



HAL
open science

An Unforeseen Equivalence Between Uncertainty and Entropy

Tim Muller

► **To cite this version:**

Tim Muller. An Unforeseen Equivalence Between Uncertainty and Entropy. 13th IFIP International Conference on Trust Management (IFIPTM), Jul 2019, Copenhagen, Denmark. pp.57-72, 10.1007/978-3-030-33716-2_5. hal-03182607

HAL Id: hal-03182607

<https://inria.hal.science/hal-03182607>

Submitted on 26 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

An Unforeseen Equivalence between Uncertainty and Entropy

Tim Muller¹

University of Nottingham tim.muller@nottingham.ac.uk

Abstract. Uncertainty and entropy are related concepts, so we would expect there to be some overlap, but the equality that is shown in this paper is unexpected. In Beta models, interactions between agents are evidence used to construct Beta distributions. In models based on the Beta Model, such as Subjective Logic, uncertainty is defined to be inversely proportional to evidence. Entropy measures how much information is lacking in a distribution. Uncertainty was neither intended nor expected to be an entropy measure. We discover that a specific entropy measure we call *EDRB* coincides with uncertainty whenever uncertainty is defined. *EDRB* is the expected Kullback-Leibler divergence between two Bernoulli trials with parameters randomly selected from the distribution. *EDRB* allows us to apply the notion of uncertainty to other distributions that may occur in the context of Beta models.

Keywords: Uncertainty · Entropy · Information Theory · Beta model · Subjective Logic

1 Introduction

The Beta model paradigm is a powerful formal approach to studying trust. Bayesian logic is at the core of the Beta model: “agents with high integrity behave honestly” becomes “honest behaviour evidences high integrity”. Its simplest incarnation is to apply Beta distributions naively, and this approach has limited success. However, more powerful and sophisticated approaches are widespread (e.g. [13,3,17]). A commonality among many approaches, is that more evidence (in the form of observing instances of behaviour) yields more certainty of an opinion. Uncertainty is inversely proportional to the amount of evidence.

Evidence is often used in machine learning. It is no surprise that there is a close link between trust models and machine learning, since the goal is to automatically create a model, based on observed data. The Beta model is based a simple Bayesian technique found in machine learning. More involved techniques may introduce hidden variables [13] or hidden Markov models [3,18]. Uncertainty as the inverse of (or lack of) evidence makes sense in this context.

We have obtained successful results applying information theory to analyse trust ratings [15,16]. Informative ratings are more useful than uninformative ones. Others have applied information theory to trust modelling in different ways, e.g. [1,2]. However, these approaches contrast the evidence-based approaches –

they were not considered to be equivalent approaches. In fact, we have studied the possibility of combining uncertainty and entropy, to understand their interplay, in [12] – and we had not expected that they would turn out to coincide.

The purpose of this paper is to demonstrate a surprising equivalence. The uncertainty used in this paper is fundamentally different from entropy in information theory. There are various entropy measures that one can define, but the standard measures do not yield an equivalence to uncertainty. However, we formulate a specific entropy measure – that we call expected Kullback-Leibler divergence of random-parameter Bernoulli trials (EDRB) – which does equate to uncertainty. The proof is based on a specific properties of functions related to Beta distributions, and does not seem provide insight in why the two are equivalent.

The main motivation for this paper, is to present this surprising result. However, there are possible practical applications too. First, EDRB allows us to compute uncertainty of a given Beta distribution with unknown parameters. Secondly, EDRB can provide the uncertainty of other distributions than the Beta distribution, generalising uncertainty. Thirdly, using EDRB, we can apply techniques from information theory on uncertainty (e.g. apply MAXENT on uncertainty).

The paper is organised as follows: In Section 2, we introduce and shortly discuss existing definitions and properties. In Section 3, we discuss the general relation between uncertainty and entropy in the setting of the Beta model. In Section 4, we present our main result, Theorem 1. Finally, in Section 5, we look at the application of Theorem 1 on more general opinions.

2 Preliminaries

In this section, we introduce the existing definitions and formalisms that are relevant to our work. The definitions can be grouped into two types, definitions surrounding the Beta model and related models (Section 2.1), and information-theoretic definitions (Section 2.2).

2.1 Beta models

The Beta models are a paradigm, and whether a specific model is a Beta model is up to debate. The core idea behind Beta models is a specific Bayesian approach to evidence [4]. Interactions with agents form evidence, and they are used to construct an opinion. The interactions correspond to Bernoulli trials [5]:

Definition 1. *A Bernoulli trial is has two outcomes, “success” and “failure”, and the probability of success is the same every time the trial is performed. A Bernoulli distribution is a discrete distribution with two outcomes, 0 and 1. Its probability mass function $f_B(p)$ has $f_B(0;p) = 1 - p$ and $f_B(1;p) = p$. A random variable B_i from a Bernoulli trial is distributed according to the Bernoulli distributions, so $P(B_i=1) = p$ and $P(B_i=0) = 1 - p$.*

There are agents $A \in \mathcal{A}$. Each agent A has an unknown parameter x_A , called its *integrity*. An agent may betray another agent, or the agent may cooperate. Which choice an agent makes is assumed to be a Bernoulli trial, where the probability of cooperating is equal to its integrity. A series of interactions, therefore, is a series of Bernoulli trials. Let $B_{A,i}$ be the random variable corresponding to the i^{th} interaction with agent A , then $P(B_{A,i} = 1) = x_A$. We refer to outcome 1 as *success* and 0 as *failure*. However, x_A is not a known quantity, so we apply the Bayesian idea of introducing a random variable X_A for the integrity of agent A . An *opinion* about an agent can be denoted as the probability density function $p_{X_A}(x_A|B_{A,1}, B_{A,2}, \dots)$.

We assume that the opinion without evidence is the uniform distribution – so $p_{X_A}(x_A) = 1$. One reason to select this prior distribution, is the principle of maximum entropy, which essentially dictates that we should pick the distribution with the highest entropy, if we want to model that we do not have any evidence – and this distribution is the uniform distribution. Another reason to select this prior distribution, is that it simplifies the notion of combining opinions. Most importantly, the prior can be changed to any arbitrary probability density function f , simply by multiplying $f(x_A) \cdot p_{X_A}(x_A|B_{A,1}, B_{A,2}, \dots) \cdot NF$.

The reason for the name “Beta model” comes from a special relationship to Beta distributions. The Beta distribution is defined as:[5,8]

Definition 2. *The Beta distribution is a continuous distribution with support in the range $[0, 1]$, with a probability density function $f_\beta(x; \alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}$, where B is the Beta function, $B(\alpha, \beta) = \int_0^1 x^{\alpha-1}(1-x)^{\beta-1} dx$, which acts as a normalisation factor. Its cumulative distribution function is $\frac{\int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt}{B(\alpha, \beta)}$, which is also known as the regularised incomplete Beta function $I_x(\alpha, \beta)$.*

We are using important properties of the Beta function and the regularised incomplete Beta function (see [8]):

Proposition 1. *The following two equalities hold:*

$$\frac{B(\alpha + 1, \beta)}{B(\alpha, \beta)} = \frac{\alpha}{\alpha + \beta} \quad \text{and} \quad I_x(\alpha + 1, \beta) = I_x(\alpha, \beta) - \frac{x^\alpha(1-x)^\beta}{\alpha B(\alpha, \beta)}$$

Given the relations between the random variables, we find that any opinion $p_{X_A}(x_A|B_{A,1}, B_{A,2}, \dots)$ is a Beta distribution. In fact, if the outcomes of the Bernoulli trials $B_{A,1}, \dots, B_{A,n}$ contain n_s success and $n - n_s = n_f$ failures, then the opinion $p_{X_A}(x_A|B_{A,1}, \dots, B_{A,n}) = f_\beta(x_A; n_s + 1, n_f + 1)$ [10].

Proposition 2. *Let $b_{A,1}, \dots, b_{A,n}$ be a list with all elements 0 or 1, $\sum_{1 \leq i \leq n} b_{A,i} = n_s$ and $n_f = n - n_s$. $p_{X_A}(x_A|B_{A,1}=b_{A,1}, \dots, B_{A,n}=b_{A,n}) = f_\beta(x_A; n_s + 1, n_f + 1)$*

We can define a fusion operator \oplus , as $p_1(x) \oplus p_2(x) = \frac{p_1(x)p_2(x)}{\int_0^1 p_1(y)p_2(y) dy}$. The fusion operator simply merges the evidence [5]:

Proposition 3. *For any series of outcomes of Bernoulli trials $b_{A,1}, \dots, b_{A,n}$ and $b'_{A,1}, \dots, b'_{A,n'}$, $p_{X_A}(x_A|B_{A,1}=b_{A,1}, \dots, B_{A,n}=b_{A,n}, B'_{A,1}=b'_{A,1}, \dots, B'_{A,n'}=b'_{A,n'}) \propto p_{X_A}(x_A|B_{A,1}=b_{A,1}, \dots, B_{A,n}=b_{A,n}) \cdot p_{X_A}(x_A|B'_{A,1}=b'_{A,1}, \dots, B'_{A,n'}=b'_{A,n'})$.*

In particular $f_\beta(x; \alpha, \beta) \oplus f_\beta(x; \alpha', \beta') = f_\beta(x; \alpha + \alpha' - 1, \beta + \beta' - 1)$.

Using a distribution to denote an opinion is a feasible approach, based on Bayesian logic, but the results are not intuitively obvious to people that may use the opinions. Subjective Logic is a formalism within the Beta model paradigm, which is developed with the purpose of being understandable to non-experts [7]¹. A Subjective Logic opinion is defined[6]:

Definition 3. *An opinion is a triple of components (b, d, u) , for positive real b, d, u with $b + d + u = 1$. The first component is belief, the second is disbelief, and the third is uncertainty.*

Subjective logic also has a fusion operator, denoted $(b, d, u) \oplus (b', d', u')$. The purpose of fusion in Subjective Logic is the same as fusion of distributions, namely to merge evidence. See [6].

Definition 4.

$$(b, d, u) \oplus (b', d', u') = \left(\frac{bu' + b'u}{u + u' - uu'}, \frac{du' + d'u}{u + u' - uu'}, \frac{uu'}{u + u' - uu'} \right).$$

That there is an isomorphism between fusion of Beta distributions and Subjective Logic fusion, is a known result [7]. In fact, this isomorphism is the primary argument in favour of the shape of Definition 4. It turns out that there is a family of isomorphisms between the two:

Proposition 4. *Let \mathbb{B}, \mathbb{S} be the groups of Beta distributions with fusion, and of SL opinions with SL fusion. Let f_r be a function $f_r : \mathbb{B} \rightarrow \mathbb{S}$ with $f_r(\alpha, \beta) = \left(\frac{\alpha-1}{\alpha+\beta-2+r}, \frac{\beta-1}{\alpha+\beta-2+r}, \frac{r}{\alpha+\beta-2+r} \right)$. For $r > 0$, f_r is an isomorphism between \mathbb{B} and \mathbb{S} .*

Proof. Keep in mind Proposition 3, so fusion simply adds α 's and β 's. The inverse of f_r is $f_r^{-1} = \left(\frac{br}{u} + 1, \frac{dr}{u} + 1 \right)$, since (w.l.o.g. for α):

$$f_r^{-1}(f_r(\alpha, \beta)) = \frac{\alpha - 1}{\alpha + \beta - 2 + r} \frac{\alpha + \beta - 2 + r}{r} + 1, \dots = r \frac{\alpha - 1}{r} + 1, \dots = \alpha, \beta.$$

Remains to prove that f_r and f_r^{-1} are homomorphisms between \mathbb{B} and \mathbb{S} :

$$\begin{aligned} & f_r^{-1} \left(\frac{bu' + b'u}{u + u' - uu'}, \frac{du' + d'u}{u + u' - uu'}, \frac{uu'}{u + u' - uu'} \right) \\ &= \left(r \frac{bu' + b'u}{uu'} + 1, r \frac{du' + d'u}{uu'} + 1 \right) \\ &= \left(r \frac{b}{u} + r \frac{b'}{u'} + 1, r \frac{d}{u} + r \frac{d'}{u'} + 1 \right) \\ &= \left(r \frac{b}{u} + 1, r \frac{d}{u} + 1 \right) \oplus \left(r \frac{b'}{u'} + 1, r \frac{d'}{u'} + 1 \right) \\ &= f_r^{-1}(b, d, u) \oplus f_r^{-1}(b', d', u') \end{aligned}$$

¹ For the purpose of this paper, we restrict ourselves to so-called *binary* Subjective Logic opinions, which corresponds to the notion that the integrity parameter determines a Bernoulli trial.

Let $D = \alpha + \beta - 2 + r$ and $D' = \alpha' + \beta' - 2 + r$. Then:

$$\begin{aligned}
 & f_r(\alpha + \alpha' - 1, \beta + \beta' - 1) \\
 = & \left(\frac{\alpha + \alpha' - 2}{\alpha + \alpha' + \beta + \beta' - 4 + r}, \frac{\beta + \beta' - 2}{\alpha + \alpha' + \beta + \beta' - 4 + r}, \frac{r}{\alpha + \alpha' + \beta + \beta' - 4 + r} \right) \\
 = & \left(\frac{\alpha + \alpha' - 2}{D + D' - r}, \frac{\beta + \beta' - 2}{D + D' - r}, \frac{r}{D + D' - r} \right) \\
 = & \left(\frac{\frac{\alpha-1}{D} \frac{r}{D'} + \frac{\alpha'-1}{D'} \frac{r}{D}}{\frac{r}{D} + \frac{r}{D'} - \frac{r^2}{DD'}}, \frac{\frac{\beta-1}{D} \frac{r}{D'} + \frac{\beta'-1}{D'} \frac{r}{D}}{\frac{r}{D} + \frac{r}{D'} - \frac{r^2}{DD'}}, \frac{\frac{r^2}{DD'}}{\frac{r}{D} + \frac{r}{D'} - \frac{r^2}{DD'}} \right) \\
 = & \left(\frac{\alpha-1}{D}, \frac{\beta-1}{D}, \frac{r}{D} \right) \oplus \left(\frac{\alpha'-1}{D'}, \frac{\beta'-1}{D'}, \frac{r}{D'} \right) \\
 = & f_r(\alpha, \beta) \oplus f_r(\alpha', \beta')
 \end{aligned}$$

Since Beta distributions and Subjective Logic are isomorphic w.r.t. fusion, we can apply notions of Subjective Logic directly to Beta distributions. So we can say that the uncertainty of $f_\beta(x; \alpha, \beta)$ is $\text{unc}_r(f_\beta(x; \alpha, \beta)) = \frac{r}{\alpha + \beta - 2 + r}$. Unless we explicitly state which isomorphism f_r we use, we assume that f_1 was used – so $\text{unc} = \text{unc}_1$. Observe that a Beta distribution based on $n = n_s + n_f$ pieces of evidence has uncertainty $\text{unc}(f_\beta(x; n_s + 1, n_f + 1)) = \frac{1}{n_s + n_f + 1}$, so the inverse of uncertainty is equal to the amount of evidence (plus 1, to avoid divide-by-zero).

2.2 Information Theory

A core notion in information theory, is the notion of *surprisal*, also known as self-information or information content. The symbol I_X is often used, but it is also used for the regularised incomplete Beta function, so we denote the surprisal of X with J_X instead. The surprisal is defined: $J_X(x) = -\log(P(X = x))$ or $J_X(x) = -\log(p_X(x))$ for discrete and continuous random variable X , respectively.

Shannon entropy is used to measure the expected amount of information carried in a random variable, which is determined by the uncertainty of the random variable [9]:

Definition 5. *The Shannon entropy of a discrete random variable X is given:*

$$H(X) = \mathbb{E}_x(J_X(x)) = - \sum_{x_i \in X} P(X=x_i) \cdot \log(P(X=x_i))$$

The Shannon entropy is maximal when all possible outcomes are equiprobable. It means that our expected surprisal is maximal, which is a common way to express we know nothing about the random variable.

Shannon entropy can be generalised for continuous random variables, to differential entropy. Differential entropy does not provide absolute values – values can go below 0 – but is useful for measuring the difference in information present in distributions.

Definition 6. *The differential entropy of a continuous random variable X is given:*

$$H(X) = \mathbb{E}_x(J_X(x)) = - \int_X p_X(x) \cdot \log(p_X(x)) \, dx$$

Kullback-Leibler divergence, also known as relative entropy, measures the distance from one distribution to another.

Definition 7. *For discrete random variables X, Y , the Kullback-Leibler divergence from X to Y is:*

$$D_{\text{KL}}(X||Y) = \mathbb{E}_x(J_X(x) - J_Y(x)) = \sum_{x_i \in X} P(X=x_i) \cdot \log\left(\frac{P(X=x_i)}{P(Y=x_i)}\right)$$

For continuous random variables X, Y , it is:

$$D_{\text{KL}}(X||Y) = \mathbb{E}_x(J_X(x) - J_Y(x)) = \int_X p_X(x) \cdot \log\left(\frac{p_X(x)}{p_Y(x)}\right) \, dx$$

Note that in general, $D_{\text{KL}}(X||Y) \neq D_{\text{KL}}(Y||X)$. Typically, X is the “true” random variable and Y is a model, in which case $D_{\text{KL}}(X||Y)$ tells us how far the model is from the truth. A divergence of 0 implies that the two random variables are identically distributed.

3 Beta Models and Entropy

In this section, we discuss different entropy measures that can be applied to a Beta distribution. We formally state each of these measures, we discuss their intuitive meaning, their application, and how they differ from uncertainty. The measure of entropy that does match uncertainty will be introduced in the next section. This section helps appreciate why that measure of entropy is the way it is.

3.1 Integrity Parameter Entropy

The most obvious measure of entropy that can be applied, is the (differential) entropy of the integrity parameter. To be precise, the entropy measure is:

$$H(X) = - \int_0^1 p_X(x) \cdot \log(p_X(x)) \, dx.$$

The standard intuition of differential entropy applies. In the case of differential entropy, values are negative and the absolute quantity tells you how much information is gained relative to the uniform distribution. The information that is gained is about the precise value of the integrity parameter. Differently put, it measures how far away from the uniform distribution, values in the distribution tend to be. Figure 1 provides two examples of graphs, Figure 1a depicts a distribution with less information about integrity than Figure 1b.

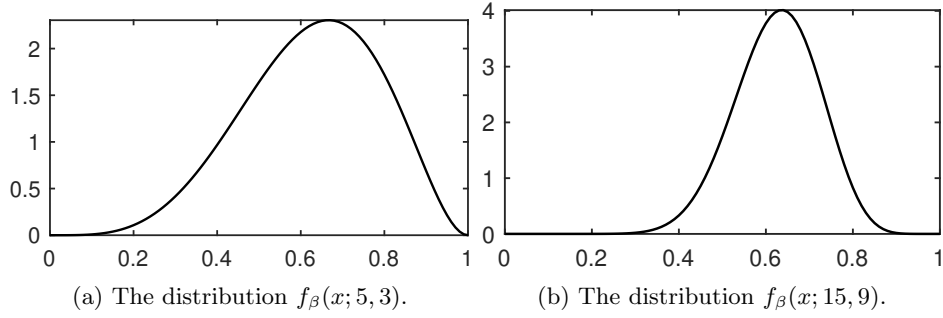


Fig. 1: Two Beta distributions equal in expected value, but not uncertainty.

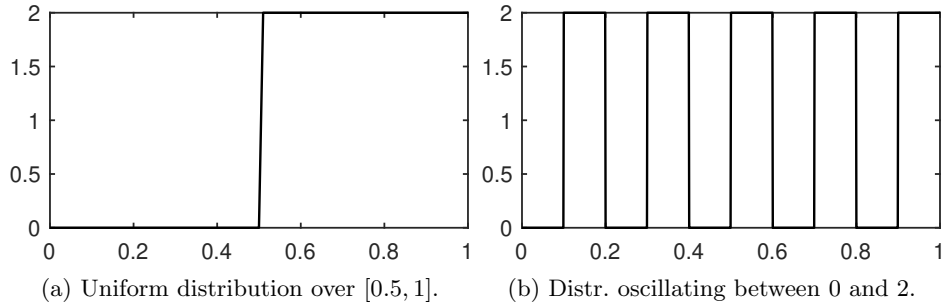


Fig. 2: Two distributions with uniform support on half the interval.

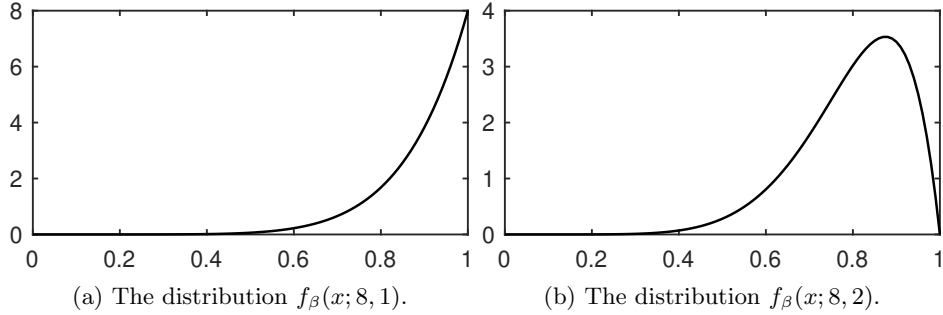


Fig. 3: Two Beta distributions with entropy increasing when adding evidence.

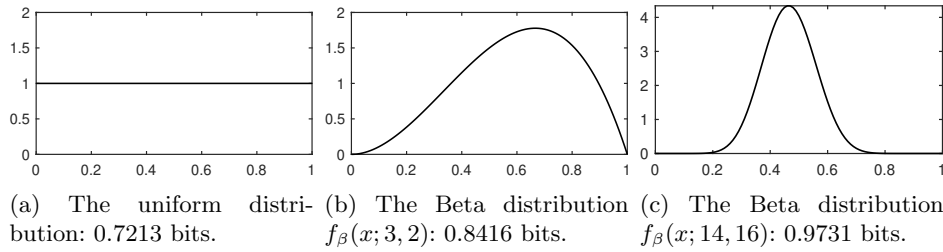


Fig. 4: More evidence continues to decrease information.

In reality, it is not important whether the integrity value is exactly 0.7, or say 0.705. For the purpose of measuring the entropy of the integrity value, these two values are considered to be completely different. For graphs such as the ones depicted in Figure 1, this is not a major issue, since the probabilities of similar integrity values tend to be similar too. However, in more extreme cases, such as in the graph depicted in Figure 2, it becomes an issue for our intuition. The graphs in Figure 2a and 2b are identical through the lens of the information measure, since both distributions have support on half the interval, and are uniformly distributed over the part with support. In both cases, the information gained over the uniform distribution is 1 bit – since we can exclude exactly half the possibilities. However, if we want to know whether we are dealing with a reliable person, the distribution in Figure 2a is likely to be helpful, but the one in Figure 2b is not.

Uncertainty is inversely proportional to the amount of evidence (i.e., the sum of the parameters of the Beta distribution). Adding evidence tends to increase the information about the integrity parameter too, as the peak tends to become narrower, as illustrated in Figure 1. However, it is not necessarily the case that adding evidence decreases the entropy, as illustrated in Figure 3. The distribution $f_\beta(8, 1)$ has an entropy of 1.7376 bits, whereas the distribution $f_\beta(8, 2)$ has an entropy of 1.1468 bits. Therefore, entropy about the integrity parameter fails to meet the basic criterion of uncertainty, which is being that it is monotonically decreasing as evidence is added.

3.2 Bernoulli Trial Entropy

An ingredient that was missing from integrity parameter entropy, was to take into account the values of integrity parameter, rather than just its probability density. Arguably, we are not necessarily interested in the exact integrity of other agents, but we are interested in knowing whether they will betray us or not. Whether an agent would betray us is determined by a Bernoulli trial based on the integrity parameter. In other words, an agent will not betray us with a probability equal to its integrity parameter. Since the Beta distribution is the estimate of that integrity parameter, the expected entropy of the Bernoulli trial is:

$$\begin{aligned} H(B) &= -E_x(x \log(x) + (1 - x) \log(1 - x)) \\ &= - \int_0^1 p_X(x) \cdot (x \log(x) + (1 - x) \log(1 - x)) dx \end{aligned}$$

Although we are computing the expectation of the entropy, the standard intuition of entropy applies: how much information about the outcome of the Bernoulli trial do we (expect to) have. The entropy of a Bernoulli trial is between 0 and 1 bits, where values close to 0 bits mean near certainty about whether we will be betrayed or not. The Beta distribution with maximal uncertainty – the uniform distribution – has an entropy of 0.7213 bits in this measure; strictly less than 1.

It can certainly be useful to measure how much you about the Bernoulli trial, but this measure has barely any connection to uncertainty. Consider a user with an integrity parameter of 0.5. A reasonable progression of Beta distributions as more evidence is accumulated is depicted in Figure 4. What we see in Figure 4 is that we are increasingly certain that the integrity parameter must be near 0.5. If the integrity parameter is 0.5, then we have 1 bit entropy of the Bernoulli trial, whereas the values near the extremes have near 0 bits entropy. As the evidence accumulates, this measure converges to 1 bit entropy. Again, this breaks the most basic requirement that entropy decreases as uncertainty decreases.

3.3 KL-Divergence from Truth

The problem with Bernoulli trial entropy as a measure for uncertainty, is that as evidence is added, it provides a value that is closer to the true Bernoulli entropy of that agent, rather than a smaller value. Assume that, somehow, we have access to the true integrity parameter of an agent, then we can measure the information-theoretic distance to that value. The standard technique is to use Kullback-Leibler divergence. Given a true integrity parameter of value x , we can apply KL-divergence to the Bernoulli trial entropy as:

$$E_y(D_{\text{KL}}(f_B(x)||f_B(y))) = \int_0^1 p_X(y)(x \log(\frac{x}{y}) + (1-x) \log(\frac{1-x}{1-y})) dy$$

As an example, say we measure 6 successes and 1 failure with an agent with parameter 0.85, then we get the KL-divergence from the truth as: $\int_0^1 f_\beta(y; 6, 1) \cdot (0.85 \log(\frac{0.85}{y}) + 0.15 \log(\frac{0.15}{1-y})) dy = 0.1247$ bits. However, it is possible that we measure 6 successes and 1 failure with an agent with parameter 0.4, in which case the distance is 1.2460. The measure does not just depend on the distribution itself.

This measure cannot be applied to compute the entropy, given an arbitrary Beta distribution, since the true integrity parameter is an unknown. Notice that the shape of the equation is such that the formula for the expectation of what behaviour will be observed is similar to the equation for the expectation of the integrity parameter given the observed behaviour. By applying Bayes' theorem, we can alter this term to talk about the expected true integrity parameter given the observed behaviour: $E_{x,y}(D_{\text{KL}}(f_B(x)||f_B(y)))$. This formula turns out to be EDRB, as we see in the next section.

4 Entropy-Uncertainty Equivalence

It may not be immediately obvious what it means for entropy measures and uncertainty measures to be equivalent. Both uncertainty and EDRB (expected KL-divergence of random Bernoulli trials) are actually families of measures, rather than a singular measure. Recall that if n_e is the amount of evidence, the general expression for uncertainty is $\frac{r}{n_e+r}$. EDRB provides different outcomes, depending on the choice of the base of the logarithm b , we will prove that it is $\frac{\log(b)}{n_e+2}$. In the case $r = 2, b = e^2$, the two formulas are equal. However, we argue that

the equivalence is stronger, since every member of the two families shares the crucial property that its inverse is a linear function of the amount of evidence.

Our goal, therefore, is to prove that $E_{x,y}(D_{\text{KL}}(f_B(x)||f_B(y))) = \frac{\log(b)}{n_e+2}$. Note that if we have s successes and f failures, our Beta distribution is $f_\beta(x; \alpha, \beta)$, with $\alpha = s + 1$ and $\beta = f + 1$. Therefore, $\alpha + \beta = s + f + 2 = n_e + 2$. Therefore, we can state our Theorem as the following equation:

Theorem 1.

$$\int_0^1 \int_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) \left((x \log_b \left(\frac{x}{y} \right) + (1-x) \log_b \left(\frac{1-x}{1-y} \right) \right) dy dx = \frac{\log(b)}{n_e + 2}$$

Proof. We will prove that :

$$\int_0^1 \int_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) x \log_b \left(\frac{x}{y} \right) dy dx = \frac{\log(b)\beta}{(\alpha + \beta)^2}.$$

Swapping α and β while substituting x for $1-x$ and y for $1-y$, it follows:

$$\int_0^1 \int_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) (1-x) \log_b \left(\frac{1-x}{1-y} \right) dy dx = \frac{\log(b)\alpha}{(\alpha + \beta)^2}.$$

This suffices to prove the theorem, since $\frac{\log(b)\beta}{(\alpha+\beta)^2} + \frac{\log(b)\alpha}{(\alpha+\beta)^2} = \frac{\log(b)}{\alpha+\beta}$.

$$\begin{aligned} & \int_0^1 \int_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) x \log_b \left(\frac{x}{y} \right) dy dx \\ &= \{ \log_b(x/y) = \log_b(x) - \log_b(y) \} \\ & \iint_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) x \log_b(x) dy dx - \iint_0^1 f_\beta(x; \alpha, \beta) f_\beta(y; \alpha, \beta) x \log_b(y) dy dx \\ &= \{ \text{Erase } y \text{ left. Isolate } x \text{ right, replace with } B(\alpha+1, \beta)/B(\alpha, \beta). \text{ Rename } y \text{ to } x. \} \\ & \int_0^1 x f_\beta(x; \alpha, \beta) \log_b(x) - \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} f_\beta(x; \alpha, \beta) \log_b(x) dx \\ &= \{ \text{Integration by parts on both summands. Recall Definition 2.} \} \\ & \left[\log_b(x) \int_0^x \frac{t^\alpha (1-t)^{\beta-1}}{B(\alpha, \beta)} dt \right]_0^1 - \int_0^1 \log(b) \frac{1}{x} \left(\int_0^x \frac{t^\alpha (1-t)^{\beta-1}}{B(\alpha, \beta)} dt \right) dx \\ & - \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} \left(\left[\log_b(x) \int_0^x \frac{t^{\alpha-1} (1-t)^{\beta-1}}{B(\alpha, \beta)} dt \right]_0^1 - \int_0^1 \log(b) \frac{1}{x} \left(\int_0^x \frac{t^{\alpha-1} (1-t)^{\beta-1}}{B(\alpha, \beta)} dt \right) dx \right) \\ &= \{ \text{Simplify to regularised incomplete Beta functions (Definition 2).} \} \\ & \left[\log_b(x) I_x(\alpha+1, \beta) \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} \right]_0^1 - \int_0^1 \log(b) \frac{1}{x} I_x(\alpha+1, \beta) \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} dx \\ & - \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} \left(\left[\log_b(x) I_x(\alpha, \beta) \right]_0^1 - \int_0^1 \log(b) \frac{1}{x} I_x(\alpha, \beta) dx \right) \\ &= \{ \text{Terms in square brackets evaluate to 0 at 0 and 1. Simplify formula.} \} \\ & \int_0^1 \log(b) \frac{1}{x} \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} (I_x(\alpha, \beta) - I_x(\alpha+1, \beta+1)) dx \\ &= \{ \text{Apply Proposition 1.} \} \\ & \int_0^1 \log(b) \frac{1}{x} \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} \frac{x^\alpha (1-x)^\beta}{\alpha B(\alpha, \beta)} dx \\ &= \{ \text{Use } 1/x \text{ to subtract 1 from the exponent } \alpha, \text{ apply Proposition 1.} \} \\ & \log(b) \frac{\alpha}{\alpha + \beta} \cdot \frac{\beta}{\alpha(\alpha + \beta)} = \frac{\log(b)\beta}{(\alpha + \beta)^2} \end{aligned}$$

There are two ways to interpret the theorem. Firstly, we can use the intuition from Section 3.3, and say that $f_\beta(x; \alpha, \beta)$ is the Bayesian estimate of the true integrity parameter that generated the history, and we measure the expected KL-divergence between the Bernoulli trial with the true integrity parameter and a new randomly selected parameter (y). Simply put, we reuse the measure from Section 3.3, but substituting the true integrity for the expected integrity. KL-divergence is an oft-used way to measure the quality of a model distribution, compared to the real one. EDRB measures the expectation of the distance between the KL-divergence of the Bernoulli trail based on an estimated true one and an estimated model. Of course, taking the expectation of the true integrity used for the Bernoulli trial is intuitively dubious.

The alternative intuition does not involve true integrities for this reason. EDRB can be interpreted to say, given two agents with the same history, how much do we learn about one agent, if we observe a new interaction with the other. As more evidence accumulates, the possible choices for the parameter for the Bernoulli trial becomes more centered around a specific value. If the probability that two Bernoulli trials use similar parameters increases, then the KL-divergence between the two decreases. This intuition is a more direct reading of the actual formula, as we are taking the expectation of a pair of integrity parameters, distributed along the same Beta distribution. The weakness of this intuition is that KL-divergence is an asymmetric measure, where one distribution represents the true distribution and the other one the model distribution, whereas this intuition is measuring the distance between two model distributions.

While both intuitions are imperfect, they do offer an explanation why we might expect uncertainty and EDRB to be related. The fact that they are indeed equivalent is non-obvious, however. The proof does not provide us with an insight as to why they are indeed equivalent – other than the fact that they are. Based on the fact that the intuitions are imperfect, and the proof does not provide any intuition either, we consider the equivalence to be surprising.

Uncertainty is a useful concept and a basic tenet of Subjective Logic. To compute the uncertainty, from a Beta distribution $f_\beta(x; \alpha, \beta)$, simply take $u = \frac{1}{\alpha + \beta - 1}$. However, this definition uses the parameters of the distribution, rather than the probability density function. Given a probability density function f , that happens to represent a Beta distribution, there is no elegant way to compute the uncertainty. For example if $f = 6(x(1-x)^5 + (1-x)^6)$, then how to determine its uncertainty – given it may not be trivial to realise $f = f_\beta(6, 1)$. Alternatively, we can compute $E_{X,Y}(D_{\text{KL}}(f_B X || f_B Y))$ for $X, Y \sim f$, and obtain $\frac{1}{\alpha + \beta}$ without knowing α and β .

The fact that we can use our new measure as an alternative way to compute the uncertainty of a Beta distribution without explicitly using the parameters, is interesting in itself. More interesting, however, is the fact that the input probability density function need not be a Beta distribution at all, for it to work. As we more rigorously argue in the next section, there are cases where it does not make sense to use a Beta distribution as opinions. These cases have been recognised implicitly in the literature (e.g. [14]), but are not typically explicitly addressed.

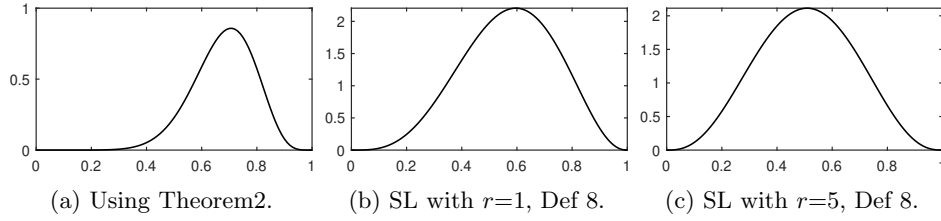


Fig. 5: Three differently computed conjunctions of $f_\beta(x; 8, 4)$ and $f_\beta(x; 9, 2)$.

We can now reason about the uncertainty present in more esoteric distributions that may pop up. In the next section, we present some of the implications to these generalised distributions.

5 Generalised Opinions

That models of information fusion found in Subjective Logic are isomorphic to Beta distributions is not surprising. After all, these models are created with this purpose in mind. Subjective Logic further incorporates logical operations, and transitive trust operations. As shown in [10] and [11] respectively, the resulting distribution of these operations is not a Beta distribution (using the assumptions of the Beta model). In other words, the isomorphism does not hold if we add the new operations. In this section, we show examples of distributions resulting from logical or transitive operations, discuss why they are not Beta distributions, and extend the result from the previous section to these distributions.

5.1 Opinion Logic

Consider performing logic on the opinions. For example, we have a distribution for A and for A' , but in order to obtain a success, we need both A and A' to succeed. In the case that A and B are independent agents, the probability that $A \wedge A'$ succeeds is a Bernoulli trial with parameter $x_A \times x_{A'}$ [10]. According to [10], if we want to obtain our opinion on $A \wedge A'$, based on our opinions on A and A' , then we need to take their product distribution:

Theorem 2. *If $B_{A,k}$ and $B'_{A',k'}$ are independent Bernoulli trials, then let C be 1 iff $B_{A,k} = 1$ and $B'_{A',k'} = 1$. Let $p(C = 1 | X_C = x_C) = x_C$ then the distribution $p_{X_C}(x_C | B_{A,1} \dots B_{A,k-1}, B'_{A',1} \dots B'_{A',k'-1}) = \int_{x_C}^1 \frac{1}{y} p_{X_C}(\frac{x_A}{y} | B_{A,1} \dots B_{A,k-1}) \cdot p_{X_{A'}}(y | B_{A',1} \dots B_{A',k'-1}) dy$; the product distribution of the opinions on A, A' .*

Proof. Theorem 4 in [10].

In Subjective Logic, conjunction is defined as:

Definition 8. *Conjunction $(b, d, u) \wedge (b', d', u') = (bb', d + d' - dd', bu' + b'u + uu')$.*

In Figure 5, we see the conjunction of $f_\beta(x; 8, 4)$ and $f_\beta(x; 9, 2)$ as derived from the product distribution, as well as the results computed using the Subjective Logic conjunction definition under f_1 and f_5 . We can see that neither f_1 nor f_5 are isomorphisms w.r.t. conjunction, since the graphs differ. In fact, for no choice of r , or even for any other Subjective Logic definition for conjunction, will f_r be an isomorphism. The reason is that all opinions in Subjective Logic are isomorphic to a Beta distribution, but the result of the product distribution is not generally (in fact, almost never) a Beta distribution. Therefore, no isomorphism can exist.

Although the resulting opinion is not a Beta distribution, we can compute the uncertainty via its equivalence to EDRB. The uncertainty of $f_\beta(x; 8, 4)$ and $f_\beta(x; 9, 2)$ is, therefore, equal to 0.0775.

5.2 Transitive Trust

Transitive trust is a fiercely debated topic. Using the assumptions of the Beta model, an issue arises. The formula contains a term χ , which is the attacker's strategy. In other words: how to use the advice of another agent, depends on how the agent would act, if he were malicious. See [11] for more details. The attacker strategy is not a topic for this paper, so we will assume the simplest attack strategy: random behaviour.

If an advisor A is honest with probability x_A , and the advisor gives us the opinion p_{X_c} , then our resulting opinion is simply $x_A \cdot p_{X_c}(x_c) + (1 - x_A)$. This can be derived from Theorem 2 in [11]. However, the intuition behind it is also clear, namely that if the advisor speaks the truth, we should listen, and if he lies, we know nothing. We do not typically know x_A , but we can use our opinion p_{X_A} to estimate this value.

The result of obtaining an opinion from advice, therefore, is not a Beta distribution, but a weighted sum of Beta distributions². However, like in Section 5.1, Subjective Logic must return a Beta distribution as the result of transitive trust. In fact, Subjective Logic defines transitive trust:

Definition 9. *Propagation* $(b, d, u) \cdot (b', d', u') = (bb', bd', bu' + d + u)$.

In Figure 6, we see the propagation of $f_\beta(x; 8, 4)$ and $f_\beta(x; 9, 2)$ as derived from the summing Beta distributions, as well as the results computed using the Subjective Logic propagation definition under f_1 and f_5 . Compared to conjunction, we see that the difference between the two approaches is even larger. In particular, we notice that Figure 6a has raised flat tails. These raised flat tails are a consequence of the fact that, no matter what malicious agents say, if they are lying, then extremely high/low integrity values remain probable.

Although the resulting opinion is not a Beta distribution, we can compute the uncertainty via its equivalence to EDRB. The uncertainty of the opinion resulting from hearing $f_\beta(x; 9, 2)$ from an agent that we have the opinion $f_\beta(x; 8, 4)$ about, is equal to 0.5354. In this case, the uncertainty is (significantly) larger than

² This is true also when the attack strategy is not trivial.

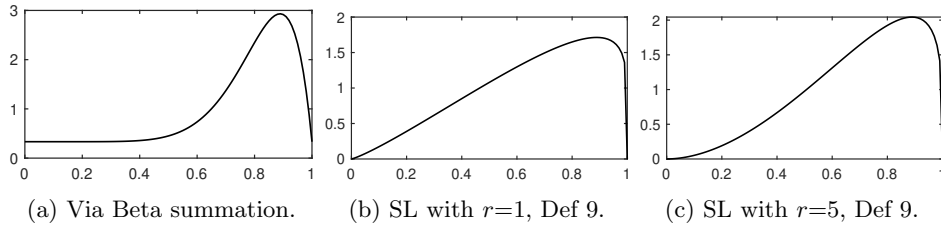


Fig. 6: When an agent with $f_\beta(x; 8, 4)$ claims to have opinion $f_\beta(x; 9, 2)$.

the uncertainty of $f_\beta(x; 9, 2)$ (which is 0.1000). However, it need not be the case that summing Beta distributions changes the uncertainty in a meaningful way. In particular, the uncertainty of $1/3(f_\beta(x; 3, 1) + f_\beta(x; 2, 2) + f_\beta(x; 1, 3))$ is the maximum: 1, even though the individual distributions have far smaller uncertainty. There may be a more subtle pattern in the EDRB entropy of a sum of Beta distributions, but this is future work.

A reasonable approach to selecting a strategy for the attacker, is to select the strategy that is the least informative. Typically, that means the strategy that gives the highest entropy. No closed formula has been found that maximises either the integrity entropy or the Bernoulli trial entropy. An open question is whether this approach can be more fruitful when using EDRB as the measure of entropy.

6 Conclusion

Theorem 1 is our main result. It states that uncertainty (the inverse of amount of evidence) is equal to a specific measure of entropy that we introduce: expected Kullback-Leibler divergence of random-parameter Bernoulli trials (EDRB). The intuition behind EDRB is that it measures the expected distance between two Bernoulli trials selected from a distribution – a more narrow distribution will have less distance between the Bernoulli trials.

While both entropy and uncertainty can be used to describe lack of knowledge. Any entropy measure is based on surprisal, whereas uncertainty is based on Bayesian evidence. Hence is surprising that they should coincide.

We discuss alternative measures of entropy in Section 3. Measures such as integrity entropy and Bernoulli trial entropy certainly have use-cases. Uncertainty simply measures something else than these two measures.

Finally, we study the implications of having EDRB on generalised opinions. These are distributions other than the Beta distributions. In Section 5, we show how these distributions arise, and why they are of interest. We plan to further study of the implications of generalised opinions under EDRB. In particular we want to explore the notion of malicious advisors maximising EDRB entropy.

References

1. Burt, D.R.: Bandwidth and echo: Trust, information, and gossip in social networks (2001)
2. Che, S., Feng, R., Liang, X., Wang, X.: A lightweight trust management based on bayesian and entropy for wireless sensor networks. *Security and Communication Networks* **8**(2), 168–175 (2015)
3. ElSalamouny, E., Sassone, V., Nielsen, M.: Hmm-based trust model. In: *International Workshop on Formal Aspects in Security and Trust*. pp. 21–35. Springer (2009)
4. Ismail, R., Jøsang, A.: The beta reputation system. *BLED 2002 proceedings* p. 41 (2002)
5. Jaynes, E.T.: *Probability theory: The logic of science*. Cambridge university press (2003)
6. Jøsang, A.: Artificial reasoning with subjective logic. In: *Proceedings of the second Australian workshop on commonsense reasoning*. vol. 48, p. 34. Citeseer (1997)
7. Jøsang, A.: *Subjective logic*. Springer (2016)
8. Kotz, S., Balakrishnan, N., Johnson, N.L.: *Continuous multivariate distributions*, vol. 1. John Wiley & Sons (2004)
9. McEliece, R.J.: *Theory of Information and Coding*. Cambridge University Press New York, USA, 2nd edition (2001)
10. Muller, T., Schweitzer, P.: A formal derivation of composite trust. In: *International Symposium on Foundations and Practice of Security*. pp. 132–148. Springer (2012)
11. Muller, T., Schweitzer, P.: On beta models with trust chains. In: *IFIP International Conference on Trust Management*. pp. 49–65. Springer (2013)
12. Muller, T., Wang, D., Jøsang, A.: Information theory for subjective logic. In: *International Conference on Modeling Decisions for Artificial Intelligence*. pp. 230–242. Springer (2015)
13. Teacy, W.L., Luck, M., Rogers, A., Jennings, N.R.: An efficient and versatile approach to trust and reputation using hierarchical bayesian modelling. *Artificial Intelligence* **193**, 149–185 (2012)
14. Teacy, W.L., Patel, J., Jennings, N.R., Luck, M.: Travos: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems* **12**(2), 183–198 (2006)
15. Wang, D., Muller, T., Irissappane, A.A., Zhang, J., Liu, Y.: Using information theory to improve the robustness of trust systems. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. pp. 791–799. International Foundation for Autonomous Agents and Multiagent Systems (2015)
16. Wang, D., Muller, T., Zhang, J., Liu, Y.: Is it harmful when advisors only pretend to be honest? In: *Thirtieth AAAI Conference on Artificial Intelligence* (2016)
17. Wu, X., Huang, J., Ling, J., Shu, L.: Bltm: Beta and lqi based trust model for wireless sensor networks. *IEEE Access* (2019)
18. Xiao, S., Dong, M.: Hidden semi-markov model-based reputation management system for online to offline (o2o) e-commerce markets. *Decision Support Systems* **77**, 87–99 (2015)