July 17, 2012

The Art of Probability Assignment

Vesselin I. Dimitrov

Abstract

The problem of assigning probabilities when little is known is analized in the case where the quanities of interest are physical observables, i.e. can be measured and their values expressed by numbers. It is pointed out that the assignment of probabilities based on observation is a process of inference, involving the use of Bayes' theorem and the choice of a probability prior. When a lot of data is available, the resulting probability are remarkable insensitive to the form of the prior. In the oposite case of scarse data, it is suggested that the probabilities are assigned such that they are the least sensitive to specific variations of the probability prior. In the continuous case this results in a probability assignment rule wich calls for minimizing the Fisher information subject to constraints reflecting all available information. In the discrete case, the corresponding quantity to be minimized turns out to be a Renyi distance between the original and the shifted distribution.

I. Introduction

The problem of probability assignment has been stirring debates and controversy ever since Laplace introduced the notion of indifference as an argument in specifying prior distributions. He thus started a quest for a statistical Holy Grail: prior distributions reflecting ignorance. Today, more than two centuries later, a satisfactory solution to this problem is still elusive. In what follows we offer a physicist's take on the somewhat narrower problem of assigning probabilities for measurable quantities, or, as physicists call them, observables. Strict space limitations will force the exposé to be much more concise than it should have been, but hopefully the main message will be able to come through.

II. Probabilities as opinions: an objective take on subjectivity

When we state "A has a probability p of being true", what we really mean is "We don't know whether A is true or false, yet we *believe* that, if our world existed together with a number of its replicas, A would be true in pN out of N of them when $N \to \infty$". Now, it should be evident that, because of the implied limit procedure, there is no practical way of verifying this statement. One cannot possibly reproduce a given physical situation down to its ever minute details several, let alone infinite, number of times - hence "we believe". Without this leap of faith no rational science would be possible. An example of this sort of belief can be found in Mechanics - we know that material points do not exist, but we *believe* that if they did, they would behave according to the Fist, Second and Third Newton's laws. The source of our faith in this case are countless observations of the behavior of real objects from afar. It is practice that sorts out "good" from "bad" beliefs. Different people, however, have different experiences, so beliefs are subjective and may differ significantly from one person to another. It is, therefore, of significant interest to inquire what is it that makes it possible for rational agents to agree among themselves on what exactly they observe. To that end, let us try to walk in Laplace's "inverse probabilities" footsteps in analyzing how opinions are formed from observations. The following builds on [1].

1. The Anatomy of a Measurement

For our purposes, we shall simplistically call "a measurement" a well-defined procedure to put a real number in correspondence with a physical phenomenon. Usually we have a good idea what the range $\mathcal{R} = [a, b]$ of this number is, but the practicalities of the particular procedure prevent it from being precise. Thus, instead of a real number $\in [a, b]$ the outcome of a single measurement is rather a pointer (index) i to a subinterval $D_i \subset [a, b]$ where $[a, b] = \cup_{i=1}^{n} D_i$

and $D_i \cap D_{j\neq i} = \varnothing$. Repeating[1] the measurement m times we end up with a histogram of n bins where each bin i contains s_i - the number of times the measurement fell in that bin. Obviously, $\sum_{i=1}^{n} s_i = m$. Now, if a result in bin i had an assigned probability p_i in a single measurement, probability theory teaches us that the probability of a set $\{s\}$ is of the multinomial form $P(\mathbf{s}|\mathbf{p}) = \frac{m!}{\Pi_{i=1}^{n} s_i!} p_1^{s_1} p_2^{s_2} \cdots p_n^{s_n}$. We, however, are interested in the opposite situation - the results of the measurements $\{s\}$ are known, and we want to assign probabilities. In this case we recognize $P(\mathbf{s}|\mathbf{p})$ as the likelihood and apply the Bayes theorem to obtain the probability of an assignment $\{p\}$ given the measurements $\{s\}$

$$P(\mathbf{p}|\mathbf{s}) = \mathcal{N}^{-1}\ p_1^{s_1} p_2^{s_2} \cdots p_n^{s_n} \pi(\mathbf{p}) \delta(\Sigma_{i=1}^{n} p_i - 1)$$

where $\mathcal{N} = \int d^n p p_1^{s_1} p_2^{s_2} \cdots p_n^{s_n} \pi(\mathbf{p}) \delta(\Sigma_{i=1}^{n} p_i - 1)$ is a normalization factor. $\pi(\mathbf{p})$ is a probability prior which originates in whatever knowledge we have about the phenomenon in question, the measuring procedure and the structure of the domain's decomposition $[a, b] = \cup_{i=1}^{n} D_i$. For example, one might find it reasonable to assign prior probability proportional to the measure (length) of D_i etc.

2. The Rôle of the Probability Prior

With (1) the most natural way to assign the individual probabilities is as the expectations

$$< p_k >= \mathcal{N}^{-1} \int d^n p p_1^{s_1} p_2^{s_2} \cdots p_k^{s_k+1} \cdots p_n^{s_n} \pi(\mathbf{p}) \delta(\Sigma_{i=1}^{n} p_i - 1)$$

where the integration is over the unit hypercube $p_. \in [0, 1]$. For a uniform prior $\pi(\mathbf{p}) = 1$ the integration [Tertsoff&Bayer] produces $< p_k >= \frac{s_k+1}{m+n} = \frac{1}{1+\frac{n}{m}} \left(f_k + \frac{1}{m}\right)$ where $f_i = s_i/m$ are the "sample frequencies". The variances of this assignment are easily calculated to be $< (\Delta p_k)^2 >= \frac{1}{m+n+1} < p_k > (1- < p_k >)$. For a different prior – uniform on a quadrant of the hypersphere defined by $p_i = \omega_i^2$ – the integrals have been evaluated in [Sykora] as $< p_k >= \frac{1}{1+\frac{n}{2m}}(f_k + \frac{1}{2m})$ and $< (\Delta p_k)^2 >= \frac{1}{m+1+n/2} < p_k > (1- < p_k >)$. For a general prior we use the average value theorem from Analysis to obtain

$$\begin{aligned} < \quad p_k >&= \frac{\pi(\varsigma_k')}{\pi(\boldsymbol{\varsigma})} < p_k >_0 \\ < \quad (\Delta p_k)^2 >&= \frac{\pi(\varsigma_k'')}{\pi(\boldsymbol{\varsigma})} < (\Delta p_k)^2 >_0 + \left[\frac{\pi(\varsigma_k'')}{\pi(\boldsymbol{\varsigma})} - \left(\frac{\pi(\varsigma_k')}{\pi(\boldsymbol{\varsigma})}\right)^2\right] < p_k >^2 \end{aligned}$$

where $\boldsymbol{\varsigma}, \boldsymbol{\varsigma}_k'$ and $\boldsymbol{\varsigma}_k''$ are points in the unit hypercube close to the maxima of $\Pi_{i=1}^{n} p_i^{s_i} \delta(1 - \sum_{i=1}^{n} p_i)$, $p_k \Pi_{i=1}^{n} p_i^{s_i} \delta(1 - \sum_{i=1}^{n} p_i)$ and $p_k^2 \Pi_{i=1}^{n} p_i^{s_i} \delta(1 - \sum_{i=1}^{n} p_i)$,

[1] "Repeating" here is a misnomer - what is meant is an "ensamble" of replicas of the world with one measurement performed in each of its members.

correspondingly, and the zero-subscript quantities are those corresponding to uniform prior. Assuming *abundance of data* (large s_i, and, correspondingly, m) and smooth prior, it can be shown that $\boldsymbol{\varsigma}'_k - \boldsymbol{\varsigma} \simeq \frac{\mathbf{n}_k}{m}$ and $\boldsymbol{\varsigma}''_k - \boldsymbol{\varsigma} \simeq \frac{2\mathbf{n}_k}{m}$ where $(\mathbf{n}_k)_i = \delta_{ki}$. Hence, expanding to the lowest non-trivial order of $1/m$

$$
\begin{aligned}
< \quad p_k > &= < p_k >_0 \left(1 + \frac{1}{m}\frac{\partial_k \pi}{\pi}\right) + O(\frac{1}{m^2}) \\
< \quad (\Delta p_k)^2 > &= < (\Delta p_k)^2 >_0 \left(1 + \frac{2}{m}\frac{\partial_k \pi}{\pi}\right) + \frac{< p_k >_0^2}{m^2}\left(\frac{3}{2}\frac{\partial_k^2 \pi}{\pi} - (\frac{\partial_k \pi}{\pi})^2\right) + O(\frac{1}{m^3})
\end{aligned}
$$

Thus, we recognize that the arbitrariness of the probability prior induces *multiplicative noise* in the assigned probabilities, and affects their variances both by rescaling and shifting. It is also worthwhile noticing that the only instance of assigning zero probability would be due to the choice of the prior; measurements alone, no matter how numerous, cannot force us to assign strictly vanishing probabilities.

In the other extreme - no ($m = 0$) data available - the probability assignment derives through (2) strictly from the prior:

$$< p_k > = \mathcal{N}^{-1} \int d^n p \; p_k \pi(\mathbf{p}) \delta(\Sigma_{i=1}^n p_i - 1) = \overline{p_k}$$

For one preformed measurement ($m = 1$) that produced a result in bin i

$$< p_k > = \mathcal{N}^{-1} \int d^n p \; p_i p_k \pi(\mathbf{p}) \delta(\Sigma_{i=1}^n p_i - 1) = \frac{\overline{p_k p_i}}{\overline{p_i}}$$

and analogously for higher values of m. Probabilities are most useful when little or no data is available, and it is seen that such "ignorance" probability assignments for measurable quantities are, not surprisingly, entirely determined by the choice of the prior $\pi(p)$.

An interesting result is obtained when we go to the continuum limit $n \to \infty$. With $p_k = \int_{D_k} dx p(x) = \int_{x_k}^{x_k + \Delta x_k} dx p(x) = \Delta x_k p(x_k) + \frac{1}{2}\Delta x_k^2 p'(x_k) + \cdots$, the usual identification $p_k = \Delta x_k p(x_k)$ for $\Delta x_k \to 0$ only makes sense when the probability density $p(x)$ is everywhere differentiable in $[a, b]$. In order to avoid handling ugly continual integrals, we perform the $n \to \infty$ limit at the stage where, with $\mu_k \equiv \frac{\pi(\varsigma'_k)}{\pi(\boldsymbol{\varsigma})}$ and $\sigma(x_k) \equiv \lim_{n\to\infty} \frac{1}{n\Delta x_k}$,

$$< p(x_k) > = \mu(x_k) \lim_{n\to\infty} \frac{< p_k >_0}{\Delta x_k} = \mu(x_k) \lim_{n\to\infty} \left[\frac{m}{m+n} f(x_k) + \frac{n}{m+n}\sigma(x_k)\right]$$

We observe that, for any finite amount of data ($m < \infty$) the assigned probability density $< p(x) > = \mu(x)\sigma(x)$ depends on the metrics σ and the prior but not on the data, while for $m = \infty$ the result depends on the order in which the limits are taken. Only for $m \to \infty$ before $n \to \infty$ is the result proportional to the "sample frequency" density $f(x)$.

To summarize, in order to relate probabilities (opinions) to the real world (sample frequencies), we need the help of the Bayes theorem where a probability prior enters the game. Hence, even when a lot of data is available, the probability assignments are not unambiguous - the arbitrariness of the prior manifests itself as a multiplicative noise in the probabilities and in their variances. When little or no data is available the assignments derive directly from the chosen prior. Let us also emphasize an important lesson from the above: the widely held opinion that a probability distribution represents a "state of knowledge" is wrong. It is rather the sample frequencies, coming from observations, which constitute "knowledge". Probabilities are necessarily inferred, and thus represent only a "state of belief". The importance of this subtle distinction will become apparent in what follows.

III. Assigning Probabilities

The most intellectually appealing scheme for assigning probabilities, in our opinion, was put forward by Jaynes around the middle of the last century, under the name "Maximum Entropy" (MaxEnt) principle. It is very difficult for a rational person to argue with its general formulation, which simply calls for inferential coherence by prescribing the assignment of the *least committed* probability distribution consistent with *all available* information. However, opinions rapidly start to diverge when it comes to specifying how exactly the "least committed" distribution is defined and what exactly constitutes "available information". On the first point, Jaynes itself maintained that the "least committed" distribution is the one with maximal Shannon entropy. His, and many others, affinity to Shannon's entropy was based on a number of appealing properties it possesses. During the years a tremendous amount of effort was invested into trying to prove that it is "the one and only" reasonable criterion to use. Eventually, however, two things were, or should have been, understood: 1) The Shannon's entropy is but a particular instance of a larger class of equally reasonable Renyi's entropies; and 2) The use of Jaynes procedure as a probability *assignment* rule is untenable, so it gradually evolved into probability *updating* rule - leaving us where we started, with the necessity of assigning an ignorance prior. On the second point, the available information is most often presented as a number of prescribed expectation values. Jaynes himself was aware of the conflict between the expectations being characteristics of probability distributions, and as such, essentially opinions, and actual information obtained by measurements, but he took the position that the available information entered in the form of constraint(s) on the probability distribution even if "It might ... be only the guess of an idiot" [2]. Before we embark on the ambitious task of trying to clarify these points, let us briefly address the question of "once assigned, how can probabilities be used?".

3. What Use are Opinions?

Probabilities being subjective, it is not immediately obvious how practical use can be made of them. In statistical sense, probabilities are the best estimators of sample frequencies, and this is about the only guiding principle for their use. Hence, it appears that plugging probabilities in place of sample frequencies in various statistical estimators would allow us to infer *predictions* about the results of measurements not yet performed. Such statistical estimators are the Kolmogorov-Nagumo averages [3], defined as $< A >_\phi = \phi^{-1}\left(\sum p_i \phi(A_i)\right)$ where $\phi(x)$ is continuous and strictly monotonic function, A is an observable, and A_i is the value of A corresponding to bin i. Different functions ϕ in general produce different values of $< A >_\phi$. When measuring physical observables, we can use rulers in different units and origin of the scale. Without an appropriate behaviour of the predictions for the results of measurements upon rescaling and shifts they would be useless. Therefore, an important requirement to be imposed on an useful estimator is that $< \alpha A + \beta >_\phi = \alpha < A >_\phi + \beta$, where α and β are arbitrary constants. It is an elementary exercise to show that this forces $\phi(x) = x$ and thus singles out $< A >= \sum_i p_i A_i$ as the rule for predicting the result of a measurement of A given the probabilities $\{p\}$[2]. The result of an actual measurement will most likely differ from the prediction, yet this is still the best we can do with a probability assignment $\{p\}$.

4. The Constraint Rule

Let us first try to make the MaxEnt principle formulation more explicit in its "using all avilable information" part. The physical problem under consideration can be envisioned as the one of studying a set of observables of a system, which we will refer to as "the primary observables". This could be, e.g. the three coordinates $\mathbf{x}$ of a material point etc. We seek to assign a probability distribution $p(\mathbf{x})$ for these observables, which would allow us to a) Predict the results of future measurements of these observables as their expectations $< \mathbf{x} >= \int dx \mathbf{x} p(\mathbf{x})$, which is of primary interest, and b) Predicting the result of future measurements of any additional obervable $Q(\mathbf{x})$ as $< Q >= \int dx Q(\mathbf{x}) p(\mathbf{x})$, which is of secondary interest. In doing this, we are generally ignorant, except possibly for the results $\{a\}$ of previous measurements of some m observables $A_r, r = 1, 2, \cdots, m$. Then the constraint rule of the MaxEnt principle can be regarded as a requirement that the asigned probability distribution correctly "predicts" the results of the already performed measurements as $a_r = \int dx A_r(\mathbf{x}) p(\mathbf{x}), r = 1, 2, \cdots, m$. In other words, the constraint rule simply forces the probability assignment, which is to be used to predict the results of future measurements, to be consistent with the results of measurements already performed. Let us stress that what is involved here are single measurements and their results $\{a\}$, and not multiple measurements from which

[2] One might be tempted to argue in favor of the most probable value instead, but one immediate indication that this is not a good rule is that it cannot produce any prediction for uniform probabilities.

the $a-s$ are obtained as sample averages, as is too often implied in the context of the MaxEnt. Indeed, if the results of, say, 10 measurements of, e.g., A_1 were known as $a_1(i), i = 1, 2, \cdots, 10$ and a_1 was taken as $\overline{a_1} = \frac{1}{10}\sum_{i=1}^{10} a_i(10)$ to be used as a constraint, this would be in a blatant violation of the "using all available information" principle, since the set of measured values of A_1 clearly contains information also about a_1's variance: $\overline{\Delta a^2} = \frac{1}{9}\sum_{i=1}^{10}[a_1(i) - \overline{a_1}]^2$ and, similarly, for its higher moments as well.

5. The Expectation as (sort of) a Parameter

Before we embark on the problem of assigning probabilities, we need to shortly discuss the parameterization of our probability distributions in terms of the expectations of their primary observables. For simplicity we will assume one primary parameter x, the case with multiple such parameters being a straight-forward generalization. In fact, we don't need to consider a full-fledged parameterization in which the value of the parameter is equal to the expectation of x, but just one that would allow us to *independently* vary the expectation of x. Thus, we are interested in a parametrization $p(x; x_e)$ such that, for any $|\varepsilon| << 1$, we have $\int dx x p(x; x_e + \varepsilon) =< x > +\varepsilon + O(\varepsilon^3)$ while the normalization of the probability distribution as well as all other cumulants $C_n(x)$ of x are preserved

$$\frac{\partial}{\partial x_e}\int dx p(x; x_e) = \frac{\partial}{\partial x_e}\int dx C_n(x) p(x; x_e) = 0 \qquad n = 2, 3, \cdots$$

We formulate the following **Conjecture**[3]: A parameterization with the above properties is only possible if the probability distribution fulfills certain conditions at the border of its domain, and in this case it is given by $p(x; x_e) = p(x + x_e)$. Obviously, with such a parameterization we always have $\frac{\partial p(x;x_e)}{\partial x_e} = \frac{\partial p(x;x_e)}{\partial x}$, which is the property we are mainly interested in. Establishing this, we finally can address the "most uncommitted" element of the general MaxEnt principle.

6. Assigning Robust Probabilities

We have shown above that probability assignments based on observations have inherent indeterminacy due to the arbitrariness of the probability prior. Therefore, a natural question to ask is whether an assignment exsists that is, in some sense, robust against variations of the prior. As already demonstrated, the latter cause multiplicative noise in the probabilities. Hence we try to formulate a robustness requirement in tems of a probability distance of the Ali-Silvey type $D(p; p + \delta p) \underset{\{p\}}{\rightarrow} \min$, where δp is the probability noise. As well known, for normalization-preserving $\delta p(\mathbf{x})$

$$D(p; p + \delta p) = \frac{\alpha}{2}\int dx p(\mathbf{x}) \left(\frac{\delta p(\mathbf{x})}{p(\mathbf{x})}\right)^2 + O(\delta p^3)$$

[3] The space restrictions do not allow us to formulate this as a theorem here.

where the constant coefficient $\alpha\tilde{}1$ depends on the particular distance used. With a general multiplicative $\delta p(\mathbf{x}) = \varepsilon(\mathbf{x})p(\mathbf{x})$ the norm-preserving variation of this with respect to $p(\mathbf{x})$ does not produce a solution. Hence, for the most general probability noise our robustness requirement is not selective enough to single out a particular distribution. However, upon some reflection, we realize that not all possible perturbations in the distribution are of equal importance: we are mainly interested in the robustness of the probabilities with regard to the perturbations which would have maximal effect on the primary observables, that is, choose the multiplicative noise such that $\delta p(\mathbf{x}) = \boldsymbol{\varepsilon}(\mathbf{x})p(\mathbf{x}) = \boldsymbol{\varepsilon}'(\mathbf{x}) \cdot \frac{\partial p(\mathbf{x};\mathbf{x}_e)}{\partial \mathbf{x}_e}$. With this noise

$$\begin{aligned} D(p; p+\delta p) &= \frac{\alpha}{2} \int dx p^{-1}(\mathbf{x};\mathbf{x_e}) \boldsymbol{\varepsilon}'(\mathbf{x}) \cdot \frac{\partial p(\mathbf{x};\mathbf{x}_e)}{\partial \mathbf{x}_e} \frac{\partial p(\mathbf{x};\mathbf{x}_e)}{\partial \mathbf{x}_e} \cdot \boldsymbol{\varepsilon}'(\mathbf{x}) + O(\varepsilon'^3) \leq \\ &\leq \frac{\alpha}{2} \int dx \boldsymbol{\varepsilon}'^2(\mathbf{x}) \mathtt{Tr} I_F(\mathbf{x}_e) \end{aligned}$$

where $I_F(\mathbf{x}_e) = \int dx p^{-1}(\mathbf{x};\mathbf{x_e}) \frac{\partial p(\mathbf{x};\mathbf{x}_e)}{\partial \mathbf{x}_e} \frac{\partial p(\mathbf{x};\mathbf{x}_e)}{\partial \mathbf{x}_e}$ is the Fisher information matrix with respect to $\mathbf{x}_e$ and the inequality follows from its postivedefinitness. Hence, for an arbitrary noise factor $\boldsymbol{\varepsilon}'(\mathbf{x})$ the tightest bound on the distance results from the distribution with minimal trace of $I_F(\mathbf{x}_e)$. Using the interchangeability of the derivatives derived above, we arrive at the final form of the robustness condition where $\mathbf{x}_e$ does not play a rôle any more and is therefore dropped

$$\int dx p^{-1}(\mathbf{x}) \left(\frac{\partial p(\mathbf{x})}{\partial \mathbf{x}} \right)^2 \to \min \quad \cap \quad \int dx p(\mathbf{x}) = 1$$

When results of measurements of some observables are known, the above minimization is constrained such that the resulting probabilities reproduce these observables. Is there any sense in which the so characterized distribution could be considered "the least committed"? The Kramer-Rao result for the most efficient estimator of $\mathbf{x}_e$ in the form $\mathtt{Tr}[cov^{-1}(\mathbf{x}_e)] = \mathtt{Tr} I_F(\mathbf{x}_e)$ indicates that in the situation where $p(\mathbf{x})$ is the one with minimal trace of $I_F(\mathbf{x}_e)$ an invariant measure of the magnitude of the primary observables' covariance is maximal. This can be formulated as "The distribution with minimal trace of the Fisher information is the one for which the *most efficient* estimator of the primary observables (whether it actually exists or not) has the *worst possible* performance". Thus the extremal property of $p(\mathbf{x})$ can indeed be interpreted as the distribution being "the least committed" with regard to the primary observables.

IV. Discussion

Previously [4] we have derived the same condition (in the one-dimensional case) for assigning uniformative probabilities from the requirement that they be the least sensitive to coarse-graining. The rational for this requirement was that coarse-graining decreased the "information content" – if such a thing could be meaningfully defined – and the distribution with minimal information content

to start with would be the one least affected by it. The approach is, in a sense, complementary to Bernardo's reference priors, where information is gained and the effect of this gain - maximized. However, in contrast to Bernardo's, our result does not depend on which particular distance is used to measure the sensitivity of the probabilities. That the same assignment rule would result from the present, quite different, considerations may bear some yet unidentified significance. Fisher information-like constructs appear almost universally in physics [5] and one cannot help but wonder to what extent physics laws could be explained as information processing rules, and answer Toffoli's question "Where does Nature shop for its Lagrangians".

Of significant interest is also whether/how the same considerations apply to probabilities on discrete domains. In physics, discrete domains are usually obtained by coarse-graining of continuous ones, and thus are "loaded" with properties inherited from the topology and the metrics of the original continuum. Such remnants could be, for example, various nearest, second nearest etc. neighbour hierarchies. Choosing the simplest case of a coarse-grained segment of the real line, the disrete domain is $\{1, 2, ..., n\}$ and the relevant observable is $< i >= \mathtt{nint}(\sum_{i=1}^{n} ip_i)$, where the function returns the nearest integer to its argument. It can be conjectured as in the continuous case that the only possible way to perturb the probabilities while best preserving their normalization and higher cumulants of i, again subject to certain conditions on p_1 and p_n, is equivalent to successive application of $p'_i = p_i - \varepsilon p_i$, $p'_{i+1} = p_{i+1} + \varepsilon p_i$, where the multiplicativity of the noise is explicitely used. Then $D(p; p') = \frac{\alpha}{2}\varepsilon^2 \sum_{i=1}^{n} p_i^2 \left(\frac{1}{p_i} + \frac{1}{p_{i+1}}\right) + O(\varepsilon^3) = \frac{\alpha}{2}\varepsilon^2 \left(1 + \sum_{i=1}^{n} p_i \frac{p_i}{p_{i+1}}\right) + O(\varepsilon^3)$. The maximal robustness with respect to $< i >$ is achieved for a distribution for which the distance is minimal, hence

$$\sum_{i=1}^{n} p_i \frac{p_i}{p_{i+1}} \to \min \quad \cap \quad \sum_{i=1}^{n} p_i = 1$$

Here the role of the Fisher information is played by the Renyi's distance of order -1.

Acknowledgment This work was partially supported by the DTRA grant HDTRA1-11-1-0035.

References

1. S. Sykora, *J. Stat. Phys.* **11,** #1 (1974) p.17
2. E.T.Jaynes, MaxEnt Formalism Conference, MIT, May 2-4, 1978
3. A. Kolmogorov, *Atti della R. Accademia Nazionale dei Lincei* **12**, (1930) p.388; M. Nagumo, *Japan. Journ. Math.* **7**, (1930) p,71
4. V.I.Dimitrov, *AIP Conference Proceedings* Vol. **954** Issue 1, (2007) p.143

5. B. Roy Frieden, "*Physics from Fisher Information*", Cambridge Univ. Press, 1999