

# Inferring Diversity: Life After Shannon

**Adom Giffin**

Department of Physics  
University at Albany–SUNY  
Albany, NY 12222,USA

The diversity of a community that cannot be fully counted must be inferred. The two preeminent inference methods are the MaxEnt method, which uses information in the form of constraints and Bayes' rule which uses information in the form of data. It has been shown that these two methods are special cases of the method of Maximum (relative) Entropy (ME). We demonstrate how this method can be used as a measure of diversity that not only reproduces the features of Shannon's index but exceeds them by allowing more types of information to be included in the inference. A specific example is solved in detail. Additionally, the entropy that is found is the same form as the thermodynamic entropy.

## 1 Introduction

Diversity is a concept that is used in many fields to describe the variability of different entities in a group. In ecology, the Shannon entropy [1] and Simpson's index [2] are the predominate measures of diversity. In this paper we focus on the Shannon entropy for two reasons: First, it has been shown that Simpson's index is an approximation of Shannon's [3]. Second, Shannon's entropy is closely tied to many other areas of research, such as information theory and physics.

It is often the case that the species in a community cannot be fully counted. In this case, when one has incomplete information, one must rely on methods of inference. The two preeminent inference methods are the MaxEnt [4] method, which has evolved to a more general method, the method of Maximum (relative) Entropy (ME) [5, 6, 7] and Bayes' rule. The choice between the two methods has traditionally been dictated by the nature of the information being processed

(either constraints or observed data). However, it has been shown that one can accommodate both types of information in one method, ME [8]. The purpose of this paper is to demonstrate how the ME method can be used as a measure of diversity that is able to include more information that Shannon's measure allows.

Traditionally when confronted with a community whose count is incomplete, the frequency of the species that are counted are used to calculate the diversity. The frequency is used because it represents an estimate of the probability of finding a particular species in the community. However, the frequency is not equivalent to the probability [9] and as such is a poor estimate. Fortunately, there are much better methods for estimating or inferring the probability such as MaxEnt and Bayes. Even more fortunate is that the new ME method can reproduce every aspect of Bayesian and MaxEnt inference *and* tackle problems that the two methods alone could not address.

We start by showing a general example of the ME method by inferring a probability with two different forms of information: expected values<sup>1</sup> and data, *simultaneously*. The solution resembles Bayes' Rule. In fact, if there are no moment constraints then the method produces Bayes rule *exactly*. If there is no data, then the MaxEnt solution is produced.

Finally we solve a toy ecological problem and discuss the diversity calculated by using Shannon's entropy and the diversity calculated by the ME method. This illustrates the many advantages to using the ME method.

## 2 Simultaneous updating

Our first concern when using the ME method to update from a prior to a posterior distribution<sup>2</sup> is to define the space in which the search for the posterior will be conducted. We wish to infer something about the values of one or several quantities,  $\theta \in \Theta$ , on the basis of three pieces of information: prior information about  $\theta$  (the prior), the known relationship between  $x$  and  $\theta$  (the model), and the observed values of the data  $x \in \mathcal{X}$ . Since we are concerned with both  $x$  and  $\theta$ , the relevant space is neither  $\mathcal{X}$  nor  $\Theta$  but the product  $\mathcal{X} \times \Theta$  and our attention must be focused on the joint distribution  $P(x, \theta)$ . The selected joint posterior  $P_{\text{new}}(x, \theta)$  is that which maximizes the entropy,

$$S[P, P_{\text{old}}] = - \int dx d\theta P(x, \theta) \log \frac{P(x, \theta)}{P_{\text{old}}(x, \theta)} , \quad (1)$$

---

<sup>1</sup>For simplicity we will refer to these expected values as *moments* although they can be considerably more general.

<sup>2</sup>In Bayesian inference, it is assumed that one always has a prior probability based on some prior information. When new information is attained, the old probability (the prior) is *updated* to a new probability (the posterior). If one has no prior information, then one uses an *ignorant* prior [10].

subject to the appropriate constraints.  $P_{\text{old}}(x, \theta)$  contains our prior information which we call the *joint prior*. To be explicit,

$$P_{\text{old}}(x, \theta) = P_{\text{old}}(\theta)P_{\text{old}}(x|\theta) , \quad (2)$$

where  $P_{\text{old}}(\theta)$  is the traditional Bayesian prior and  $P_{\text{old}}(x|\theta)$  is the likelihood. It is important to note that they *both* contain prior information. The Bayesian prior is defined as containing prior information. However, the likelihood is not traditionally thought of in terms of prior information. Of course it is reasonable to see it as such because the likelihood represents the model (the relationship between  $\theta$  and  $x$ ) that has already been established. Thus we consider both pieces, the Bayesian prior and the likelihood to be *prior* information.

The new information is the *observed data*,  $x'$ , which in the ME framework must be expressed in the form of a constraint on the allowed posteriors. The family of posteriors that reflects the fact that  $x$  is now known to be  $x'$  is such that

$$C_1 : P(x) = \int d\theta P(x, \theta) = \delta(x - x') . \quad (3)$$

This amounts to an *infinite* number of constraints: there is one constraint on  $P(x, \theta)$  for each value of the variable  $x$  and each constraint will require its own Lagrange multiplier  $\lambda(x)$ . Furthermore, we impose the usual normalization constraint,

$$\int dx d\theta P(x, \theta) = 1 , \quad (4)$$

and include additional information about  $\theta$  in the form of a constraint on the expected value of some function  $f(\theta)$ <sup>3</sup>,

$$C_2 : \int dx d\theta P(x, \theta) f(\theta) = \langle f(\theta) \rangle = F . \quad (5)$$

We emphasize that constraints imposed at the level of the prior need not be satisfied by the posterior. What we do here differs from the standard Bayesian practice in that we *require* the constraint to be satisfied by the posterior distribution.

Maximize (1) subject to the above constraints,

$$\delta \left\{ \begin{array}{l} S + \alpha [\int dx d\theta P(x, \theta) - 1] \\ + \beta [\int dx d\theta P(x, \theta) f(\theta) - F] \\ + \int dx \lambda(x) [\int d\theta P(x, \theta) - \delta(x - x')] \end{array} \right\} = 0 , \quad (6)$$

yields the joint posterior,

$$P_{\text{new}}(x, \theta) = P_{\text{old}}(x, \theta) \frac{e^{\lambda(x) + \beta f(\theta)}}{Z} , \quad (7)$$

---

<sup>3</sup>Including an additional constraint in the form of  $\int dx d\theta P(x, \theta) g(x) = \langle g \rangle = G$  could only be used when it does not contradict the data constraint (3). Therefore, it is redundant and the constraint would simply get absorbed when solving for  $\lambda(x)$ .

where  $Z$  is determined by using (4),

$$Z = e^{-\alpha+1} = \int dx d\theta e^{\lambda(x)+\beta f(\theta)} P_{\text{old}}(x, \theta) \quad (8)$$

and the Lagrange multipliers  $\lambda(x)$  are determined by using (3)

$$e^{\lambda(x)} = \frac{Z}{\int d\theta e^{\beta f(\theta)} P_{\text{old}}(x, \theta)} \delta(x - x') . \quad (9)$$

The posterior now becomes

$$P_{\text{new}}(x, \theta) = P_{\text{old}}(x, \theta) \delta(x - x') \frac{e^{\beta f(\theta)}}{\zeta(x, \beta)} , \quad (10)$$

where  $\zeta(x, \beta) = \int d\theta e^{\beta f(\theta)} P_{\text{old}}(x, \theta)$ .

The Lagrange multiplier  $\beta$  is determined by first substituting the posterior into (5),

$$\int dx d\theta \left[ P_{\text{old}}(x, \theta) \delta(x - x') \frac{e^{\beta f(\theta)}}{\zeta(x, \beta)} \right] f(\theta) = F . \quad (11)$$

Integrating over  $x$  yields,

$$\frac{\int d\theta e^{\beta f(\theta)} P_{\text{old}}(x', \theta) f(\theta)}{\zeta(x', \beta)} = F , \quad (12)$$

where  $\zeta(x, \beta) \rightarrow \zeta(x', \beta) = \int d\theta e^{\beta f(\theta)} P_{\text{old}}(x', \theta)$ . Now  $\beta$  can be determined by

$$\frac{\partial \ln \zeta(x', \beta)}{\partial \beta} = F . \quad (13)$$

The final step is to marginalize the posterior,  $P_{\text{new}}(x, \theta)$  over  $x$  to get our updated probability,

$$P_{\text{new}}(\theta) = P_{\text{old}}(x', \theta) \frac{e^{\beta f(\theta)}}{\zeta(x', \beta)} \quad (14)$$

Additionally, this result can be rewritten using the product rule as

$$P_{\text{new}}(\theta) = P_{\text{old}}(\theta) P_{\text{old}}(x'|\theta) \frac{e^{\beta f(\theta)}}{\zeta'(x', \beta)} , \quad (15)$$

where  $\zeta'(x', \beta) = \int d\theta e^{\beta f(\theta)} P_{\text{old}}(\theta) P_{\text{old}}(x'|\theta)$ . The right side resembles Bayes theorem, where the term  $P_{\text{old}}(x'|\theta)$  is the standard Bayesian likelihood and  $P_{\text{old}}(\theta)$  is the prior. The exponential term is a *modification* to these two terms. Notice when  $\beta = 0$  (no moment constraint) we recover Bayes' rule. For  $\beta \neq 0$  Bayes' rule is modified by a "canonical" exponential factor.

It must be noted that MaxEnt has been traditionally used for obtaining a prior for use in Bayesian statistics. When this is the case, the updating is sequential. This is not the case here where both types of information are processed simultaneously. In the sequential updating case, the multiplier  $\beta$  is chosen so that the posterior  $P_{\text{new}}$  only satisfies  $C_2$ . In the simultaneous updating case the multiplier  $\beta$  is chosen so that the posterior  $P_{\text{new}}$  satisfies both  $C_1$  and  $C_2$  or  $C_1 \wedge C_2$  [8].

### 3 Inference in Ecology

In the following sections we will discuss the traditional way diversity is measured and the way it is measured using ME. This will be done by examining a simple example and comparing the two methods. In addition, we will show how the ME method could include information that the traditional method cannot.

The general information for the example is as follows: There are  $k$  types of plants in a forest. A portion of the forest is examined and the amount of each species is counted where  $m_1, m_2 \dots m_k$  represents the counts of each species and  $n$  represents the total count so that  $n = \sum_i^k m_i$ . Additionally, we know from biological examination that one species,  $s_2$  and another species,  $s_5$  are codependent. Perhaps they need each others pollen in such supply that they cannot exist unless there are on the average, twice the number of  $s_2$  as compared to  $s_5$ .

#### 3.1 Traditional Diversity

We calculate the Shannon diversity by using Shannon's entropy as follows,

$$S_{Traditional} = - \sum_i^k p_i \log p_i , \quad (16)$$

where  $p_i = m_i/n$ . The problem with using this method is not in the method itself but with the reason it is being used. If the purpose of using this method was to measure the diversity of the portion that was counted then the method is acceptable. However, if the purpose of the method is to estimate or infer the diversity of the *whole* forest, then it is a poor estimate. First,  $p_i$  is meant to represent the probability of finding the  $i$ th species in the forest. As previously stated, the frequency of the sample is not equivalent to the probability. In fact, it is the expected value of the frequency that is equivalent to the probability,  $\langle F \rangle = p$  [9]. It would only make sense to use the frequency as an estimate of the probability when  $n$  is very large (i.e.  $n \rightarrow \infty$ ) but this is not usually the case. Second, the diversity of two samples that have the same ratio of frequencies will be the same. Therefore this measure does not reflect the abundance of the species. This might be a desirable feature [3]. Third, there is no clear way to process the information about the codependence using Shannon's entropy.

#### 3.2 ME Diversity

Here we intend to use a better method to estimate or infer  $p_i$  and that method is the ME method. The first task is to realize that the correct mathematical model for the probability of getting a particular species where the information that we have is the number of species counted is a multinomial distribution. The probability of finding  $k$  species in  $n$  counts which yields  $m_i$  instances for the  $i^{th}$

species is

$$P_{\text{old}}(m|p, n) = P_{\text{old}}(m_1 \dots m_k | p_1 \dots p_k, n) = \frac{n!}{m_1! \dots m_k!} p_1^{m_1} \dots p_k^{m_k} , \quad (17)$$

where  $m = (m_1, \dots, m_k)$  with  $\sum_{i=1}^k m_i = n$ , and  $p = (p_1, \dots, p_k)$  with  $\sum_{i=1}^k p_i = 1$ . The general problem is to infer the parameters  $p$  on the basis of information about the data,  $m'$ . Here we see the first advantage with using the ME diversity; we allow for fluctuations in our inference by looking at a distribution of  $p$ 's as opposed to claiming that we know the "true"  $p$ .

Additionally we can include information about the codependence by using the following general constraint,

$$\langle f(p) \rangle = F \quad \text{where} \quad f(p) = \sum_i^k f_i p_i , \quad (18)$$

where  $f_i$  is used to represent the codependence. For our example, on the average, we will find twice the number of  $s_2$  as compared to  $s_5$  thus, *on the average*, the probability of finding one of the species will be twice that of the other,  $\langle p_2 \rangle = 2 \langle p_5 \rangle$ . In this case,  $f_2 = 1$ ,  $f_5 = -2$  and  $f_{i \neq (2,5)} = F = 0$ .

Next we need to write the data (counts) as a constraint which in general is

$$P(m|n) = \delta_{mm'} , \quad (19)$$

where  $m' = \{m'_1, \dots, m'_k\}$ . Finally we write the appropriate entropy to use,

$$S[P, P_{\text{old}}] = - \sum_m \int dp P(m, p|n) \log \frac{P(m, p|n)}{P_{\text{old}}(m, p|n)} , \quad (20)$$

where

$$\sum_m = \sum_{m_1 \dots m_k=0}^n \delta(\sum_{i=1}^k m_i - n) , \quad (21)$$

and

$$\int dp = \int dp_1 \dots dp_k \delta\left(\sum_{i=1}^k p_i - 1\right) , \quad (22)$$

and where  $P_{\text{old}}(m, p|n) = P_{\text{old}}(p|n)P_{\text{old}}(m|p, n)$ . The prior  $P_{\text{old}}(p)$  is not important for our current purpose so for the sake of definiteness we can choose it flat for our example (there are most likely better choices for priors). We then maximize this entropy with respect to  $P(m, p|n)$  subject to normalization and our constraints which after marginalizing over  $m'$  yields,

$$P(p) = P_{\text{old}}(m'|p, n) \frac{e^{\beta f(p)}}{\zeta} , \quad (23)$$

where

$$\zeta = \int dp e^{\beta f(p)} P_{\text{old}}(m'|p, n) \quad \text{and} \quad F = \frac{\partial \log \zeta}{\partial \beta} . \quad (24)$$

The probability distribution  $P(p)$  has sometimes been criticized for being too strange. The idea of getting a probability of a probability may seem strange at first but makes absolute sense. We do not know the "true" distribution of species,  $p_i$ . Therefore it seems natural to express our knowledge with some uncertainty in the form of a distribution. Notice that if one has no information relating the species then  $\beta = 0$ .

Finally by substituting (23) into (20), and using our constraints (18) and (19) we introduce our new general measure for diversity,

$$S_{ME} = \log \zeta - \beta F . \quad (25)$$

## 4 Conclusions

Diversity is an important concept in many fields. In this paper we provided a toy example of how ME would be used as a measure of diversity that may simulate real world situations. By using the multinomial, we not only properly infer  $p$  so that fluctuations are represented, we get the additional bonus of having the abundance of the species represented in the measure. It is critical to note that our diversity,  $S_{ME}$  satisfies all of Pielou's axioms [11].

This of course could all be done with only using Bayes to infer  $p$ . However, by using the ME method we can include additional information allowing to go beyond what Bayes' rule and MaxEnt methods alone could do. Therefore, we would like to emphasize that anything one can do with Bayesian or MaxEnt methods, one can now do with ME. Additionally, in ME one now has the ability to apply additional information that Bayesian or MaxEnt methods could not. Further, any work done with Bayesian techniques can be implemented into the ME method directly through the joint prior.

Although Shannon had discovered the entropy that bears his name quite independently of thermodynamic considerations, it nevertheless is directly proportional to the thermodynamic entropy. The realization that the ME diversity is of the exact same form as the thermodynamic entropy<sup>4</sup> is of no small consequence. All of the concepts that thermodynamics utilizes can now also be utilized in ecology, whether it be energy considerations or equilibrium conditions, etc.

To see a detailed method for calculating  $\zeta$ , see [8], for a numeric example, see [12] and for an example of what to do when one knows that there are species in the forest but simply have not been counted (perhaps they are rare), see [13].

**Acknowledgements:** We would like to acknowledge many valuable discussions with A. Caticha

---

<sup>4</sup>The thermodynamic entropy is actually,  $S = \log \zeta + \beta F$ . The fact that our entropy (25) has a  $-\beta F$  is a reflection of our choice to add our Lagrange multipliers in (6) as opposed to subtracting them as is the case in thermodynamics. However, this is trivial because when one solves for  $\beta$  in (13) the sign will be accounted for. Thus, if the Lagrange multiplier was subtracted, the solution to (13) would be  $-F$  and the entropy would have a  $+\beta F$ .

## Bibliography

- [1] C. E. Shannon, "A Mathematical Theory of Communication", *Bell System Technical Journal*, **27**, 379 (1948).
- [2] E. H. Simpson, "Measurement of Diversity", *Nature*, **163**, 688 (1949).
- [3] R. Guiasu and S. Guiasu, "Conditional and Weighted Measures of Ecological Diversity", *International J. of Uncertainty*, **11**, 283 (2002).
- [4] E. T. Jaynes, Phys. Rev. **106**, 620 and **108**, 171 (1957); R. D. Rosenkrantz (ed.), *E. T. Jaynes: Papers on Probability, Statistics and Statistical Physics* (Reidel, Dordrecht, 1983); E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, 2003).
- [5] J. E. Shore and R. W. Johnson, IEEE Trans. Inf. Theory **IT-26**, 26 (1980); IEEE Trans. Inf. Theory **IT-27**, 26 (1981).
- [6] J. Skilling, "The Axioms of Maximum Entropy", *Maximum-Entropy and Bayesian Methods in Science and Engineering*, G. J. Erickson and C. R. Smith (eds.) (Kluwer, Dordrecht, 1988).
- [7] A. Caticha and A. Giffin, "Updating Probabilities", *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. by Ali Mohammad-Djafari (ed.), AIP Conf. Proc. **872**, 31 (2006) (<http://arxiv.org/abs/physics/0608185>).
- [8] A. Giffin and A. Caticha, "Updating Probabilities with Data and Moments", *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. by Kevin Knuth, et al., AIP Conf. Proc. **954**, 74 (2007) (<http://arxiv.org/abs/0708.1593>).
- [9] D. S. Sivia, *Data Analysis: A Bayesian Tutorial* (Oxford U. Press, 1996).
- [10] A. Gelman, et al., *Bayesian Data Analysis, 2nd edition* (CRC Press, 2004).
- [11] E. Pielou. *Ecological Diversity* (Wiley, New York 1975).
- [12] A. Giffin, "Updating Probabilities with Data and Moments: An Econometric Example", presented at the *3rd Econophysics Colloquium*, Ancona, Italy, 2007 (<http://arxiv.org/abs/0710.2912>).
- [13] A. Giffin, "Updating Probabilities with Data and Moments: A Complex Agent Based Example", presented at the *7th International Conference on Complex Systems*, Boston, 2007.