

Updating Probabilities with Data and Moments

Cite as: AIP Conference Proceedings **954**, 74 (2007); <https://doi.org/10.1063/1.2821302>

Published Online: 15 November 2007

Adom Giffin, and Ariel Caticha



View Online



Export Citation

ARTICLES YOU MAY BE INTERESTED IN

Updating Probabilities

AIP Conference Proceedings **872**, 31 (2006); <https://doi.org/10.1063/1.2423258>

Entropic Inference

AIP Conference Proceedings **1305**, 20 (2011); <https://doi.org/10.1063/1.3573619>

Relative Entropy and Inductive Inference

AIP Conference Proceedings **707**, 75 (2004); <https://doi.org/10.1063/1.1751358>

Lock-in Amplifiers
up to 600 MHz



Updating Probabilities with Data and Moments

Adom Giffin and Ariel Caticha

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

Abstract. We use the method of Maximum (relative) Entropy to process information in the form of observed data and moment constraints. The generic “canonical” form of the posterior distribution for the problem of simultaneous updating with data and moments is obtained. We discuss the general problem of non-commuting constraints, when they should be processed sequentially and when simultaneously. As an illustration, the multinomial example of die tosses is solved in detail for two superficially similar but actually very different problems.

INTRODUCTION

The original method of Maximum Entropy, MaxEnt [1], was designed to assign probabilities on the basis of information in the form of constraints. It gradually evolved into a more general method, the method of Maximum relative Entropy (abbreviated ME) [2]-[6], which allows one to update probabilities from arbitrary priors unlike the original MaxEnt which is restricted to updates from a uniform background measure.

The realization [5] that ME includes not just MaxEnt but also Bayes’ rule as special cases is highly significant. First, it implies that ME is *capable of reproducing every aspect of orthodox Bayesian inference* and proves the complete compatibility of Bayesian and entropy methods. Second, it opens the door to tackling problems that could not be addressed by either the MaxEnt or orthodox Bayesian methods individually. The main goal of this paper is to explore this latter possibility: the problem of processing data plus additional information in the form of expected values.¹

When using Bayes’ rule it is quite common to impose constraints on the prior distribution. In some cases these constraints are also satisfied by the posterior distribution, but these are special cases. In general, constraints imposed on priors do not “propagate” to the posteriors. Although Bayes’ rule can handle *some* constraints, we seek a procedure capable of enforcing *any* constraint on the posterior distributions.

After a brief review of how ME processes data and reproduces Bayes’ rule, we derive our main result, the general “canonical” form of the posterior distribution for the problem of simultaneous updating with data and moment constraints. The final result is deceptively simple: Bayes’ rule is modified by a “canonical” exponential factor. Although this result is very simple, it should be handled with caution: once we consider several sources of information such as multiple constraints we must confront the problem of non-commuting constraints. We discuss the question of whether they should be

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

processed simultaneously, or sequentially, and in what order. Our general conclusion is that these different alternatives correspond to different states of information and accordingly we expect that they will lead to different inferences.

As an illustration, the multinomial example of die tosses is solved in some detail for two problems. They appear superficially similar but are in fact very different. The first die problem requires that the constraints be processed sequentially. This corresponds to the familiar situation of using MaxEnt to derive a prior and then using Bayes to process data. The second die problem, which requires that the constraints be processed simultaneously, provides a clear example that lies beyond the reach of Bayes' rule.

UPDATING WITH DATA USING THE ME METHOD

Our first concern when using the ME method to update from a prior to a posterior distribution is to define the space in which the search for the posterior will be conducted. We wish to infer something about the value of a quantity $\theta \in \Theta$ on the basis of three pieces of information: prior information about θ (the prior), the known relationship between x and θ (the model), and the observed values of the data $x \in \mathcal{X}$.² Since we are concerned with both x and θ , the relevant space is neither \mathcal{X} nor Θ 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,

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

subject to the appropriate constraints. All prior information is codified into the *joint prior* $P_{\text{old}}(x, \theta) = P_{\text{old}}(\theta)P_{\text{old}}(x|\theta)$. Both $P_{\text{old}}(\theta)$ (the familiar Bayesian prior distribution) and $P_{\text{old}}(x|\theta)$ (the likelihood) contain prior information.³ 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 $P(x, \theta)$ that reflects the fact that x is now known to be x' is such that

$$P(x) = \int d\theta P(x, \theta) = \delta(x - x'). \quad (2)$$

This amounts to an *infinite* number of constraints on $P(x, \theta)$: for each value of x there is one constraint and one Lagrange multiplier $\lambda(x)$.

Maximizing S , (1), subject to the constraints (2) plus normalization,

$$\delta \{ S + \alpha [\int dx d\theta P(x, \theta) - 1] + \int dx \lambda(x) [\int d\theta P(x, \theta) - \delta(x - x')] \} = 0, \quad (3)$$

² We use the concise notation θ and x to represent one or many unknown variables, $\theta = (\theta_1, \theta_2 \dots)$, and one or multiple experiments, $x = (x_1, x_2 \dots)$.

³ The notion that the likelihood function contains prior information may sound unfamiliar from the point of view of standard Bayesian practice. It should be clear that the likelihood is *prior* information in the sense that its functional form is known *before* the actual data is known, or at least before it can be processed.

yields the joint posterior,

$$P_{\text{new}}(x, \theta) = P_{\text{old}}(x, \theta) \frac{e^{\lambda(x)}}{z}, \quad (4)$$

where z is a normalization constant, and $\lambda(x)$ is determined from (2),

$$\int d\theta P_{\text{old}}(x, \theta) \frac{e^{\lambda(x)}}{z} = P_{\text{old}}(x) \frac{e^{\lambda(x)}}{z} = \delta(x - x'). \quad (5)$$

The final expression for the joint posterior is

$$P_{\text{new}}(x, \theta) = \frac{P_{\text{old}}(x, \theta) \delta(x - x')}{P_{\text{old}}(x)} = \delta(x - x') P_{\text{old}}(\theta|x), \quad (6)$$

and the marginal posterior distribution for θ is

$$P_{\text{new}}(\theta) = \int dx P_{\text{new}}(x, \theta) = P_{\text{old}}(\theta|x'), \quad (7)$$

which is the familiar Bayes' conditionalization rule.

To summarize: $P_{\text{old}}(x, \theta) = P_{\text{old}}(x) P_{\text{old}}(\theta|x)$ is updated to $P_{\text{new}}(x, \theta) = P_{\text{new}}(x) P_{\text{new}}(\theta|x)$ with $P_{\text{new}}(x) = \delta(x - x')$ fixed by the observed data while $P_{\text{new}}(\theta|x) = P_{\text{old}}(\theta|x)$ remains unchanged. We see that in accordance with the minimal updating philosophy that drives the ME method *one only updates those aspects of one's beliefs for which corrective new evidence (in this case, the data) has been supplied.*

SIMULTANEOUS UPDATING WITH MOMENTS AND DATA

Here we generalize the previous section to include additional information about θ in the form of a constraint on the expected value of some function $f(\theta)$,

$$\int dx d\theta P(x, \theta) f(\theta) = \langle f(\theta) \rangle = F. \quad (8)$$

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.

Maximizing the entropy (1) subject to normalization, the data constraint (2), and the moment constraint (8) yields the joint posterior,

$$P_{\text{new}}(x, \theta) = P_{\text{old}}(x, \theta) \frac{e^{\lambda(x) + \beta f(\theta)}}{z}, \quad (9)$$

where z is a normalization constant,

$$z = \int dx d\theta e^{\lambda(x) + \beta f(\theta)} P_{\text{old}}(x, \theta). \quad (10)$$

The Lagrange multipliers $\lambda(x)$ are determined from the data constraint, (2),

$$\frac{e^{\lambda(x)}}{z} = \frac{\delta(x-x')}{ZP_{\text{old}}(x')} \quad \text{where} \quad Z(\beta, x') = \int d\theta e^{\beta f(\theta)} P_{\text{old}}(\theta|x'), \quad (11)$$

so that the joint posterior becomes

$$P_{\text{new}}(x, \theta) = \delta(x-x') P_{\text{old}}(\theta|x') \frac{e^{\beta f(\theta)}}{Z}. \quad (12)$$

The remaining Lagrange multiplier β is determined by imposing that the posterior $P_{\text{new}}(x, \theta)$ satisfy (8). This yields an implicit equation for β ,

$$\frac{\partial \log Z}{\partial \beta} = F. \quad (13)$$

Note that since $Z = Z(\beta, x')$ the resultant β will depend on the observed data x' . Finally, the new marginal distribution for θ is

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

For $\beta = 0$ (no moment constraint) we recover Bayes' rule. For $\beta \neq 0$ Bayes' rule is modified by a ‘‘canonical’’ exponential factor.

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. In what order should they be processed? Or should they be processed at the same time? The answer depends on the nature of the constraints and the question being asked.

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 multiple experiments: for the purpose of inferring θ it is well-known that the order in which the observed data $x' = \{x'_1, x'_2, \dots\}$ is processed does not matter. The proof that ME is completely compatible with Bayes' rule implies that data constraints implemented through δ functions, as in (2), commute. It is useful to see how this comes about.

When an experiment is repeated it is common to refer to the value of x in the first experiment and the value of x in the second experiment. This is a dangerous practice because it obscures the fact that we are actually talking about *two* separate variables. We do not deal with a single x but with a composite $x = (x_1, x_2)$ and the relevant space is $\mathcal{X}_1 \times \mathcal{X}_2 \times \Theta$. After the first experiment yields the value x'_1 , represented by the constraint $c_1 : P(x_1) = \delta(x_1 - x'_1)$, we can perform a second experiment that yields x'_2 and is represented by a second constraint $c_2 : P(x_2) = \delta(x_2 - x'_2)$. These constraints




FIGURE 1. Illustrating the difference between processing two constraints C_1 and C_2 sequentially ($P_{\text{old}} \rightarrow P_1 \rightarrow P_{\text{new}}^{(a)}$) and simultaneously ($P_{\text{old}} \rightarrow P_{\text{new}}^{(b)}$ or $P_{\text{old}} \rightarrow P_1 \rightarrow P_{\text{new}}^{(b)}$).

c_1 and c_2 commute because they refer to *different* variables x_1 and x_2 . An experiment, once performed and its outcome observed, cannot be *un-performed* and its result cannot be *un-observed* by a second experiment. Thus, imposing one constraint does not imply a revision of the other.

In general constraints need not commute and when this is the case the order in which they are processed is critical. For example, suppose the prior is P_{old} and we receive information in the form of a constraint, C_1 . To update we maximize the entropy $S[P, P_{\text{old}}]$ subject to C_1 leading to the posterior P_1 as shown in Figure 1. Next we receive a second piece of information described by the constraint C_2 . At this point we can proceed in essentially two different ways:

(a) Sequential updating. Having processed C_1 , we use P_1 as the current prior and maximize $S[P, P_1]$ subject to the new constraint C_2 . This leads us to the posterior $P_{\text{new}}^{(a)}$.

(b) Simultaneous updating. Use the original prior P_{old} and maximize $S[P, P_{\text{old}}]$ subject to both constraints C_1 and C_2 simultaneously. This leads to the posterior $P_{\text{new}}^{(b)}$.⁴

To decide which path (a) or (b) is appropriate, we must be clear about how the ME method treats constraints. The ME machinery interprets a constraint such as C_1 in a very mechanical way: all distributions satisfying C_1 are in principle allowed and all distributions violating C_1 are ruled out.

Updating to a posterior P_1 consists precisely in revising those aspects of the prior P_{old} that disagree with the new constraint C_1 . However, there is nothing final about the distribution P_1 . It is just the best we can do in our current state of knowledge and we fully expect that future information may require us to revise it further. Indeed, when new information C_2 is received we must reconsider whether the original C_1 remains valid or not. Are *all* distributions satisfying the new C_2 really allowed, even those that violate C_1 ? If this is the case then the new C_2 takes over and we update from P_1 to $P_{\text{new}}^{(a)}$.

⁴ At first sight it might appear that there exists a third possibility of simultaneous updating: (c) use P_1 as the current prior and maximize $S[P, P_1]$ subject to both constraints C_1 and C_2 simultaneously. Fortunately, and this is a valuable check for the consistency of the ME method, it is easy to show that case (c) is equivalent to case (b). Whether we update from P_{old} or from P_1 the selected posterior is $P_{\text{new}}^{(b)}$.

The constraint C_1 may still retain some lingering effect on the posterior $P_{\text{new}}^{(a)}$ through P_1 , but in general C_1 has now become obsolete.

Alternatively, we may decide that the old constraint C_1 retains its validity. The new C_2 is not meant to revise C_1 but to provide an additional refinement of the family of allowed posteriors. In this case the constraint that correctly reflects the new information is not C_2 but the more restrictive $C_1 \wedge C_2$. The two constraints should be processed simultaneously to arrive at the correct posterior $P_{\text{new}}^{(b)}$.

To summarize: sequential updating is appropriate when old constraints become obsolete and are superseded by new information; simultaneous updating is appropriate when old constraints remain valid. The two cases refer to different states of information and therefore *we expect* that they will result in different inferences. These comments are meant to underscore the importance of understanding what information is being processed; failure to do so will lead to errors that do not reflect a shortcoming of the ME method but rather a misapplication of it.

SEQUENTIAL UPDATING: A LOADED DIE EXAMPLE

This is a loaded die example illustrating the appropriateness of sequential updating. The background information is the following: A certain factory makes loaded dice. Unfortunately because of poor quality control, the dice are not identical and it is not known how each die is loaded. It is known, however, that the dice produced by this factory are such that face 2 is on the average twice as likely to come up as face number 5.

The mathematical representation of this situation is as follows. The fact that we deal with dice is modelled in terms of multinomial distributions. The probability that casting a k -sided die n times yields m_i instances for the i^{th} face is

$$P_{\text{old}}(m|\theta) = P_{\text{old}}(m_1 \dots m_k | \theta_1 \dots \theta_k, n) = \frac{n!}{m_1! \dots m_k!} \theta_1^{m_1} \dots \theta_k^{m_k}, \quad (15)$$

where $m = (m_1, \dots, m_k)$ with $\sum_{i=1}^k m_i = n$, and $\theta = (\theta_1, \dots, \theta_k)$ with $\sum_{i=1}^k \theta_i = 1$. The generic problem is to infer the parameters θ on the basis of information about moments of θ and data m' . The additional information about how the dice are loaded is represented by the constraint $\langle \theta_2 \rangle = 2 \langle \theta_5 \rangle$. Note that this piece of information refers to the factory as a whole and not to any individual die. The constraint is of the general form of (8)

$$C_1 : \langle f(\theta) \rangle = F \quad \text{where} \quad f(\theta) = \sum_{i=1}^k f_i \theta_i. \quad (16)$$

For this particular factory $F = 0$, and all $f_i = 0$ except for $f_2 = 1$ and $f_5 = -2$. Now that the background information has been given, here is our first example.

We purchase a die. On the basis of our general knowledge of dice we are led to write down a joint prior

$$P_{\text{old}}(m, \theta) = P_{\text{old}}(\theta) P_{\text{old}}(m|\theta). \quad (17)$$

(The particular form of $P_{\text{old}}(\theta)$ is not important for our current purpose so for the sake of definiteness we can choose it flat.) At this point the only information we have is that

we have a die and it came from a factory described by C_1 . Accordingly, we use ME to update to a new joint distribution. This is shown as P_1 in Figure 1. The relevant entropy is

$$S[P, P_{\text{old}}] = - \sum_m \int d\theta P(x, \theta) \log \frac{P(x, \theta)}{P_{\text{old}}(x, \theta)}, \quad (18)$$

where

$$\sum_m = \sum_{m_1 \dots m_k=0}^n \delta(\sum_{i=1}^k m_i - n) \quad \text{and} \quad \int d\theta = \int d\theta_1 \dots d\theta_k \delta(\sum_{i=1}^k \theta_i - 1),$$

Maximizing S subject to normalization and C_1 gives the P_1 posterior

$$P_1(m, \theta) = \frac{e^{\lambda f(\theta)}}{Z_1} P_{\text{old}}(m, \theta), \quad (19)$$

where the normalization constant Z_1 and the Lagrange multiplier λ are determined from

$$Z_1 = \int d\theta e^{\lambda f(\theta)} P_{\text{old}}(\theta) \quad \text{and} \quad \frac{\partial \log Z_1}{\partial \lambda} = F. \quad (20)$$

The joint distribution $P_1(m, \theta) = P_1(\theta)P_1(m|\theta)$ can be rewritten as

$$P_1(m, \theta) = P_1(\theta)P_{\text{old}}(m|\theta) \quad \text{where} \quad P_1(\theta) = P_{\text{old}}(\theta) \frac{e^{\lambda f(\theta)}}{Z_1}. \quad (21)$$

To find out more about this particular die we toss it n times and obtain data $m' = (m'_1, \dots, m'_k)$ which we represent as a new constraint

$$C_2 : P(m) = \delta(m - m'). \quad (22)$$

Our goal is to infer the θ that apply to our particular die. The original constraint C_1 applies to the whole factory while the new constraint C_2 refers to the actual die of interest and thus takes precedence over C_1 . As $n \rightarrow \infty$ we expect C_1 to become less and less relevant. Therefore the two constraints should be processed sequentially.

Using ME, that is (6), we impose C_2 and update from $P_1(m, \theta)$ to a new joint distribution (shown as $P_{\text{new}}^{(a)}$ in Figure 1)

$$P_{\text{new}}^{(a)}(m, \theta) = \delta(m - m') P_1(\theta|m). \quad (23)$$

Marginalizing over m and using (21) the final posterior for θ is

$$P_{\text{new}}^{(a)}(\theta) = P_1(\theta|m') = P_1(\theta) \frac{P_1(m'|\theta)}{P_1(m')} = \frac{1}{Z_2} e^{\lambda f(\theta)} P_{\text{old}}(\theta) P_{\text{old}}(m'|\theta). \quad (24)$$

where

$$Z_2 = \int d\theta e^{\lambda f(\theta)} P_{\text{old}}(\theta) P_{\text{old}}(m'|\theta). \quad (25)$$

The readers will undoubtedly recognize that (24) is precisely the result obtained by using MaxEnt to obtain a prior, in this case $P_1(\theta)$ given in (21), and then using Bayes' theorem to take the data into account. This familiar result has been derived in some detail for two reasons: first, to reassure the readers that ME does reproduce the standard solutions to standard problems and second, to establish a contrast with the example discussed next.

SIMULTANEOUS UPDATING: A LOADED DIE EXAMPLE

Here is a different problem illustrating the appropriateness of simultaneous updating. The background information is the same as in the previous example. The difference is that the factory now hires a quality control engineer who wants to learn as much as he can about the factory. His initial knowledge is described by the same prior $P_{\text{old}}(m, \theta)$, (17). After some inquiries he is told that the only available information is $C_1 : \langle \theta_2 \rangle = 2 \langle \theta_5 \rangle$. Not satisfied with this limited information he decides to collect data that reflect the production of the whole factory. Randomly chosen dice are tossed n times yielding data $m' = (m'_1, \dots, m'_k)$ which is represented as a constraint,

$$C_2 : P(m) = \delta(m - m') . \quad (26)$$

The apparent resemblance with (22) may be misleading: (22) refers to a single die, while (26) now refers to the whole factory. The goal here is to infer the distribution of θ that describes the overall population of dice produced by the factory. The new constraint C_2 is information in addition to, rather than instead of, the old C_1 : the two constraints should be processed simultaneously. From (12) the joint posterior is ⁵

$$P_{\text{new}}^{(b)}(m, \theta) = \delta(m - m') P_{\text{old}}(\theta | m') \frac{e^{\beta f(\theta)}}{Z} . \quad (27)$$

Marginalizing over m the posterior for θ is

$$P_{\text{new}}^{(b)}(\theta) = P_{\text{old}}(\theta | m') \frac{e^{\beta f(\theta)}}{Z} = \frac{1}{\zeta} e^{\beta f(\theta)} P_{\text{old}}(\theta) P_{\text{old}}(m' | \theta) , \quad (28)$$

where the new normalization constant is

$$\zeta = \int d\theta e^{\beta f(\theta)} P_{\text{old}}(\theta) P_{\text{old}}(m' | \theta) \quad \text{and} \quad \frac{\partial \log \zeta}{\partial \beta} = F . \quad (29)$$

This looks like the sequential case, (24), but there is a crucial difference: $\beta \neq \lambda$ and $\zeta \neq Z_2$. In the sequential updating case, the multiplier λ is chosen so that the intermediate P_1 satisfies C_1 while the posterior $P_{\text{new}}^{(a)}$ only satisfies C_2 . In the simultaneous updating case

⁵ As mentioned in the previous footnote, whether we update from P_{old} or from P_1 we obtain the same posterior $P_{\text{new}}^{(b)}$.

the multiplier β is chosen so that the posterior $P_{\text{new}}^{(b)}$ satisfies both C_1 and C_2 or $C_1 \wedge C_2$. Ultimately, the two distributions $P_{\text{new}}(\theta)$ are different because they refer to different problems: $P_{\text{new}}^{(a)}(\theta)$ refers to a single die, while $P_{\text{new}}^{(b)}(\theta)$ applies to all the dice produced by the factory.⁶

SUMMARY AND FINAL REMARKS

The realization that the ME method incorporates Bayes' rule as a special case has allowed us to go beyond Bayes' rule to process both data and expected value constraints simultaneously. To put it bluntly, anything one can do with Bayes can also be done with ME with the additional ability to include information that was inaccessible to Bayes alone. This raises several questions and we have offered a few answers.

First, it is not uncommon to claim that the non-commutability of constraints represents a *problem* for the ME method. Processing constraints in different orders might lead to different inferences and this is said to be unacceptable. We have argued that, on the contrary, the information conveyed by a particular sequence of constraints is not the same information conveyed by the same constraints in different order. Since different informational states should in general lead to different inferences, the way ME handles non-commuting constraints should not be regarded as a *shortcoming* but rather as a *feature* of the method.

Second, we are capable of processing both data and moments. Is this kind of information of purely academic interest or is it something we might encounter in real life? At this early stage our answer must be tentative: we have given just one example – the die factory – which we think is fairly realistic. However, we feel that other applications (e.g. in econometrics and ecology) can be handled in this way as well.[7, 8]

Finally, is it really true that this type of problem lies beyond the reach of Bayesian methods? After all, we can always interpret an expected value as a sample average in a sufficiently large number of trials. True. We can always construct a large imaginary ensemble of experiments. Entropy methods then become in principle *superfluous*; all we need is probability. The problem with inventing *imaginary* ensembles to do away with entropy in favor of mere probabilities, or to do away with probabilities in favor of more intuitive frequencies, is that the ensembles are just what they are claimed to be, imaginary. They are purely artificial constructions invented for the purpose of handling incomplete information. It seems to us that a safer way to proceed is to handle the available information directly as given (i.e., as expected values) without making additional assumptions about an imagined reality.

Acknowledgements: We would like to acknowledge valuable discussions with C. Cafaro, K. Knuth, and C. Rodríguez.

⁶ For the sake of completeness, we note that, because of the peculiarities of δ functions, had the constraints been processed sequentially but in the opposite order, first the data C_2 , and then the moment C_1 , the resulting posterior would be the same as for simultaneous update to $P_{\text{new}}^{(b)}$.

REFERENCES

1. 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).
2. J. E. Shore and R. W. Johnson, IEEE Trans. Inf. Theory **IT-26**, 26 (1980); IEEE Trans. Inf. Theory **IT-27**, 26 (1981).
3. 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).
4. A. Caticha, "Relative Entropy and Inductive Inference", *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, G. J. Erickson and Y. Zhai (eds.), AIP Conf. Proc. **707**, 75 (2004) (arXiv.org/abs/physics/0311093).
5. 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>).
6. A. Caticha, "Information and Entropy", presented at the *27th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, Saratoga Springs, NY, 2007.
7. A. Giffin, "Updating Probabilities with Data and Moments: an Econometric Example", to be presented at the *3rd Econophysics Colloquium*, Ancona, Italy, 2007.
8. A. Giffin, "Updating Probabilities with Data and Moments: an Ecological Example", to be presented at the *7th International Conference on Complex Systems*, Boston, 2007.

APPENDIX: MORE ON THE MULTINOMIAL PROBLEM

Here we pursue the calculation of the posterior (28) in more detail. To be specific we choose a flat prior, $P_{\text{old}}(\theta) = \text{constant}$. Then, dropping the superscript (b),

$$P_{\text{new}}(\theta) = \frac{1}{\zeta_e} \delta(\sum_i^k \theta_i - 1) \prod_{i=1}^k e^{\beta f_i \theta_i} \theta_i^{m'_i}. \quad (30)$$

where ζ_e differs from ζ in (29) only by a combinatorial coefficient,

$$\zeta_e = \int \delta(\sum_i^k \theta_i - 1) \prod_{i=1}^k d\theta_i e^{\beta f_i \theta_i} \theta_i^{m'_i}, \quad (31)$$

and β is determined from (13) which in terms of ζ_e now reads $\partial \log \zeta_e / \partial \beta = F$. A brute force calculation gives ζ_e as a nested hypergeometric series,

$$\zeta_e = e^{\beta f_k} I_1(I_2(\dots(I_{k-1}))), \quad (32)$$

where each I is written as a sum of Γ functions,

$$I_j = \Gamma(b_j - a_j) \sum_{q_j=0}^{\infty} \frac{\Gamma(a_j + q_j)}{\Gamma(b_j + q_j)} \frac{t_j^{q_j}}{q_j!} I_{j+1} \quad \text{with} \quad I_k = 1. \quad (33)$$

The index j takes all values from 1 to $k-1$ and the other symbols are defined as follows: $t_j = \beta (f_{k-j} - f_k)$, $a_j = m'_{k-j} + 1$, and

$$b_j = n + j + 1 + \sum_{i=0}^{j-1} q_i - \sum_{i=0}^{k-j-1} m'_i, \quad (34)$$

with $q_0 = m'_0 = 0$. The terms that have indices ≤ 0 are equal to zero (i.e. $b_0 = q_0 = 0$, etc.). A few technical details are worth mentioning: First, one can have singular points when $t_j = 0$. In these cases the sum must be evaluated in the limit as $t_j \rightarrow 0$. Second, since a_j and b_j are positive integers the gamma functions involve no singularities. Lastly, the sums converge because $a_j > b_j$. The normalization for the first die example, (25), can be calculated in a similar way. Currently, for small values of k (less than 10) it is feasible to evaluate the nested sums numerically; for larger values of k it is best to evaluate the integral for ζ_e using sampling methods. A more detailed version of the multinomial example is worked out in [7].