Maximum-entropy from the probability calculus: exchangeability, sufficiency

P.G.L. Porta Mana <pgl@portamana.org>

7 June 2017; updated 19 January 2018

Dedicato alla mia fantastica sorellina Marianna per il suo compleanno

The classical maximum-entropy principle method (Jaynes 1963) appears in the probability calculus as an approximation of a particular *model by exchangeability* or a particular *model by sufficiency*.

The approximation from the exchangeability model can be inferred from an analysis by Jaynes (1996) and to some extent from works on entropic priors (Rodríguez 1989; 2002; Skilling 1989a; 1990). I tried to show it explicitly in a simple context (Porta Mana 2009). The approximation from the sufficiency model can be inferred from Bernardo & Smith (2000 § 4.5) and Diaconis & Freedman (1981) in combination with the Koopman-Pitman-Darmois theorem (see references in § 3).

In this note I illustrate how either approximations arises, in turn, and then give a heuristic synopsis of both. At the end I discuss some questions: Prediction or retrodiction? Which of the two models is preferable? (the exchangeable one.) How good is the maximum-entropy approximation? Is this a “derivation” of maximum-entropy?

I assume that you are familiar with: the maximum-(relative)-entropy method (Jaynes 1957a; much clearer in Jaynes 1963; Sivia 2006; Hobson et al. 1973), especially the mathematical form of its distributions and its prescription “expectations = empirical averages”; the probability calculus (Jaynes 2003; Hailperin 1996; Jeffreys 2003; Lindley 2014); the basics of models by exchangeability and sufficiency (Bernardo et al. 2000 ch. 4), although I’ll try to explain the basic ideas behind them – likely you’ve often worked with them even if you’ve never heard of them under these names.

1 Context and notation

We have a potentially infinite set of measurements, each having K possible outcomes. Dice rolls and their six outcomes are a typical example. I use the terms “measurement” and “outcome” to lend concreteness to the discussion, but the formulae below apply to much more general contexts.

The proposition that the n th measurement has outcome k is denoted $E_k^{(n)}$. The relative frequencies of the K possible outcomes in a set of measurements are denoted $f := (f_k)$. It may happen that in a measurement we observe not directly an outcome but an “observable” having values $(O_k) =: \mathbf{O}$ for the K outcomes. This observable may be vector-valued. The *empirical average* of the observable in a set of N measurements with outcomes $\{k_1, \dots, k_N\}$ is $\sum_{n=1}^N O_{k_n}/N$, equivalent to $\sum_k O_k f_k$.

Probabilities have propositions as arguments (for good definitions of what a proposition is – it isn’t a sentence, for example – see Strawson 1964; Copi 1979; Barwise et al. 2003). Johnson’s definition remains one of the simplest and most beautiful: “Probability is a magnitude to be attached to any possibly true or possibly false proposition; not, however, to the proposition in and for itself, but in reference to another proposition the truth of which is supposed to be known” (Johnson 1924 Appendix, § 2). See also Hailperin’s (1996; 2011) formalization, sadly neglected in the literature. The assumptions or knowledge underlying our probabilities – our “model” – will be generically denoted by I , with subscripts denoting specific assumptions. We will sometimes let a quantity stand as abbreviation for a proposition, for example f for “the observed relative frequencies in N measurements are f ”. In such cases the probability symbol will be in lower-case to remind us of our notational sins.

Lest this note become an anthill of indices let’s use the following notation: for positive K -tuples $x := (x_i)$, $y := (y_i)$, and number a ,

$$\begin{aligned} ax &:= (ax_i), & x/y &:= (x_i/y_i), & xy &:= (x_i y_i), & x^y &:= (x_i^{y_i}), \\ \exp x &:= (\exp x_i), & \ln x &:= (\ln x_i), & x! &:= (x_i!), \\ \sum x &:= \sum_k x_k, & \prod x &:= \prod_k x_k, & \binom{a}{ax} &:= \frac{a!}{\prod (ax)!}. \end{aligned} \tag{1}$$

The symbol δ indicates a Dirac delta (Lighthill 1964; even better: Egorov 1990; 2001) or a characteristic function (cf. Knuth 1992), depending on the context.

The Shannon entropy $H(x) := -\sum x \ln x$, and the relative Shannon entropy or negative discrimination information $H(x; y) := -\sum x \ln(x/y)$. Let’s keep in mind the important properties

$$H(x; y) \leq 0, \quad H(x; y) = 0 \Leftrightarrow x = y. \tag{2}$$

The problem typically addressed by maximum-entropy is this: given that in a large number N of measurements we have observed an average having value in a *convex* set A (which can consist of a single number),

$$\sum \mathbf{O}f \in A, \quad (3)$$

what is the probability of having outcome k in an $(N+1)$ th measurement? In symbols,

$$\boxed{P[E_k^{(N+1)} \mid \sum \mathbf{O}f \in A, I] = ?} \quad (4)$$

where I denotes our state of knowledge. The maximum-entropy answer (Mead et al. 1984; Fang et al. 1997; Boyd et al. 2009) has the form

$$\frac{r_k \exp(\lambda O_k)}{\sum r \exp(\lambda O)} \quad (5)$$

where r is a reference distribution and λ is determined by the constraints in a way that we don't need to specify here. The convexity of A ensures the uniqueness of this answer.

2 Maximum-entropy from a model by exchangeability

Let's assume that in our state of knowledge I_x we deem the measurements to be infinitely exchangeable (Bernardo et al. 2000 § 4.2); that is, there can be a potentially unlimited number of them and their indices are irrelevant for our inferences. De Finetti's theorem (1930; 1937; Heath et al. 1976) states that this assumption *forces* us to assign probabilities of this form:

$$\begin{aligned} P[E_{k_1}^{(1)}, \dots, E_{k_N}^{(N)} \mid I_x] &= \int q_{k_1} \cdots q_{k_N} p(q \mid I_x) dq \\ &\equiv \int \left(\prod q^{N_f} \right) p(q \mid I_x) dq, \end{aligned} \quad (6)$$

where the distribution q can be interpreted as the relative frequencies in the long run,¹ and integration is over the $(K-1)$ -dimensional simplex (Grünbaum 2003) of such distributions, $\{q \in \mathbf{R}_{\geq 0}^K \mid \sum q = 1\}$. The term

¹"But this *long run* is a misleading guide to current affairs. *In the long run* we are all dead." (Keynes 2013 § 3.I, p. 65)

$p(q|I_x) dq$ can be interpreted as the prior probability density of observing the long-run frequencies q in an infinite number of measurements. This probability is not determined by the theorem.

Let's call the expression above an *exchangeability model* (Bernardo et al. 2000 § 4.3).

We assume that our state of knowledge I_x is also expressed by a particular prior density for the long-run frequencies:

$$p(q|I_x) dq = \kappa(L, r) \binom{L}{Lq} \prod r^{Lq} dq, \quad L \geq 1, \quad (7)$$

which we can call “multinomial prior” because it is a sort of continuous interpolation of the multinomial distribution (Johnson et al. 1996 ch. 35). In the multinomial distribution each q_k assumes discrete values in $\{0, 1/L, \dots, 1\}$ and the normalizing constant is unity; for this reason the normalizing constant $\kappa(L, r) \approx L$ in eq. (7). The results that follow also hold for any other prior density that is asymptotically equal to the multinomial prior for L large, for example proportional to $\exp[LH(q; r)]$, which appears in Rodríguez's (1989; 2002) entropic prior and in Skilling's (1989a; 1990) prior for “classical” and “quantified” maximum-entropy.

To find the probability (4) queried by maximum-entropy we need the probability for each possible frequency distribution in the N measurements, which by combinatorial arguments is

$$p(f|I_x) = \int \binom{N}{Nf} (\prod q^{Nf}) p(q|I_x) dq. \quad (8)$$

There are $\binom{N+K-1}{K-1}$ possible frequency distributions (Csiszár et al. 2004).

By marginalization over the subset of frequencies consistent with our data, the probability for the empirical average is

$$P(\sum O f \in A | I_x) = \sum_f \delta(\sum O f \in A) \int \binom{N}{Nf} (\prod q^{Nf}) p(q|I_x) dq. \quad (9)$$

Finally using Bayes's theorem with the probabilities (6)–(9) we find

$$P[E_k^{(N+1)} | \sum O f \in A, I_x] = \frac{\int q_k \sum_f \delta(\sum O f \in A) \binom{N}{Nf} (\prod q^{Nf}) p(q|I_x) dq}{\int \sum_f \delta(\sum O f \in A) \binom{N}{Nf} (\prod q^{Nf}) p(q|I_x) dq}, \quad (10)$$

where the density $p(q|I_x)dq$ is specified in eq. (7), even though the formula above holds as well with any other prior density.

I have graphically emphasized this formula because it is the *exact* answer given to the question (4) by a general exchangeability model: it holds for all numbers K of possible outcomes, all numbers N of observations, and all sets A – even *non-convex* ones.

If N and L are large we can use the bounds of the multinomial (Csiszár et al. 1981 Lemma 2.3)

$$\binom{N}{Nf} = \epsilon(N, f) \exp[N H(f)], \quad (N+1)^{-K} \leq \epsilon(N, f) \leq 1, \quad (11)$$

analogously for $\binom{L}{Lr}$.

From the bounds above it can be shown that the exact probability expression (10) has the asymptotic form

$$P[E_k^{(N+1)} | \sum \mathbf{O}f \in A, I_x] \simeq \kappa(N, L, r) \int q_k \sum_f \delta(\sum \mathbf{O}f \in A) \exp[NH(f; q) + LH(q; r)] dq, \\ N, L \text{ large.} \quad (12)$$

I prefer the symbol “ $\simeq$ ”, “is asymptotically equal to” (ISO 2009; IEEE 1993; IUPAC 2007), to the limit symbol “ $\rightarrow$ ” because the latter may invite to think about a *sequence*, but no such sequence exists. In each specific problem N has one, fixed, possibly unknown value, and cannot be increased at will. The symbol “ $\simeq$ ” says that the right side differs from the left side by an error that may be negligible. It is our duty to check whether this error is really negligible for our purposes.

The asymptotic expression above shows an interesting interplay of two relative entropies. The two exponential terms give rise to two Dirac deltas. The delta in f requires some mathematical care owing to the discreteness of this quantity; see Csiszár (1984; 1985), and Berk (1966; 1970) for an analysis in a slightly different setting. In particular, if $N < K$ the discrete set of $\binom{N+K-1}{K-1}$ possible frequency distributions lies within the $(N-1)$ -dimensional facets of the $(K-1)$ -dimensional simplex of distributions q ; it does not “fill” the simplex. In this case the frequency sum $\sum_f$ cannot be meaningfully approximated by an integral. The approximations below are valid if the number N of observations is much larger than the number K of possible outcomes.

If L/N is also large, taking limits in the proper order gives

$$P[E_k^{(N+1)} | \sum \mathbf{O}f \in A, I_x] \simeq r_k, \quad N, L, L/N \text{ large.} \quad (13)$$

Note how the data about the average (3) are practically discarded in this (L/N) -large case. Compare with Skilling's remark that the parameter L (his α) shouldn't be "particularly large" (cf. Skilling 1998 p. 2).

The asymptotic case that interests us is N/L large: the exponential in N dominates the integral of eq. (12), which becomes

$$\kappa(L, r) \sum_f f_k \delta[\sum \mathbf{O}f \in A] \exp[LH(f; r)] \simeq \arg \sup_f^{\sum \mathbf{O}f \in A} H(f; r), \quad (14)$$

so that, finally,

$$P[E_k^{(N+1)} | \sum \mathbf{O}f \in A, I_x] \simeq f_k^*, \quad N, L, L/N \text{ large,} \\ \text{with } f^* \text{ maximizing } H(f; r) \text{ under constraints } \sum \mathbf{O}f \in A, \quad (15)$$

which is the maximum-entropy recipe, giving the distribution (5).

3 Maximum-entropy from a model by sufficiency

Consider the following assumption or working hypothesis, denoted I_s : To predict the outcome of an $(N + 1)$ th measurement given knowledge of the outcomes of N measurements, all we need to know is the average $\sum \mathbf{O}f$ of an observable $\mathbf{O}$ in those N measurements, no matter the value of N . In other words, any data about known measurements, besides the empirical average of $\mathbf{O}$, is irrelevant for our prediction. The average $\sum \mathbf{O}f$ is then called a *minimal sufficient statistics* (Bernardo et al. 2000 § 4.5; Lindley 2008 § 5.5). In symbols,

$$P[E_k^{(N+1)} | E_{k_1}^{(1)}, \dots, E_{k_N}^{(N)}, I_s] = P[E_k^{(N+1)} | \sum \mathbf{O}f, N, I_s]. \quad (16)$$

Note that the data $\{E_{k_n}^{(n)}\}$ determine the data $\{\sum \mathbf{O}f, N\}$ but not vice versa, so some data have effectively been discarded in the conditional.

The Koopman-Pitman-Darmois theorem (Koopman 1936; Pitman 1936; Darmois 1935; see also later analyses: Hipp 1974; Andersen 1970; Denny 1967; Fraser 1963; Barankin et al. 1963) states that this assumption

forces us to assign probabilities of this form:

$$\begin{aligned} P[E_{k_1}^{(1)}, \dots, E_{k_N}^{(N)} | I_s] &= \int p(k_1 | \lambda, \mathbf{r}, I_s) \cdots p(k_N | \lambda, \mathbf{r}, I_s) p(\lambda | I_s) d\lambda, \\ &\equiv \int [\prod p(k | \lambda, \mathbf{r}, I_s)^{N_f}] p(\lambda | I_s) d\lambda, \end{aligned} \quad (17a)$$

$$\text{with } p(k | \lambda, \mathbf{r}, I_s) := \mathbf{r} \frac{\exp(\lambda \mathbf{O})}{Z(\lambda)}, \quad Z(\lambda) := \sum \mathbf{r} \exp(\lambda \mathbf{O}), \quad (17b)$$

and we have defined $p(k | \dots) := (p(1 | \dots), \dots, p(K | \dots))$. The integration of the parameter λ is over $\mathbf{R}^M$, with M the dimension of the vector-valued observable $\mathbf{O}$, and $\mathbf{r}$ is a K -dimensional distribution. Neither $\mathbf{r}$ or the distribution $p(\lambda | I_s)$ are determined by the theorem.

Let's call the expression above a *sufficiency model* (Bernardo et al. 2000 § 4.5). A sufficiency model can be viewed as a mixture, with weight density $p(\lambda | I_s) d\lambda$, of distributions having maximum-entropy form (5) with multipliers λ .

To find the probability (4) we calculate, as in the previous section, the probabilities for the frequencies:

$$p(f | I_s) = \int \binom{N}{N_f} [\prod p(k | \lambda, \mathbf{r}, I_s)^{N_f}] p(\lambda | I_s) d\lambda, \quad (18)$$

and for the empirical average by marginalization:

$$\begin{aligned} P(\sum \mathbf{O} f \in A | I_s) &= \\ \sum_f \delta(\sum \mathbf{O} f \in A) \int \binom{N}{N_f} [\prod p(k | \lambda, \mathbf{r}, I_s)^{N_f}] p(\lambda | I_s) d\lambda. \end{aligned} \quad (19)$$

From these using Bayes's theorem we finally find

$$\begin{aligned} P[E_k^{(N+1)} | \sum \mathbf{O} f \in A, I_s] &= \\ \frac{\int p(k | \lambda, \mathbf{r}, I_s) \sum_f \delta(\sum \mathbf{O} f \in A) \binom{N}{N_f} [\prod p(k | \lambda, \mathbf{r}, I_s)^{N_f}] p(\lambda | I_s) d\lambda}{\int \sum_f \delta(\sum \mathbf{O} f \in A) \binom{N}{N_f} [\prod p(k | \lambda, \mathbf{r}, I_s)^{N_f}] p(\lambda | I_s) d\lambda}. \end{aligned} \quad (20)$$

This is the *exact* answer given to the maximum-entropy question by a sufficiency model *if the constraints used in maximum-entropy are considered to be a sufficient statistics*. This proviso has serious consequences discussed in § 5.2. The expression above holds for all N and all sets A , even *non-convex* ones.

The asymptotic analysis for large N uses again the multinomial's bounds (11). We find

$$P[E_k^{(N+1)} | \sum \mathbf{O} f \in A, I_s] \simeq \kappa(N, \mathbf{r}) \int p(k | \lambda, \mathbf{r}, I_s) \times \\ \sum_f \delta(\sum \mathbf{O} f \in A) \exp\{NH[f; p(k | \lambda, \mathbf{r}, I_s)]\} p(\lambda | I_s) d\lambda, \\ N \text{ large.} \quad (21)$$

A rigorous analysis of this limit can be done using “information projections” (Berk 1966; 1970; Csiszár 1984; 1985); here is a heuristic summary. Consider the sum in f for fixed λ . We have two cases. (1) If λ is such that $\sum \mathbf{O} p(k | \lambda, \mathbf{r}, I_s) \in A$, there exists a unique f in the sum for which the relative entropy in the exponential reaches its maximum, zero, making the exponential unity. For all other f the relative entropy is negative and the exponential asymptotically vanishes for large N . The integral therefore doesn't vanish asymptotically. (2) If λ is such that $p(k | \lambda, \mathbf{r}, I_s)$ doesn't satisfy the constraints, the relative entropy in the exponential will be negative for all f in the sum, making the exponential asymptotically vanish for all f . The integral therefore vanishes asymptotically. The distinction between these two cases actually requires mathematical care owing to the discreteness of the sum. The f sum then acts as a delta or characteristic function (depending on whether A has measure zero or not):

$$\sum_f \delta(\sum \mathbf{O} f \in A) \exp\{NH[f; p(k | \lambda, \mathbf{r}, I_s)]\} \simeq \delta[\sum \mathbf{O} p(k | \lambda, \mathbf{r}, I_s) \in A]. \\ (22)$$

Thus asymptotically we have, using the explicit expression (17b) for $p(k | \lambda, \mathbf{r}, I_s)$:

$$P[E_k^{(N+1)} | \sum \mathbf{O} f \in A, I_s] \simeq \\ \int \delta\left[\sum \mathbf{O} \mathbf{r} \frac{\exp(\lambda \mathbf{O})}{Z(\lambda)} \in A\right] r_k \frac{\exp(\lambda O_k)}{Z(\lambda)} p(\lambda | I_s) d\lambda, \quad N \text{ large.} \quad (23)$$

This result can also be found first integrating λ and then summing f , using a heuristic argument similar to the one above. This is a mixture, with weight density $p(\lambda | I_s) d\lambda$, of maximum-relative-entropy distributions f^* that satisfy the individual constraints $\sum O f^* = a$, $a \in A$. The final distribution thus differs from the maximum-entropy one if the set A is not a singleton: maximum-entropy would pick up only one distribution. But if the constraint set is a singleton, $A = \{a\}$, we do obtain the same answer (5) as the maximum-entropy recipe:

$$P[E_k^{(N+1)} | \sum O f = a, I_x] \simeq f_k^*, \quad N \text{ large,} \\ \text{with } f^* \text{ maximizing } H(f; r) \text{ under constraints } \sum O f = a. \quad (24)$$

4 Heuristic explanation of both asymptotic approximations

First of all let's note that both the exchangeability (6) and sufficiency (17) models have the parametric form

$$P[E_{k_1}^{(1)}, \dots, E_{k_N}^{(N)} | I] = \int p(k_1 | \nu, I) \cdots p(k_N | \nu, I) p(\nu | I) d\nu \\ \equiv \int [\prod p(k | \nu, I)^{N_f}] p(\nu | I) d\nu. \quad (25)$$

The final probability distribution p for the K outcomes of the $(N + 1)$ th measurement belongs to a $(K - 1)$ -dimensional simplex $\{p \in \mathbf{R}_{\geq 0}^K \mid \sum p = 1\}$. The expression above first selects, within this simplex, a family of distributions $\{p(k | \nu, I)\}$ parametrized by ν ; then it delivers the distribution p as a mixture of the distributions of this family, with weight density $p(\nu | I) d\nu$. In the exchangeability model this family is actually the whole simplex (that's why it's sometimes called a "non-parametric" model). In the sufficiency model it is an *exponential family* (Bernardo et al. 2000 § 4.5.3; Barndorff-Nielsen 2014).

When we conditionalize on data D , the weight density is determined by the mutual modulation of two weights: that of the probability of the

data $p(D|\nu, I)$ and the initial weight $p(\nu|I)$. Pictorially, if $K = 3$:

$$p(D|\nu, I) \quad \times \quad p(\nu|I) \quad = \quad \kappa p(\nu|D, I) \quad (26)$$

the final p is given by the mixture with the weight density $p(\nu|D, I) d\nu$ ensuing from this modulation. The mathematical expression of the data weight $p(D|\nu, I)$ is typically exponentiated to the number of measurements N from which the data originate; compare with eqs (19), (20). If N is large this weight is very peaked on the subset of distributions that give highest probability to the data, that is, that have expectations very close to the empirical averages. It effectively restricts the second weight $p(\nu|I) d\nu$ to such “data subset”. In our case the data subset consists of all distributions satisfying the constraints.

The mechanism described so far is common to the exchangeability and the sufficiency model. Their difference lies in how they choose the final distribution from the data subset.

In the exchangeability model (6) the choice is made by the weight density $p(\nu|I) d\nu$, i.e. the multinomial prior (7). It is extremely peaked owing to the large parameter L , and its level curves are isentropics. Once it’s restricted to the data subset by the data weight $p(D|\nu, I)$, it gives highest weight to the distribution p lying on the highest isentropic curve, which is unique if the data subset is convex; compare with fig.-eq. (26). Hence this is a maximum-entropy distribution satisfying the data constraints. For this mechanism to work it’s necessary that the dominance of the data weight comes first, and the dominance of the multinomial prior comes second. This is the reason why the correct asymptotic limit (15) has N , L , and N/L large.

In the sufficiency model (17) the choice is made by the family of distributions $\{p(k|\nu, I)\}_{\nu}$. These distributions have by construction a maximum-entropy form for the particular observable O . This family intersects the data subset in only one point if the constraint has the form $\sum O f = a$. This point is therefore the maximum-entropy distribution satisfying the data constraints.

The mechanism above also explains why these two models still work if the data subset is non-convex and touches the highest isentropics (exchangeability model) or the exponential family (sufficiency model) in multiple points, bringing the maximum-entropy recipe to an impasse. The final distribution will simply be an equal mixture of such tangency points; it may well lie outside of the data subset.

5 Discussion

5.1 Prediction or retrodiction? An essential aspect of the maximum-entropy method is surprisingly often disregarded in the literature. If we have data from N measurements, we can ask two questions:

“Prediction”: what is the outcome of a *further* similar measurement?

“Retrodiction”: what is the outcome of the first of the N measurements?

Note that despite the literal meaning of these terms *the distinction is not between future and past, but between unknown and partially known*.

Some literature rarely makes clear whether the maximum-entropy probabilities refer to the predictive or to the retrodictive question. Yet these two questions are fundamentally different; their answers rely on very different principles.

To answer the retrodictive question we can – but need not – fully rely on symmetry principles in the discrete case. It is a matter of combinatorics and equal probabilities; a drawing-from-an-urn problem. Most derivations of the maximum-entropy method (e.g. Jaynes 1963; Shore et al. 1980; van Campenhout et al. 1981; Csiszár 1985) address this question only, as often betrayed by the presence of “ $p(x_1)$ ” or similar expressions in their final formulae.

To answer the predictive question, symmetry and combinatorics alone are no use: additional principles are needed. This is the profound philosophical question of *induction*, with its ocean of literature; my favourite sample are the classic Hume (1896 book I, § III.VI), Johnson (1922 esp. chs VIII ff; 1924 Appendix; 1932), de Finetti (1937; 1959), Jeffreys (1955; 1973 ch. I; 2003 § 1.0), Jaynes (2003 § 9.4). De Finetti, foreshadowed by Johnson, was probably the one who expressed most strongly, and explained brilliantly, that the probability calculus does not and cannot *explain* or *justify* our inductive reasoning; it only *expresses* it in a quantitative way. This shift in perspective was very much like

Galilei's shift from *why* to *how* in the study of physical phenomena.² We do inductive inferences in many different ways (Jaynes 2003 § 9.4). The notion of exchangeability (de Finetti 1937; Johnson 1924 Appendix; 1932) captures one of the most intuitive and expresses it mathematically.

The calculations of the previous sections and the final probabilities (10), (20) for our two models pertain the predictive question, as clear from the $E^{(N+1)}$ in their arguments. The two models can also be used to answer the retrodictive question. The resulting formulae are different; they can again be found applying the rules of the probability calculus and Bayes's theorem. The retrodictive formula for the exchangeability model is (proof in Porta Mana 2009 § B):

$$P[E_k^{(n)} | \sum O f \in A, I_x] = \frac{\int \sum_f f_k \delta(\sum O f \in A) \binom{N}{N_f} (\prod q^{N_f}) p(q | I_x) dq}{\int \sum_f \delta(\sum O f \in A) \binom{N}{N_f} (\prod q^{N_f}) p(q | I_x) dq} dq, \\ n \in \{1, \dots, N\}. \quad (27)$$

Graphically it differs from the predictive one (10) only in the replacement of q_k by f_k . An analogous replacement appears in the retrodictive formula for the sufficiency model. But this graphically simple replacement leads to a mechanism very different from the one of § 4 in delivering the final probability: it's a mixture on the data subset rather than on the whole simplex. Predictive and retrodictive probabilities can therefore be very different for small N . See for example figs 1 and 2 below and their accompanying discussion.

This means that the goodness of the maximum-entropy distribution as an approximation of our two models can depend on whether we are asking a *predictive* or a *retrodictive* question. This fact is very important in every application.

²"According to credible traditions it was in the sixteenth century, an age of very intense spiritual emotions, that people gradually ceased trying, as they had been trying all through two thousand years of religious and philosophic speculation, to penetrate into the secrets of Nature, and instead contented themselves, in a way that can only be called superficial, with investigations of its surface. The great Galileo, who is always the first to be mentioned in this connection, did away with the problem, for instance, of the intrinsic reasons why Nature abhors a vacuum, so that it will cause a falling body to enter into and occupy space after space until it finally comes to rest on solid ground, and contented himself with a much more general observation: he simply established the speed at which such a body falls, what course it takes, what time it takes, and what its rate of acceleration is. The Catholic Church made a grave mistake in threatening this man with death and forcing him to recant, instead of exterminating him without more ado." (Musil 1979 vol. 1, ch. 72)

5.2 Which of the two models is preferable? A maximum-entropy distribution can be seen as an approximation of the distribution obtained from the exchangeability model or the sufficiency one (*repetita iuvant*). The two inferential models are not equivalent though, and there are reasons to prefer the exchangeability one – despite the frequent association, in the literature, of maximum-entropy with exponential families. The most important and quite serious difference is this:

Suppose that we have used either model to assign a predictive distribution conditional on the empirical average a of the observable O , obtained from N measurements. If N is large the distributions obtained from either model will be approximately equal, and equal to the maximum-entropy one. Now someone gives us a new empirical average a' of a different observable O' , obtained from the *same* N measurements. This observable turns out to be complementary to the previous one, in the sense that in general from knowing the value of $\sum O f$ we cannot deduce the value of $\sum O' f$, and vice versa. These new data therefore reveal more about the outcomes of our N measurements and of possible further measurements.

The new empirical average a' can be incorporated in the exchangeability model; the resulting predictive and retrodictive distributions conditional on (a', a) will be numerically different from the ones conditional on a only. They will be approximated by a maximum-entropy one based on the old *and* new constraints.

If we incorporate the new average in the sufficiency model, however, *the resulting predictive conditional distribution will be unchanged: knowledge of the new data has no effect in the prediction of new measurements*. The reason is simple: the sufficiency model expresses by construction that the average of the old observable O is all we need for our inferences about further measurements. Any other observable is irrelevant. The new average automatically drops out under predictive conditioning. The only way to obtain a different predictive conditional distribution would be to *discard* the sufficiency model based on O , and use a new one based on (O, O') . But that would be cheating!

This shows how dramatically absolute and categorical the assumption of the existence of a sufficient statistics is. The difficulty above doesn't happen for the retrodictive distribution; the proof is left as an exercise for you.

Since the maximum-entropy method is meant to always employ new constraints, we deduce that it's more correct to interpret it as an approximation of the exchangeability model than of the sufficiency model.

5.3 How good is the maximum-entropy approximation? How does maximum-entropy compare with the exchangeability model (6) with multinomial prior (7) away from the asymptotic approximation?

Their distributions are compared in the classic example of dice rolling in figs 1 and 2 for empirical averages of 5 and 6 (see Porta Mana 2009 for the calculations). The **maximum-entropy distribution** (red) is at the top; the distribution of the exchangeability model with $L = 1$ (blue) and $L = 50$ (bluish purple) is shown underneath for the cases $N = 2$, $N = 12$, $N = \infty$, and for the retrodiction of an “old roll” $E_k^{(n)}$, $n \in \{1, \dots, N\}$, and the prediction of a “new roll” $E_k^{(N+1)}$. The charts also report the Shannon entropies H of the distributions.

The exchangeability model gives very reasonable and even “logical” probabilities for small N . For example, if you obtain an average of 5 in two rolls, it's impossible that either of them was 6 – unless, of course, you own a six-sided die with nine pips on one face. The exchangeability model logically gives zero probability in this case (fig. 1 bottom left). Maximum-entropy gives an erroneous non-zero probability. And having obtained an average of 5 or 6 in two rolls, would you really give a much higher probability to 6 or 5 for a *third* roll? I'd still give 1/6. The exchangeability model reasonably gives an almost uniform distribution, especially for large L (both figures bottom right). The maximum-entropy distribution is unreasonably biased towards high values. If we observe a high average in twelve rolls we start to suspect that the die/dice or the roll technique are biased. The exchangeability model expresses this bias, but more conservatively than maximum-entropy.

In fact the predictive exchangeability-model distribution can have *higher* entropy than the maximum-entropy one! This happens because, when N is small compared to L , the maximum-entropy prescription “what you've seen in N measurements = what you should expect in an $(N + 1)$ th measurement” is silly (MacKay 2003 Exercise 22.13). The exchangeability model intelligently doesn't respect this prescription

Figure 1 Maximum-entropy and exchangeability model, empirical average $a = 5$

Figure 2 Maximum-entropy and exchangeability model, empirical average $a = 6$

strictly, if N isn't large.³ See Porta Mana (2009) for comparisons under other values of the empirical average and of the number of measurements.

When is N large enough for the prescription to become reasonable? In other words, when is maximum-entropy a good approximation of the exchangeability model with multinomial prior? The answer depends on the interplay among the number of measurements N , the number of possible outcomes K , the parameter L , the reference distribution r , and the value a (or range A) of the observed average. The first three ingredients determine the maximum heights of the densities involved in the integral and sum of eq. (10); the last three ingredients determine the size of the effective integration and sum region relative to the integration simplex, and the distance between the peaks of the data weights and the prior weights of fig.-eq. (26). All five ingredients determine how good are the delta approximations in the integral and sum of eq. (10). We saw in § 2, p. 5, after eq. (12), that N needs to be much larger than K for the integral and delta approximations of the frequency sum to be meaningful. Maximum-entropy approximations are not meaningful if the number of possible outcomes is much larger than the number of observations.

It would be very useful to have explicit estimates of the maximum-entropy-approximation error as a function of the four quantities above. I hope to analyse them in a future note, and promise it would be a shorter note.

5.4 Is this a “derivation” of maximum-entropy? The heuristic explanation of § 4 shows that the maximum-entropy distributions appear asymptotically owing to our specific choices of a multinomial prior in the exchangeability model, and of an exponential family with observable O in the sufficiency model. They are therefore not derived only from first principles or from some sort of universal limit. This is why I don't call the asymptotic analysis discussed in this note a “derivation” of the maximum-entropy “principle”. In my opinion this analysis shows that it is not a principle at all.

The information-theoretic arguments – or should we say incentives – behind the standard maximum-entropy recipe can be lifted to a *meta*⁴

³“Obedience is no longer a virtue.” (Milani 1965)

⁴“This is an expression used to hide the absence of any mathematical idea [...]. Personally, I never use this expression in front of children.” (Girard 2001 p. 446)

level and used for priors asymptotically equivalent to the multinomial prior (7), as done by Rodríguez (1989; 2002) for the entropic prior (see also Skilling 1989a; 1990). Such arguments don't determine the parameters L and r , though. They seem to be prone to an infinite regress; Jaynes was aware of this (Jaynes 2003 § 11.1, p. 344).

It would be useful if the multinomial or entropic priors could be uniquely determined by intuitive inferential assumptions, as for example is the case with the Johnson-Dirichlet prior, proportional to $q^L dq$: this prior *must* be used if we believe (denote this by I_j) that the frequencies of other outcomes are irrelevant for predicting a particular one:

$$p[E_k^{(N+1)} | f, N, I_j] = p[E_k^{(N+1)} | f_k, N, I_j], \quad k \in \{1, \dots, K\}, \quad (28)$$

a condition called “sufficientness” (Johnson 1924; 1932; Good 1965 ch. 4; Zabell 1982; Jaynes 1996). Asymptotically it leads to a maximum-entropy distribution with Burg's (1975) entropy $\sum \ln x$ (see Jaynes 1996; Porta Mana 2009).

But, after all, the logical calculus doesn't tell us which truths to choose at the beginning of a logical deduction. Why should the probability calculus tell us which probabilities to choose at the beginning of a probabilistic induction?

5.5 Conclusion Interpreting the maximum-entropy method as an approximation of the exchangeable model (6) with multinomial prior (7) has many advantages:

- it clears up the meaning of the “expectation = average” prescription of the maximum-entropy method;
- it identifies the range of validity of such prescription;
- it quantifies the error of the maximum-entropy approximation;
- it gives a more sensible solution when this approximation doesn't hold;
- it clearly differentiates between prediction and retrodiction;
- it can be backed up by information-theoretic incentives (Rodríguez 1989; 2002) if you're into those.

Disadvantages:

- It can't be used to answer the question “Where did the cat go?”. But this question lies forever beyond the reach of the probability calculus. That's all (Hanshaw 1928).

Thanks

... to Philip Goyal, Moritz Helias, Vahid Rostami, Jakob Jordan, Alper Yegenoglu, Emiliano Torre for many insightful discussions about maximum-entropy. To Mari & Miri for continuous encouragement and affection, and to Buster Keaton and Saitama for filling life with awe and inspiration. To the developers and maintainers of L^AT_EX, Emacs, AUC_TE_X, Open Science Framework, PhilSci, Hal archives, biorXiv, Python, Inkscape, Sci-Hub for making a free and unfiltered scientific exchange possible. ☒

Bibliography

(“van *X*” is listed under *V*; similarly for other prefixes, regardless of national conventions.)

- Andersen, E. B. (1970): *Sufficiency and exponential families for discrete sample spaces*. J. Am. Stat. Assoc. **65**³³¹, 1248–1255.
- Barankin, E. W., Maitra, A. P. (1963): *Generalization of the Fisher-Darmois-Koopman-Pitman theorem on sufficient statistics*. Sankhyā A **25**³, 217–244.
- Barnard, G. A., Jaynes, E. T., Seidenfeld, T., Polasek, W., Csiszár, I. (1985): *Discussion [An extended maximum entropy principle and a Bayesian justification] and Reply*. In: Bernardo, DeGroot, Lindley, Smith (1985), 93–98. See Csiszár (1985).
- Barndorff-Nielsen, O. E. (2014): *Information and Exponential Families: In Statistical Theory*, reprint. (Wiley, New York). First publ. 1978.
- Barwise, J., Etchemendy, J. (2003): *Language, Proof and Logic*. (CSLI, Stanford). Written in collaboration with Gerard Allwein, Dave Barker-Plummer, Albert Liu. First publ. 1999.
- Berk, R. H. (1966): *Limiting behavior of posterior distributions when the model is incorrect*. Ann. Math. Stat. **37**¹, 51–58.
- (1970): *Consistency a posteriori*. Ann. Math. Stat. **41**³, 894–906.
- Bernardo, J.-M., DeGroot, M. H., Lindley, D. V., Smith, A. F. M., eds. (1985): *Bayesian Statistics 2*. (Elsevier and Valencia University Press, Amsterdam and Valencia).
- Bernardo, J.-M., Smith, A. F. (2000): *Bayesian Theory*, reprint. (Wiley, New York). First publ. 1994.
- Boyd, S., Vandenberghe, L. (2009): *Convex Optimization*, 7th printing with corrections. (Cambridge University Press, Cambridge). <http://www.stanford.edu/~boyd/cvxbook/>. First publ. 2004.
- Burg, J. P. (1975): *Maximum entropy spectral analysis*. PhD thesis. (Stanford University, Stanford). <http://sepwww.stanford.edu/data/media/public/oldreports/sep06/>.
- Copi, I. M. (1979): *Symbolic Logic*, 5th ed. (Macmillan, New York). First publ. 1954.
- Csiszár, I. (1984): *Sanov property, generalized I-projection and a conditional limit theorem*. Ann. Prob. **12**³, 768–793.
- (1985): *An extended maximum entropy principle and a Bayesian justification*. In: Bernardo, DeGroot, Lindley, Smith (1985), 83–93. With discussion and reply Barnard, Jaynes, Seidenfeld, Polasek, Csiszár (1985).
- Csiszár, I., Körner, J. (1981): *Information Theory: Coding Theorems for Discrete Memoryless Systems*. (Academic Press, New York). Second ed. Csiszár, Körner (2011).

- Csiszár, I., Körner, J. (2011): *Information Theory: Coding Theorems for Discrete Memoryless Systems*, 2nd ed. (Cambridge University Press, Cambridge). First publ. 1981.
- Csiszár, I., Shields, P. C. (2004): *Information theory and statistics: a tutorial*. Foundations and Trends in Communications and Information Theory **1**⁴, 417–528. <http://www.renyi.hu/~csiszar/>.
- Curien, P.-L. (2001): *Preface to Locus solum*. Math. Struct. in Comp. Science **11**³, 299–300. See also Girard (2001).
- Darmois, G. (1935): *Sur les lois de probabilité à estimation exhaustive*. Comptes rendus hebdomadaires des séances de l'Académie des sciences **200**, 1265–1266.
- de Finetti, B. (1930): *Funzione caratteristica di un fenomeno aleatorio*. Atti Accad. Lincei: Sc. Fis. Mat. Nat. **IV**⁵, 86–133. <http://www.brunodefinetti.it/Opere.htm>.
- (1937): *La prévision : ses lois logiques, ses sources subjectives*. Ann. Inst. Henri Poincaré **7**¹, 1–68. Transl. in Kyburg, Smokler (1980), pp. 53–118, by Henry E. Kyburg, Jr.
- (1959): *La probabilità e la statistica nei rapporti con l'induzione, secondo i diversi punti di vista*. In: de Finetti (2011), 1–115. Transl. in de Finetti (1972), ch. 9, pp. 147–227.
- (1972): *Probability, Induction and Statistics: The art of guessing*. (Wiley, London).
- ed. (2011): *Induzione e statistica*, reprint. (Springer, Berlin). First publ. 1959.
- Demidov, A. S. (2001): *Generalized Functions in Mathematical Physics: Main Ideas and Concepts*. (Nova Science, Huntington, USA). With an addition by Yu. V. Egorov.
- Denny, J. L. (1967): *Sufficient conditions for a family of probabilities to be exponential*. Proc. Natl. Acad. Sci. (USA) **57**⁵, 1184–1187.
- Diaconis, P., Freedman, D. (1981): *Partial exchangeability and sufficiency*. In: Ghosh, Roy (1981), 205–236. <http://statweb.stanford.edu/~cgates/PERSI/year.html>. Also publ. 1982 as technical report https://www.stat.berkeley.edu/~aldous/206-Exch/Papers/diaconis_freedman_PES.pdf.
- Egorov, Yu. V. (1990): *A contribution to the theory of generalized functions*. Russ. Math. Surveys (Uspekhi Mat. Nauk) **45**⁵, 1–49.
- (2001): *A new approach to the theory of generalized functions*. In: Demidov (2001), 117–123.
- Erickson, G. J., Rychert, J. T., Smith, C. R., eds. (1998): *Maximum Entropy and Bayesian Methods*. (Springer, Dordrecht).
- Fang, S.-C., Rajasekera, J. R., Tsao, H.-S. J. (1997): *Entropy Optimization and Mathematical Programming*, reprint. (Springer, New York).
- Ford, K. W., ed. (1963): *Statistical Physics*. (Benjamin, New York).
- Fougère, P. F., ed. (1990): *Maximum Entropy and Bayesian Methods: Dartmouth, U.S.A., 1989*. (Kluwer, Dordrecht).
- Fraser, D. A. S. (1963): *On sufficiency and the exponential family*. J. Roy. Stat. Soc. B **25**¹, 115–123.
- Ghosh, J. K., Roy, J., eds. (1981): *Statistics: Applications and New Directions*. (Indian Statistical Institute, Calcutta).
- Girard, J.-Y. (2001): *Locus solum: From the rules of logic to the logic of rules*. Math. Struct. in Comp. Science **11**³, 301–506. <http://iml.univ-mrs.fr/~girard/Articles.html>. See also Curien (2001).
- Good, I. J. (1965): *The Estimation of Probabilities: An Essay on Modern Bayesian Methods*. (MIT Press, Cambridge, USA).
- Grünbaum, B. (2003): *Convex Polytopes*, 2nd ed. (Springer, New York). Prep. by Volker Kaibel, Victor Klee, and Günter M. Ziegler. First publ. 1967.
- Hailperin, T. (1996): *Sentential Probability Logic: Origins, Development, Current Status, and Technical Applications*. (Associated University Presses, London).

- Hailperin, T. (2011): *Logic with a Probability Semantics: Including Solutions to Some Philosophical Problems*. (Lehigh University Press, Plymouth, UK).
- Hanshaw, A. (1928): *My Blackbirds are Bluebirds now*. (Velvet Tone, Washington, D.C.). With her Sizzlin' Syncopators; written by Cliff Friend, composed by Irving Caesar.
- Heath, D., Sudderth, W. (1976): *De Finetti's theorem on exchangeable variables*. *American Statistician* 30⁴, 188–189.
- Hipp, C. (1974): *Sufficient statistics and exponential families*. *Ann. Stat.* 2⁶, 1283–1292.
- Hobson, A., Cheng, B.-K. (1973): *A comparison of the Shannon and Kullback information measures*. *J. Stat. Phys.* 7⁴, 301–310.
- Hume, D. (1896): *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*, reprint. (Oxford University Press, London). Ed., with an analytical index, by L. A. Selby-Bigge. https://archive.org/details/treatiseofhumann00hume_0. First publ. 1739–1740.
- IEEE (1993): *ANSI/IEEE Std 260.3-1993: American National Standard: Mathematical signs and symbols for use in physical sciences and technology*. Institute of Electrical and Electronics Engineers.
- ISO (2009): *ISO 80000:2009: Quantities and units*. International Organization for Standardization. First publ. 1993.
- IUPAC (2007): *Quantities, Units and Symbols in Physical Chemistry*, 3rd ed. (RSC, Cambridge). Prepared for publication by E. Richard Cohen, Tomislav Cvitaš, Jeremy G. Frey, Bertil Holmström, Kozo Kuchitsu, Roberto Marquardt, Ian Mills, Franco Pavese, Martin Quack, Jürgen Stohner, Herbert L. Strauss, Michio Takami, Anders J Thor. First publ. 1988.
- Jaynes, E. T. (1957a): *Information theory and statistical mechanics*. *Phys. Rev.* 106⁴, 620–630. <http://bayes.wustl.edu/etj/node1.html>, see also Jaynes (1957b).
- (1957b): *Information theory and statistical mechanics. II*. *Phys. Rev.* 108², 171–190. <http://bayes.wustl.edu/etj/node1.html>, see also Jaynes (1957a).
- (1963): *Information theory and statistical mechanics*. In: Ford (1963), 181–218. Repr. in Jaynes (1989), ch. 4, 39–76. <http://bayes.wustl.edu/etj/node1.html>.
- (1989): *E. T. Jaynes: Papers on Probability, Statistics and Statistical Physics*, reprint. (Kluwer, Dordrecht). Ed. by R. D. Rosenkrantz. First publ. 1983.
- (1996): *Monkeys, kangaroos, and N*. <http://bayes.wustl.edu/etj/node1.html>. First publ. 1986. (Errata: in equations (29)–(31), (33), (40), (44), (49) the commas should be replaced by gamma functions, and on p. 19 the value 0.915 should be replaced by 0.0915).
- (2003): *Probability Theory: The Logic of Science*. (Cambridge University Press, Cambridge). Ed. by G. Larry Bretthorst; <http://www-biba.inrialpes.fr/Jaynes/prob.html>, <http://omega.albany.edu:8008/JaynesBook.html>, <http://omega.albany.edu:8008/JaynesBookPdf.html>. First publ. 1994.
- Jeffreys, H. (1955): *The present position in probability theory*. *Brit. J. Phil. Sci.* 5²⁰, 275–289.
- (1973): *Scientific Inference*, 3rd ed. (Cambridge University Press, Cambridge). First publ. 1931.
- (2003): *Theory of Probability*, 3rd ed. (Oxford University Press, London). First publ. 1939.
- Johnson, N. L., Kotz, S., Balakrishnan, N. (1996): *Discrete Multivariate Distributions*. (Wiley, New York). First publ. 1969 in chapter form.
- Johnson, W. E. (1922): *Logic. Part II: Demonstrative Inference: Deductive and Inductive*. (Cambridge University Press, Cambridge).

- Johnson, W. E. (1924): *Logic. Part III: The Logical Foundations of Science*. (Cambridge University Press, Cambridge). <https://archive.org/details/logic03john>, <https://archive.org/details/johnsonslogic03johnuoft>.
- (1932): *Probability: the deductive and inductive problems*. *Mind* **41**¹⁶⁴, 409–423. With some notes and an appendix by R. B. Braithwaite.
- Keynes, J. M. (2013): *A Tract on Monetary Reform*, repr. of second ed. (Cambridge University Press, Cambridge). First publ. 1923.
- Knuth, D. E. (1992): *Two notes on notation*. *Am. Math. Monthly* **99**⁵, 403–422. [arXiv: math/9205211](https://arxiv.org/abs/math/9205211).
- Koopman, B. O. (1936): *On distributions admitting a sufficient statistic*. *Trans. Am. Math. Soc.* **39**³, 399–409.
- Kyburg Jr., H. E., Smokler, H. E., eds. (1980): *Studies in Subjective Probability*, 2nd ed. (Robert E. Krieger, Huntington, USA). First publ. 1964.
- Lighthill, M. J. (1964): *Introduction to Fourier Analysis and Generalised Functions*. (Cambridge University Press, London). First publ. 1958.
- Lindley, D. V. (2008): *Introduction to Probability and Statistics from a Bayesian Viewpoint. Part 2: Inference*, reprint. (Cambridge University Press, Cambridge). First publ. 1965.
- (2014): *Understanding Uncertainty*, rev. ed. (Wiley, Hoboken, USA). First publ. 2006.
- MacKay, D. J. C. (2003): *Information Theory, Inference, and Learning Algorithms*. (Cambridge University Press, Cambridge). <http://www.inference.phy.cam.ac.uk/mackay/itila/>. First publ. 1995.
- Mead, L. R., Papanicolaou, N. (1984): *Maximum entropy in the problem of moments*. *J. Math. Phys.* **25**⁸, 2404–2417. <http://bayes.wustl.edu/Manual/MeadPapanicolaou.pdf>.
- Milani, L. (1965): *L'obbedienza non è più una virtù*. (Libreria Editrice Fiorentina, Florence). <https://cleliabartoli.files.wordpress.com/2015/09/lobbedienza-non-c3a8-pi-c3b9-una-virtc3b9.pdf>.
- Musil, R. (1979): *The Man Without Qualities*. (Picador, London). Transl. by E. Wilkins and E. Kaiser. First publ. in German 1930 as Musil (2000).
- (2000): *Der Mann ohne Eigenschaften*. (Rowohlt, Reinbek bei Hamburg). Herausgegeben von Adolf Frisé. First publ. 1930. Transl. as Musil (1979).
- Pitman, E. J. G. (1936): *Sufficient statistics and intrinsic accuracy*. *Math. Proc. Camb. Phil. Soc.* **32**⁴, 567–579.
- Porta Mana, P. G. L. (2009): *On the relation between plausibility logic and the maximum-entropy principle: a numerical study*. [arXiv:0911.2197](https://arxiv.org/abs/0911.2197). Presented as invited talk at the 31st International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering “MaxEnt 2011”, Waterloo, Canada.
- Rodríguez, C. C. (1989): *The metrics induced by the Kullback number*. In: *Skilling (1989b)*, 415–422.
- (2002): *Entropic priors for discrete probabilistic networks and for mixtures of Gaussians models*. *Am. Inst. Phys. Conf. Proc.* **617**, 410–432. [arXiv:physics/0201016](https://arxiv.org/abs/physics/0201016).
- Shore, J. E., Johnson, R. W. (1980): *Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy*. *IEEE Trans. Inform. Theor.* **IT-26**¹, 26–37. See also comments and correction Shore, Johnson (1983).
- (1983): *Comments on and correction to “axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy”*. *IEEE Trans. Inform. Theor.* **IT-29**⁶, 942–943.
- Sivia, D. S. (2006): *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, Oxford). Written with J. Skilling. First publ. 1996.

- Skilling, J. (1989a): *Classic maximum entropy*. In: Skilling (1989b), 45–52.
- ed. (1989b): *Maximum Entropy and Bayesian Methods: Cambridge, England, 1988*. (Kluwer, Dordrecht).
- (1990): *Quantified maximum entropy*. In: Fougère (1990), 341–350.
- (1998): *Massive inference and maximum entropy*. In: Erickson, Rychert, Smith (1998), 1–14. <http://www.maxent.co.uk/documents/massinf.pdf>.
- Strawson, P. F. (1964): *Introduction to Logical Theory*. (Methuen, London). First publ. 1952.
- van Campenhout, J. M., Cover, T. M. (1981): *Maximum entropy and conditional probability*. IEEE Trans. Inform. Theor. **IT-27**⁴, 483–489.
- Zabell, S. L. (1982): *W. E. Johnson's "sufficientness" postulate*. Ann. Stat. **10**⁴, 1090–1099. Repr. in Zabell (2005 pp. 84–95).
- (2005): *Symmetry and Its Discontents: Essays on the History of Inductive Probability*. (Cambridge University Press, Cambridge).