

Article

Entropy, Information, and the Updating of Probabilities

Ariel Caticha

Physics Department, University at Albany-SUNY, Albany, NY 12222, USA; acaticha@albany.edu;
Tel.: +1-(518)-442-4592

Abstract: This paper is a review of a particular approach to the method of maximum entropy as a general framework for inference. The discussion emphasizes pragmatic elements in the derivation. An epistemic notion of information is defined in terms of its relation to the Bayesian beliefs of ideally rational agents. The method of updating from a prior to posterior probability distribution is designed through an eliminative induction process. The logarithmic relative entropy is singled out as a unique tool for updating (a) that is of universal applicability, (b) that recognizes the value of prior information, and (c) that recognizes the privileged role played by the notion of independence in science. The resulting framework—the ME method—can handle arbitrary priors and arbitrary constraints. It includes the MaxEnt and Bayes' rules as special cases and, therefore, unifies entropic and Bayesian methods into a single general inference scheme. The ME method goes beyond the mere selection of a single posterior, and also addresses the question of how much less probable other distributions might be, which provides a direct bridge to the theories of fluctuations and large deviations.

Keywords: maximum entropy; Bayesian inference; updating probabilities



Citation: Caticha, A. Entropy, Information, and the Updating of Probabilities. *Entropy* **2021**, *23*, 895. <https://doi.org/10.3390/e23070895>

Academic Editor: Petr Jizba and Jan Korbek

Received: 17 June 2021
Accepted: 10 July 2021
Published: 14 July 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Inductive inference is a framework for coping with uncertainty, for reasoning with incomplete information. The framework must include a means to represent a state of partial knowledge—this is handled through the introduction of probabilities—and it must allow us to change from one state of partial knowledge to another when new information becomes available. Indeed, any inductive method that recognizes that a situation of incomplete information is in some way unfortunate—by which we mean that it constitutes a problem in need of a solution—would be severely deficient if it failed to address the question of how to proceed in those fortunate circumstances when new information becomes available. The theory of probability, if it is to be useful at all, demands a method for assigning and updating probabilities.

The challenge is to develop updating methods that are both systematic, objective and practical. When information consists of data and a likelihood function, Bayesian updating is the uniquely natural method of choice. Its foundation lies in recognizing the value of prior information: whatever was learned in the past is valuable and should not be disregarded, which amounts to requiring that beliefs ought to be revised but only to the extent required by the new data. This immediately raises a number of questions: How do we update when the information is not in the form of data? If the information is not data, what else could it possibly be? Indeed, what, after all, is “information”? On a separate line of development, the method of Maximum Gibbs–Shannon Entropy (MaxEnt) allows one to process information in the form of constraints on the allowed probability distributions. This provides a partial answer to one of our questions: in addition to data, information can also take the form of constraints. However, it immediately raises several other questions: What is the interpretation of entropy? Is there unique entropy? Are Bayesian and entropic methods mutually compatible?

The purpose of this paper is to review one particular approach to entropic updating. The presentation below, which is meant to be pedagogical and self-contained, is based on work presented in a sequence of papers [1–5] and in the sets of lectures [6–8]. As we shall see below, we adopt a pragmatic approach in which entropy is a tool designed for the specific purpose of updating probabilities.

Historically, the method of maximum relative entropy (ME) is a direct descendant of the MaxEnt method, pioneered by Jaynes [9,10]. In the MaxEnt framework, entropy is interpreted through the Shannon axioms as a measure of the amount of information that is missing in a probability distribution. This approach has its limitations. The Shannon axioms refer to probabilities of discrete variables; for continuous variables, the Shannon entropy is not defined. A more serious objection is that even if we grant that the Shannon axioms do lead to a reasonable expression for entropy, to what extent do we believe the axioms themselves? Shannon's third axiom, the grouping property, is indeed sort of reasonable, but is it necessary? Is entropy the only consistent measure of uncertainty or of information? Indeed, there exist examples in which the Shannon entropy does not seem to reflect one's intuitive notion of information [8,11]. One could introduce other entropies justified by different choices of axioms (e.g., [12–14]), but this move raises problems of its own: Which entropy should one adopt? If different systems are handled using different entropies, how does one handle composite systems?

From our perspective, the problem can be traced to the fact that neither Shannon nor Jaynes were concerned with the task of updating probabilities. Shannon's communication theory aimed to characterize the sources of information, to measure the capacity of the communication channels, and to learn how to control the degrading effects of noise. On the other hand, Jaynes conceived MaxEnt as a method to assign probabilities on the basis of constraint information and a fixed underlying measure and not from an arbitrary prior distribution.

Considerations such as these motivated several attempts to develop ME directly as a method for updating probabilities without invoking questionable measures of information [1,5,15–17]. The important contribution by Shore and Johnson was the realization that one could axiomatize the updating method itself rather than the information measure. Their axioms have, however, raised criticisms [11,18–20] and counter-criticisms [2,6,8,21,22]. Despite the controversies, Shore and Johnson's pioneering papers have had an enormous influence: they identified the correct goal to be achieved.

The concept of relative entropy is introduced as a tool for updating probabilities. Hereinafter, we drop the qualifier "relative" and adopt the simpler term "entropy". The reasons for the improved nomenclature are the following: (1) The general concept should receive the general name "entropy", while the more specialized concepts should be the ones receiving a qualifier, such as "thermodynamic" or "Clausius" entropy, and "Gibbs–Shannon" entropy. (2) All entropies are relative, even if they happen to be relative to an implicit uniform prior. Making this fact explicit has tremendous pedagogical value. (3) The practice is already in use with the concept of energy: all energies are relative too, but there is no advantage in constantly referring to a "relative energy". Accordingly, ME will be read as "maximum entropy"; additional qualifiers are redundant.

As with all tools, entropy too is *designed* to perform a certain function, and its performance must meet certain *design criteria* or *specifications*. There is no implication that the method is "true", or that it succeeds because it achieves some special contact with reality. Instead, the claim is that the method succeeds in the pragmatic sense that it works as designed—this is satisfactory because when properly deployed, it leads to empirically adequate models. In this approach, *entropy needs no interpretation* whether it be in terms of heat, multiplicity of states, disorder, uncertainty, or even in terms of an amount of information. Incidentally, this may explain why the search for the meaning of entropy has proved so elusive: we need not know what "entropy" means—we only need to know how to use it.

Since our topic is the updating of probabilities when confronted with new information, our starting point is to address the question, “what is information?”. In Section 2, we develop a concept of information that is both pragmatic and Bayesian. “Information” is defined in terms of its effects on the beliefs of rational agents. The design of entropy as a tool for updating is the topic of Section 3. There, we state the design specifications that define what function entropy is supposed to perform, and we derive its functional form. To streamline the presentation, some of the mathematical derivations are left to the appendices.

To conclude, we present two further developments. In Section 4, we show that Bayes’ rule can be derived as a special case of the ME method. An earlier derivation of this important result following a different line of argument was given by Williams [23] before a sufficient understanding of entropy as an updating tool had been achieved. It is not, therefore, surprising that Williams’ achievement has not received the widespread appreciation it deserves. Thus, within the ME framework, entropic and Bayesian methods are unified into a single consistent theory of inference. One advantage of this insight is that it allows a number of generalizations of Bayes’ rule [2,8]. Another is that it provides an important missing piece for the old puzzles of quantum mechanics concerning the so-called collapse of the wave function and the quantum measurement problem [24,25].

There is yet another function that the ME method must perform in order to fully qualify as a method of inductive inference. Once we have decided that the distribution of maximum entropy is to be preferred over all others, the following question arises immediately: the maximum of the entropy functional is never infinitely sharp, so are we confident that distributions that lie very close to the maximum are completely ruled out? In Section 5, the ME method is deployed to assess quantitatively the extent to which distributions with lower entropy are ruled out. The significance of this result is that it provides a direct link to the theories of fluctuations and large deviations. Concluding remarks are given in Section 6.

2. What Is Information?

The term “information” is used with a wide variety of different meanings [10,26,27]. There is the Shannon notion of information that is meant to measure an amount of information and is quite divorced from semantics. There is also an algorithmic notion of information that captures a notion of complexity and originates in the work of Solomonov, Kolmogorov, and Chaitin [26]; there is a related notion of entropy as a minimum description length [28]. Furthermore, in the general context of the thermodynamics of computation, it is said that “information is physical” because systems “carry” or “contain” information about their own physical state [29–31] (see also [32,33]).

Here, we follow a different path [3,4]. We seek an epistemic notion of information that is closer to the everyday colloquial use of the term—roughly, information is what we request when we ask a question. In a Bayesian framework, this requires an explicit account of the relation between information and the beliefs of ideally rational agents. We emphasize that our concern here is with *idealized rational agents*. Our subject is not the psychology of actual humans who often change their beliefs by processes that are neither fully rational nor fully conscious. We adopt a Bayesian interpretation of probability as a degree of credibility: the degree to which we *ought* to believe that a proposition is true if only we were ideally rational. For a discussion of a decision theory that might be relevant to the economics and psychology of partially rational agents see [34–36]. An entropic framework for modelling economies that bypasses all issues of bounded rationality is described in [37].

It is implicit in the recognition that most of our beliefs are held on the basis of incomplete information that not all probability assignments are equally good; some beliefs are preferable to others in the very pragmatic sense that they enhance our chances to successfully navigate this world. Thus, a theory of probability demands a theory of updating probabilities in order to improve our beliefs.

We are now ready to address the question: What, after all, is “information”? The answer is pragmatic. *Information is what information does*. Information is defined by its effects: (a) it restricts our options as to what we are honestly and rationally allowed to believe; and (b) it induces us to update from prior beliefs to posterior beliefs. This, I propose, is a defining characteristic of information:

Information is that which induces a change from one state of rational belief to another.

One aspect of this notion is that for a rational agent, the identification of what constitutes information—as opposed to mere noise—already involves a judgment, an evaluation. Another aspect is that the notion that information is directly related to changing our minds does not involve any reference to *amounts* of information, but it nevertheless allows precise quantitative calculations. Indeed, constraints on the acceptable posterior probabilities are precisely the kind of information that the method of maximum entropy is designed to handle. In short,

Information constrains probability distributions. The constraints are the information.

To the extent that the probabilities are Bayesian, this definition captures the Bayesian notion that information is directly related to changing our minds, that it is the driving force behind the process of learning. It also incorporates an important feature of rationality: being rational means accepting that “not everything goes”, and that our beliefs must be constrained in very specific ways. However, the indiscriminate acceptance of any arbitrary constraint does not qualify as rational behavior. To be rational, an agent must exercise some judgment before accepting a particular piece of information as a reliable basis for the revision of its beliefs, which raises questions about what judgments might be considered sound. Furthermore, there is no implication that the information must *be true*; only that we *accept it as true*. False information is information too, at least to the extent that we are prepared to accept it and allow it to affect our beliefs.

The paramount virtue of the definition above is that it is useful; it allows precise quantitative calculations. The constraints that constitute information can take a wide variety of forms. They can be expressed in terms of expected values, they can specify the functional form of a distribution, or be imposed through various geometrical relations. Examples are given in Section 5 and in [38].

Concerning the act of updating, it may be worthwhile to point out an analogy with dynamics. In Newtonian mechanics, the state of motion of a system is described in terms of momentum, and the change from one state to another is said to be “caused” by an applied force or impulse. Bayesian inference is analogous in that a state of belief is described in terms of probabilities, and the change from one state to another is “caused” by information. Just as a force is that which induces a change from one state of motion to another, so *information is that which induces a change from one state of belief to another*. Updating is a form of dynamics. In [39], the analogy is taken seriously: the logic is reversed and quantum mechanics is derived as an example of the entropic updating of probabilities.

3. The Pragmatic Design of Entropic Inference

Once we have decided, as a result of the confrontation of new information with old beliefs, that our beliefs require revision, the problem becomes one of deciding how precisely this ought to be done. First, we identify some general features of the kind of belief revision that one might count as rational. Then, we design a method—a systematic procedure—that implements those features. To the extent that the method performs as desired, we can claim success. The point is not that success derives from our method having achieved some intimate connection to the inner wheels of reality; success simply means that the method seems to be working.

The one obvious requirement is that the updated probabilities ought to agree with the newly acquired information. Unfortunately, this requirement, while necessary, is not sufficiently restrictive: we can update in many ways that preserve both internal consistency

and consistency with the new information. Additional criteria are needed. What rules would an ideally rational agent choose?

3.1. General Criteria

The rules are motivated by the same pragmatic criteria that motivate the design of probability theory itself [8]—universality, consistency, and practical utility. However, this is admittedly too vague; we must be very specific about the precise way in which the criteria are implemented.

3.1.1. Universality

In principle, different systems and different situations could require different problem-specific induction methods. However, in order to be useful in practice, the method we seek must be of *universal* applicability. Otherwise, it would fail us when most needed, for we would not know which method to choose when not much is known about the system. To put in different words, what we want to design is a general-purpose method that captures what all the other problem-specific methods might have in common. The idea is that the peculiarities of a particular problem will be captured by the specific constraints that describe the information that is relevant to the problem at hand.

The analogy with mechanics can be found here as well. The possibility of a science of mechanics hinges on identifying a law of motion of universal applicability (e.g., the Schrödinger equation), while the specifics of each system are introduced through initial conditions and the choice of potentials or forces. Here, we shall design an entropy of universal applicability, while the specifics of each problem are introduced through prior probabilities and the choice of constraints.

3.1.2. Parsimony

To specify the updating, we adopt a very conservative criterion that recognizes the value of information: what has been laboriously learned in the past is valuable and should not be disregarded unless rendered obsolete by new information. The only aspects of one's beliefs that should be updated are those for which new evidence has been supplied. Thus, we adopt the following.

Principle of Minimal Updating (PMU): *Beliefs should be updated only to the minimal extent required by the new information.*

The special case of updating in the absence of new information deserves a comment. The PMU states that when there is no new information, ideally, rational agents should not change their minds. In fact, it is difficult to imagine any notion of rationality that would allow the possibility of changing one's mind for no apparent reason.

Minimal updating offers yet another pragmatic advantage. As we shall see below, rather than identifying what features of a distribution are singled out for updating and then specifying the detailed nature of the update, we will adopt design criteria that stipulate what is not to be updated. The practical advantage of this approach is that it enhances objectivity—there are many ways to change something but only one way to keep it the same. The analogy with mechanics can be pursued even further: if updating is a form of dynamics, then minimal updating is the analogue of inertia. Rationality and objectivity demand a considerable amount of inertia.

3.1.3. Independence

The next general requirement turns out to be crucially important because without it, the very possibility of scientific theories would be compromised. The point is that every scientific model, whatever the topic, if it is to be useful at all, must assume that all relevant variables have been taken into account and that whatever was left out—the rest of the universe—should not matter. To put it another way, in order to do scientific work, we must be able to understand parts of the universe without having to understand the universe as a

whole. Granted, a pragmatic understanding need not be complete and exact; it must be merely adequate for our purposes.

The assumption, then, is that it is possible to focus our attention on a suitably chosen system of interest and neglect the rest of the universe because the system and the rest of the universe are “sufficiently independent”. Thus, in any form of science, the notion of statistical independence must play a central and privileged role. This idea—that some things can be neglected and that not everything matters—is implemented by imposing a criterion that tells us how to handle independent systems. The chosen criterion is quite natural: *whenever two systems are a priori believed to be independent and we receive information about just one, it should not matter if the other is included in the analysis or not*. This is an example of the PMU in action; it amounts to requiring that independence to be preserved unless information about correlations is explicitly introduced.

Again, we emphasize that none of these criteria are imposed by nature. They are desirable for pragmatic reasons; they are imposed by design.

3.2. Entropy as a Tool for Updating Probabilities

Consider a set of propositions $\{x\}$ about which we are uncertain. The proposition x can be discrete or continuous, in one or in several dimensions. It could, for example, represent the microstate of a physical system, a point in phase space, or an appropriate set of quantum numbers. The uncertainty about x is described by a probability distribution $q(x)$. The goal is to update from the prior distribution $q(x)$ to a posterior distribution $p(x)$ when new information—by which we mean a set of constraints—becomes available. The question is, which distribution among all those that satisfy the constraints should we select?

Our goal is to design a method that allows a systematic search for the preferred posterior distribution. The central idea, first proposed by Skilling [16], is disarmingly simple: to select the posterior, first rank all candidate distributions in increasing order of “preference” and then pick the distribution that ranks the highest. Irrespective of what it is that makes one distribution “preferable” over another (we will get to that soon enough), it is clear that any such ranking must be transitive: if distribution p_1 is preferred over distribution p_2 , and p_2 is preferred over p_3 , then p_1 is preferred over p_3 . Transitive rankings are implemented by assigning to each p a real number $S[p]$, which is called the entropy of p in such a way that if p_1 is preferred over p_2 , then $S[p_1] > S[p_2]$. The selected distribution (one or possibly many, for *there may be several equally preferred distributions*) is that which maximizes the entropy functional.

The importance of Skilling’s strategy of ranking distributions cannot be overestimated: it answers the questions “why entropy?” and “why a maximum?”. The strategy implies that the updating method will take the form of a variational principle—the method of maximum entropy (ME)—involving a certain functional that maps distributions to real numbers. These features are not imposed by nature; they are all imposed by design. They are dictated by the function that the ME method is supposed to perform. (Thus, it makes no sense to seek a generalization in which entropy is a complex number or a vector; such generalized entropies would just not perform the desired function.)

Next, we specify the ranking scheme, that is, we choose a specific functional form for the entropy $S[p]$. Note that the purpose of the method is to update *from priors to posteriors* so the ranking scheme must depend on the particular prior q and therefore, the entropy S must be a functional of both p and q . The entropy $S[p, q]$ describes a ranking of the distributions p relative to the given prior q . $S[p, q]$ is the entropy of p relative to q , and accordingly, $S[p, q]$ is commonly called *relative entropy*. This is appropriate and sometimes we will follow this practice. However, since all entropies are relative, even when relative to a uniform distribution, the qualifier “relative” is redundant and can be dropped.

The functional $S[p, q]$ is designed by a process of elimination—this is a process of *eliminative induction*. First, we state the desired design criteria; this is the crucial step that defines what makes one distribution preferable over another. Candidate functionals that

fail to satisfy the criteria are discarded—hence, the qualifier “eliminative”. As we shall see, the criteria adopted below are so constraining that there is a single entropy functional $S[p, q]$ that survives the process of elimination.

This approach has a number of virtues. First, to the extent that the design criteria are universally desirable, the single surviving entropy functional will also be of universal applicability. Second, the reason why alternative entropy candidates are eliminated is quite explicit—at least one of the design criteria is violated. Thus, *the justification behind the single surviving entropy is not that it leads to demonstrably correct inferences, but rather, that all other candidates demonstrably fail to perform as desired.*

3.3. Specific Design Criteria

Consider a lattice of propositions generated by a set \mathcal{X} of atomic propositions that are mutually exclusive and exhaustive and are labeled by a discrete index $i = 1, 2, \dots, n$. The extension to infinite sets and to continuous labels turns out to be straightforward. The index i might, for example, label the microstates of a physical system but, since the argument below is supposed to be of general validity, we shall not assume that the labels themselves carry any particular significance. We can always permute labels; this should have no effect on the updating of probabilities.

We adopt design criteria that reflect the structure of the lattice of propositions—the propositions are related to each other by disjunctions (OR) and conjunctions (AND) and the consistency of the web of beliefs is implemented through the sum and product rules of the probability theory. Our criteria refer to the two extreme situations of propositions that are mutually exclusive and of propositions that are mutually independent. At one end, we deal with the probabilities of propositions that are highly correlated (if one proposition is true, the other is false and vice versa); at the other end, we deal with the probabilities of propositions that are totally uncorrelated (the truth or falsity of one proposition has no effect on the truth or falsity of the other). One extreme is described by the simplified sum rule, $p(i \vee j) = p(i) + p(j)$, and the other extreme by the simplified product rule, $p(i \wedge j) = p(i)p(j)$. (For an alternative approach to the foundations of inference that exploits the various symmetries of the lattice of propositions see [40,41].)

The two design criteria and their consequences for the functional form of entropy are given below. Detailed proofs are deferred to the appendices.

3.3.1. Mutually Exclusive Subdomains

DC1: *Probabilities that are conditioned on one subdomain are not affected by information about other non-overlapping subdomains.*

Consider a subdomain $\mathcal{D} \subset \mathcal{X}$ composed of atomic propositions $i \in \mathcal{D}$ and suppose the information to be processed refers to some other subdomain $\mathcal{D}' \subset \mathcal{X}$ that does not overlap with \mathcal{D} , $\mathcal{D} \cap \mathcal{D}' = \emptyset$. In the absence of any new information about \mathcal{D} , the PMU demands, we do not change our minds about probabilities that are conditional on \mathcal{D} . Thus, we design the inference method so that $q(i|\mathcal{D})$, the prior probability of i conditioned on $i \in \mathcal{D}$, is not updated. Thus, the selected conditional posterior is

$$P(i|\mathcal{D}) = q(i|\mathcal{D}) . \quad (1)$$

We adopt the following notation: priors are denoted by q , candidate posteriors by lower case p , and the selected posterior by upper case P . We shall write either $p(i)$ or p_i . Furthermore, we adopt the notation, standard in physics where the probabilities of x and θ are written $p(x)$ and $p(\theta)$ but there is no implication that p refers to the same mathematical function.

We emphasize that the point is not that we make the unwarranted assumption that keeping $q(i|\mathcal{D})$ unchanged is guaranteed to lead to correct inferences. It need not; induction is risky. The point is, rather, that in the absence of any evidence to the contrary, there is no reason to change our minds and the prior information takes priority.

The consequence of DC1 is that non-overlapping domains of i contribute additively to the entropy,

$$S(p, q) = \sum_i F(p_i, q_i), \tag{2}$$

where F is some unknown function of two arguments. The proof is given in Appendix A.

Comment 1: It is essential that DC1 refers to *conditional* probabilities—local information about a domain \mathcal{D}' can (via normalization) have a non-local effect on the probability of another domain \mathcal{D} .

Comment 2: An important special case is the “update” from a prior $q(i)$ to a posterior $P(i)$ in a situation in which no new information is available. The criterion DC1 applied to a situation where the subdomain \mathcal{D} covers the whole space of i s, $\mathcal{D} = \mathcal{X}$, requires that *in the absence of any new information, the prior conditional probabilities are not to be updated:* $P(i|\mathcal{X}) = q(i|\mathcal{X})$ or $P(i) = q(i)$.

Comment 3: The criterion DC1 implies Bayesian conditionalization as a special case. Indeed, if the information is given through the constraint $p(\bar{\mathcal{D}}) = 0$, where $\bar{\mathcal{D}}$ is the complement of \mathcal{D} , then $P(i|\mathcal{D}) = q(i|\mathcal{D})$, which is referred to as Bayesian conditionalization. More explicitly, if θ is the variable to be inferred on the basis of prior information about a likelihood function $q(i|\theta)$ and observed data i' , then the update from the prior q to the posterior P , is

$$q(i, \theta) = q(i)q(\theta|i) \rightarrow P(i, \theta) = P(i)P(\theta|i), \tag{3}$$

consists of updating $q(i) \rightarrow P(i) = \delta_{ii'}$ to agree with the new information and invoking the PMU so that $P(\theta|i') = q(\theta|i')$ remains unchanged. Therefore,

$$P(i, \theta) = \delta_{ii'}q(\theta|i') \quad \text{so that} \quad P(\theta) = q(\theta|i') = q(\theta) \frac{q(i'|\theta)}{q(i')}, \tag{4}$$

which is Bayes’ rule. Thus, *entropic inference is designed to include Bayesian inference as a special case.* Note, however, that imposing DC1 is not identical to imposing Bayesian conditionalization: DC1 is not restricted to information in the form of absolute certainties, such as $p(\mathcal{D}) = 1$.

Comment 4: If the label i is turned into a continuous variable x , the criterion DC1 requires that information that refers to points infinitely close but just outside the domain \mathcal{D} will have no influence on probabilities conditional on \mathcal{D} . This may seem surprising, as it may lead to updated probability distributions that are discontinuous, but it is not a problem. In situations where we have explicit reasons to believe that conditions of continuity or differentiability hold, then such conditions should be imposed explicitly. The inference process should not be expected to discover and replicate information with which it was not supplied.

3.3.2. Subsystem Independence

DC2: *When two systems are a priori believed to be independent and the information we receive about one of them makes no reference to the other, then it should not matter whether the latter is included in the analysis of the former or not.*

Consider a system of propositions labeled by a composite index, $i = (i_1, i_2) \in \mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$. For example, $\{i_1\} = \mathcal{X}_1$ and $\{i_2\} = \mathcal{X}_2$ might describe the microstates of two separate physical systems. Assume that all prior evidence led us to believe the two subsystems are independent, that is, any two propositions $i_1 \in \mathcal{X}_1$ and $i_2 \in \mathcal{X}_2$ are believed to be independent. This belief is reflected in the prior distribution: if the individual subsystem priors $q_1(i_1)$ and $q_2(i_2)$, then the prior for the whole system is $q_1(i_1)q_2(i_2)$. Next, suppose that new information is acquired such that $q_1(i_1)$ would by itself be updated to $P_1(i_1)$, and that $q_2(i_2)$ would by itself be updated to $P_2(i_2)$. DC2 requires that $S[p, q]$ be

such that the joint prior $q_1(i_1)q_2(i_2)$ updates to the product $P_1(i_1)P_2(i_2)$ so that inferences about one subsystem do not affect inferences about the other.

The consequence of DC2 is to fully determine the unknown function F in (2) so that probability distributions $p(i)$ should be ranked relative to the prior $q(i)$ according to the relative entropy,

$$S[p, q] = -\sum_i p(i) \log \frac{p(i)}{q(i)}. \tag{5}$$

Comment 1: We emphasize that the point is not that when we have no evidence for correlations, we draw the firm conclusion that the systems must necessarily be independent. Induction involves risk; the systems might, in actual fact, be correlated through some unknown interaction potential. The point is rather that if the joint prior reflected independence and the new evidence is silent on the matter of correlations, then the evidence we actually have—namely, the prior—takes precedence, and there is no reason to change our minds. As before, the PMU requires that a feature of the probability distribution—in this case, independence—will not be updated unless the evidence requires it.

Comment 2: We also emphasize that DC2 is *not a consistency requirement*. The argument we deploy is *not* that both the prior *and* the new information tell us the systems are independent in which case consistency requires that it should not matter whether the systems are treated jointly or separately. DC2 refers to a situation where the new information does not say whether the systems are independent or not. Rather, the updating is being *designed*—through the PMU—so that the independence reflected in the prior is maintained in the posterior by default.

Comment 3: The generalization to continuous variables $x \in \mathcal{X}$ is approached as a Riemann limit from the discrete case. A continuous probability density $p(x)$ or $q(x)$ can be approximated by the discrete distributions. Divide the region of interest \mathcal{X} into a large number N of small cells. The probabilities of each cell are as follows:

$$p_i = p(x_i)\Delta x_i \quad \text{and} \quad q_i = q(x_i)\Delta x_i, \tag{6}$$

where Δx_i is an appropriately small interval. The discrete entropy of p_i relative to q_i is as follows:

$$S_N = -\sum_{i=1}^N \Delta x_i p(x_i) \log \left[\frac{p(x_i)\Delta x_i}{q(x_i)\Delta x_i} \right], \tag{7}$$

and in the limit as $N \rightarrow \infty$ and $\Delta x_i \rightarrow 0$ we get the Riemann integral

$$S[p, q] = -\int dx p(x) \log \left[\frac{p(x)}{q(x)} \right]. \tag{8}$$

(To simplify the notation, we include multi-dimensional integrals by writing $d^n x = dx$.) It is easy to check that the ranking of distributions induced by $S[p, q]$ is invariant under coordinate transformations. The insight that coordinate invariance could be derived as a consequence of the requirement of subsystem independence first appeared in [5].

3.4. The ME Method

We can now summarize the overall conclusion.

The ME method: *The goal is to update from a prior distribution q to a posterior distribution when there is new information in the form of constraints \mathcal{C} that specify a family $\{p\}$ of candidate posteriors. The preferred posterior P is that which maximizes the relative entropy,*

$$S[p, q] = -\sum_i p_i \log \frac{p_i}{q_i} \quad \text{or} \quad S[p, q] = -\int dx p(x) \log \left[\frac{p(x)}{q(x)} \right], \tag{9}$$

within the family $\{p\}$ specified by the constraints \mathcal{C} .

This extends the method of maximum entropy beyond its original purpose as a rule to assign probabilities from a given underlying measure (MaxEnt) to a method for updating probabilities from any arbitrary prior (ME). Furthermore, the logic behind the updating procedure does not rely on any particular meaning assigned to the entropy whether in terms of information, or heat, or disorder. Entropy is merely a tool for inductive inference. *No interpretation for $S[p, q]$ is given and none is needed.*

The derivation above has singled out a unique $S[p, q]$ to be used in inductive inference. Other “entropies” (such as the one-parameter families of entropies proposed in [12–14]) might turn out to be useful for other purposes—perhaps as measures of some kind of “information”, as measures of discrimination or distinguishability among distributions, of ecological diversity, or for some altogether different function—but they are unsatisfactory for the purpose of updating because they fail to perform the functions stipulated by the design criteria DC1 and DC2. They induce correlations that are unwarranted by the information in the priors or the constraints.

4. Bayes’ Rule as a Special Case of ME

Back in Section 3.3.1, we saw that ME is designed to include Bayes’ rule as a special case. Here, we wish to verify this explicitly [2]. The goal is to update our beliefs about $\theta \in \Theta$ (θ represents one or many parameters) on the basis of three pieces of information: (1) the prior information codified into a prior distribution $q(\theta)$; (2) the new information conveyed by data $x \in \mathcal{X}$ (obtained in one or many experiments); and (3) the known relation between θ and x given by a model defined by the sampling distribution or likelihood, $q(x|\theta)$. The updating will result in replacing the *prior* probability distribution $q(\theta)$ by a *posterior* distribution $P(\theta)$ that applies after the data information has been processed.

The crucial element that will allow the Bayes’ rule to be smoothly integrated into the ME scheme is the realization that before the data are collected, not only do we not know θ , but we do not know x either. Thus, the relevant space for inference is not the space Θ but the product space $\Theta \times \mathcal{X}$, and the relevant joint *prior* is $q(x, \theta) = q(\theta)q(x|\theta)$. Let us emphasize two points: first, the likelihood function is an integral part of the *prior* distribution; second, the prior information about how x is related to θ is contained in the *functional form* of the distribution $q(x|\theta)$ and not in the numerical values of the arguments x and θ , which, at this point, are still unknown.

Next, data are collected and the observed values turn out to be x' . We must update to a posterior that lies within the family of distributions $p(x, \theta)$ that reflect the fact that the previously unknown x is now known to be x' , that is,

$$p(x) = \int d\theta p(\theta, x) = \delta(x - x') . \tag{10}$$

The information in this data constrains but is not sufficient to fully determine the joint distribution,

$$p(x, \theta) = p(x)p(\theta|x) = \delta(x - x')p(\theta|x') . \tag{11}$$

Any choice of $p(\theta|x')$ is, in principle, possible. So far, the formulation of the problem parallels Section 3.3.1 exactly. We are, after all, solving the same problem. The next step is to apply the ME method.

According to the ME method, the selected joint posterior $P(x, \theta)$ is that which maximizes the entropy,

$$S[p, q] = - \int dx d\theta p(x, \theta) \log \frac{p(x, \theta)}{q(x, \theta)} , \tag{12}$$

subject to the data constraints. Note that Equation (10) represents an *infinite* number of constraints on the family $p(x, \theta)$: there is one constraint and one Lagrange multiplier $\lambda(x)$ for each value of x . Maximizing S , (12), subject to (10) and normalization,

$$\delta \{ S + \alpha [\int dx d\theta p(x, \theta) - 1] + \int dx \lambda(x) [\int d\theta p(x, \theta) - \delta(x - x')] \} = 0 , \tag{13}$$

yields the joint posterior

$$P(x, \theta) = q(x, \theta) \frac{e^{\lambda(x)}}{Z}, \quad (14)$$

where Z is a normalization constant, and the multiplier $\lambda(x)$ is determined from (10) as follows:

$$\int d\theta q(x, \theta) \frac{e^{\lambda(x)}}{Z} = q(x) \frac{e^{\lambda(x)}}{Z} = \delta(x - x'), \quad (15)$$

so that the joint posterior is

$$P(x, \theta) = q(x, \theta) \frac{\delta(x - x')}{q(x)} = \delta(x - x') q(\theta|x). \quad (16)$$

The corresponding marginal posterior probability $P(\theta)$ is

$$P(\theta) = \int dx P(\theta, x) = q(\theta|x') = q(\theta) \frac{q(x'|\theta)}{q(x')}, \quad (17)$$

which is Bayes' rule. Thus, Bayes' rule is derivable from, and therefore consistent with, the ME method.

To summarize, the prior $q(x, \theta) = q(x)q(\theta|x)$ is updated to the posterior $P(x, \theta) = P(x)P(\theta|x)$, where $P(x) = \delta(x - x')$ is fixed by the observed data while $P(\theta|x') = q(\theta|x')$ remains unchanged. Note that in accordance with the PMU philosophy that drives the ME method, *one only updates those aspects of one's beliefs for which corrective new evidence has been supplied*. In [2,8,42], further examples are given that show how ME allows generalizations of Bayes' rule to situations where the data itself are uncertain, there is information about moments of x or moments of θ , or even in situations where the likelihood function is unknown. In conclusion, the ME method of maximum entropy can fully reproduce and then go beyond the results obtained by the standard Bayesian methods.

5. Deviations from Maximum Entropy

The basic ME problem is to update from a prior $q(x)$ given information specified by certain constraints. The constraints specify a family of candidate distributions as follows:

$$p_\theta(x) = p(x|\theta) \quad (18)$$

which can be conveniently labeled with a finite number of parameters θ^a , $a = 1 \dots n$. (The generalization to an infinite number of parameters poses technical but not insurmountable difficulties.) Thus, the parameters θ are coordinates on the statistical manifold specified by the constraints. The distributions in this manifold are ranked according to their entropy,

$$S[p_\theta, q] = - \int dx p(x|\theta) \log \frac{p(x|\theta)}{q(x)} = S(\theta), \quad (19)$$

and the selected posterior is the distribution $p(x|\theta_0)$ that maximizes the entropy $S(\theta)$. (The notation indicates that $S[p_\theta, q]$ is a functional of p_θ while $S(\theta)$ is a function of θ .)

The question we now address concerns the extent to which $p(x|\theta_0)$ should be preferred over other distributions with lower entropy or, to put it differently, to what extent is it rational to believe that the selected value ought to be the entropy maximum θ_0 rather than any other value θ [1]? This is a question about the probability $p(\theta)$ of various values of θ . The original problem which led us to design the maximum entropy method was to assign a probability to the quantity x ; we now see that the full problem is to assign probabilities to both x and θ . We are concerned not just with $p(x)$, but rather with the joint distributions which we denote as $\pi(x, \theta)$; the universe of discourse has been expanded from \mathcal{X} (the space of x s) to the product space $\mathcal{X} \times \Theta$ (Θ is the space of parameters θ).

To determine the joint distribution $\pi(x, \theta)$, we make use of essentially the only (universal) method at our disposal—the ME method itself—but this requires that we address

the standard two preliminary questions: First, what is the prior distribution? What do we know about x and θ before we receive information about the constraints? Second, what is the new information that constrains the allowed joint distributions $\pi(x, \theta)$?

This first question is the more subtle one: when we know absolutely nothing about the θ s, we know neither their physical meaning nor whether there is any relation to the x s. A joint prior that reflects this lack of correlations is a product, $q(x, \theta) = q(x)q(\theta)$. We will assume that the prior $q(x)$ is known—it is the same prior we had used when we updated from $q(x)$ to $p(x|\theta_0)$ using (19).

However, we are not totally ignorant about the θ s: we know that they label distributions $\pi(x|\theta)$ on some as yet unspecified statistical manifold Θ . Then there exists a natural measure of distance in the space Θ . It is given by the information metric $d\ell^2 = g_{ab}d\theta^a d\theta^b$ [8,43], where

$$g_{ab} = \int dx p(x|\theta) \frac{\partial \log p(x|\theta)}{\partial \theta^a} \frac{\partial \log p(x|\theta)}{\partial \theta^b}, \tag{20}$$

and the corresponding volume elements are given by $g^{1/2}(\theta)d^n\theta$, where $g(\theta)$ is the determinant of the metric. The uniform prior for θ , which assigns equal probabilities to equal volumes, is proportional to $g^{1/2}(\theta)$, and therefore we choose $q(\theta) = g^{1/2}(\theta)$. Therefore, the joint prior is $q(x, \theta) = q(x)g^{1/2}(\theta)$.

Next, we tackle the second question: what are the constraints on the allowed joint distributions $\pi(x, \theta)$? Consider the space of all joint distributions. To each choice of the functional form of $\pi(x|\theta)$ (for example, whether we talk about Gaussians, Boltzmann–Gibbs distributions, or something else), there corresponds a different subspace defined by distributions of the form $\pi(x, \theta) = \pi(\theta)\pi(x|\theta)$. The crucial constraint is that which specifies the subspace by imposing that $\pi(x|\theta)$ takes the particular functional form given by the constraint (18), $\pi(x|\theta) = p(x|\theta)$. This defines the meaning to the θ s and also fixes the prior $g^{1/2}(\theta)$ on the relevant subspace.

The preferred joint distribution, $P(x, \theta) = P(\theta)p(x|\theta)$, is the distribution, $\pi(x, \theta) = \pi(\theta)p(x|\theta)$, that maximizes the joint entropy,

$$\begin{aligned} S[\pi, q] &= - \int dx d\theta \pi(\theta)p(x|\theta) \log \frac{\pi(\theta)p(x|\theta)}{g^{1/2}(\theta)q(x)} \\ &= - \int d\theta \pi(\theta) \log \frac{\pi(\theta)}{g^{1/2}(\theta)} + \int d\theta \pi(\theta)S(\theta), \end{aligned} \tag{21}$$

where $S(\theta)$ is given in (19). Varying (21) with respect to $\pi(\theta)$ with $\int d\theta \pi(\theta) = 1$ and $p(x|\theta)$ fixed yields the posterior probability that the value of θ lies within the small volume $g^{1/2}(\theta)d^n\theta$,

$$P(\theta)d^n\theta = \frac{1}{\zeta} e^{S(\theta)}g^{1/2}(\theta)d^n\theta \quad \text{with} \quad \zeta = \int d^n\theta g^{1/2}(\theta) e^{S(\theta)}. \tag{22}$$

Equation (22) is the result we seek. It tells us that, as expected, the preferred value of θ is the value θ_0 that maximizes the entropy $S(\theta)$, Equation (19), because this maximizes the scalar density $\exp S(\theta)$. However, it also tells us the degree to which values of θ away from the maximum are ruled out. (Note that the density $\exp S(\theta)$ is a scalar function and the presence of the Jacobian factor $g^{1/2}(\theta)$ makes Equation (22) manifestly invariant under changes of the coordinates θ in the space Θ .)

This discussion allows us to refine our understanding of the ME method. ME is not an all-or-nothing recommendation to pick the single distribution that maximizes entropy and reject all others. The ME method is more nuanced: in principle, all distributions within the constraint manifold ought to be included in the analysis; they contribute in proportion to the exponential of their entropy and this turns out to be significant in situations where the entropy maximum is not particularly sharp.

Going back to the original problem of updating from the prior $q(x)$, given information that specifies the manifold $\{p(x|\theta)\}$, the preferred update within the family $\{p(x|\theta)\}$ is $p(x|\theta_0)$, but to the extent that other values of θ are not totally ruled out, a better update is obtained marginalizing the joint posterior $P(x, \theta) = P(\theta)p(x|\theta)$ over θ ,

$$P(x) = \int d^n \theta P(\theta) p(x|\theta) = \int d^n \theta g^{1/2}(\theta) \frac{e^{S(\theta)}}{\zeta} p(x|\theta). \quad (23)$$

In situations where the entropy maximum at θ_0 is very sharp, we recover the old result,

$$P(x) \approx p(x|\theta_0). \quad (24)$$

When the entropy maximum is not very sharp a more honest update is Equation (23), which, incidentally, is a form of superstatistics.

One of the limitations of the standard MaxEnt method is that it selects a single “posterior” $p(x|\theta_0)$ and strictly rules out all other distributions. The result (22) overcomes this limitation and finds many applications. For example, it extends the Einstein theory of thermodynamic fluctuations beyond the regime of small fluctuations; it provides a bridge to the theory of large deviations; and, suitably adapted for Bayesian data analysis, it leads to the notion of entropic priors [44].

6. Discussion

Consistency with the law of large numbers.

Entropic methods of inference are of general applicability but there exist special situations—for example, those involving large numbers of independent subsystems—where inferences can be made by purely probabilistic methods without ever invoking the concept of entropy. In such cases, one can check (see, for example, [6,45]) that the two methods of calculation are consistent with each other. It is significant, however, that alternative entropies, such as those proposed in [12–14], do not pass this test [46,47], which rules them out as tools for updating. Some probability distributions obtained by maximizing the alternative entropies have, however, turned out to be physically relevant. It is, therefore, noteworthy that those successful distributions can also be derived through a more standard application of MaxEnt or ME, as advocated in this review [8,48–51]. In other words, what is being ruled out are not the distributions themselves, but the alternative entropies from which they were inferred.

On priors.

Choosing the prior density $q(x)$ can be tricky. Sometimes, symmetry considerations can be useful but otherwise, there is no fixed set of rules to translate information into a probability distribution except, of course, for Bayes’ rule and the ME method themselves.

What if the prior $q(x)$ vanishes for some values of x ? $S[p, q]$ can be infinitely negative when $q(x)$ vanishes within some region \mathcal{D} . This means that the ME method confers an infinite preference on those distributions $p(x)$ that vanish whenever $q(x)$ does. One must emphasize that this is as it should be. A similar situation also arises in the context of Bayes’ theorem, where assigning a vanishing prior represents a tremendously serious commitment because no amount of data to the contrary would allow us to revise it. In both ME and Bayes updating, we should recognize the implications of assigning a vanishing prior. Assigning a very low but non-zero prior represents a safer and possibly less prejudiced representation of one’s prior beliefs.

Commuting and non-commuting constraints.

The ME method allows one to process information in the form of constraints. When we are confronted with several constraints, we must be particularly cautious. Should they be processed simultaneously or sequentially? And, if the latter, in what order? The answer depends on the problem at hand [42].

We refer to constraints as *commuting* when it makes no difference whether they are handled simultaneously or sequentially. The most common example is that of Bayesian updating on the basis of data collected in several independent experiments. In this case, the order in which the observed data $x' = \{x'_1, x'_2, \dots\}$ are processed does not matter for the purpose of inferring θ . In general, however, constraints need not commute and when this is the case, the order in which they are processed is critical.

To decide whether constraints are to be handled sequentially or simultaneously, one must be clear about how the ME method handles constraints. The ME machinery interprets a constraint in a very mechanical way: all distributions satisfying the constraint are, in principle, allowed, while all distributions violating it are ruled out. Therefore, sequential updating is appropriate when old constraints become obsolete and are superseded by new information, while simultaneous updating is appropriate when old constraints remain valid. The two cases refer to different states of information, and therefore, it is to be expected that they will result in different inferences. These comments are meant to underscore the importance of understanding what information is and how it is processed by the ME method; failure to do so will lead to errors that do not reflect a shortcoming of the ME method but rather a misapplication of it.

Pitfalls?

Entropy is a tool for reasoning and—as with all tools for reasoning or otherwise—it can be misused, leading to unsatisfactory results [52]. Should that happen, the inevitable questions are “what went wrong?” and “how do we fix it?” It helps to first ask what components of the analysis can be trusted so that the possible mistakes can be looked for elsewhere. The answers proposed by the ME method are radically conservative: problems always arise through a wrong choices of variables, priors, or constraints. Indeed, one should not blame the entropic method for not having discovered and taken into account relevant information that was not explicitly introduced into the analysis. Indeed, just as one would be very reticent about questioning the basic rules of arithmetic, or the basic rules of calculus, one should not question the basic sum and product rules of the probability calculus and, taking this one step farther, one should not question the applicability of entropy as the updating tool. The adoption of this conservative approach leads us to reject alternative entropies and quantum probabilities. Fortunately, those constructs are not actually needed—as mentioned above, those Tsallis distributions that have turned out to be useful can be derived with standard entropic methods [8,48–51], and quantum mechanics can be handled within standard probability theory without invoking exotic probabilities [39,53].

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: I would like to acknowledge many valuable discussions on probability and entropy with N. Caticha, A. Giffin, K. Knuth, R. Preuss, C. Rodríguez, J. Skilling, and K. Vanslette.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. DC1—Mutually Exclusive Subdomains

In these appendices, we establish the consequences of the two criteria DC1 and DC2, leading to the final result: Equation (9). The details of the proofs are important not just because they lead to our final conclusions, but also because the translation of the verbal statement of the criteria into precise mathematical form is a crucial part of unambiguously specifying what the criteria actually say.

First, we prove that criterion DC1 leads to the expression Equation (2) for $S[p, q]$. Consider the case of a discrete variable, p_i with $i = 1 \dots n$, so that $S[p, q] = S(p_1 \dots p_n, q_1 \dots q_n)$. Suppose the space of states \mathcal{X} is partitioned into two non-overlapping domains \mathcal{D} and $\tilde{\mathcal{D}}$ with $\mathcal{D} \cup \tilde{\mathcal{D}} = \mathcal{X}$, and that the information to be processed is in the form of a constraint that refers to the domain $\tilde{\mathcal{D}}$,

$$\sum_{j \in \tilde{\mathcal{D}}} a_j p_j = A. \tag{A1}$$

DC1 states that the constraint on $\tilde{\mathcal{D}}$ does not have an influence on the *conditional* probabilities $p_{i|\mathcal{D}}$. It may, however, influence the probabilities p_i within \mathcal{D} through an overall multiplicative factor. To deal with this complication, consider then a special case where the overall probabilities of \mathcal{D} and $\tilde{\mathcal{D}}$ are also constrained:

$$\sum_{i \in \mathcal{D}} p_i = P_{\mathcal{D}} \quad \text{and} \quad \sum_{j \in \tilde{\mathcal{D}}} p_j = P_{\tilde{\mathcal{D}}}, \tag{A2}$$

with $P_{\mathcal{D}} + P_{\tilde{\mathcal{D}}} = 1$. Under these special circumstances, constraints on $\tilde{\mathcal{D}}$ will not influence p_i s within \mathcal{D} , and vice versa.

To obtain the posterior, maximize $S[p, q]$ subject to these three constraints,

$$0 = \left[\delta S - \lambda \left(\sum_{i \in \mathcal{D}} p_i - P_{\mathcal{D}} \right) + \right. \\ \left. - \tilde{\lambda} \left(\sum_{j \in \tilde{\mathcal{D}}} p_j - P_{\tilde{\mathcal{D}}} \right) + \mu \left(\sum_{j \in \tilde{\mathcal{D}}} a_j p_j - A \right) \right],$$

leading to

$$\frac{\partial S}{\partial p_i} = \lambda \quad \text{for } i \in \mathcal{D}, \tag{A3}$$

$$\frac{\partial S}{\partial p_j} = \tilde{\lambda} + \mu a_j \quad \text{for } j \in \tilde{\mathcal{D}}. \tag{A4}$$

Equations (A1)–(A4) are $n + 3$ equations; we must solve for the p_i s and the three Lagrange multipliers, λ , $\tilde{\lambda}$, and μ . Since $S = S(p_1 \dots p_n, q_1 \dots q_n)$ its derivative

$$\frac{\partial S}{\partial p_i} = f_i(p_1 \dots p_n, q_1 \dots q_n)$$

could, in principle, also depend on all $2n$ variables. However, this violates the DC1 criterion because any arbitrary change in a_j within $\tilde{\mathcal{D}}$ would influence the p_i s within \mathcal{D} . The only way that probabilities conditioned on \mathcal{D} can be shielded from arbitrary changes in the constraints pertaining to $\tilde{\mathcal{D}}$ is that for any $i \in \mathcal{D}$, the function f_i depends only on p_j s with $j \in \mathcal{D}$. Furthermore, this must hold not just for one particular partition of \mathcal{X} into domains \mathcal{D} and $\tilde{\mathcal{D}}$, but it must hold for *all conceivable partitions*, including the partition into atomic propositions. Therefore, f_i can depend only on p_i ,

$$\frac{\partial S}{\partial p_i} = f_i(p_i, q_1 \dots q_n). \tag{A5}$$

The power of the criterion DC1 is not exhausted yet. The information that affects the posterior can enter not just through constraints, but also through the prior. Suppose that the local information about domain $\tilde{\mathcal{D}}$ is altered by changing the prior within $\tilde{\mathcal{D}}$. Let $q_j \rightarrow q_j + \delta q_j$ for $j \in \tilde{\mathcal{D}}$. Then (A5) becomes

$$\frac{\partial S}{\partial p_i} = f_i(p_i, q_1 \dots q_j + \delta q_j \dots q_n),$$

which shows that p_i with $i \in \mathcal{D}$ will be influenced by information about $\tilde{\mathcal{D}}$ unless f_i with $i \in \mathcal{D}$ is independent of all the q_j s for $j \in \tilde{\mathcal{D}}$. Again, this must hold for all possible partitions into \mathcal{D} and $\tilde{\mathcal{D}}$, and therefore,

$$\frac{\partial S}{\partial p_i} = f_i(p_i, q_i) \quad \text{for all } i \in \mathcal{X} .$$

The choice of the functions $f_i(p_i, q_i)$ can be restricted further. If we maximize $S[p, q]$, subject to constraints

$$\sum_i p_i = 1 \quad \text{and} \quad \sum_i a_i p_i = A$$

we obtain

$$\frac{\partial S}{\partial p_i} = f_i(p_i, q_i) = \lambda + \mu a_i \quad \text{for all } i \in \mathcal{X} ,$$

where λ and μ are Lagrange multipliers. Solving for p_i gives a posterior,

$$P_i = g_i(q_i, \lambda, \mu, a_i)$$

for some functions g_i . As stated in Section 3.3 we do not assume that the labels i themselves carry any particular significance. This means, in particular, that for any proposition labeled i , we want the selected posterior P_i to depend only on the numbers q_i , λ , μ , and a_i . We do not want to have different updating rules for different propositions: two different propositions i and i' with the same $q_i = q_{i'}$ and the same $a_i = a_{i'}$ should be updated to the same posteriors, $P_i = P_{i'}$. In other words, the functions g_i and f_i must be independent of i . Therefore,

$$\frac{\partial S}{\partial p_i} = f(p_i, q_i) \quad \text{for all } i \in \mathcal{X} . \tag{A6}$$

Integrating, one obtains

$$S[p, q] = \sum_i F(p_i, q_i) + \text{constant} .$$

for some still undetermined function F . The constant has no effect on the entropy maximization and can be dropped.

The corresponding expression for a continuous variable x is obtained replacing i by x , and the sum over i by an integral over x leading to Equation (2),

$$S[p, q] = \int dx F(p(x), q(x)) .$$

Appendix B. DC2—Independent Subsystems

Here, we show that DC2 leads to Equation (9). Let the microstates of a composite system be labeled by $(i_1, i_2) \in \mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$. We shall consider two special cases.

Case (a)

First, we treat the two subsystems separately. Suppose that for subsystem 1, we have the extremely constraining information that updates $q_1(i_1)$ to be $P_1(i_1)$, and for subsystem 2 we have no new information at all. For subsystem 1, we maximize $S_1[p_1, q_1]$ subject to the constraint $p_1(i_1) = P_1(i_1)$ and the selected posterior is, of course, $p_1(i_1) = P_1(i_1)$. For subsystem 2, we maximize $S_2[p_2, q_2]$ subject only to normalization and there is no update, $P_2(i_2) = q_2(i_2)$.

When the systems are treated jointly, however, the inference is not nearly as trivial. We want to maximize the entropy of the joint system,

$$S[p, q] = \sum_{i_1, i_2} F(p(i_1, i_2), q_1(i_1)q_2(i_2)) ,$$

subject to the constraint on subsystem 1,

$$\sum_{i_1} p(i_1, i_2) = P_1(i_1) .$$

Notice that this is not just one constraint: we have one constraint for each value of i_1 , and each constraint must be supplied with its own Lagrange multiplier, $\lambda_1(i_1)$. Then,

$$\delta [S - \sum_{i_1} \lambda_1(i_1) (\sum_{i_2} p(i_1, i_2) - P_1(i_1))] = 0 .$$

The independent variations $\delta p(i_1, i_2)$ yield the following:

$$f(p(i_1, i_2), q_1(i_1)q_2(i_2)) = \lambda_1(i_1) ,$$

where f is given in (A6),

$$\frac{\partial S}{\partial p} = \frac{\partial}{\partial p} F(p, q_1q_2) = f(p, q_1q_2) .$$

Next, we impose that the selected posterior is the product $P_1(i_1)q_2(i_2)$. The function f must be such that

$$f(P_1q_2, q_1q_2) = \lambda_1 .$$

Since the RHS is independent of the argument i_2 , the f function must be such that the i_2 -dependence cancels out, and this cancellation must occur for all values of i_2 and all choices of the prior q_2 . Therefore, we impose that for any value of x the function $f(p, q)$ must satisfy

$$f(px, qx) = f(p, q) .$$

Choosing $x = 1/q$, we obtain

$$f\left(\frac{p}{q}, 1\right) = f(p, q) \quad \text{or} \quad \frac{\partial F}{\partial p} = f(p, q) = \phi\left(\frac{p}{q}\right) . \tag{A7}$$

Thus, the function $f(p, q)$ has been reduced to a function $\phi(p/q)$ of a single argument.

Case (b)

Next, we consider a situation in which both subsystems are updated by extremely constraining information: when the subsystems are treated separately, $q_1(i_1)$ is updated to $P_1(i_1)$ and $q_2(i_2)$ is updated to $P_2(i_2)$. When the systems are treated jointly, we require that the joint prior for the combined system $q_1(i_1)q_2(i_2)$ be updated to $P_1(i_1)P_2(i_2)$.

First we treat the subsystems separately. Maximize the entropy of subsystem 1,

$$S[p_1, q_1] = \sum_{i_1} F(p_1(i_1), q_1(i_1)) \quad \text{subject to} \quad p_1(i_1) = P_1(i_1) .$$

To each constraint—one constraint for each value of i_1 —we must supply one Lagrange multiplier, $\lambda_1(i_1)$. Then, we obtain

$$\delta [S - \sum_{i_1} \lambda_1(i_1) (p(i_1) - P_1(i_1))] = 0 .$$

Using Equation (A7),

$$\frac{\partial S}{\partial p_1} = \frac{\partial}{\partial p_1} F(p_1, q_1) = \phi\left(\frac{p_1}{q_1}\right) ,$$

and, imposing that the selected posterior be $P_1(i_1)$, we find that the function ϕ must obey

$$\phi\left(\frac{P_1(i_1)}{q_1(i_1)}\right) = \lambda_1(i_1) . \tag{A8}$$

Similarly, for system 2 we find the following:

$$\phi\left(\frac{P_2(i_2)}{q_2(i_2)}\right) = \lambda_2(i_2). \quad (\text{A9})$$

Next, we treat the two subsystems jointly. Maximize the entropy of the joint system as follows:

$$S[p, q] = \sum_{i_1, i_2} F(p(i_1, i_2), q_1(i_1)q_2(i_2)),$$

subject to the following constraints on the joint distribution $p(i_1, i_2)$:

$$\sum_{i_2} p(i_1, i_2) = P_1(i_1) \quad \text{and} \quad \sum_{i_1} p(i_1, i_2) = P_2(i_2).$$

Again, there is one constraint for each value of i_1 and of i_2 and we introduce Lagrange multipliers, $\eta_1(i_1)$ or $\eta_2(i_2)$. Then,

$$\delta[S - \sum_{i_1} \eta_1(i_1)(\sum_{i_2} p(i_1, i_2) - P_1(i_1)) - \{1 \leftrightarrow 2\}] = 0,$$

where $\{1 \leftrightarrow 2\}$ indicates a third term, similar to the second, with 1 and 2 interchanged. The independent variations $\delta p(i_1, i_2)$ yield

$$\phi\left(\frac{p(i_1, i_2)}{q_1(i_1)q_2(i_2)}\right) = \eta_1(i_1) + \eta_2(i_2),$$

and we impose that the selected posterior be the product $P_1(i_1)P_2(i_2)$. Therefore, the function ϕ must be such that

$$\phi\left(\frac{P_1P_2}{q_1q_2}\right) = \eta_1 + \eta_2.$$

To solve this equation, we take the exponential of both sides, let $\xi = \exp \phi$, and rewrite as

$$\xi\left(\frac{P_1P_2}{q_1q_2}\right)e^{-\eta_2(i_2)} = e^{\eta_1(i_1)}. \quad (\text{A10})$$

This shows that for any value of i_1 , the dependences of the LHS on i_2 through P_2/q_2 and η_2 must cancel each other out. In particular, if for some subset of i_2 s, the subsystem 2 is updated so that $P_2 = q_2$, which amounts to no update at all, the i_2 dependence on the left is eliminated but the i_1 dependence remains unaffected,

$$\xi\left(\frac{P_1}{q_1}\right)e^{-\eta'_2} = e^{\eta_1(i_1)}.$$

where η'_2 is some constant independent of i_2 . A similar argument with $\{1 \leftrightarrow 2\}$ yields

$$\xi\left(\frac{P_2}{q_2}\right)e^{-\eta'_1} = e^{\eta_2(i_2)},$$

where η'_1 is a constant. Taking the exponential of (A8) and (A9) leads to the following:

$$\xi\left(\frac{P_1}{q_1}\right)e^{-\eta'_2} = e^{\lambda_1 - \eta'_2} = e^{\eta_1} \quad \text{and} \quad \xi\left(\frac{P_2}{q_2}\right)e^{-\eta'_1} = e^{\lambda_2 - \eta'_1} = e^{\eta_2}.$$

Substituting back into (A10), we obtain

$$\xi\left(\frac{P_1P_2}{q_1q_2}\right) = \xi\left(\frac{P_1}{q_1}\right)\xi\left(\frac{P_2}{q_2}\right),$$

where a constant factor $e^{-(\eta'_1+\eta'_2)}$ is absorbed into a new function ζ . The general solution of this functional equation is a power,

$$\zeta(xy) = \zeta(x)\zeta(y) \implies \zeta(x) = x^a,$$

so that

$$\phi(x) = a \log x + b,$$

where a and b are constants. Finally, integrate (A7),

$$\frac{\partial F}{\partial p} = \phi\left(\frac{p}{q}\right) = a \log \frac{p}{q} + b,$$

to obtain

$$F[p, q] = ap \log \frac{p}{q} + b'p + c$$

where b' and c are constants.

At this point, the entropy takes the general form

$$S[p, q] = \sum_i \left(ap_i \log \frac{p_i}{q_i} + b'p_i + c \right).$$

The additive constant c may be dropped: it contributes a term that does not depend on the probabilities and has no effect on the ranking scheme. Furthermore, since $S[p, q]$ will be maximized subject to constraints that include normalization, the b' term has no effect on the selected distribution and can also be dropped. Finally, the multiplicative constant a has no effect on the overall ranking, except in the trivial sense that inverting the sign of a will transform the maximization problem to a minimization problem or vice versa. We can, therefore, set $a = -1$ so that maximum S corresponds to maximum preference, which gives us Equation (9) and concludes our derivation.

References and Notes

- Caticha, A. Relative Entropy and Inductive Inference. *AIP Conf. Proc.* **2004**, *707*, 75.
- Caticha, A.; Giffin, A. Updating Probabilities. *AIP Conf. Proc.* **2006**, *872*, 31.
- Caticha, A. Information and Entropy. *AIP Conf. Proc.* **2007**, *954*, 11.
- Caticha, A. Towards an Informational Pragmatic Realism. *Mind Mach.* **2014**, *24*, 37. [CrossRef]
- Vanslette, K. Entropic Updating of Probabilities and Density Matrices. *Entropy* **2017**, *19*, 664. [CrossRef]
- Caticha, A. *Lectures on Probability, Entropy, and Statistical Physics*; MaxEnt: São Paulo, Brazil, 2008.
- Caticha, A. *Entropic Inference and the Foundations of Physics*; EBEB: São Paulo, Brazil, 2012. Available online: <https://www.albany.edu/physics/faculty/ariel-caticha> (accessed on 12 July 2021).
- Caticha, A. Entropic Physics: Probability, Entropy, and the Foundations of Physics. 2021. Available online: <https://www.albany.edu/physics/faculty/ariel-caticha> (accessed on 12 July 2021)
- Jaynes, E.T. Information Theory and Statistical Mechanics, I and II. *Phys. Rev.* **1957**, *106*, 620. [CrossRef]
- Rosenkrantz, R.D. (Ed.) *E. T. Jaynes: Papers on Probability, Statistics and Statistical Physics*; Reidel: Dordrecht, The Netherlands, 1983.
- Uffink, J. Can the Maximum Entropy Principle be explained as a consistency requirement? *Stud. Hist. Philos. Mod. Phys.* **1995**, *26*, 223. [CrossRef]
- Rényi, A. On measures of entropy and information. In Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, CA, USA, 20 June–30 July 1961; p. 547.
- Aczél, J.; Daróczy, Z. *On Measures of Information and Their Characterizations*; Academic Press: New York, NY, USA, 1975.
- Tsallis, C. Possible Generalization of Boltzmann-Gibbs Statistics. *J. Stat. Phys.* **1988**, *52*, 479. [CrossRef]
- Shore, J.E.; Johnson, R.W. Axiomatic derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy. *IEEE Trans. Inf. Theory* **1980**, *26*, 26–37. [CrossRef]
- Skilling, J. The Axioms of Maximum Entropy. In *Maximum-Entropy and Bayesian Methods in Science and Engineering*; Erickson, G.J., Smith, C.R., Eds.; Kluwer: Dordrecht, The Netherlands, 1988; pp. 173–187.
- Skilling, J. Classic Maximum Entropy. In *Maximum Entropy and Bayesian Methods*; Skilling, J., Ed.; Kluwer: Dordrecht, The Netherlands, 1989; pp. 45–52.

18. Karbelkar, S.N. On the axiomatic approach to the maximum entropy principle of inference. *Pramana J. Phys.* **1986**, *26*, 301–310. [[CrossRef](#)]
19. Tsallis, C. Conceptual Inadequacy of the Shore and Johnson Axioms for Wide Classes of Complex systems. *Entropy* **2015**, *17*, 2853. [[CrossRef](#)]
20. Jizba, P.; Korbel, J. Maximum Entropy Principle in Statistical Inference: Case for Non-Shannonian Entropies. *Phys. Rev. Lett.* **2019**, *122*, 120601. [[CrossRef](#)]
21. Pressé, S.; Ghosh, K.; Lee, J.; Dill, K.A. Nonadditive Entropies Yield Probability Distributions with Biases not Warranted by the Data. *Phys. Rev. Lett.* **2013**, *111*, 180604. [[CrossRef](#)]
22. Pressé, S.; Ghosh, K.; Lee, J.; Dill, K.A. Reply to Tsallis' Conceptual inadequacy of the Shore and Johnson axioms for wide classes of complex systems. *Entropy* **2015**, *17*, 5043. [[CrossRef](#)]
23. Williams, P.M. Bayesian Conditionalization and the Principle of Minimum Relative Information. *Br. J. Philos. Sci.* **1980**, *31*, 131. [[CrossRef](#)]
24. Johnson, D.T.; Caticha, A. Entropic dynamics and the quantum measurement problem. *AIP Conf. Proc.* **2012**, *1443*, 104.
25. Vanslette, K.; Caticha, A. Quantum measurement and weak values in entropic quantum dynamics. *AIP Conf. Proc.* **2017**, *1853*, 090003.
26. Cover, T.; Thomas, J. *Elements of Information Theory*; Wiley: New York, NY, USA, 1991.
27. Golan, A. *Foundations of Info-Metrics: Modeling, Inference, and Imperfect Information*; Oxford University Press: New York, NY, USA, 2018.
28. Rissanen, J. Modeling by shortest data description. *Automatica* **1978**, *14*, 465–471. [[CrossRef](#)]
29. Landauer, R. Information is Physical. *Phys. Today* **1991**, *44*, 23–29. [[CrossRef](#)]
30. Bennett, C. The thermodynamics of computation—A review. *Int. J. Theor. Phys.* **1982**, *21*, 905–940. [[CrossRef](#)]
31. Bennett, C. Notes on Landauer's principle, reversible computation, and Maxwell's demon. *Stud. Hist. Philos. Mod. Phys.* **2003**, *34*, 501–510. [[CrossRef](#)]
32. Norton, J.D. Waiting for Landauer. *Stud. Hist. Philos. Mod. Phys.* **2011**, *36*, 184–198. [[CrossRef](#)]
33. Norton, J.D. The End of the Thermodynamics of Computation: A No-Go Result. *Philos. Sci.* **2013**, *80*, 1182–1192. [[CrossRef](#)]
34. Binmore, K. On the foundations of decision theory. *Homo Oecon.* **2017**, *34*, 259. [[CrossRef](#)]
35. Harre, M.S. Information Theory for Agents in Artificial Intelligence, Psychology, and Economics. *Entropy* **2021**, *23*, 310. [[CrossRef](#)] [[PubMed](#)]
36. Evans, B.P.; Prokopenko, M. A Maximum Entropy Model of Bounded Rational Decision-Making with Prior Beliefs and Market Feedback. *Entropy* **2021**, *23*, 669. [[CrossRef](#)]
37. Caticha, A.; Golan, A. An Entropic framework for Modeling Economies. *Physica A* **2014**, *408*, 149. [[CrossRef](#)]
38. Pessoa, P.; Costa, F.X.; Caticha, A. Entropic dynamics on Gibbs statistical manifolds. *Entropy* **2021**, *23*, 494. [[CrossRef](#)]
39. Caticha, A. The Entropic Dynamics approach to Quantum Mechanics. *Entropy* **2019**, *21*, 943. [[CrossRef](#)]
40. Knuth, K.H. Lattice duality: The origin of probability and entropy. *Neurocomputing* **2005**, *67*, 245–274. [[CrossRef](#)]
41. Knuth, K.H.; Skilling, J. Foundations of Inference. *Axioms* **2012**, *1*, 38–73. [[CrossRef](#)]
42. Giffin, A.; Caticha, A. Updating Probabilities with Data and Moments. *AIP Conf. Proc.* **2007**, *954*, 74.
43. Amari, S. *Differential-Geometrical Methods in Statistics*; Springer: Berlin, Germany, 1985.
44. Caticha, A.; Preuss, R. Maximum entropy and Bayesian data analysis: Entropic prior distributions. *Phys. Rev. E* **2004**, *70*, 046127. [[CrossRef](#)]
45. Grendar, M., Jr.; Grendar, M. Maximum Probability and Maximum Entropy Methods: Bayesian interpretation. *AIP Conf. Proc.* **2004**, *707*, 490.
46. La Cour, B.R.; Schieve, W.C. Tsallis maximum entropy principle and the law of large numbers. *Phys. Rev. E* **2000**, *62*, 7494. [[CrossRef](#)] [[PubMed](#)]
47. Nauenberg, M. Critique of q -entropy for thermal statistics. *Phys. Rev. E* **2003**, *67*, 036114. [[CrossRef](#)] [[PubMed](#)]
48. Plastino, A.R.; Plastino, A. From Gibbs microcanonical ensemble to Tsallis generalized canonical distribution. *Phys. Lett. A* **1994**, *193*, 140–143. [[CrossRef](#)]
49. Beck, C. Dynamical Foundations of nonextensive Statistical Mechanics. *Phys. Rev. Lett.* **2001**, *87*, 180601. [[CrossRef](#)]
50. Beck, C.; Cohen, E.G.D. Superstatistics. *Physica A* **2003**, *322*, 267–275. [[CrossRef](#)]
51. Crooks, G.E. Beyond Boltzmann-Gibbs statistics: Maximum entropy hyperensembles out of equilibrium. *Phys. Rev. E* **2007**, *75*, 041119. [[CrossRef](#)]
52. Caticha, A. Entropic inference: Some pitfalls and paradoxes we can avoid. *AIP Conf. Proc.* **2013**, *1553*, 176.
53. Bohm, D.; Hiley, B.J. *The Undivided Universe: An Ontological Interpretation on Quantum Theory*; Routledge: New York, NY, USA, 1993.