap also considered inductive logic defined on a first-order predicate language, but in the special case in which the premisses are all categorical, i.e., no uncertainty attaches to the premiss sentences. Carnap's approach can be thought of as providing semantics for entailment relationships of the form:

$$\varphi_1,\dots,\varphi_k \mathrel{|\!\!\approx}_\lambda \psi^Y,$$

where $\lambda \in [0,\infty]$ is a parameter that selects one out of a continuum of entailment relations. Carnap was particularly interested in inference from a sample $\pm Ut_1,\dots,\pm Ut_k$ of past observations. For Carnap, an entailment relationship of the form

$$\pm Ut_1,\dots,\pm Ut_k \mathrel{|\!\!\approx}_\lambda Ut_{k+1}^Y$$

holds just when

$$c_\lambda(Ut_{k+1}|\pm Ut_1,\dots,\pm Ut_k) \stackrel{\text{df}}{=} \frac{\#U + \lambda/2}{k+\lambda} \in Y,$$

where $\#U$ is the number of positive instances of U in the sample.[8]

There are a number of key differences between OBIL and Carnap's approach. Firstly, OBIL is more general, as it is not restricted to categorical premisses. Because Carnap's approach appeals to conditional probabilities, the premisses that are conditioned on must be categorical. Second, Carnap's approach yields a continuum of inductive logics, and, unfortunately, he gave little concrete guidance as to which logic to choose from this continuum. Thus, inductive logic is somewhat underdetermined in Carnap's framework.

The third and perhaps the most fundamental difference is that Carnap's approach tries to embed learning from experience within the logic, while OBIL takes this to be a separate, statistical question. The extent to which Carnap's logics exhibit learning from experience varies with the parameter λ. If $\lambda = 0$, then the probability of the next outcome being positive is set to the observed frequency

[8]Here, c_λ is one of a continuum of probability functions that is defined by the above assignment of conditional probabilities.

$\#U/k$ of positive outcomes in the sample. As λ increases, the influence of the sample decreases until, when $\lambda = \infty$, the probability of the next outcome being positive remains at the value $\frac{1}{2}$ regardless of the proportion of positive outcomes in the sample.

OBIL, on the other hand, takes the question of the influence of the sample to have been determined when formulating the premisses. Recall that objective Bayesianism requires degrees of belief to be appropriately calibrated to empirical probabilities. The way this is fleshed out by Williamson (2017, Chapter 7) is by appeal to confidence intervals. Given a sample $\pm Ut_1, \ldots, \pm Ut_k$ of outcomes, consider the narrowest confidence interval $[l,u]$ within which one is prepared to infer that the empirical probability of a positive outcome lies. If there is no more pertinent information about the next outcome, the Evidential norm would say that one should believe to some degree in the interval $[l,u]$ that the next outcome will be positive. One thus adds $Ut_{k+1}{}^{[l,u]}$ to the list of premisses. This is because the premisses are supposed to represent all the relevant constraints imposed by the Evidential norm on degrees of belief. Therefore, the objective Bayesian needs to consider as premisses $\pm Ut_1, \ldots, \pm Ut_k, Ut_{k+1}{}^{[l,u]}, Ut_{k+2}{}^{[l,u]}, \ldots$. The maximal entropy principle would then yield, for example,

$$\pm Ut_1, \ldots, \pm Ut_k, Ut_{k+1}{}^{[.7,.8]}, Ut_{k+2}{}^{[.7,.8]}, \ldots \mathrel{\overset{\circ}{\approx}} Ut_{k+1}{}^{.7}.$$

Thus, OBIL admits a clear demarcation between entailment in inductive logic and the 'knowledge representation' process of formulating suitable premisses, which can itself involve statistical inference. Theoretical work on OBIL often focusses on entailment, though it presupposes adequate knowledge representation. This demarcation accords well with the analogous case of deductive logic, where considerable work is required in formulating implicit as well as explicit premisses, but where the logical theory tends to presuppose that this work has been done and focusses on entailment. Here, constraints imposed by a sample, such as $Ut_{k+1}{}^{[.7,.8]}$, can be thought of as implicit premisses.

There are advantages to keeping these two aspects of induction distinct, rather than conflating them as Carnap does. As Williamson (2017, Chapter 4) argues, problems can arise for Carnap's approach because the probabilities are always treated as 'exchangeable', i.e., invariant under permutations of the constants. Exchangeability is sometimes appropriate (e.g., when drawing inferences from a sample from independent and identically distributed random variables), but not always. OBIL's separation of these statistical concerns from the logic allows for more flexibility with respect to statistical methods: exchangeability can be invoked where appropriate but need not always be adhered to.

§6
Relation to subjective Bayesianism

The subjective Bayesian approach to inductive inference has some points in common with Carnap's approach. In particular, it invokes conditional probabilities and thus requires categorical premisses. There are two key variants.

The *precise* subjectivist approach appeals to a prior probability function P_0 and holds that:

$$\varphi_1, \ldots, \varphi_k \approx_{P_0} \psi^Y \text{ iff } P_0(\psi | \varphi_1, \ldots, \varphi_k) \in Y.$$

Howson (2000) was a notable advocate of this sort of approach. It can be criticised as being radically underdetermined: according to the subjectivist any probability function is admissible as a prior probability function and there are no rational grounds for choosing one prior over another. Rather, the task is to elicit an agent's subjective degrees of belief and represent these using the prior probability function.

In order to address the problem of underdetermination, one could, of course, argue for stronger constraints on rational degrees of belief, but this arguably leads us back to objective Bayesianism. Alternatively one could adopt an *imprecise* subjectivist approach (Nilsson, 1986):

$$\varphi_1,\ldots,\varphi_k \approx \psi^Y \text{ iff } P(\psi|\varphi_1,\ldots,\varphi_k)\in Y \text{ for all } P\in\mathbb{P}.$$

This can be generalised to what sometimes known as the *standard semantics* for inductive logic (Williamson, 2017, §3.5):

$$\varphi_1^{X_1},\ldots,\varphi_k^{X_k} \approx \psi^Y \text{ iff } P(\psi)\in Y \text{ for all } P\in\mathbb{E}.$$

Unfortunately, however, the imprecise approach arguably trades one problem (underdetermination) for another, namely weakness: it is often the case that the most one can infer from given premisses is that the probability of conclusion sentence ψ lies in the unit interval. OBIL usually motivates much stronger inferences.

A key point of difference between subjective and objective Bayesianism concerns conditional probabilities. Conditional probabilities are absolutely central to subjective Bayesianism, because evidence is taken into account by conditionalising, i.e., by ensuring that all probabilities are conditional on total evidence (see, e.g., Howson and Urbach, 1989, §3.h). This stands in marked contrast to objective Bayesianism as developed above, which takes evidence into account by means of calibration to empirical probabilities and applying the maximal entropy principle.[9]

The question nevertheless remains as to the connection between the two ways of handling evidence. Landes et al. (2023, Theorem 34) show that updating in OBIL captures Bayesian conditionalisation as a special case:[10]

Theorem 11 (*Bayesian conditionalisation*). *Given categorical premisses* $\varphi_1,\ldots,\varphi_k$*, if* $P_=(\varphi_1\wedge\cdots\wedge\varphi_k)>0$ *then*

$$\mathrm{maxent}\,\mathbb{E}=\{P_=(\cdot|\varphi_1,\ldots,\varphi_k)\}=\{P_=(\cdot|\check{\varphi}_1,\ldots,\check{\varphi}_k)\}.$$

Landes et al. (2023, Theorem 41) go on to show that a generalisation of Bayesian conditionalisation, namely Jeffrey conditionalisation (Jeffrey, 2004, §3.2), also holds in OBIL:

Theorem 12 (*Jeffrey conditionalisation*). *Given* $c\in(0,1)$ *and premiss* φ^c*, if* $P_=(\varphi)\in(0,1)$ *then*

$$\mathrm{maxent}\,\mathbb{E}=\{c\cdot P_=(\cdot|\varphi)+(1-c)\cdot P_=(\cdot|\neg\varphi)\}=\{c\cdot P_=(\cdot|\check{\varphi})+(1-c)\cdot P_=(\cdot|\neg\check{\varphi})\}.$$

§7
Inference

In this section, we turn to questions of inference in OBIL.

[9] Indeed, Williamson (2023) argues that coherent calibration to empirical probabilities requires abandoning conditionalisation in favour of maximising entropy.

[10] This extends a result of Seidenfeld (1986, §2.1) from the finite case to the case of a first-order predicate language.

§7.1. Decidability

As Landes et al. (2024, Corollary 1) observe, inference in inductive logic under the standard semantics is undecidable. However, there is a surprisingly large decidable class of entailment relationships in OBIL—decidable in the sense that there is an effective procedure for deciding whether a given entailment relationship lies in the class and, if so, whether the entailment is inductively valid. To obtain decidability results, we need to suppose that all intervals $X_1,\ldots,X_k,Y$ are expressible in the sense that they have endpoints that are expressed as terminating decimal fractions to some fixed number of decimal places. First, Landes et al. (2024, Theorem 1) show:

Theorem 13 (*Quantifier-free decidability*). *The class of all quantifier-free entailment relationships is decidable in OBIL.*

Then, in virtue of the fact that there is an effective procedure for determining the support of any sentence, and this support is quantifier-free (Landes et al., 2024, Theorem 5),

Theorem 14 (*Support-satisfiable decidability*). *The class of all support-satisfiable entailment relationships is decidable in OBIL.*

§7.2. Truth tables

Given that such a large class of entailment relationships is decidable, the question arises as to how one might determine whether a given support-satisfiable entailment relationship is inductively valid. In this subsection we see that truth tables can be used for inference, while in the next, that probabilistic graphical models can be used for inference.

We will illustrate these methods by providing a running example—see Landes et al. (2024) for more detail on the methods themselves. Consider the following entailment relationship:

$$\exists x V t_1 x,\ (Ut_2 \vee \forall x Rx)^{0.9},\ Ut_1 \to Vt_1t_3,\ (Ut_1 \vee (\exists x Vxt_3 \to Ut_2))^{[0.95,1]} \mathrel{\mathring{\approx}} Vt_1t_3{}^{[.5,1]}$$

We can enumerate the atomic propositions as follows:

$$a_1 : Ut_1, \quad a_2 : Rt_1, \quad a_3 : Vt_1t_1,$$

$$a_4 : Ut_2, \quad a_5 : Rt_2, \quad a_6 : Vt_1t_2, \quad a_7 : Vt_2t_1, \quad a_8 : Vt_2t_2,$$

$$a_9 : Ut_3, \quad a_{10} : Rt_3, \quad a_{11} : Vt_1t_3, \quad a_{12} : Vt_3t_1, \quad a_{13} : Vt_2t_3, \quad \cdots$$

The support of each premiss sentence is:

i	φ_i	$\check{\varphi}_i$
1	$\exists x V t_1 x$	$a_1 \vee \neg a_1$
2	$Ut_2 \vee \forall x Rx$	a_4
3	$Ut_1 \to Vt_1t_3$	$a_1 \to a_{11}$
4	$Ut_1 \vee (\exists x Vxt_3 \to Ut_2)$	$a_1 \vee a_4$

Note in particular for φ_1, i.e., $\exists x V t_1 x$, we have that $\breve{\varphi}_1$ is $a_1 \vee \neg a_1$ because φ_1 mentions no atomic propositions and $P_=(a_1 \wedge \exists x V t_1 x) = P_=(\neg a_1 \wedge \exists x V t_1 x) = 1/2 > 0$. Intuitively, since $\exists x V t_1 x$ is an inductive tautology, it is treated as if it were a deductive tautology when determining $\breve{\varphi}_1$. In determining $\breve{\varphi}_2$, the disjunct $\forall x R x$ is an inductive contradiction and so ignored. Strictly speaking, $\breve{\varphi}_3$ is defined as $(a_1 \wedge a_{11}) \vee (\neg a_1 \wedge a_{11}) \vee (\neg a_1 \wedge \neg a_{11})$ but we abbreviate this sentence by the logically equivalent sentence $a_1 \to a_{11}$. Similarly for $\breve{\varphi}_4$. Note that $\exists x V x t_3$ is an inductive tautology and hence treated as a deductive tautology when determining $\breve{\varphi}_4$. The support of the conclusion is simply a_{11}.

The support inference is thus:

$$a_1 \vee \neg a_1,\ a_4^{.9},\ a_1 \to a_{11},\ a_1 \vee a_4^{[0.95,1]} \mathrel{|\!\!\!\approx}^{\circ} a_{11}^{[0.5,1]}$$

The states of the atomic propositions that occur in the inference are $\Xi = \{\pm a_1 \wedge \pm a_4 \wedge \pm a_{11}\}$. These correspond to rows of the truth table for the support sentences:

$P^\dagger$	a_1	a_4	a_{11}	$a_1 \vee \neg a_1$	a_4	$a_1 \to a_{11}$	$a_1 \vee a_4$	a_{11}
0.3	T	T	T	T	T	T	T	T
0	T	T	F	T	T	F	T	F
0.05	T	F	T	T	F	T	T	T
0	T	F	F	T	F	F	T	F
0.3	F	T	T	T	T	T	T	T
0.3	F	T	F	T	T	T	T	F
0.025	F	F	T	T	F	T	F	T
0.025	F	F	F	T	F	T	F	F

The first column of the truth table shows the probability given to the state corresponding to each row of the truth table by the maximal entropy function $P^\dagger$. This is found by maximising the entropy $-\sum_{\xi \in \Xi} P(\xi) \log P(\xi)$ subject to the constraints imposed by the premisses. Standard numerical methods such as gradient ascent can be used to straightforwardly obtain the values that maximise entropy (see, e.g., Boyd and Vandenberghe, 2004). The probability of each statement is the sum of the probabilities at the rows of the truth table at which it is true. We see then that $P^\dagger(a_{11}) = 0.3 + 0.05 + 0.3 + 0.025 = 0.675 \in [0.5, 1]$, so the entailment relationship is indeed inductively valid in OBIL.

§7.3. Objective Bayesian Nets

The truth-table method is simple and powerful, but does not scale well to large and complex entailment relationships, because the number of rows in a truth table increases exponentially in the number of atomic sentences that feature in it. However, there is an alternative method for inference that is more tractable in many cases. This is the graphical modelling approach of *objective Bayesian networks*. This approach was originally applied to finite propositional inductive logic (Williamson, 2005b, 2008; Haenni et al., 2011), but has been extended to first-order OBIL (Landes et al., 2024).

An objective Bayesian network (OBN) is a Bayesian network that represents the maximal entropy probability function $P^\dagger$. It consists of a directed acyclic graph whose vertices are the atomic sentences that feature in the support inference, together with the probability distribution of each atomic sentence conditional on each of its parents in the graph. From the OBN one can calculate the probability of the conclusion sentence to determine whether the entailment relationship holds. The distinctive feature of an OBN is that, unlike other kinds of Bayesian net, the construction of the directed acyclic graph is computationally trivial: there is no need for any computationally expensive probabilistic independence tests, for example.

Consider again the entailment relationship of the previous section. To construct an OBN we first construct an undirected graph $\mathscr{G}$ by taking atomic sentences that feature in the supports of the premisses as vertices and linking two atomic sentences if they occur together in one of the premiss supports. Thus, we take a_1, a_4 and a_{11} as vertices because they are the atomic sentences that feature in the premiss supports. We include an edge between a_1 and a_{11} because they both feature in $\check{\varphi}_3$, and an edge between a_1 and a_4 because they both feature in $\check{\varphi}_4$:

$$a_4 \;\frac{\qquad}{}\; a_1 \;\frac{\qquad}{}\; a_{11}$$

Maximal entropy probability functions render atomic propositions probabilistically independent in the absence of any evidence linking them. More precisely, one can prove that separation in $\mathscr{G}$ implies conditional probabilistic independence of the maximal entropy function $P^\dagger$ (Williamson, 2002). In our example, $P^\dagger$ renders a_4 and a_{11} probabilistically independent conditional on a_1.

We next use a standard algorithm to transform $\mathscr{G}$ into a directed acyclic graph $\mathscr{H}$ that preserves as many of the conditional independencies of $\mathscr{G}$ as possible. For example, we can set $\mathscr{H}$ to be:

$$a_4 \longrightarrow a_1 \longrightarrow a_{11}$$

D-separation in $\mathscr{H}$ implies that $P^\dagger$ renders a_4 and a_{11} probabilistically independent conditional on a_1.[11]

Finally, we determine the conditional probability distributions by finding the values of the following parameters that maximise entropy:

$$P(a_4),\ P(a_1|a_4),\ P(a_1|\neg a_4),\ P(a_{11}|a_1),\ P(a_{11}|\neg a_1).$$

A numerical optimisation yields:

$$P(a_4) = 0.9,\ P(a_1|a_4) = 1/3,\ P(a_1|\neg a_4) = 1/2,\ P(a_{11}|a_1) = 1,\ P(a_{11}|\neg a_1) = 1/2.$$

Standard Bayesian network inference algorithms can then be used to determine the probability of a conclusion sentence of interest. For example,

$$\exists x V t_1 x,\ (Ut_2 \vee \forall x Rx)^{0.9},\ Ut_1 \rightarrow V t_1 t_3,\ (Ut_1 \vee (\exists x V x t_3 \rightarrow U t_2))^{[0.95,1]} \\ \overset{\circ}{\mid\!\approx} (\neg(Ut_1 \vee Ut_3) \wedge \exists x V xx \wedge V t_1 t_3)^{0.1625}$$

[11]Subset Z D-separates subsets X from Y of nodes if each path between a node in X and a node in Y contains either (i) some node a_i in Z at which the arrows on the path meet head-to-tail ($\longrightarrow a_i \longrightarrow$) or tail-to-tail ($\longleftarrow a_i \longrightarrow$), or (ii) some node a_j at which the arrows on the path meet head-to-head ($\longrightarrow a_j \longleftarrow$) and neither a_j nor any of its descendants are in Z. The key result is that if Z D-separates X from Y in $\mathscr{H}$ then the maximal entropy function renders X and Y probabilistically independent conditional on Z (Williamson, 2005a, Theorem 5.3).

To see this, note that the support of the conclusion sentence is $\neg a_1 \wedge \neg a_7 \wedge a_{11}$ and that a_7 is not mentioned by any of the premisses so $P^\dagger$ renders a_7 probabilistically independent of a_1 and a_{11} and $P^\dagger(a_7) = 1/2$. Hence, $P^\dagger(\neg a_1 \wedge \neg a_7 \wedge a_{11}) = 1/2 \times P(a_{11}|\neg a_1)(P(\neg a_1|a_4)P(a_4) + P(\neg a_1|\neg a_4)P(\neg a_4)) = 1/2 \times 1/2(2/3 \times 9/10 + 1/2 \times 1/10) = 0.1625$.

A key advantage of OBNs over the truth table method is a reduction in the number of parameters required to specify the maximal entropy function $P^\dagger$. In this example, the support problem involves three atomic propositions and the truth table requires 8 parameters while the OBN requires only 5. Typically, this reduction in the number of parameters becomes very marked as the number of atomic propositions in the premisses increases.

§8
Logical properties

Here we briefly survey some of the logical properties of OBIL.

Some logical properties are common to all standard probabilistic logics (Landes, 2023; Landes et al., 2024). For example:

Preservation of Deductive Entailment. If $\varphi \models \psi$ then $\varphi \mathrel{\overset{\circ}{|\!\approx}} \psi$.

Reflexivity. $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \varphi^X$.

Landes et al. (2023, Corollary 24) prove the following basic fact:

Theorem 15 (*Zero-One Law*). ***Every constant-free sentence is either an inductive contradiction or an inductive tautology.***

Moreover, Landes et al. (2024) show the following:

Preservation of Inductive Tautologies (PIT). If $\mathrel{\overset{\circ}{|\!\approx}} \psi$ and $\varphi_1^{X_1}, \ldots, \varphi_k^{X_k}$ have satisfiable support then $\varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \mathrel{\overset{\circ}{|\!\approx}} \psi$.

Note that PIT implies that inductive contradictions are also preserved by premisses with satisfiable support. Hence, by the zero-one law,

Corollary 16 (*Extended Zero-One Law*). ***Every constant-free sentence is given either probability 1 or 0 by premisses with satisfiable support.***

Landes (2023, §3) derives several logical properties, including:[12]

Modus Ponens. If $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \psi^Y$ and $\mathrel{\overset{\circ}{|\!\approx}} \varphi^X$ then $\mathrel{\overset{\circ}{|\!\approx}} \psi^Y$.

Modus Tollens. If $\varphi \mathrel{\overset{\circ}{|\!\approx}} \psi$ and $\mathrel{\overset{\circ}{|\!\approx}} \neg\psi$ then $\mathrel{\overset{\circ}{|\!\approx}} \neg\varphi$.

Implication. If $\varphi \mathrel{\overset{\circ}{|\!\approx}} \psi$ then $\mathrel{\overset{\circ}{|\!\approx}} \varphi \rightarrow \psi$.

Weak Cautious Monotonicity. If $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \psi^Y$, $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \theta^Z$ and $\mathrel{\overset{\circ}{|\!\not\approx}} \neg\varphi$ then $\varphi^X, \psi^Y \mathrel{\overset{\circ}{|\!\approx}} \theta^Z$, as long as $X \neq \{0\}$.

Cautious Cut. If $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \psi^Y$ and $\varphi^X, \theta^Z \mathrel{\overset{\circ}{|\!\approx}} \psi^Y$ then $\varphi^X \mathrel{\overset{\circ}{|\!\approx}} \theta^Z$.

[12] Note that in deriving these results, Landes assumed that any conclusion follows from unsatisfiable premisses ('explosion').

Very Cautious Or. If $\varphi \overset{\circ}{\approx}\psi$, $\theta \overset{\circ}{\approx}\psi$, $\overset{\circ}{\not\approx}\neg\varphi$ and $\overset{\circ}{\not\approx}\neg\theta$ then $\varphi \vee \theta \overset{\circ}{\approx}\psi$.

However, versions of some standard rules of nonmonotonic logic do not hold in general in OBIL (Landes, 2023, §§4–5). For example,

Rational Monotonicity. If $\varphi^X \approx \psi^Y$ and $\varphi^X \not\approx \neg\theta$ then $\varphi^X, \theta \approx \psi^Y$.

Cautious Monotonicity. If $\varphi^X \approx \psi^Y$ and $\varphi^X \approx \theta^Z$ then $\varphi^X, \psi^Y \approx \theta^Z$.

And. If $\varphi^X \approx \psi^Y$ and $\varphi^X \approx \theta^Y$ then $\varphi^X \approx (\psi \wedge \theta)^Y$.

Transitivity. If $\varphi^X \approx \psi^Y$ and $\psi^Y \approx \theta^Z$ then $\varphi^X \approx \theta^Z$.

Interesting open questions remain. For example, is there a sound and complete set of rules for OBIL?

§9
Language invariance

An inductive logic is *language invariant* if, whenever $\varphi_1, \ldots, \varphi_k, \psi$ are sentences that can be formulated in both $\mathcal{L}^1$ and $\mathcal{L}^2$,

$$\varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \approx^1 \psi^Y \text{ if and only if } \varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \approx^2 \psi^Y,$$

where $\approx^1, \approx^2$ are the entailment relations defined on $\mathcal{L}^1, \mathcal{L}^2$ respectively. An inductive logic may be language invariant on a specific class of entailment relationships, even if it is not language invariant simpliciter.

What is known about language invariance is generally positive:

Theorem 17. *OBIL is language invariant:*

- *for languages that differ only with respect to predicate symbols, on the class of finitely generated entailment relationships. (Williamson, 2017, Theorem 5.9.)*
- *for languages whose constant symbols are permuted, as long as the permutation σ is such that $\{t_1, \ldots, t_n\} = \{t_{\sigma(1)}, \ldots, t_{\sigma(n)}\}$ for sufficiently large n. (Williamson, 2010, Proposition 5.10.)*
- *for languages whose constant symbols are permuted, as long as the permutation σ preserves the set $\mathbb{E} = [\varphi_1^{X_1}, \ldots, \varphi_k^{X_k}]$. (Landes et al., 2023, Proposition 49.)*

A key question for further research concerns the extent to which these results can be generalised.

§10
The entropy-limit conjecture

The maximal entropy principle is an extension of the maximum entropy principle to the infinite predicate language $\mathcal{L}$. It is not the only such extension, however: Barnett and Paris (2008) put forward another approach to generalising the maximum entropy principle from a finite domain to an infinite predicate language $\mathcal{L}$. The question thus arises as to the relationship between the two approaches.

The core idea of Barnett and Paris (2008) is to consider a finite sublanguage $\mathcal{L}_n$ in which the premisses are expressible, and to re-interpret the premisses as saying something about a finite domain of $\mathscr{L}_n$, rather than as propositions about the infinite domain of $\mathscr{L}$. Thus $\forall x\theta(x)$ is re-expressed as $\theta(t_1) \wedge \cdots \wedge \theta(t_{r_n})$ and $\exists x\theta(x)$ is re-expressed as $\theta(t_1) \vee \cdots \vee \theta(t_{r_n})$. Then one can find the function $P^{(n)}$ that maximises n-entropy subject to the constraints imposed by the re-expressed premisses. One can consider the behaviour of $P^{(n)}$ as n increases by taking the limit $P^{(\infty)}(\psi) \stackrel{\text{df}}{=} \lim_{n\to\infty} P^{(n)}(\psi)$, where it exists. Finally, one can provide *entropy-limit* semantics for inductive logic on $\mathscr{L}$ itself: $\varphi_1^{X_1}, \ldots, \varphi_k^{X_k} \approx^{\infty} \psi^Y$ if and only if $P^{(\infty)}(\psi) \in Y$.

As it stands, this inductive logic is only partially defined, as the limit function $P^{(\infty)}$ does not always exist—for example, in certain cases where there are categorical Σ_2 or Π_2 premisses. However, Williamson (2017, p. 191) put forward the following conjecture:

Entropy-limit conjecture. Entropy-limit entailment $\approx^{\infty}$ agrees with maximal-entropy entailment $\stackrel{\circ}{\approx}$ wherever the limit $P^{(\infty)}$ exists and is in $\mathbb{E}$.

This conjecture has been explored in a number of cases and found to hold in each:

Theorem 18. *The entropy-limit conjecture has been verified to hold in the following special cases:*

- *Inferences with categorical monadic premisses.* *(Barnett and Paris, 2008; Rafiee Rad, 2009, Theorem 29; Rafiee Rad, 2021)*
- *Inferences with categorical Σ_1 premisses.* *(Rafiee Rad, 2018)*
- *Inferences with categorical Π_1 premisses.* *(Landes et al., 2021)*
- *Inferences from a premiss of the form φ^c, where the entropy-limit conjecture holds for both φ and $\neg\varphi$.* *(Landes et al., 2021)*
- *Certain inferences in which $P^{(\infty)}$ is a limit in entropy of the $P^{(n)}$.* *(Landes et al., 2021)*

If the entropy-limit conjecture were to hold generally, it would support the claim that there is a canonical inductive logic and that OBIL captures that logic. Thus, the status of the entropy-limit conjecture is one of the central open questions in this field.

§11
Conclusions

As we have seen, the maximal entropy principle generalises Jaynes' maximum entropy principle and underpins an inductive logic defined on a first-order predicate language. This logic, objective Bayesian inductive logic, suffers neither from the underdetermination of the Carnapian and precise Bayesian approaches nor from the weak inferences of an imprecise Bayesian approach. A large class of entailment relationships (including all those that satisfy closure in entropy) induce a unique maximal entropy function, and a large subclass of those entailment relationships (the class of those whose premisses have satisfiable support) is known to be decidable. These entailment relationships can be verified by means of truth tables or objective Bayesian nets. OBIL has been found to be language independent in key cases, and to agree with the alternative approach of Barnett and Paris (2008) in key cases. Thus OBIL is now a mature and promising theory.

But there remain many open questions, not just concerning language independence and the entropy-limit conjecture, but also, for example, the question of whether OBIL can be characterised by logical rules, the extent to which limits in entropy exist, and the behaviour of OBIL in cases where $\mathbb{E}$ is not closed in entropy. Thus, there are many opportunities to advance the theory and develop our understanding of OBIL. Enough has been done to demonstrate the potential of OBIL, but not so much as to rob us of exciting new results.

Acknowledgements

Most of the results discussed here are the product of joint research with Jürgen Landes and Soroush Rafiee Rad. This research was supported by Leverhulme Trust grant RPG-2022-336.

Bibliography

Barnett, O. and Paris, J. (2008). Maximum entropy inference with quantified knowledge. *Logic Journal of the IGPL*, 16(1):85–98.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization.* Cambridge University Press, Cambridge.

Carnap, R. (1952). *The continuum of inductive methods.* University of Chicago Press, Chicago IL.

de Finetti, B. (1937). Foresight. Its logical laws, its subjective sources. In Kyburg, H. E. and Smokler, H. E., editors, *Studies in subjective probability*, pages 53–118. Robert E. Krieger Publishing Company, Huntington, NY, second (1980) edition.

Haenni, R., Romeijn, J.-W., Wheeler, G., and Williamson, J. (2011). *Probabilistic logics and probabilistic networks.* Synthese Library. Springer, Dordrecht.

Howson, C. (2000). *Hume's problem: induction and the justification of belief.* Clarendon Press, Oxford.

Howson, C. and Urbach, P. (1989). *Scientific reasoning: the Bayesian approach.* Open Court, Chicago IL, third (2006) edition.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *The Physical Review*, 106(4):620–630.

Jaynes, E. T. (1979). Where do we stand on maximum entropy? In Levine, R. D. and Tribus, M., editors, *The maximum entropy formalism*, pages 15–118. MIT Press, Cambridge MA.

Jaynes, E. T. (2003). *Probability theory: the logic of science.* Cambridge University Press, Cambridge.

Jeffrey, R. (2004). *Subjective probability: the real thing.* Cambridge University Press, Cambridge.

Keynes, J. M. (1921). *A treatise on probability.* Macmillan, London, 1973 edition.

Landes, J. (2023). Rules of proof for maximal entropy inference. *International Journal of Approximate Reasoning*, 153:144–171.

Landes, J., Rafiee Rad, S., and Williamson, J. (2021). Towards the entropy-limit conjecture. *Annals of Pure and Applied Logic*, 172(2):102870.

Landes, J., Rafiee Rad, S., and Williamson, J. (2023). Determining maximal entropy functions for objective Bayesian inductive logic. *Journal of Philosophical Logic*, 52:555–608.

Landes, J., Rafiee Rad, S., and Williamson, J. (2024). A decidable class of inferences in first-order objective Bayesian inductive logic. *Under review*.

Landes, J. and Williamson, J. (2013). Objective Bayesianism and the maximum entropy principle. *Entropy*, 15(9):3528–3591.

Landes, J. and Williamson, J. (2015). Justifying objective Bayesianism on predicate languages. *Entropy*, 17(4):2459–2543.

Nilsson, N. J. (1986). Probabilistic logic. *Artificial Intelligence*, 28:71–87.

Paris, J. B. (1994). *The uncertain reasoner's companion*. Cambridge University Press, Cambridge.

Rafiee Rad, S. (2009). *Inference processes for probabilistic first order languages*. PhD thesis, Department of Mathematics, University of Manchester. Available at http://www.maths.manchester.ac.uk/~jeff/theses/srthesis.pdf.

Rafiee Rad, S. (2018). Maximum Entropy Models for Σ_1 Sentences. *Journal of Applied Logics - IfCoLoG Journal of Logics and their Applications*, 5(1):287–300.

Rafiee Rad, S. (2021). Probabilistic characterisation of models of first-order theories. *Annals of Pure and Applied Logic*, 172(1):102875.

Ramsey, F. P. (1926). Truth and probability. In Kyburg, H. E. and Smokler, H. E., editors, *Studies in subjective probability*, pages 23–52. Robert E. Krieger Publishing Company, Huntington, New York, second (1980) edition.

Rosenkrantz, R. D. (1977). *Inference, method and decision: towards a Bayesian philosophy of science*. Reidel, Dordrecht.

Seidenfeld, T. (1986). Entropy and uncertainty. *Philosophy of Science*, 53(4):467–491.

Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423 and 623–656.

Topsøe, F. (1979). Information theoretical optimization techniques. *Kybernetika*, 15:1–27.

Williamson, J. (2002). Maximising entropy efficiently. *Electronic Transactions in Artificial Intelligence Journal*, 6. www.etaij.org.

Williamson, J. (2005a). *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.

Williamson, J. (2005b). Objective Bayesian nets. In Artemov, S., Barringer, H., d'Avila Garcez, A. S., Lamb, L. C., and Woods, J., editors, *We Will Show Them! Essays in Honour of Dov Gabbay*, volume 2, pages 713–730. College Publications, London.

Williamson, J. (2008). Objective Bayesian probabilistic logic. *Journal of Algorithms in Cognition, Informatics and Logic*, 63:167–183.

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

Williamson, J. (2017). *Lectures on inductive logic*. Oxford University Press, Oxford.

Williamson, J. (2018). Justifying the principle of indifference. *European Journal for the Philosophy of Science*, 8:559–586.

Williamson, J. (2023). Bayesianism from a philosophical perspective and its application to medicine. *International Journal of Biostatistics*, doi 10.1515/ijb-2022-0043.