



Kent Academic Repository

Leisen, Fabrizio, Villa, Cristiano and Walker, Stephen G. (2019) *On a Class of Objective Prior from Scoring Rules*. Bayesian Analysis . ISSN 1936-0975.

Downloaded from

<https://kar.kent.ac.uk/77889/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1214/19-BA1187>

This document version

Publisher pdf

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

On a Class of Objective Priors from Scoring Rules

Fabrizio Leisen^{*}, Cristiano Villa[†], and Stephen G. Walker[‡]

Abstract. Objective prior distributions represent an important tool that allows one to have the advantages of using a Bayesian framework even when information about the parameters of a model is not available. The usual objective approaches work off the chosen statistical model and in the majority of cases the resulting prior is improper, which can pose limitations to a practical implementation, even when the complexity of the model is moderate. In this paper we propose to take a novel look at the construction of objective prior distributions, where the connection with a chosen sampling distribution model is removed. We explore the notion of defining objective prior distributions which allow one to have some degree of flexibility, in particular in exhibiting some desirable features, such as being proper, or log-concave, convex etc. The basic tool we use are proper scoring rules and the main result is a class of objective prior distributions that can be employed in scenarios where the usual model based priors fail, such as mixture models and model selection via Bayes factors. In addition, we show that the proposed class of priors is the result of minimising the information it contains, providing solid interpretation to the method.

Keywords: calculus of variation, differential entropy, Euler–Lagrange equation, Fisher information, invariance, objective Bayes, proper scoring rules.

1 Introduction

With the ever increasing popularity of Bayesian methods, attributable largely to the advent of Markov chain Monte Carlo methods and other sampling techniques, the need for default, otherwise known as objective or noninformative, priors is also in demand. Model based objective priors, such as the reference prior (Berger et al., 2009) and Jeffreys prior (Jeffreys, 1961), are commonly used when available. However, as models become larger and more complex, so it is that such priors are becoming more difficult to obtain, if not altogether unavailable. Indeed, it is our contention that model based objective priors have now reached their natural ceiling with little progress or advances in recent years.

Our observation is that limits to the progress in the research on objective priors is connected to their improperness. In fact, with very few exceptions, objective priors are improper. Although this may not represent a problem, as long as the posterior is

^{*}School of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, UK, F.Leisen@kent.ac.uk

[†]School of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, UK, C.Villa-88@kent.c.uk

[‡]Department of Mathematics, University of Texas at Austin, USA, s.g.walker@math.utexas.edu

proper, it causes severe limitations to the use of objective prior distributions. Indeed, the impropriety of objective priors is the main motivation which brought us to investigate a novel approach to derive objective priors. A thorough discussion of the problems that improper priors cause can be found in Kass and Wasserman (1996), where they illustrate issues such as incoherence, strong inconsistencies and nonconglomerability, dominating effect of the prior, inadmissibility, marginalisation paradoxes and improper posteriors. Although all points are undoubtedly important, it is probably the last issue that requires careful consideration. The main concern is that, as of today, general results that allow one to assess if a given improper prior yields a proper posterior are yet to be found. Research has progressed on a case by case basis; for example, see Ibrahim and Laud (1991) for the use of Jeffreys prior in generalised linear models, extended to overdispersed models of the same kind by Dey et al. (1993), Natarajan and McCulloch (1995), Berger and Strawderman (1993), and Yand and Chen (1995). More recently, Rubio and Steel (2018) describe general conditions to use improper priors for linear mixed models with longitudinal and survival data. As one would expect, the task of assessing posterior propriety becomes more onerous the more complex the model is. But, even for simple models, the risk is high; as, for example, the discussion in Vallejos and Steel (2013) about the use of the Jeffreys rule prior for the Student- t regression model, derived in Fonseca et al. (2008), shows.

In this paper, we investigate constructing objective prior distributions that are not model dependent and based on the sole knowledge of the parameter space, say, Θ . As such, the connection between the prior distribution, $p(\theta)$, and the likelihood function, $f(\cdot|\theta)$, is limited to the common parameter space only.

Therefore, the prior $p(\theta)$ using only Θ loses the connection with the subjective component $f(\cdot|\theta)$ and could be argued as a consequence to be *more objective*. Conversely, model based priors, such as the Jeffreys prior (Jeffreys, 1946, 1961) or the reference prior (Berger et al., 2009), necessarily include the subjective choice of the model. In fact, models are by and large misspecified and, consequently, model based priors are propagating this misspecification. So, while a model based prior reinforces the connection between the misspecified model and the prior itself, a prior that depends on the parameter space loses only the connection.

The method we propose to derive a novel class of objective priors is based on defining a scoring rule as a combination of the *log-score* and of the *Hyvärinen* score (Parry et al., 2012). We then seek a prior such that the above score, say $S(\theta, p)$, is constant, that is

$$S(\theta, p) = \text{constant} \quad \forall \theta \in \Theta. \quad (1)$$

We show in Section 3 that the density p satisfying (1) identifies a class of objective priors. Furthermore, we show in Section 4 that the density $p(\theta)$ solving (1) minimises the information in the prior as measured by a combination of the Shannon information and the Fisher information. As a result of these equivalent approaches, we show that the objective prior $p(\theta)$ is obtained by solving the following differential equation:

$$\frac{p''(\theta)}{p(\theta)} - \frac{1}{2} \left\{ \frac{p'(\theta)}{p(\theta)} \right\}^2 = 1 + \log p(\theta). \quad (2)$$

Our prior is the class of solutions to the differential equation (2). The class allows us to include constraints such as properness, convexity and being decreasing, among others.

The development of objective priors is thoroughly reviewed in Consonni et al. (2018). The idea is that in a scenario where prior elicitation is not feasible, or not desirable, a prior distribution can be formed through structural or formal rules (Kass and Wasserman, 1996). The most popular objective prior is Jeffreys prior (Jeffreys, 1946, 1961), who proposed a prior distribution for continuous parameter spaces which is invariant for one-to-one transformations of the parameter space. Although in scenarios where there is only one parameter of interest, Jeffreys prior yields sensible posterior distributions. However, in cases where the parameter space has a dimension of two or more, the prior is known to yield posteriors with poor performance (sometimes giving paradoxical results, such as the marginalisation paradox, see Stone and Dawid (1972)). In such cases the priors are taken to be independent. Although other more general invariance priors have been proposed, such as in Dawid (1983), Hartigan (1964) and Jaynes (1968), the reference prior of Bernardo (Berger et al., 2009) represents an alternative to Jeffreys prior. A limitation of reference priors is sensitivity to the order of importance of parameters; this issue and possible solutions have been discussed in Berger et al. (2015). Other objective priors proposed include that of Box and Tiao (1973), based on data-translated likelihoods, and maximum entropy priors, see, for example, Jaynes (1957, 1968). As discussed in Kass (1990), these priors turn out to be very restrictive. Another important class of objective priors are the probability matching priors, first proposed in Welch and Peers (1963). The aim is to obtain a prior distribution under which the posterior probabilities of certain regions coincide with their coverage probabilities, either exactly or approximately. Recent developments of this method can be found in Sweeting et al. (2006) and Sweeting (2008). A different method, based on information theoretical concepts, has been proposed by Zellner and Min (1993), giving the so called maximal data information prior. Possibly, the most recent development in defining prior distributions, although not strictly in an objective sense, is discussed in Simpson et al. (2017). The idea is to identify the parts in a complex model that require subjective input, while the remaining parts can be associated with non-informative priors. It is important to point out that objective priors derived from scoring rules have been proposed in the recent work of Giummolè et al. (2018). A final consideration is reserved for discrete parameter spaces, whose systematic discussion can be seen to be generated by the paper of Rissanen (1983). The lack of general methods, due to the challenges that discreteness imposes, has been filled by Berger et al. (2012) first, and by Villa and Walker (2015) later.

It is known that for a mixture model Jeffreys prior can only be found under specific conditions, as studied in Grazian and Robert (2018), for reasons identified, in part, by Titterton et al. (1985). Also, standard objective priors are problematic in model comparison since Bayes factors depend on the arbitrary normalizing constants of improper prior. Though, special solutions have been proposed, for example, by O'Hagan (1995) and Berger and Pericchi (1996). Further discussion of these contexts is contained in Sections 6.1 and 6.2, where solutions provided by our approach are explored. There are other examples where the use of improper priors is challenging. For example, Kass and Wasserman (1996) discuss some issues related to their use in hierarchical modelling. Also, paradoxical results may appear, such as the marginalisation paradox (Stone and

Dawid, 1972) or the Stein’s paradox (Bernardo and Smith, 1994). For more examples, see Stone (1976) and Syversveen (1998).

The paper is structured as follows. In Section 2 we give an outline of the idea and discuss a simple introductory example. In Section 3 we introduce the foundations of the proposed prior on the basis of scoring rules and their properties. An interesting aspect of the prior based on scoring rules is its interpretation in terms of the information content carried by the prior itself. This aspect is explored in Section 4. In Section 5 we present the objective priors concentrating on $\Theta = (0, 1)$, $(0, \infty)$, $(-\infty, +\infty)$ and $(-M, M)$, for a finite M , as well as on a multidimensional space. The implementation of the prior for some specific applications is presented in Section 6 where, in addition, we explore some properties of the proposed prior. Finally, Section 7 is dedicated to some concluding remarks.

2 Outline of the idea

Here we illustrate two key motivating examples, in particular in relation to the use of improper priors, and give an intuitive description of the core ideas of this work.

The key to our idea is to consider a loss function $l(\theta, p(\theta))$ which penalizes for each $\theta \in \Theta$ a choice of a prior density $p(\theta)$. The objective criterion is then based on the idea of finding the class of p which makes $l(\theta, p(\theta))$ constant. For obvious reasons, the loss function should have the following property

$$\int l(\theta, p(\theta))q(\theta) d\theta \geq \int l(\theta, q(\theta))q(\theta) d\theta, \quad (3)$$

for all q ’s representing a density for the θ . In other words, if a “true” density for θ exists, the expected loss should be minimised when such a density is chosen. The condition in (3) identifies a particular class of loss functions, known as *proper scoring rules*. One way of interpreting (proper) scoring rules is as loss functions that measure the quality of a quoted density p for an uncertain quantity θ ; see, for example, Parry et al. (2012). We indicate a proper scoring rule by $S(\theta, p)$, and we ask it to be constant for all $\theta \in \Theta$. So we set

$$S(\theta, p) = \text{constant} \quad \forall \theta \in \Theta,$$

and the densities satisfying the above equality identify a class of objective priors. We set the constant to 1 and show later that this choice is without loss of generality. The criterion defining this class of priors is clearly objective, for if the scoring rule were not constant, some parts of the space Θ would be given preference above others.

As discussed in Parry et al. (2012) a scoring rule is defined as *local* if it depends on the density function $p(\cdot)$ only through its value at θ , that is $p(\theta)$. Holding the above definition, we have that any *proper local* scoring rule is equivalent to the *log score*, $-\log p(\theta)$, also known as the self information loss function. However, the above scoring rule is not suitable to us, for if we set $-\log p(\theta) = \text{constant}$, we only achieve $p(\theta) \propto 1$, and the result is not interesting. Hence, we extend the score function to include the first and second derivatives of $p(\theta)$; a natural extension, in our opinion.

Thus, we consider additionally the Hyvärinen scoring rule (Hyvärinen, 2005) which makes use of the first two derivatives of p , written as p' and p'' . We then have a scoring

rule $S(\theta, p)$ which has two components; the log score and the Hyvärinen score. Finding solutions to $S(\theta, p) = 1$ will now involve solving a second order differential equation and we obtain the class of prior through the two constants connected with the two derivatives. Multivariate versions of the scoring rule criterion we are proposing, and corresponding solutions, are discussed in Section 3.

Parry et al. (2012) describe a larger class of score based on the first two derivatives; see (39) in their paper. In the second order case these are based on an Euler–Lagrange equation which itself is based on a measure of information of the form $\int [p'(\theta)]^k / [p(\theta)]^{k-1} d\theta$ for $k = 2, 3, \dots$. The only widely recognised information occurs with $k = 2$; the Fisher information and the corresponding score coincides with the Hyvärinen score. Hence, we use this. A connection between our scoring rule criterion and information is made through an alternative formulation through variational methods, which is discussed in Section 4.

3 Priors from scoring rules

Let us consider a quantity of interest, θ , which can take values in the space $\Theta \subset \mathbb{R}^k$. The fundamental argument behind objective prior distributions is that they should represent a state of actual or alleged prior ignorance about the true value of θ . Several criteria have been proposed to select such a prior, all of which assume that a probabilistic model generating the data (given θ) has been chosen. What we propose is to avoid this choice and derive a prior depending on Θ only. The idea is to measure the quality of the prior p with a proper scoring function, say $S(\theta, p)$, and require it to be constant, as discussed in the Introduction.

Definition 1. *A density p with respect to the Lebesgue measure on Θ , is objective (in accordance with commonly accepted meaning of the expression) if $S(\theta, p) = \text{constant}$ for all $\theta \in \Theta$, where S is a proper scoring rule.*

The constant here is unimportant and we can take it as 0. Any constant works – effectively it becomes 0 once the normalizing constant for p has been established.

Before proceeding we provide a brief discussion on scoring rules. Scoring rules are *proper* if $\int_{\Theta} S(\theta, p) q(\theta) d\theta$ is minimized at $p = q$ and *local* if it depends on p only through the value $p(\theta)$. The unique proper local scoring rule is the log score, defined as

$$S_L(\theta, p) = -\log p(\theta).$$

Parry et al. (2012) extend the local property to m -local, in that now $S(\theta, p)$ depends also on the l -derivative $p^{(l)}(\theta)$, for $0 \leq l \leq m$. In particular, for $m = 2$, there is the Hyvärinen scoring rule, (Hyvärinen, 2005), given by

$$S_H(\theta, p) = \Delta \log p(\theta) + \frac{1}{2} |\nabla \log p(\theta)|^2.$$

In Section 4 we will illustrate the connection of the Hyvärinen scoring rule to Fisher information. Our choice of scoring rule following our reasoning in Section 2, and with

the weighting factor, is

$$\begin{aligned} S(\theta, p) &= w S_L(\theta, p) + S_H(\theta, p) \\ &= -w \log p(\theta) + \sum_{j=1}^k \frac{\partial^2 p / \partial \theta_j^2}{p}(\theta) - \frac{1}{2} \sum_{j=1}^k \left(\frac{\partial p / \partial \theta_j}{p}(\theta) \right)^2. \end{aligned}$$

That this is a proper scoring rule is derived from the fact that it is the sum of two proper scoring rules. It is also clearly 2-local. Previously, priors have been sought based solely on $\log p$; for example, the reference prior, and the math becomes unnatural as a consequence. On the other hand, including higher derivatives yields well defined solutions to optimization procedures. That we set this score to 1 for all θ is done without loss of generality, as we shall see later on. That we understand this to be an objective procedure is evident from the fact that no part of Θ is being given preference; the *loss* at θ for our choice of $p(\theta)$ is the same for all θ . For, if $S(\theta, p)$ did depend on θ then we argue that this could only be driven by *information*; i.e. parts of Θ space are preferential to others.

Predominantly, throughout the paper, we will be using the choice of $w = 1$. After all the value of w is a calibration issue between the two scores; i.e. to put them on a comparable scale. The reason for $w = 1$ is that for the benchmark standard normal density function, i.e. $p(\theta) \propto e^{-\frac{1}{2}\theta^2}$, the difference between the scores S_L and S_H is a constant (i.e. does not depend on θ), and so one does not end up dominating the other, only for $w = 1$.

Hence, we see that the objective prior, $p(\theta) \propto \exp\{-u(\theta)\}$, is obtained by solving the following differential equation:

$$\sum_{j=1}^k \frac{\partial^2 u}{\partial \theta_j^2} - \frac{1}{2} \sum_{j=1}^k \left(\frac{\partial u}{\partial \theta_j} \right)^2 = wu. \quad (4)$$

To derive the solution, we have the following result.

Theorem 3.1. *The solution to (4) is given by, for $j = 1, \dots, k$,*

$$\frac{\partial u}{\partial \theta_j} = \pm \sqrt{c_j e^{u(\theta)} - 2w(1 + u(\theta))}/k, \quad (5)$$

for some suitable constants $c = (c_j)$ and a specified value of u at some point; e.g. $u(0, \dots, 0)$.

Proof. Solving the differential equation (4) is equivalent to solving the following differential equation;

$$\sum_{j=1}^k v_j \frac{\partial v_j}{\partial u} = \frac{1}{2} \sum_{j=1}^k v_j^2 + uw, \quad (6)$$

having defined $v_j = \partial u / \partial \theta_j$. It is now seen the solution is given by (5). This follows since

$$v_j \frac{\partial v_j}{\partial u} = v_j \frac{1}{2} \frac{1}{v_j} (c_j e^u - 2w/k)$$

and note that $c_j e^u - 2w/k = v_j^2 + 2wu/k$. Therefore (6) holds. \square

The missing pieces in (5) are $c = (c_j)$ and say $u(0)$, the constants of integration. Note that, the initial value $u(0)$ (together with the constant c) is required to ensure the existence and uniqueness of the solution. We will see how to complete these when we look at illustrations in Section 4. In general, as the solution depends on the above arbitrary constants, our method provides a class of solutions, where some are proper and some are improper and, more general, where the priors will have some assigned properties via specification of $(c, u(0))$.

We also note here that we do not need the normalizing constant for p and neither do we need to find an explicit solution for u , and p , beyond (5). The reason for this is that we can find an accurate solution via numerical methods; i.e. if we have $u(\theta)$ at a particular θ value, then we can evaluate

$$u(\theta + \varepsilon) = u(\theta) + \varepsilon' \frac{\partial u}{\partial \theta}(\theta) + \frac{1}{2} \varepsilon' \frac{\partial^2 u}{\partial \theta^2}(\theta) \varepsilon + o(|\varepsilon|^2)$$

for small ε , and the $\partial u / \partial \theta$ and $\partial^2 u / \partial \theta^2$ are available explicitly, combined with the ease of obtaining higher derivatives if needed. From here we can evaluate $p(\theta)$.

To ease the reader into the proposed prior, we illustrate the following simple example, where the parameter space is $\Theta = (-M, +M)$. Here we discuss an explicit solution to the equation $u'(\theta) = \pm \sqrt{ce^{u(\theta)} - 2(1 + u(\theta))}$, which is (5) where we have set $w = 1$ and, as the parameter space is unidimensional, $k = 1$. If we set $c = 0$ then for a solution to exist we must have $1 + u$ to be negative. Consider a prior on $\Theta = (-M, +M)$ for some finite M . With $c = 0$ we have the solution for u in the form $1 + u(\theta) = -\frac{1}{2}(\theta - \mu)^2$ for some μ . Hence, $p(\theta) \propto \exp\{\frac{1}{2}(\theta - \mu)^2\}$, which will provide a proper density on Θ .

A more general solution $p(\theta) \propto \exp\{\frac{1}{2}w(\theta - \mu)^2\}$ arises when we take the more general form of score function; i.e. $S(\theta, p) = w S_L(\theta, p) + S_H(\theta, p)$, providing interpretation for the score weighting parameter in this case. Plots of such $p(\theta)$ depending on w are presented in Figure 1, with $M = 2$.

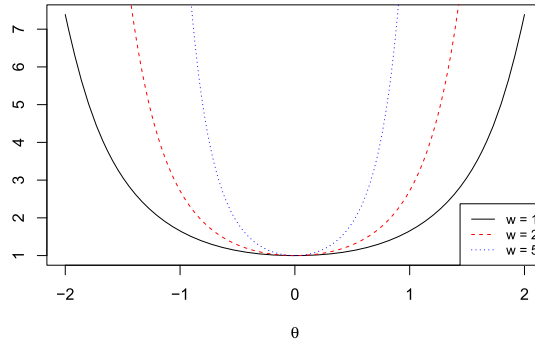


Figure 1: Prior for the parameter space $(-2, 2)$ with weights $w = 1$ (continuous black line), $w = 2$ (dashed red line) and $w = 5$ (dotted blue line).

4 Variational problems and solutions

Here we provide an alternative derivation of (4) using information theory, specifically entropy information and Fisher information. We show that the p solving (4) can also be regarded as a density carrying minimal local information. This material then is to provide support for the solution to (4) being an objective prior.

The entropy information (negative entropy) of a density function p is given by $I_E(p) = \int p(\theta) \log p(\theta) d\theta$, which is related to Shannon's entropy and is equal to negative the expected self-information loss. In addition to $I_E(p)$, we consider a measure of the information in the density p known as Fisher information, given by

$$I_F(p) = \int \frac{p'(\theta)^2}{p(\theta)} d\theta = \int p(\theta) \left\{ \frac{\partial}{\partial \theta} \log p(\theta) \right\}^2 d\theta.$$

See, for example, Bobkov et al (2014).

Now consider $I(p) = I_E(p) + \frac{1}{2}I_F(p)$ and the aim is to find the p which minimizes $I(p)$. Recalling variational methods (Rustagi, 1976), if we wish to minimise $\int_a^b L(\theta, p, p') d\theta$, a necessary condition for a local extremum of the integral of the Lagrangian $L(\theta, p, p')$ is that

$$\frac{\partial L}{\partial p} = \frac{d}{d\theta} \frac{\partial L}{\partial p'}. \quad (7)$$

Minimising $\int_a^b L(\theta, p, p') d\theta$ reduces to the classical calculus of variation problem where we want to extremize the integral of the function

$$L(\theta, p, p') = \frac{1}{2} \frac{p'(\theta)^2}{p(\theta)} + p(\theta) \log p(\theta).$$

The solution to the extremal problem, if it exists, is obtained from the Euler–Lagrange equation, given by (7). According to page 44 of Rustagi (1976), if $L(p, p')$ is strictly convex on $(0, \infty) \times (-\infty, +\infty)$, and p satisfies the Euler equation, then p is a minimum of $\int_a^b L(\theta, p, p') d\theta$. Now $L(p, p')$ is strictly convex if the matrix

$$H = \begin{pmatrix} \partial^2 L / \partial p^2 & \partial^2 L / \partial p \partial p' \\ \partial^2 L / \partial p \partial p' & \partial^2 L / \partial (p')^2 \end{pmatrix}$$

is positive definite.

Theorem 4.1. *A minimum satisfying the Euler–Lagrange equations is given by the p solving the differential equation $p' = \pm p \sqrt{c/(ep) + 2 \log p}$, for some suitable c .*

Proof. Calculations give

$$H = \frac{1}{p} \begin{pmatrix} \kappa^2 + 1 & -\kappa \\ -\kappa & 1 \end{pmatrix},$$

where $\kappa = p'/p$. This is easily seen to be a positive definite matrix; the eigenvalues are given by

$$\frac{1}{p} \left[\frac{1}{2}(2 + \kappa^2) \pm \sqrt{\frac{1}{4}(2 + \kappa^2)^2 - 1} \right],$$

which are positive.

Then (7), after some elementary algebra and differentiation, leads to the differential equation (2), which we report here for convenience,

$$\frac{p''(\theta)}{p(\theta)} - \frac{1}{2} \left\{ \frac{p'(\theta)}{p(\theta)} \right\}^2 = 1 + \log p(\theta),$$

which is the same as (4). □

This differential equation has the solution derived in the previous section. It is interesting that the Euler–Lagrange equations are solved by precisely the same p solving (4).

It might be thought that we would need the constraint $\int p(\theta) d\theta = 1$, that is to consider $L(\theta, p, p') + \lambda(p)$, i.e. to include a Lagrange multiplier. However, any ensuing differences are covered by our note in Section 5.

5 Illustrations

Before proceeding with some illustrations and applications, it is important that we discuss two aspects of the proposed class of priors within the boundaries of *Objective Bayes*: i.e. uniqueness and invariance.

It is important to discuss uniqueness and flexibility associated with objective priors. It is widely acknowledged that a prior representing total ignorance is elusive, and it might not even be possible to obtain in principle, see Bernardo and Smith (1994). As a consequence, any prior distribution, objective or not, must out of necessity provide some knowledge about something, and this “something” is not necessarily unique. For example, given a particular problem, the corresponding objective prior over a given parameter space could be proper or improper; differentiable everywhere or not; convex; log-concave; etc. In other words, a prior can be objective and exhibit desirable features of choice without impinging on subjective components relating to information.

So while we will be introducing a Bayesian objective prior criterion, it does not lead to a unique prior, rather to a class of priors, where some desirable features may or may not be included. We believe that this level of flexibility is a point of strength of the proposed approach, making it adaptable to different scenarios, including those where model based priors do not work.

Another fundamental point of discussion about prior distributions and, in particular, objective prior distributions, is invariance. Indeed, Jeffreys’ rule to derive a prior distribution for the parameters of a given model is based on an invariance requirement,

in particular on invariance under one-to-one reparameterisations. Also, other common objective priors, such as reference priors, have been shown to be invariant and the same apply, for example, to the priors in Simpson et al. (2017).

Here we discuss invariance from two opposite perspectives: that it is not important, and that it is important. Before discussing this apparent contradiction, we need to point out that we define the objective prior by setting the scoring rule equal to a constant, that is $S(\theta, p(\theta)) = \text{constant}$, is invariant under location transformations.

The question is whether lack of invariance has any practical implications given that only one parameterisation will be used. Current model based objective procedures are bound to throw away some coherence properties to achieve invariance, see Kass and Wasserman (1996). However, our point is that there is no practical consequence of any relevance arising from the lack of invariance, given that, as mentioned above, a single parameterisation will be used. For example, in the case the chosen model is the normal density, one either considers the precision parameter or the variance parameter, not both. And whichever parameterisation is used, our claim is that the corresponding objective prior is adequate for the purpose to which it has been assigned.

The above points of discussion are concerned with the perspective that invariance is not important. To consider the opposite point of view let us assume that there is a canonical parameterisation for the model $f(\cdot|\theta)$. Certainly, for most models the set of parameters for which priors would be assigned is obvious. For example, the exponential family has

$$f(x|\theta) = h(x) \exp \left\{ \sum_{j=1}^p \theta_j T_j(x) - A(\theta) \right\}$$

with $\theta = (\theta_1, \dots, \theta_p)$ being the canonical parameterisation. We can then define the canonical objective prior for statistical model $f(\cdot|\theta)$, $\theta \in \Theta$, as $p_{\Theta}(\theta) = \prod_{j=1}^p p_{\Theta_j}(\theta_j)$, where $\Theta = \otimes_{j=1}^p \Theta_j$. Then, any transformed prior can be obtained in the usual way involving variable transformations; that is $p(\phi) = |J| p_{\Theta}(\theta(\phi))$, where J is the Jacobian matrix for the transformation.

To illustrate the proposed method we consider three common parameter spaces in the unidimensional case and one in the multidimensional case. In particular, we consider the space for a parameter representing a probability, that is $\Theta = (0, 1)$, the space $\Theta = (0, \infty)$, usually representing the support of scale parameters, and the support for (location) parameters $\Theta = (-\infty, +\infty)$. The space $\Theta = (-M, M)$, for some finite M , has been illustrated in Section 3. As an illustration for the multidimensional case, we consider the bidimensional parameter space $(0, \infty)^2$.

5.1 One dimensional parameter space

The aim here is to solve (7) for particular motivated choices of $(c, u(0))$, equivalently, $(c, p(0))$ or $(p'(0), p(0))$. In fact, the solutions to the Euler–Lagrange equations are many, and the choice of the two constants $(c, u(0))$ will then determine a unique solution.

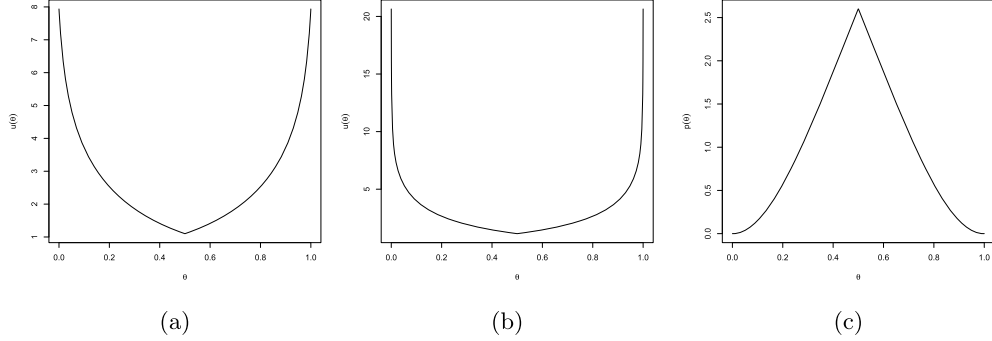


Figure 2: Plot of numerical solutions for u , panels (a) and (b) with, respectively, $u(\frac{1}{2}) = 1.1$ and $u(\frac{1}{2}) = 1.14$; plot of the normalised prior density obtained by setting $u(\frac{1}{2}) = 1.14$ (c).

Now there is the flat solution for all Θ in (4) given by $p(\theta) \propto 1$. This is achieved by setting $c = 2$ and $u(0) = 0$. However, in each of the settings of Θ considered we can find alternate priors with particular features. So, e.g. for $\Theta = (0, 1)$ we ask that $p(0) = p(1) = 0$ and for $\Theta = (0, \infty)$ we ask that p is convex and decreasing.

Case $\Theta = (0, 1)$ Here we consider the u function, recall $p \propto e^{-u}$, and so for $p(0) = p(1) = 0$ we require $u(0) = u(1) = \infty$. For additional symmetry, we can take $u(\frac{1}{2}) > 0$ and taking $c = 2$ as the extremal value, we have

$$\begin{aligned} u' &= \sqrt{2}\sqrt{e^u - 1 - u} & \theta > \frac{1}{2}, \\ u' &= -\sqrt{2}\sqrt{e^u - 1 - u} & \theta < \frac{1}{2}. \end{aligned}$$

Note there is a discontinuity in the derivative of u at $\theta = \frac{1}{2}$. As $u(\frac{1}{2})$ increases it is that $u(0)$ and $u(1)$ got to ∞ . A plot of u is given in Figure 2a for $u(\frac{1}{2}) = 1.1$ and for $u(\frac{1}{2}) = 1.14$ in Figure 2b. For these figures we used a grid of 1000 either side of $\theta = \frac{1}{2}$ to obtain the numerical solutions. In the latter case, the corresponding density for p is presented in the right plot of Figure 2c.

It is also possible to obtain a prior that mimics Jeffreys'; that is, a distribution that has spikes at $\theta = 0$ and $\theta = 1$ with the lowest value at $\theta = \frac{1}{2}$. This is done by simply inverting u : i.e. set $u = -u$ in the above prior.

Case $\Theta = (0, \infty)$ For a prior defined on the space $(0, \infty)$ we require a specific shape property for p (convex and decreasing) and then take extremal values for the c and $u(0)$. This property is common to most objective priors on $(0, \infty)$. Thus, since $p' < 0$ we require $u' > 0$ and so $u' = \sqrt{ce^u - 2(1+u)}$, and for u' to exist for all u we must have $c \geq 2$. Thus, as an extremal value, we take $c = 2$.

For the prior to be convex we require $p'' \geq 0$. Now $p' = -u'p$ implying $p'' = p\{(u')^2 - u''\}$. Therefore, p is convex when $(u')^2 \geq u''$. From (6) and (5) we have $(u')^2 = ce^u - 2(1+u)$ and $u'' = \frac{1}{2}ce^u - 1$, and hence we are interested in the $u(0)$ for

which $c = 2 \geq 2(1 + 2u)e^{-u}$ for all u ; i.e. $(1 + 2u)e^{-u} \leq 1$ for all u . Given that the function $(1 + 2u)e^{-u}$ is maximum, with a value of 1.31, at $u = \frac{1}{2}$, and u is increasing, since $u' > 0$, we need $u(0) > \frac{1}{2}$ and $(1 + 2u(0))e^{-u(0)} = 1$, again as an extremal value. Solving this gives a value for $u(0)$ of approximately 1.31.

In the next result we show that u' is bounded away from 0, and this will have important consequences for the properness of p .

Lemma 1. *It is that u' is bounded away from 0.*

Proof. To show this we need to show that $e^u - 1 - u$ is bounded away from 0 for $u \geq u(0)$. This follows trivially since $e^u - 1 - u \geq \frac{1}{2}u^2 \geq \frac{1}{2}u(0)^2$. \square

The result of Lemma 1 has also the implication that p is a proper density function. To show this, we require Gronwall's inequality (Gronwall, 1919). This inequality states that, if f and g are real valued functions on $\Theta = (0, \infty)$, g is differentiable on $\text{int}(\Theta)$, and $g'(t) \leq g(t) f(t)$ for all $t \in \Theta$ then $g(t) \leq g(0) \exp\{\int_0^t f(s) ds\}$.

Lemma 2. *If $p(0) < \infty$, it is that $p(\theta) \leq p(0) e^{-\epsilon\theta}$, and hence p is proper.*

Proof. Since $p' = -u'p$ and we have $u' \geq \epsilon$ for some $\epsilon > 0$, it is that $p' \leq -\epsilon p$. From Gronwall's lemma, with $f(t) = -\epsilon$ and $g = p$, we have that

$$p(\theta) \leq p(0) \exp\left\{-\int_0^\theta \epsilon ds\right\},$$

and hence the proof is complete. \square

To have a graphical image, in Figure 3a we plot the prior using the approximation available via a numerical solution to the differential equation for p . Note that this is the unnormalised p .

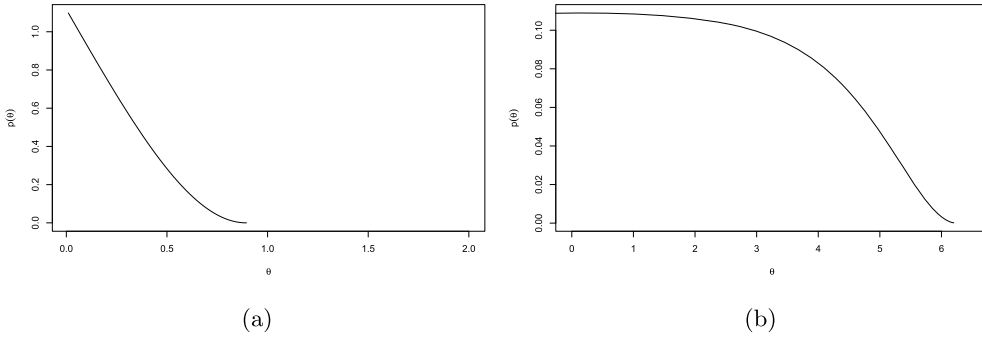


Figure 3: Plots of the prior on $(0, \infty)$ with $c = 2$ and $u(0) = 1.31$ (a). Plot of the positive half of the prior on $(-\infty, \infty)$, obtained by symmetrising the first one and make it differentiable at the origin (b). Both densities are normalised.




Figure 4: Surface plot of the prior for the bidimensional parameter space $\Theta = (0, \infty)^2$.

Case $\Theta = (-\infty, +\infty)$ A solution here is a symmetric version of the case $\Theta = (0, +\infty)$, which will represent a proper prior. On the other hand, if we ask that p is smooth at the origin, i.e. $p'(0) = 0$, then we need $u'(0) = 0$ and hence we must take $(c, u(0))$ to satisfy $ce^{u(0)} = 2 + 2u(0)$. If now we take $c = 2$, then $u(\theta)$ is a constant, resulting in a flat (improper) prior for p .

For a proper prior here one could take $u(0)$ to be small, say $u(0) = 0.01$ and then to take $c = 2\{1 + u(0)\}/\exp\{u(0)\}$. We computed numerically the right side; i.e. the $(0, \infty)$ side, for p , which is shown in Figure 3b.

To make the value of $u(0)$ more diverted, we could equally set a motivated choice for $p(0) = 1/(\sigma\sqrt{2\pi})$, corresponding to a normal density with zero mean and variance σ^2 .

5.2 Multidimensional parameter space

Here we consider the bidimensional parameter space $\Theta = (0, \infty)^2$, thus $\boldsymbol{\theta} = (\theta_1, \theta_2)$. As discussed in Section 3, the prior will have the solution $p(\boldsymbol{\theta}) \propto \exp\{-u(\boldsymbol{\theta})\}$, with the partial derivatives as in (5) and additional parameters c_1 and c_2 . It is sensible to expect the marginal priors $p_1(\theta_1)$ and $p_2(\theta_2)$ to be as the prior for the $\theta = (0, \infty)$ case discussed above. We can then set $u(0, 0) = (1.31, 1.31)$ and $c_1 = c_2 = 2$. The bidimensional prior is obtained using the bivariate Taylor expansion and, by applying the algorithm described in Appendix A of the Supplementary Material (Leisen et al., 2019), we obtain the prior for $(0, \infty)^2$, which surface is plotted in Figure 4.

6 Applications

The proposed class of priors is illustrated through a simulation study and the discussion of some practical implementation. Besides two initial simple examples, this section is dedicated to show the behaviour of the prior for cases where improper priors cannot be used (i.e. mixture models and model comparison via Bayes factors), while simula-

tion studies and a real data example are illustrated, respectively, in Appendix B and Appendix C in the Supplementary Material.

Although we do not have an explicit form for $p(\theta)$, we can use (5) to calculate it numerically quite easily. In particular, if we know $p(\theta)$ then we calculate $p(\theta + \delta\theta)$ for small $\delta\theta$, hence setting up the possibility of a posterior estimation process via Metropolis–Hastings sampling. The algorithm employed is detailed in Appendix A of the Supplementary Material.

To be specific, suppose we are currently at θ and the proposal value is θ' . The acceptance probability is

$$\alpha = \min \left\{ 1, \frac{l(\theta') p(\theta') q(\theta|\theta')}{l(\theta) p(\theta) q(\theta'|\theta)} \right\}, \quad (8)$$

where $l(\theta)$ is the likelihood function, and $q(\theta'|\theta)$ is the proposal density. The evaluation of $p(\theta')/p(\theta)$ in (8) does not represent any particular challenge. In fact, we have $p(\theta')/p(\theta) = \exp\{-[u(\theta') - u(\theta)]\}$, where u is the solution of the differential equation

$$u' = \sqrt{ce^u - 2(1 + u)}. \quad (9)$$

Equation (9) allows us to evaluate $u(\theta') - u(\theta)$ numerically, via

$$u(\theta') \approx u(\theta) + (\theta' - \theta)u'(\theta) + \frac{(\theta' - \theta)^2}{2}u''(\theta) + \frac{(\theta' - \theta)^3}{6}u'''(\theta),$$

where the derivatives are $u'(\theta) = \sqrt{ce^{u(\theta)} - 2(1 + u(\theta))}$, $u''(\theta) = \frac{1}{2}ce^{u(\theta)} - 1$, and $u'''(\theta) = \frac{1}{2}ce^{u(\theta)}\sqrt{ce^{u(\theta)} - 2(1 + u(\theta))}$, and so on. Depending on how far θ' is from θ we can either use the direct approximation just given or otherwise get from θ to θ' using smaller step sizes.

The frequentist performances are illustrated in a thorough simulation study, which is presented in Appendix B of the Supplementary Material. There, we also show the complete analysis for two single i.i.d. samples, that is where we obtain data from a Poisson distribution and from a normal density with unknown mean and known variance.

6.1 Mixture models

In this section we discuss the application of the proposed method to a scenario where objective priors have been notoriously challenging, namely mixture models. Due to their flexibility, mixtures of probability distributions allow models suitable for complex data by building on simple components. As an example, consider a mixture of normal densities,

$$f(x) = \sum_{j=1}^k \omega_j N(x|\mu_j, \sigma_j^2), \quad (10)$$

where k is a positive integer, including ∞ , and the $(\omega_j, \mu_j, \sigma_j)$ are the collection of parameters. Even under the scenario when k is known, the reference prior for model (10)

has yet to be derived, and Jeffreys prior can only be obtained under specific conditions; see Grazian and Robert (2018). Furthermore, this type of model is subject to other issues related to non-identifiability and unbounded likelihoods, among others. The issues mainly arise from the fact that improper priors may not be appropriate as we might not observe outcomes from every component of the mixture (Titterington et al., 1985). For example, Grazian and Robert (2018) show that Jeffreys prior is suitable for mixtures of normal densities only in certain circumstances; that is, when the unknown parameters are the weights. If the unknown parameters are the means or the variances, then using Jeffreys prior may lead to improper posteriors. In particular, if the unknown parameters are the means only, proper posteriors exist only when the number of mixture components is at most two; while, if the unknown parameters are the variance, or the mean and the variances, then Jeffreys prior is not suitable for inference. The above issues can be generalised to apply to any type of mixture model.

Given that the objective prior we propose is proper, it allows to make inference on the parameters of a mixture density as the yielded posteriors will be proper. As an illustration, we consider a mixture of three normal densities, where the weights and the parameters of the components are unknown. In particular, we sample from the following model

$$f(y|\omega_1, \omega_2, \omega_3, \mu_1, \mu_2, \mu_3, \sigma_1^2, \sigma_2^2, \sigma_3^2) = \sum_{i=1}^3 \omega_i N(\mu_i, \sigma_i^2), \quad (11)$$

with weights $\omega_1 = 0.25$, $\omega_2 = 0.35$ and $\omega_3 = 0.40$, means $\mu_1 = -3.5$, $\mu_2 = 0$ and $\mu_3 = 2.5$, and variances $\sigma_1^2 = 0.5$, $\sigma_2^2 = 0.1$ and $\sigma_3^2 = 1.2$. Note that we have chosen mixture components that are reasonably distant, so not to be forced to impose any constraint to overcome identifiability, as the focus of the paper is not in this sense. However, the implementation of constraints in that sense is straightforward. For the parameters we have chosen prior independence, where each prior is the prior on the space $(0, 1)$ for the weights, on the space $(-\infty, \infty)$ for the means and on the space $(0, \infty)$ for the variances, in agreement with Section 5. The prior on $(-\infty, \infty)$ is the symmetrised version from $\Theta = (0, \infty)$. To ensure properness of the priors for the means and variances, we have set $c = 2$ and $u(0) = 1.31$ (as discussed in Section 5).

We have performed the analysis on two data sets of size $n = 100$ and $n = 250$. The whole details, including histograms of the sample data, description of the algorithm implemented, as well as convergence diagnostics are reported in Appendix D of the Supplementary Material.

6.2 Model comparison

Another simple case where objective priors are problematic is in model comparison (or selection) via Bayes factors. So, if we wish to compare model $M_1 = \{f_1(x|\boldsymbol{\theta}_1), p_1(\boldsymbol{\theta}_1)\}$ to model $M_2 = \{f_2(x|\boldsymbol{\theta}_2), p_2(\boldsymbol{\theta}_2)\}$, where both $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are vector of parameters with some elements not in common, then the Bayes factor

$$BF_{12} = \frac{\int f_1(x|\boldsymbol{\theta}_1) p_1(\boldsymbol{\theta}_1) d\boldsymbol{\theta}_1}{\int f_2(x|\boldsymbol{\theta}_2) p_2(\boldsymbol{\theta}_2) d\boldsymbol{\theta}_2},$$

is, in general, meaningful if the priors assigned to non-common parameters are proper. If not, then the arbitrary multiplicative constant up to which they are defined do not cancel and the Bayes factor depends on an arbitrary constant. Solutions to the issue have been proposed, see, for example, O’Hagan (1995) and Berger and Pericchi (1996), however, the resulting procedures are still quite tedious to implement and are limited to simple models. By and large the above issue stays; however, Berger et al. (1998) give an exception of the issue. Furthermore, in Dawid and Musio (2015) and Dawid et al. (2017) the authors propose to make use of homogeneous scoring rules that circumvent the problem of using improper priors on the parameters.

An example on the application of the proposed prior in model selection is discussed in Appendix E of the Supplementary Material. There, a Poisson distribution and a geometric distribution are compared, where the priors are, respectively, of a parameter space $(0, \infty)$ and $(0, 1)$. As propriety is necessary in this context, we have set $c = 2$ and $u(0) = 1.31$.

Nested models

When models under comparison are nested, there are particular considerations which are needed to be taken into account; see, for example, Consonni et al. (2013). The point is that a diffuse type prior for the larger model will end up lacking focus so that the mass assigned to the smaller model is too much. However, our argument is that if two nested models are under comparison, it is essential, at least from a coherent point of view, to center the larger prior on the fixed part of the smaller one. Let us elaborate.

Suppose $f(y|\theta)$ for $\theta \in \Theta_1$ is the larger model and the smaller one is given by $\theta \in \Theta_0$ where $\Theta_0 \subset \Theta_1$. Typically Θ_1 will be a higher dimension to Θ_0 and to get the latter from the former one fixes a particular value in the higher dimension. To make this concrete, let us consider Example 2.1 from Consonni et al. (2013), where $M_0 : f(y|\theta_0)$ is binomial(n, θ_0), with $\theta_0 = 1/4$ fixed, and $M_1 : f(y|\theta)$ is binomial(n, θ), for which a prior for θ , $p(\theta)$, is required. Given the nature of the comparison it is our argument that $p(\theta)$ must be centered on θ_0 .

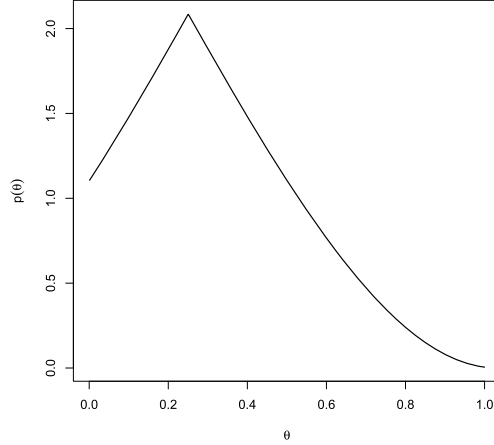
We can adapt quite easily the prior obtained in Section 4, the $\Theta = (0, 1)$, to be centered on $1/4$ rather than $\frac{1}{2}$. Without repeating the mathematics, we can take $u(1/4) = w$ and $c = 2$ and then

$$\begin{aligned} u' &= \sqrt{2}\sqrt{e^u - 1 - u} & \theta > 1/4, \\ u' &= -\sqrt{2}\sqrt{e^u - 1 - u} & \theta < 1/4. \end{aligned}$$

For the illustration of the prior $p(\theta)$, obtained numerically from u , in Figure 5 we took $w = 1.5$.

The proposed prior, centered at $\theta_0 = 0.25$, is compared with the intrinsic prior in Consonni et al. (2013), that is

$$p^I(\theta|b, t) = \sum_{x=0}^t \text{Beta}(\theta|b+x, b+t-x) \text{Bin}(x|t, \theta_0),$$

Figure 5: Numerical solution for normalized p with $w = 1.5$.

where $b = 1$ and $t = 8$. The intrinsic prior defined above is centered at $w\theta_0 + (1-w)\frac{1}{2}$, where $w = t/(2b+t)$, and has behaviour similar to the one in Figure 5, giving the Intrinsic Bayes Factor in favour of M_1

$$BF_{10}^I = \sum_{x=0}^8 \frac{B(1+x+y, 1+t-x+n-y)}{B(1+x, 1+t-x)\theta_0^y(1-\theta_0)^{n-y}},$$

where $n = 12$. A relatively small sample size allows to better capture differences in the performance of the two priors. Figure 6 shows the posterior probability for M_1 , i.e. $P(M_1|y) = (1 + 1/BF_{10}^I)^{-1}$. The priors yield model probabilities that are similar; in fact in both cases the lowest point is at $\theta = \theta_0$ and, the more θ moves away from θ_0 the higher the posterior probability for M_1 .

6.3 Further properties of the prior

It is important to illustrate how the choices of the constraints $u(0)$ and c impact the prior; in particular, in terms of mean, variance and tail behaviour.

As a general result, if we need to center the prior at a particular θ_0 , we can simply set

$$\begin{aligned} u' &= \sqrt{ce^u - 2(1+u)} & \theta > \theta_0, \\ u' &= -\sqrt{ce^u - 2(1+u)} & \theta < \theta_0. \end{aligned}$$

See Section 6.2.1. In some cases, for example, when nested models are compared, one wishes to have a prior distribution with tails that are as heaviest as it is possible; i.e. suitable for alternative priors. This is achieved by aiming to have a prior for which u' is as small as possible, thus we can set c to the smallest value for which we can solve the equation, which is $c = 2$, if $u(0) < 0$, and $c = 2(1+u(0))e^{-u(0)}$, if $u(0) > 0$.

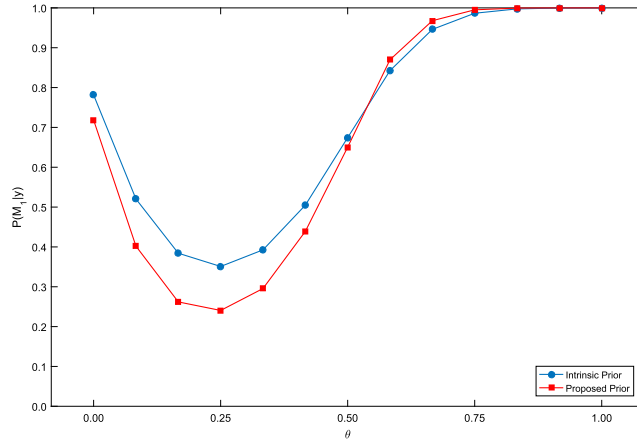


Figure 6: Small sample evidence for the Binomial example. The graph shows the posterior probability for model $f(y|\theta) = \text{Bin}(n = 12, \theta)$ using the proposed prior (squares) and the intrinsic prior (circles).

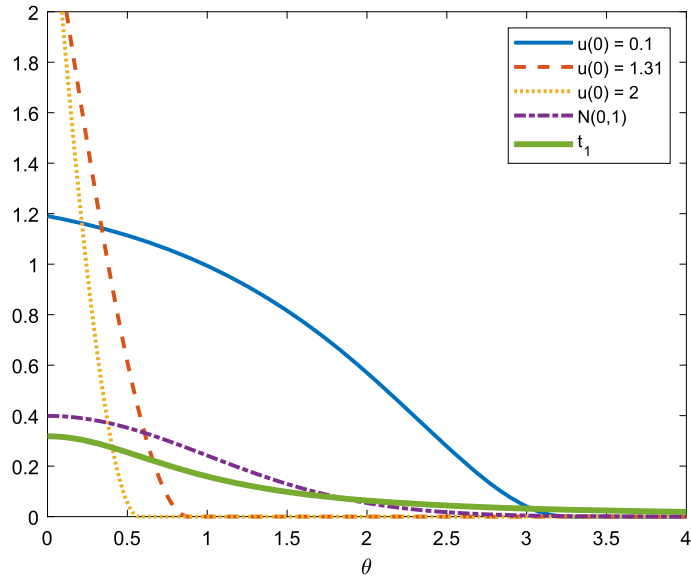


Figure 7: Comparison of the prior tail, with $c = 2$ and different values of $u(0)$, to the standard normal and the Student- t with 1 degree of freedom.

As an illustration, let us consider the case $\Theta = (-\infty, \infty)$, and compare the tail of the prior distribution to the tails of a normal density and a Student- t with 1 degree of freedom. For simplicity, we consider only the positive half of the real line. Figure 7 shows

the comparison of the prior based on scoring rules for $c = 2$ and for $u(0) = (0.1, 1.31, 2)$. When compare to both the standard normal and the Student- t with 1 degree of freedom, the proposed prior appears to have lighter tails for relatively large values of $u(0)$. In particular, the higher the value the quicker the prior drops towards 0. On the other hand, should we select a small value of the initial condition $u(0)$, then the proposed prior has a more gentle decrease to zero.

To understand the effect of the initial conditions on mean and variance of the prior, let us start by considering the case $\Theta = (0, \infty)$. We explore options where both $u(0)$ and c have increasing values; that is,

$$u(0) = (0.01, 0.05, 0.1, 0.5, 1, 1.31, 1.5, 2, 10), \quad c = (2, 2.01, 2.05, 2.1, 2.5, 3, 5, 10, 20, 50).$$

The prior mean and variance are reported in, respectively, Table 1 and Table 2. We see that when the conditions increase, mean and variance of the prior tend to 0. This is in

c	$u(0)$								
	0.01	0.05	0.1	0.5	1	1.31	1.5	2	10
2	3.54	1.82	1.07	0.23	0.09	0.05	0.04	0.02	5.59E-06
2.01	1.63	1.23	0.90	0.23	0.09	0.05	0.04	0.02	5.57E-06
2.05	0.89	0.77	0.64	0.21	0.08	0.05	0.04	0.02	5.51E-06
2.1	0.66	0.58	0.51	0.19	0.08	0.05	0.04	0.02	5.43E-06
2.5	0.27	0.26	0.23	0.12	0.06	0.04	0.03	0.02	4.46E-06
3	0.17	0.16	0.15	0.09	0.04	0.03	0.02	0.01	3.60E-06
5	0.07	0.07	0.06	0.04	0.02	0.02	0.01	0.01	2.34E-06
10	0.03	0.03	0.03	0.02	0.01	0.01	0.01	0.00	1.11E-06
20	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	5.59E-07
50	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.00E-07

Table 1: Mean of the prior distribution over $\Theta = (0, \infty)$ for different combinations of $u(0)$ and c .

c	$u(0)$								
	0.01	0.05	0.1	0.5	1	1.31	1.5	2	10
2	0.78	0.99	0.54	0.10	0.03	0.02	0.01	0.00	2.15E-08
2.01	0.88	0.63	0.45	0.10	0.03	0.02	0.01	0.00	2.14E-08
2.05	0.45	0.39	0.32	0.09	0.03	0.01	0.01	0.00	2.10E-08
2.1	0.33	0.29	0.25	0.08	0.03	0.01	0.01	0.00	2.05E-08
2.5	0.12	0.11	0.10	0.04	0.02	0.01	0.01	0.00	1.54E-08
3	0.07	0.06	0.06	0.03	0.01	0.01	0.00	0.00	1.14E-08
5	0.02	0.02	0.02	0.01	0.00	0.00	0.00	0.00	5.83E-09
10	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	2.03E-09
20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	7.69E-10
50	0.00	0.00	0.00	0.00	0.00	7.77E-05	5.87E-05	2.78E-05	2.07E-10

Table 2: Variance of the prior distribution over $\Theta = (0, \infty)$ for different combinations of $u(0)$ and c .

c	$u(0)$								
	0.01	0.05	0.1	0.5	1	1.31	1.5	2	10
2	10.50	3.44	1.59	0.17	0.04	0.02	0.01	0.01	2.16E-08
2.01	2.93	1.94	1.24	0.16	0.04	0.02	0.01	0.00	2.15E-08
2.05	1.21	0.97	0.75	0.14	0.04	0.02	0.01	0.00	2.10E-08
2.1	0.77	0.65	0.53	0.13	0.03	0.02	0.01	0.00	2.05E-08
2.5	0.21	0.19	0.17	0.06	0.02	0.01	0.01	0.00	1.55E-08
3	0.11	0.10	0.09	0.04	0.01	0.01	0.01	0.00	1.14E-08
5	0.03	0.03	0.02	0.01	0.01	0.00	0.00	0.00	5.84E-09
10	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	2.03E-09
20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	7.70E-10
50	0.00	0.00	0.00	0.00	0.00	8.07E-05	6.08E-05	2.85E-05	2.07E-10

Table 3: Variance of the prior distribution over $\Theta = (-\infty, \infty)$ for different combinations of $u(0)$ and c .

line with the tail behaviour discussed above, where the larger the value chosen for $u(0)$, the faster the prior drops to 0.

For the case $\Theta = (-\infty, \infty)$ we have chosen to center the prior at zero; although other values are easily attainable yielding similar results. Again, the prior variance (Table 3) decreases as we set the initial conditions to increasingly large values.

A final application of the proposed prior is for a Poisson regression model. The example is presented in Appendix C of the Supplementary Material, where we have worked with simulated data as well as real data. One objective of the simulation study is to show the robustness of the prior when nuisance covariates are added in the regression model; in fact, it can be seen that there is no noticeable impact on the size of the posterior credible intervals.

7 Discussion

In this paper we have introduced a new class of objective priors derived from scoring rules. A remarkable aspect is that we have been able to show that the same result can be achieved via the rigour of calculus of variations, by finding objective priors which solve the Euler–Lagrange equation for finding extremum to integrals of the type

$$\int L(\theta, p, p') d\theta.$$

If we can establish suitable choices of $L(\theta, p, p')$ which can be motivated and satisfy conditions for the existence of extremum, then new classes of objective prior can be sought. The case we have considered, which we can consider as a first step, is to use a combination of two well known measures of information in a prior density function; i.e.

$$L(\theta, p, p') = \frac{1}{2} \frac{p'(\theta)^2}{p(\theta)} + p(\theta) \log p(\theta).$$

The objective priors here defined have two desirable properties. The first is that they are somewhat detached from the choice of the sampling distribution and are dependent on the parameter space only. In other words, the information required to derive the prior is limited to the range of values that the quantity of interest can take.

The second property is that the prior can be proper. Besides the advantage of not having to check properness of the posterior, it allows to exploit the prior in scenarios where improper objective priors have been challenging. For example, as illustrated in Section 6.1, the proposed prior is used to estimate the means of a mixture of normal densities with three components. Another potential application, discussed in Section 6.2, is in model selection. In particular, the objective prior may be used to represent minimal information on the parameters that are not common to two models. In fact, the Bayes factor used to compare two models is, in general, sensitive to the proportionality constant of improper priors. While for common parameters the constant will cancel out, this is not the case if the parameter is either at the numerator or at the denominator of the ratio only. Hence, the necessity of having a proper prior assigned to this kind of parameters.

The simulation study, aimed to compare the frequentist performances of the proposed prior with the ones of the Jeffreys prior, has shown no appreciable differences, with the exception of a slightly larger Mean Squared Error (MSE) for the proposed prior; the last result is expected as it is a consequence of the smaller information used to define the proposed prior in comparison with any model based objective prior.

Future and ongoing work involves using the Fisher information alone; i.e. to minimize $\int (p')^2/p dp$ subject to p having certain constraints; for example, a zero mean or a specified variance or being log-concave. The mathematical results here would be able to provide explicit solutions including a class of non-local prior distributions (Johnson and Rossell, 2010).

Supplementary Material

On a Class of Objective Priors from Scoring Rules. Supplementary Material (DOI: [10.1214/19-BA1187SUPP](https://doi.org/10.1214/19-BA1187SUPP); .pdf). Supplement to “On a Class of Objective Priors from Scoring Rules”. The supplementary material contains the Appendixes A, B, C and D of the paper.

References

- Berger, J. O. (2006). “The case for objective Bayesian analysis.” *Bayesian Analysis*, **1**, 1–17. MR2221271. doi: <https://doi.org/10.1214/06-BA115>.
- Berger, J. O., Bernardo, J. M. and Sun, D. (2009). “The formal definition of reference priors.” *Annals of Statistics*, **37**, 905–938. MR2502655. doi: <https://doi.org/10.1214/07-AOS587>. 1, 2, 3
- Berger, J. O., Bernardo, J. M. and Sun, D. (2012). “Objective priors for discrete parameter spaces.” *Journal of the American Statistical Association*, **107**, 636–648. MR2980073. doi: <https://doi.org/10.1080/01621459.2012.682538>. 3

- Berger, J. O., Bernardo, J. M. and Sun, D. (2015). “Overall objective priors (with discussion)”. *Bayesian Analysis*, **10**, 189–221. MR3420902. doi: <https://doi.org/10.1214/14-BA915>. 3
- Berger, J. O. and Pericchi, L. R. (1996). “The intrinsic Bayes factor for model selection and prediction.” *Journal of the American Statistical Association*, **91**, 109–122. MR1394065. doi: <https://doi.org/10.2307/2291387>. 3, 16
- Berger, J. O., Pericchi, L. R., and Varshavsky, J. (1998). “Bayes factors and marginal distributions in invariant situations.” *Sankhya*, **60**, 109–122. MR1718789. 16
- Berger, J. O. and Strawderman, W. (1993). “Choice of hierarchical priors: admissibility in estimation of normal means.” *Technical report*, 93-34C, Purdue University, Dept. of Statistics. MR1401831. doi: <https://doi.org/10.1214/aos/1032526950>. 2
- Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian Theory*. John Wiley & Sons, Inc., Hoboken, NJ, USA. doi: [10.1002/9780470316870.ch1](https://doi.org/10.1002/9780470316870.ch1). 4, 9
- Bobkov, S. G., Gozlan, N., Roberto, C. and Samson, P. M. (2014). “Bounds on the deficit in the logarithmic Sobolev inequality.” *Journal of Functional Analysis*, **267**, 4110–4138. MR3269872. doi: <https://doi.org/10.1016/j.jfa.2014.09.016>. 8
- Box, G. E. P. and Tiao, G. C. (1973). *Bayesian Inference in Statistical Analysis*. Reading, MA: Addison-Wesley. MR0418321. 3
- Consonni, G., Forster, J. J. and La Rocca, L. (2013). “The Whetstone and the alum block: Balanced objective Bayesian comparison of nested models for discrete data.” *Statistical Science*, **28**, 398–423. MR3135539. doi: <https://doi.org/10.1214/13-STS433>. 16
- Consonni, G., Fouskakis, D., Liseo, B. and Ntzoufras, I. (2018). “Prior distributions for objective Bayesian analysis.” *Bayesian Analysis*, **13**, 627–679. MR3807861. doi: <https://doi.org/10.1214/18-BA1103>. 3
- Dawid, A. P. (1983). “Invariant prior distributions.” In *Encyclopedia of Statistical Sciences*, eds. S. Kotz and N. L. Johnson, New York: John Wiley. 3
- Dawid, A. P. and Musio, M. (2015). “Bayesian Model Selection based on Proper Scoring Rules (with discussion).” *Bayesian Analysis*, **10**, 479–521. MR3420890. doi: <https://doi.org/10.1214/15-BA942>. 16
- Dawid, A. P., Musio, M. and Columbu, S. (2017). “A Note on Bayesian Model Selection for Discrete Data Using Proper Scoring Rules.” *Statistics and Probability Letters*, **129**, 101–106. MR3688521. doi: <https://doi.org/10.1016/j.spl.2017.05.010>. 16
- Dey, D. K., Gelfand, A. E. and Peng, F. (1993). “Overdispersed generalized linear models.” *Technical report*, University of Connecticut, Dept. of Statistics. 2
- Fonseca, T. C. O. , Ferreira, M. A. R. and Migon, H. S. (2008). “Objective Bayesian analysis for the Student-T regression model.” *Biometrika*, **95**, 325–333. MR2521587. doi: <https://doi.org/10.1093/biomet/asn001>. 2

- Giummolè, F., Mameli, V., Ruli, E. and Ventura, L. (2019). “Objective Bayesian inference with proper scoring rules.” *TEST*, **28**, 728–755. MR3992136. doi: <https://doi.org/10.1007/s11749-018-0597-z>. 3
- Grazian, C. and Robert, C. P. (2018). “Jeffreys priors for mixture estimation: properties and alternatives.” *Computational Statistics and Data Analysis*, **121**, 149–163. MR3759204. doi: <https://doi.org/10.1016/j.csda.2017.12.005>. 3, 15
- Gronwall, T. H. (1919). “Note on the derivatives with respect to a parameter of the solutions of a system of differential equations.” *Annals of Mathematics*, **20**, 292–296. MR1502565. doi: <https://doi.org/10.2307/1967124>. 12
- Hartigan, J. A. (1964). “Invariant prior distributions.” *Annals of Mathematical Statistics*, **35**, 836–845. MR0161406. doi: <https://doi.org/10.1214/aoms/1177703583>. 3
- Hyvärinen, A. (2005). “Estimation of non-normalized statistical models by score matching.” *Journal of Machine Learning Research*, **6**, 695–709. MR2249836. 4, 5
- Ibrahim, J. G. and Laud, P. W. (1991). “On Bayesian analysis of generalized linear models using Jeffreys’s prior.” *Journal of the American Statistical Association*, **86**, 981–986. MR1146346. 2
- Kass, R. E. (1990). “Data-translate likelihood and Jeffreys’s rule.” *Biometrika*, **77**, 107–114. MR1049412. doi: <https://doi.org/10.2307/2336053>. 3
- Kass, R. E. and Wasserman, L. (1996). “The selection of prior distributions by formal rules.” *Journal of the American Statistical Association*, **91**, 1343–1370. 2, 3, 10
- Leisen, F., Villa, C. and Walker, S. G. (2019). “On a Class of Objective Priors from Scoring Rules. Supplementary Material.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1187SUPP>. 13
- Jaynes, E. T. (1957). “Information theory and statistical mechanics I, II.” *Physical Review*, **106**, 620–630; **108**, 171–190. MR0087305. 3
- Jaynes, E. T. (1968). “Prior probabilities.” *IEEE Transactions on Systems Science and Cybernetics*, **SSC-4**, 227–241. 3
- Jeffreys, H. (1946). “An invariant form for the prior probability in estimation problems.” *Proceedings of the Royal Society of London, Ser. A*, **186**, 453–461. MR0017504. doi: <https://doi.org/10.1098/rspa.1946.0056>. 2, 3
- Jeffreys, H. (1961). *Theory of Probability and Inference*, 3rd ed., Cambridge University Press, London. MR0745621. 1, 2, 3
- Johson, V. E. and Rossell, D. (2010). “On the use of non-local prior densities in Bayesian hypothesis tests.” *Journal of the Royal Statistical Society, Series B*, **72**, 143–170. MR2830762. doi: <https://doi.org/10.1111/j.1467-9868.2009.00730.x>. 21
- Natarajan, R. and McCulloch, C. E. (1995). “A note on the existence of the posterior distribution for a class of mixed models for binomial responses.” *Biometrika*, **82**, 639–643. MR1366287. doi: <https://doi.org/10.1093/biomet/82.3.639>. 2

- O'Hagan, A. (1995). "Fractional Bayes factors for model comparison." *Journal of the Royal Statistical Society, Series B*, **57**, 99–138. MR1325379. 3, 16
- Parry, M., Dawid, A. P. and Lauritzen S. (2012). "Proper local scoring rules." *Annals of Statistics*, **40**, 561–592. MR3014317. doi: <https://doi.org/10.1214/12-AOS971>. 2, 4, 5
- Rissanen, J. (1983). "A universal prior for integers and estimation by minimum description length." *Annals of Statistics*, **11**, 416–431. MR0696056. doi: <https://doi.org/10.1214/aos/1176346150>. 3
- Rubio, F. J. and Liseo, B. (2014). "On the independence Jeffreys prior for skew-symmetric models." *Statistics & Probability Letters*, **85**, 91–97. MR3157886. doi: <https://doi.org/10.1016/j.spl.2013.11.012>.
- Rubio, F. J. and Steel, M. F. J. (2018). "Flexible linear mixed models with improper priors for longitudinal and survival data." *Electronic Journal of Statistics*, **12**, 572–598. MR3769189. doi: <https://doi.org/10.1214/18-EJS1401>. 2
- Rustagi, J. S. (1976). *Variational Methods in Statistics*. Academic Press. MR0402569. 8
- Stone, M. (1972). "Strong Inconsistency from Uniform Priors." *Journal of the American Statistical Association*, **71**, 114–116. MR0415866. 4
- Stone, M. and Dawid, A. (1972). "Un-Bayesian implications of improper Bayes inference in routine statistical problems." *Biometrika*, **59**, 369–375. MR0431449. doi: <https://doi.org/10.1093/biomet/59.2.369>. 3
- Syversveen, A. R. (1998). "Noninformative Bayesian priors. Interpretation and problems with construction and applications." *Technical Report*. 4
- Simpson, D., Rue, H., Riebler, A., Martins, T. G. and Sørbye, S. H. (2017). "Penalising model component complexity: a principled, practical approach to constructing priors." *Statistical Science*, **32**, 1–28. MR3634300. doi: <https://doi.org/10.1214/16-STS576>. 3, 10
- Sweeting, T. J. (2008). "On predictive probability matching priors." *IMS Collections: Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, eds. B. Clarke and S. Ghosal. **3**, 46–59. MR2459215. doi: <https://doi.org/10.1214/074921708000000048>. 3
- Sweeting, T. J., Datta, G. S. and Ghosh, M. (2006). "Nonsubjective priors via predictive relative entropy loss." *Annals of Statistics*, **34**, 441–468. MR2275249. doi: <https://doi.org/10.1214/009053605000000804>. 3
- Titterton, D., Smith, A. and Makov, U. (1985). *Statistical Analysis of Finite Mixture Distributions*. John Wiley, New York. MR0838090. 3, 15
- Vallejos, C. A. and Steel, M. J. F. (2013). *On posterior propriety for the Student-t linear regression model under Jeffreys priors*. arxiv:1311.1454. MR2949504. doi: <https://doi.org/10.1080/00949655.2011.563239>. 2

- Villa, C. and Walker, S. G. (2015). “An objective approach to prior mass functions for discrete parameter spaces.” *Journal of the American Statistical Association*, **110**, 1072–1082. MR3420685. doi: <https://doi.org/10.1080/01621459.2014.946319>. 3
- Welch, B. L. and Peers, H. W. (1963). “On formulae for confidence points based on integrals of weighted likelihoods.” *Journal of the Royal Statistical Society, Series B*, **35**, 318–329. MR0173309. 3
- Yang, R. and Chen, M. H. (1995). “Bayesian analysis for random coefficient regression models using noninformative priors.” *Journal of Multivariate Analysis*, **55**, 283–311 MR1370405. doi: <https://doi.org/10.1006/jmva.1995.1080>. 2
- Zellner, A. and Min, C. (1993). “Bayesian analysis model selection and prediction.” In *Physics and Probability: Essays in honour of Edwin T. Jaynes*, eds. W. T. Grandy, Jr. and P. W. Milonni, Cambridge, U.K.: Cambridge University Press. MR1992316. doi: <https://doi.org/10.1017/CB09780511790423>. 3

Acknowledgments

We thank the Editor, the Associate Editor, and two Reviewers for insightful comments on a previous version of the paper. We also extend our thanks to Guido Consonni for feedback on an earlier draft of the paper. Fabrizio Leisen was supported by the European Community’s Seventh Framework Programme [FP7/2007-2013] under grant agreement no: 630677. The third author is partially supported by NSF grant DMS 1612089.