

Accepted Manuscript

On the proper treatment of improper distributions

Bo H. Lindqvist, Gunnar Taraldsen

PII: S0378-3758(17)30165-9

DOI: <https://doi.org/10.1016/j.jspi.2017.09.008>

Reference: JSPI 5594

To appear in: *Journal of Statistical Planning and Inference*



Please cite this article as: Lindqvist Bo H., Taraldsen G., On the proper treatment of improper distributions. *J. Statist. Plann. Inference* (2017), <https://doi.org/10.1016/j.jspi.2017.09.008>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

On the proper treatment of improper distributions

Bo H. Lindqvist, Gunnar Taraldsen

*Department of Mathematical Sciences, Norwegian University of Science and Technology,
N-7491 Trondheim, Norway*

Abstract

The axiomatic foundation of probability theory presented by Kolmogorov has been the basis of modern theory for probability and statistics. In certain applications it is, however, necessary or convenient to allow improper (unbounded) distributions, which is often done without a theoretical foundation. The paper reviews a recent theory which includes improper distributions, and which is related to Renyi's theory of conditional probability spaces. It is in particular demonstrated how the theory leads to simple explanations of apparent paradoxes known from the Bayesian literature. Several examples from statistical practice with improper distributions are discussed in light of the given theoretical results, which also include a recent theory of convergence of proper distributions to improper ones.

Keywords: axioms of probability, Bayesian statistics, conditional law, Gibbs sampling, intrinsic Gaussian Markov random fields, marginalization paradox

1. Introduction

Bayes' formula forms the basis of Bayesian statistics. Suppose a parameter θ is of interest, and that we have data x which is supposed to give information about θ . The idea of Bayesian inference is to first express one's prior knowledge (some would call it *uncertainty*) of θ in the form of a *prior distribution*, commonly given in the form of a density function $\pi(\theta)$, and then combine this knowledge with the new knowledge provided by the data x . The

Email addresses: bo.lindqvist@ntnu.no (Bo H. Lindqvist),
gunnar.taraldsen@ntnu.no (Gunnar Taraldsen)

influence of θ on the data x is modeled by a *statistical model*, represented by the conditional density of the data given the parameter, $f(x|\theta)$. Note that $f(x|\theta)$ will sometimes be interpreted as the 'likelihood' of θ for a given observation x , in which case the function $\theta \mapsto f(x|\theta)$ will be the *likelihood function*.

Bayes' formula is used to express the updated information about θ obtained after x is observed, given in the form of the *posterior distribution*,

$$\pi(\theta|x) = \frac{f(x|\theta)\pi(\theta)}{f(x)}. \quad (1)$$

Here $f(x) = \int f(x|\theta)\pi(\theta)d\theta$ is the *marginal density* of x . Equation (1) is one version of Bayes' theorem.

The algorithm for calculation of posterior distributions given by (1) is well defined as long as $0 < f(x) < \infty$, and in this case it will always lead to a proper probability distribution in the sense that $\pi(\theta|x)$ is a non-negative function which integrates to 1.

The usual proof of Bayes' formula is restricted to the case where $\pi(\theta)$ is a probability density, where basic rules of probability are used in the derivation. As pointed out above we may however get a proper distribution as the output of the formula even if the prior $\pi(\theta)$ is not proper. This fact is the obvious excuse for using Bayes' formula for improper priors.

A natural question is now, why should one want to use improper distributions for θ ? In practice, improper distributions often result from the search for so-called *non-informative* priors. The most prominent example of such priors is the Jeffreys' prior, which is proportional to the square root of the determinant of the Fisher information matrix and has the key property of being invariant under reparameterizations.

Intuitively, a non-informative prior should be one that does not favor any parameter values above others, suggesting "flat" priors. In practice, this often means to use proper standard probability models such as normal, gamma or uniform distributions with (very) large variances. Taking the limit as the variance tends to ∞ is then most often the excuse for using an improper prior density which is constant over the complete parameter space. Such distributions may, however, not have the invariance properties required by Jeffreys' priors. It is in fact well-known that flat priors may be very informative on non-location parameters. We refer to Irony and Singpurwalla (1997) for an interesting discussion on non-informative and improper priors.

As already indicated, posterior distributions computed by Bayes' formula are proper probability distributions only under the condition of $0 < f(x) < \infty$. Standard Bayesian calculations are, however, made by just recognizing the proportionality

$$\pi(\theta|x) \propto f(x|\theta)\pi(\theta),$$

which can be used without loss of information in the case that Bayes' formula gives a proper distribution, but gives an improper posterior distribution in case $f(x)$ is not finite. The latter case, if ignored, may lead to misleading inferences, as discussed later in this paper.

Improper distributions also appear naturally in certain non-Bayesian analyses. Lindqvist and Taraldsen (2005) considered conditional sampling of data x given a sufficient statistic $T(x)$ for the unknown parameter θ , which has numerous applications in statistical inference (Casella and Berger, 2002). The key is that this conditional distribution is independent of the value of θ . The general idea of the conditional sampling method of Lindqvist and Taraldsen (2005) was to use this fact, but instead of fixing the value of θ , to let it be a random quantity with some suitable distribution. Under certain mild restrictions, this distribution can be freely chosen, often with improper distributions giving the most efficient methods. Improper distributions appear likewise as useful ingredients in fiducial statistics, see for example Taraldsen and Lindqvist (2013) and Taraldsen and Lindqvist (2015).

The purpose of the present paper is to review and discuss some important aspects of the use of improper priors in statistical practice. Some would say that no new theory is needed, since improper priors are just approximations to proper ones. As will be discussed in the paper, this is a too simple attitude. The literature on Bayesian statistics includes a lot of paradoxes and misleading conclusions due to improper priors and posteriors. There are, however, not a lot of theoretical treatments of proper versus improper distributions in the literature. Some exceptions are, e.g., Hartigan (1983), Chang and Pollard (1997) and the more recent paper by McCullagh et al. (2011).

Our point of departure will be the paper by Taraldsen and Lindqvist (2010) which has a slightly different view than the above references. The idea is here simply to allow infinite probabilities in Kolmogorov's axioms. While this implies that all random variables have infinite mass, all conditional distributions will be proper probability distributions under a certain crucial condition which turns out to be equivalent to the above mentioned condition of finite $f(x)$. Formally, this condition is the σ -finiteness of the random

quantity that is conditioned on, here x . Details will be given in Section 2 which reviews the theoretical results of Taraldsen and Lindqvist (2010).

The above idea is not new, however. We quote from Renyi (1962), motivating the introduction of improper distributions:

“One can indeed give an axiomatic theory of probability which matches the above-mentioned requirements. This theory contains the theory of Kolmogorov as a special case. The fundamental concept of the theory is that of conditional probability; it contains cases where ordinary probabilities are not defined at all.”

In a footnote he adds:

“The idea of such a theory is due to Kolmogorov himself; he, however, did not publish anything about it.”

The theory presented in Taraldsen and Lindqvist (2010) is in fact closely related to Renyi’s theory of conditional probability spaces (Renyi, 1970). This connection is studied in more detail in Taraldsen and Lindqvist (2016).

Having introduced the basic elements of the theory of Taraldsen and Lindqvist (2010) in Section 2, we proceed to Section 3 which discusses some consequences of the theory when applied to Bayesian statistics. In particular we investigate in some detail a so-called marginalization paradox presented by Stone and Dawid (1972). Section 4 is devoted to Gibbs sampling, where a possible pitfall is the fact that posteriors may be improper even if all full conditionals are proper. A recent theoretical paper on approximation of improper priors, Bioche and Druilhet (2016), is briefly reviewed in Section 5. This is an important paper giving precise conditions for convergence of proper priors to improper ones and for convergence of the corresponding posterior distributions. Section 6 discusses a class of improper models which is popular in spatial statistics. Some concluding remarks are finally given in Section 7.

2. The theoretical framework

2.1. The modified Kolmogorov axioms

As in Kolmogorov’s axioms we consider an abstract space Ω of outcomes, where events A are represented by subsets of Ω and where the family \mathcal{E} of events is assumed to be a σ -algebra. We next let the measurable space (Ω, \mathcal{E}) be equipped with a fixed law Pr with

- $\Pr(A) \geq 0$ for all $A \in \mathcal{E}$.
- $\Pr(A_1 \cup A_2 \cup \dots) = \Pr(A_1) + \Pr(A_2) + \dots$
whenever A_1, A_2, \dots are pairwise disjoint events.

However, where Kolmogorov adds the axiom $\Pr(\Omega) = 1$, we assume only

- $\Pr(\emptyset) = 0$,

and hence allow the case $\Pr(\Omega) = \infty$. Note that the above axioms are exactly the axioms of a positive measure from standard measure theory (Royden, 1968).

2.2. Random quantities

A random quantity X with values in a measurable space $(\Omega_X, \mathcal{E}_X)$, is identified with a measurable function $X : \Omega \rightarrow \Omega_X$, i.e., such that $(X \in A) \equiv \{\omega \in \Omega \mid X(\omega) \in A\}$ is an event in \mathcal{E} for any event A in \mathcal{E}_X . The law \Pr on Ω now induces the law \Pr_X of a random quantity X by defining

$$\Pr_X(A) = \Pr(X \in A) \text{ for } A \in \mathcal{E}_X.$$

Hence the joint law of a pair (X, Y) is determined by

$$\Pr_{X,Y}(A \times B) = \Pr((X, Y) \in A \times B) \text{ for } A \in \mathcal{E}_X, B \in \mathcal{E}_Y,$$

while marginal laws are found from

$$\Pr_X(A) = \Pr_{X,Y}(A \times \Omega_Y) \text{ for } A \in \mathcal{E}_X.$$

The random quantity Y is called σ -finite if the law \Pr_Y is σ -finite, i.e., if there exist events $E_1, E_2, \dots \in \mathcal{E}_Y$ with:

$$\Omega_Y = \cup_i E_i \text{ and } \Pr_Y(E_i) < \infty \text{ for } i = 1, 2, \dots$$

2.3. Conditional distributions

A key feature of our approach is that if Y is σ -finite, then we can define a *unique proper conditional probability*

$$\Pr^y(A) \equiv \Pr(A|Y = y)$$

as a function of y for each $A \in \mathcal{E}$. The following approach equals the standard approach for definition of conditional probabilities and expectation in ordinary probability theory.

For a given event A in \mathcal{E} , conditional probabilities should satisfy, for all $B \in \mathcal{E}_Y$:

$$\begin{aligned} \Pr(A \cap (Y \in B)) &= \int_B \Pr(A|Y = y) \Pr_Y(dy) \\ &= \int_B \Pr^y(A) \Pr_Y(dy). \end{aligned} \quad (2)$$

By the assumed σ -finiteness of the measure \Pr_Y , the Radon-Nikodym theorem (Royden, 1968) states exactly that the function $g(y) = \Pr^y(A)$ exists and is uniquely (a.e.) defined by the above.

Since the measure \Pr_Y must satisfy $\Pr_Y(Y \in B) = \int_B \Pr_Y(dy)$ for all $B \in \mathcal{E}_Y$, it is seen by letting $A = \Omega$ in (2) and using uniqueness of $g(y)$, that we have

$$\Pr^y(\Omega) = 1.$$

Under regularity conditions, which we will not pursue here, we may from this conclude that conditional laws \Pr^y can always be represented as proper probability distributions, as long as Y is σ -finite. If, on the other hand, Y is not σ -finite, then \Pr^y is not defined due to the requirement of σ -finiteness in Radon-Nikodym's theorem.

Having defined the conditional law \Pr^y on Ω , we now define the conditional distribution of a random quantity X given $Y = y$ for $A \in \mathcal{E}_X$ by

$$\Pr_X^y(A) = \Pr^y(X \in A).$$

2.4. A Bayesian statistical model

A Bayesian statistical model involves an observation, represented by a random quantity $X : \Omega \rightarrow \Omega_X$, and a random parameter θ , represented as a σ -finite random quantity $\Theta : \Omega \rightarrow \Omega_\Theta$. The law \Pr_Θ of Θ is then the prior distribution.

The conditional distribution of X given $\Theta = \theta$, i.e., $\{\Pr_X^\theta : \theta \in \Omega_\Theta\}$, defines in a consistent way a *statistical model*. This follows directly from the above approach since Θ is assumed to be σ -finite and since conditional distributions are always proper.

2.5. Implications of improper prior Pr_{Θ}

So far we have not specified the value of $Pr(\Omega)$. Suppose that Θ is σ -finite with $Pr_{\Theta}(\Omega_{\Theta}) = \infty$. We claim that this implies that Pr is σ -finite and $Pr(\Omega) = \infty$. To see this, suppose $A_1, A_2, \dots \in \mathcal{E}_{\Theta}$ are such that

$$\Omega_{\Theta} = \cup_i A_i \text{ and } Pr_{\Theta}(A_i) < \infty \text{ for } i = 1, 2, \dots$$

Then

$$\Omega = (\Theta \in \cup_i A_i) = \cup_i (\Theta \in A_i),$$

which implies that Pr is also by necessity improper and σ -finite. This follows since $Pr(\Omega) = Pr_{\Theta}(\Omega_{\Theta}) = \infty$ and $Pr(\Theta \in A_i) = Pr_{\Theta}(A_i) < \infty$.

Assume now that $Pr(\Omega) = \infty$. Then every random quantity has an improper law, since

$$Pr_X(\Omega_X) = Pr\{\omega : X(\omega) \in \Omega_X\} = Pr(\Omega) = \infty.$$

On the other hand, a random quantity X is not necessarily σ -finite, even if Pr is. Namely, let X take values 0 and 1. Then

$$\infty = Pr(X \in \Omega_X) = Pr(X = 0) + Pr(X = 1)$$

and at least one of these is necessarily equal to ∞ . Hence X is not σ -finite.

2.6. Bayesian posteriors

Recall that a Bayesian model is given by a σ -finite law Pr_{Θ} , the prior distribution, and an observation X with distribution Pr_X^{θ} . For an observation x , Bayesian inference considers the *posterior law*, i.e., the conditional law of Θ given $X = x$, which in our notation is Pr_{Θ}^x . This conditional distribution is well defined if X is σ -finite, in which case it is a *proper* probability distribution. On the other hand, if X is *not* σ -finite, then the posterior is not defined. Hence, in the current theory there is nothing such as an *improper posterior*!

3. Bayesian statistics and marginalization paradoxes

3.1. The absolutely continuous case

Random quantities are said to be absolutely continuous if they can be defined by densities with respect to Lebesgue measure, for example $Pr_{X,Y} \sim$

$f(x, y)$. The *marginal density* of X is then given by the density $f(x) = \int f(x, y)dy$, where $f(x) = \infty$ is a permitted value.

It is seen that X with density $f(x)$ is σ -finite according to the definition of Section 2.2 if and only if $f(x) < \infty$ (a.e.), and the approach of the previous section can be shown to lead to the Bayes' formula (1), which corresponds to \Pr_{Θ}^x in the notation of Section 2.3.

3.2. What may “go wrong” with improper distributions?

A prior model for $\theta = (\theta_1, \theta_2)$ in Bayesian statistics is commonly given on the form of a joint density $\pi(\theta_1, \theta_2) = \pi_1(\theta_1)\pi_2(\theta_2)$ of the pair (Θ_1, Θ_2) , where $\pi_1(\theta_1), \pi_2(\theta_2)$ are two non-negative, finite-valued functions. One would then say that the parameters are given “independent priors, with marginal priors $\pi_1(\theta_1)$ and $\pi_2(\theta_2)$ ”. In practice, one might have chosen one or both of the “marginal” priors $\pi_1(\theta_1)$ and $\pi_2(\theta_2)$ as improper ones. To be concrete, suppose $\pi_1(\theta_1)$ is a proper probability density, while $\pi_2(\theta_2)$ is improper, i.e., integrates to ∞ . As indicated above, it would be tempting to call $\pi_1(\theta_1)$ and $\pi_2(\theta_2)$ the *marginal densities* of Θ_1 and Θ_2 , respectively. But are they?

By the definition given in Section 2.2, the marginal density of Θ_1 is $\int \pi(\theta_1, \theta_2)d\theta_2 = \pi_1(\theta_1) \int \pi_2(\theta_2)d\theta_2$, which however equals $\pi_1(\theta_1) \cdot \infty$ since $\pi_2(\cdot)$ is improper. Since this equals ∞ whenever $\pi_1(\theta_1) > 0$, it follows that $\pi_1(\cdot)$ is *not* the marginal density of Θ_1 ! However, integrating instead with respect to θ_1 and recalling that $\pi_1(\cdot)$ was assumed to be a probability density, we find that the marginal density of Θ_2 is $\int \pi(\theta_1, \theta_2)d\theta_1 = \pi_2(\theta_2) \int \pi_1(\theta_1)d\theta_1 = \pi_2(\theta_2)$, showing that $\pi_2(\cdot)$ is indeed the marginal density of Θ_2 .

So which interpretation can we give of $\pi_1(\theta_1)$? Since the marginal density of Θ_2 , $\pi_2(\theta_2)$, is finite (although not proper), we conclude that Θ_2 is σ -finite. Hence it has meaning to condition on it, and it can be seen that $\pi_1(\theta_1)$ is the *conditional density* of Θ_1 given $\Theta_2 = \theta_2$, using the approach of Section 2.3. In particular this shows that even if Θ_1 is not σ -finite, and has an infinite marginal density, it has a well defined proper conditional distribution given the σ -finite random quantity Θ_2 .

3.3. A marginalization paradox (Stone and Dawid, 1972)

For given parameters $\Theta = \theta, \Phi = \phi$, let X and Y be independent and exponentially distributed with hazard rates, respectively, $\theta\phi$ and ϕ . Suppose the interest is in the ratio θ of the hazard rates, which suggests consideration of the ratio $Z = Y/X$.

Let the joint prior distribution of (Θ, Φ) be given by $\pi(\theta, \phi) = \pi(\theta) \cdot 1$, where $\pi(\theta)$ is proper. (Note that by the previous subsection, $\pi(\theta)$ is not the *marginal* density of θ .) The joint density of (X, Z, Θ, Φ) is readily obtained to be

$$f(x, z, \theta, \phi) = \theta \phi^2 x e^{-\phi x(\theta+z)} \pi(\theta). \quad (3)$$

Integration with respect to (θ, ϕ) shows that (X, Z) is σ -finite, and we readily get the marginal conditional distribution of Θ given (X, Z) to be

$$f(\theta|x, z) \propto \frac{\theta \pi(\theta)}{(\theta + z)^3}. \quad (4)$$

Since this does not depend on x , it is tempting to conclude that the right hand side of (4) is also the conditional distribution of Θ given $Z = z$, i.e.,

$$f(\theta|z) = f(\theta|x, z) \propto \frac{\theta \pi(\theta)}{(\theta + z)^3}. \quad (5)$$

Starting differently, by integrating out x in (3) and conditioning with respect to (θ, ϕ) (which is obviously σ -finite), we get

$$f(z|\theta, \phi) = \int_0^\infty \theta \phi^2 x e^{-\phi x(\theta+z)} dx = \frac{\theta}{(\theta + z)^2}.$$

Since this depends only on θ , one might suggest that $f(z|\theta) = \frac{\theta}{(\theta+z)^2}$. and from this obtain

$$f(\theta|z) \propto f(z|\theta) \pi(\theta) = \frac{\theta \pi(\theta)}{(\theta + z)^2}. \quad (6)$$

But (5) and (6) contradict each other! It is therefore not clear how to proceed if one wants to do inference on θ based on Z alone. This is an example of a *marginalization paradox*.

So what is the problem? Considering the approach of Section 2, the problem is that the marginal distribution of Z is not σ -finite, so one is not allowed to condition on it. The conclusion in (5) is hence not correct. In fact, neither is the distribution of Θ σ -finite, making the conclusion (6) incorrect as well.

The clue is that we have above, in fact twice, used the generally invalid result that

$$f(\theta|x, z) \text{ does not depend on } x \Rightarrow f(\theta|z) = f(\theta|x, z).$$

This is well-known to hold for probability distributions, but holds for improper distributions only provided (X, Z) and Z both have σ -finite distributions.

3.4. Marginalization paradox revisited

In order to understand better the mechanisms of the previous example, let us redefine the problem and let the prior of (θ, ϕ) be given by

$$\pi(\theta, \phi) = \pi(\theta)h(\phi),$$

where $h(\phi)$ may be proper or improper, while $\pi(\theta)$ is proper as before, unless otherwise stated below. Multiplying (3) by $h(\phi)$ and integrating with respect to ϕ we get

$$f(x, z, \theta) = \frac{\theta\pi(\theta)}{x^2} \int_0^\infty u^2 e^{-u(\theta+z)} h\left(\frac{u}{x}\right) du. \quad (7)$$

The joint marginal distribution of (z, θ) is obtained by integrating with respect to x , which gives

$$\begin{aligned} f(z, \theta) &= \theta\pi(\theta) \int_0^\infty u^2 e^{-u(\theta+z)} \left[\int_0^\infty \frac{1}{x^2} h\left(\frac{u}{x}\right) dx \right] du \\ &= \frac{\theta\pi(\theta)}{(\theta+z)^2} \int_0^\infty h(w) dw. \end{aligned} \quad (8)$$

Hence, if h is proper, then for inference about θ we may base ourselves on Z only and use the relation (6). This case corresponds to the typical frequentist approach for this example, where one concludes that the distribution of Z depends on the parameters only via θ . In the case where h is improper, we get however $f(z, \theta) = \infty$ in (8), and neither Z nor Θ are σ -finite.

Recalling that we want to make inference about θ , let us go back to (7). It follows that

$$f(\theta|x, z) \propto \frac{\theta\pi(\theta)}{(\theta+z)^3} \int_0^\infty w^2 e^{-w} h\left(\frac{w}{x(\theta+z)}\right) dw. \quad (9)$$

Setting $h \equiv 1$ leads to the right hand side of (5). Strange enough, choosing h to be the improper density $h(\phi) = 1/\phi$, it follows that

$$f(\theta|x, z) \propto \frac{\theta\pi(\theta)}{(\theta+z)^2} \quad (10)$$

in which case we would apparently not have a marginalization paradox. Still, however, Z is not σ -finite, so we cannot conclude that (10) equals $f(\theta|z)$. It is notable that Stone and Dawid (1972) explain this apparent absence of a

marginalization paradox by the fact that we now use the prior $1/\phi$ for Φ which is the common prior for a scale parameter. In our opinion this seems to be more like a coincidence since we here use an improper h under which Z is not σ -finite.

As another comment on (10), note that the pair (X, Z) is σ -finite even if we let $\pi(\theta)$ be the improper density $\pi(\theta) = 1/\theta$, while we keep $h(\phi) = 1/\phi$. This is seen by integrating (7) with respect to θ . Hence (10) is meaningful and leads after normalization to the posterior density for θ given by

$$f(\theta|x, z) = \frac{z}{(\theta + z)^2}. \quad (11)$$

It is interesting to note that the density (11) also appears as the optimal invariant confidence distribution for θ in a frequentist approach involving the observations X and Y , where θ is the parameter of interest. The argument follows Schweder and Hjort (2016), Chapter 5.

We close the present section by returning to the original assumption where $h(\phi) \equiv 1$ and $\pi(\theta)$ is unspecified, but proper. Consider a proper density which can be seen as an approximation to the improper $h(\phi)$, e.g.,

$$h_M(\phi) = \frac{1}{M}I(0 < \phi \leq M), \quad (12)$$

where $I(\cdot)$ is the indicator function and $M > 0$ is considered to be large. From (9) we get

$$\begin{aligned} f(\theta|x, z) &\propto \frac{\theta\pi(\theta)}{(\theta + z)^3} \int_0^{xM(\theta+z)} w^2 e^{-w} dw \\ &= \frac{\theta\pi(\theta)}{(\theta + z)^3} [2 - e^{-A}(A^2 + 2A + 2)], \end{aligned} \quad (13)$$

where $A = xM(\theta + z)$.

It is seen that the limit as M tends to infinity in (13) is consistent with (4). This is in fact a consequence of a general result in Bioche and Druilhet (2016) (see Section 5), since $h_M \rightarrow h$ in their approach with $h \equiv 1$. As seen from (13), the convergence of $f(\theta|x, z)$ as M tends to infinity is not uniform in the observations (x, z) . This point was made by Akaike (1980) in his discussion of certain marginalization paradoxes. More precisely, Akaike questioned the common interpretation of an improper prior distribution as a limit of proper prior distributions, and he argued that an improper prior can

more adequately be described as the limit of certain data adaptive proper prior distributions. He concluded that a prior distribution without data adaptability may produce poor inference due to a gross misspecification of the prior. We illustrate this point in Figure 1. It is seen that, even if we set M as a large number (here 500) in (12), the posterior distribution for Θ in (13) depends rather distinctly on the value of x , at least for small x . This is despite the non-appearance of x in the right hand side of (4). For higher values of x , like $x = 1$, the posterior is however indistinguishable from the one we get when M equals infinity. The figure also includes a corresponding plot of (6) which illustrates the difference between (4) and (6).

4. An example from Gibbs-sampling

4.1. Gibbs-sampling from improper posterior distribution

Hobert and Casella (1996) gave an example showing that the output from Gibbs sampling corresponding to an improper posterior distribution may still appear perfectly reasonable. The authors' advice is thus that before implementing a Gibbs chain one should check that the posterior is proper. For this it is important to note that propriety of the conditionals of a Gibbs chain does not imply that the full posterior is proper (see example below).

Gelfand and Sahu (1999) consider similar problems with Gibbs sampling, focusing on parameter identifiability and posterior propriety. In particular, they provide rather general results for propriety of posteriors in the case of GLMs. As a simple illustration they consider in an earlier technical report (Gelfand and Sahu, 1996) the following example.

Let $Y|\theta_1, \theta_2 \sim N(\theta_1 + \theta_2, 1)$ with improper prior $\pi(\theta_1, \theta_2) = 1$. Then the joint distribution of (Y, Θ_1, Θ_2) is

$$f(y, \theta_1, \theta_2) = (1/\sqrt{2\pi})e^{-(1/2)(y-\theta_1-\theta_2)^2},$$

leading to the marginal density of Y given as

$$\int \int f(y, \theta_1, \theta_2) d\theta_1 d\theta_2 = \infty.$$

Hence Y is not σ -finite, so the posterior $f(\theta_1, \theta_2|y)$ does not exist (or, is *improper*).




Figure 1: Marginalization paradox example. Let $\pi(\theta) = \exp(-\theta)$ and suppose $z = 1$ is observed. Solid line: the (normalized) density (13) with $x = 1, M = 500$ (which is indistinguishable from (4)). Dashed line: the (normalized) density (13) with $x = 0.001, M = 500$. Dotted line: the (normalized) density (10).

On the other hand, the pairs (Y, Θ_1) and (Y, Θ_2) are both σ -finite, so the following conditional distributions exist and are proper:

$$\Theta_1 | \theta_2, y \sim N(y - \theta_2, 1), \quad (14)$$

$$\Theta_2 | \theta_1, y \sim N(y - \theta_1, 1). \quad (15)$$

Thus Gibbs-sampling of pairs (θ_1, θ_2) for given y is possible. The question is, however, how the pairs (θ_1, θ_2) will behave, knowing that the posterior $f(\theta_1, \theta_2 | y)$ does not exist. Figure 2 shows a simulation from (14)-(15). The

large fluctuations seen in the plots are due to the impropriety of the joint posterior given y .




Figure 2: Gibbs chains for θ_1 (left) and θ_2 (right), drawn from (14) and (15), respectively.




Figure 3: Simulated values of $\delta = \theta_1 + \theta_2$ using (14) and (15).

4.2. The proper embedded posterior (Gelfand and Sahu, 1999)

Gelfand and Sahu observed, however, that if one makes a 1-1 transformation

$$(\theta_1, \theta_2) \rightarrow (\delta, \rho), \quad \text{where } \delta = \theta_1 + \theta_2,$$

then the distribution $\delta|y \sim N(y, 1)$ can be recovered in the Gibbs-sampling. Indeed, the plot of $\delta = \theta_1 + \theta_2$ from (14) and (15) (Figure 3), is apparently well-behaved. Gelfand and Sahu (1999) call $\delta|y$ the unique proper embedded posterior, regarding it as embedded within the improper posterior for (θ_1, θ_2) .

A critical remark is of course appropriate here in view of the previous theory. Since Y is not σ -finite it has apparently no meaning to consider $\delta|y$. On the other hand, we clearly have from (15) that $(\Theta_1 + \Theta_2|y, \theta_1) \sim N(y, 1)$, i.e., if we let $\rho \equiv \theta_1$, we have $(\delta|y, \rho) \sim N(y, 1)$. Thus Gelfand and Sahu's conclusion is similar to the one that we deemed to be incorrect in connection with the marginalization paradox. Namely, when Y is not σ -finite, then even if the density of $(\delta|y, \rho)$ does not depend on ρ , this is not the conditional density of δ given y .

The nice behavior of δ in the simulation can be explained theoretically as follows. Suppose that the prior distribution of (Θ_1, Θ_2) is given by $\pi(\theta_1, \theta_2) = g(\theta_1)$, where $g(\theta_1)$ is a proper density. Then under the transformation $(\theta_1, \theta_2) \rightarrow (\rho, \delta)$, we have

$$f(y, \rho, \delta) = (1/\sqrt{2\pi})e^{-(1/2)(y-\delta)^2} \cdot g(\rho). \quad (16)$$

In this model Y is clearly σ -finite (in fact, the marginal density of Y is the constant 1). Thus the posterior $\pi(\rho, \delta|y)$ exists and is given by (16). The marginal posterior of δ is hence

$$\delta|y \sim N(y, 1)$$

whatever be the density $g(\rho)$, as long as it is proper. Gelfand and Sahu (1999) let $g(\rho)$ correspond to $N(0, \tau^2)$ and let $\tau^2 \rightarrow \infty$, and concluded that also in the limit will have $\delta|y \sim N(y, 1)$. This is, however, not a valid conclusion since Y is now not σ -finite.

4.3. Using proper priors for both θ_1 and θ_2

A proper posterior for (θ_1, θ_2) can of course be achieved by giving (Θ_1, Θ_2) a proper prior. Assume for example that Θ_1 and Θ_2 are independent with

$\Theta_1 \sim N(0, \tau^2)$, $\Theta_2 \sim N(0, \kappa^2)$. Then, as shown in Gelfand and Sahu (1996),

$$\begin{aligned}\Theta_1 | \theta_2, y &\sim N\left(\frac{\tau^2}{1 + \tau^2}(y - \theta_2), \frac{\tau^2}{1 + \tau^2}\right), \\ \Theta_2 | \theta_1, y &\sim N\left(\frac{\kappa^2}{1 + \kappa^2}(y - \theta_1), \frac{\kappa^2}{1 + \kappa^2}\right).\end{aligned}$$

If τ and κ are large, then the trajectories of the Gibbs chains for θ_1 and θ_2 , respectively, will still tend to drift in a way similar to the behavior in Figure 2. Thus, if we use a proper but diffuse priors for Θ_1 and Θ_2 , the posteriors will be proper but will in practice be indistinguishable from those obtained under the corresponding limiting improper prior. As concluded by Gelfand and Sahu (1996), an implicit byproduct of this observation is the infeasibility of numerical sampling based diagnostics for propriety of posteriors. A similar conclusion is expressed by Hobert and Casella (1996).

5. Convergence of priors and posteriors (Bioche and Druilhet)

5.1. q -vague convergence of measures

Bioche and Druilhet (2016) propose a convergence mode for measures allowing a sequence of probability measures to have an improper limiting measure. They also study convergence of corresponding posterior distributions.

Technically the authors study the set of positive Radon-measures on the state space Ω_Θ , i.e., the set of positive measures Π which are finite on compact subsets of Ω_Θ . Noting that the output of Bayes' formula (1) is unchanged if $\pi(\theta)$ is multiplied by a constant, they define the equivalence relation $\Pi \sim \Pi'$ to mean that there is an $\alpha > 0$ such that $\Pi = \alpha\Pi'$. Their basic space of measures is then the corresponding quotient space, equipped with the quotient topology resulting from vague convergence of positive Radon measures. Convergence in this topology has been denoted as q -vague convergence. A similar quotient topology is introduced by Taraldsen and Lindqvist (2016).

A useful way of expressing the definition of q -vague convergence is the following: A sequence of positive Radon-measures $\{\Pi_n\}$ converges q -vaguely to Π if there exists a sequence $\{a_n\}$ such that

$$a_n \Pi_n \rightarrow \Pi \text{ (vaguely),}$$

(see, e.g., Billingsley (2008) for the definition of vague convergence).

From this definition it is not difficult to prove that for any improper distribution Π there is a sequence Π_n of proper distributions such that $\Pi_n \rightarrow \Pi$ (q -vaguely). In this case, the a_n given in the above definition tend to ∞ as n increases.

As an example, consider the proper distribution with density h_M given by (12). We claim that $h_M \rightarrow h$ (q -vaguely) as $M \rightarrow \infty$, where $h \equiv 1$. To see this we need to find constants a_M such that $a_M h_M \rightarrow h$ (vaguely) as $M \rightarrow \infty$, i.e., such that

$$\int a_M h_M(\phi) f(\phi) d\phi \rightarrow \int f(\phi) d\phi$$

for each continuous function with compact support. But this is clear by the dominated convergence theorem by letting $a_M = M$ for all $M > 0$.

Bioche and Druilhet (2016) also consider convergence of posterior densities. If $f(x|\theta)$ is the likelihood of the data x and $\pi(\theta)$ is the prior density, then they define the posterior distribution as the distribution with density in the equivalence class corresponding to $\pi(\theta|x) \propto f(x|\theta)\pi(\theta)$, thus allowing also improper posterior distributions.

Their main proposition on convergence of posteriors states that if for the priors we have $\pi_n \rightarrow \pi$ (q -vaguely), and if $\theta \mapsto f(x|\theta)$ is continuous, then the posteriors converge in the sense that $\pi_n(\cdot|x) \rightarrow \pi(\cdot|x)$ (q -vaguely). We have already seen an example in Section 3.4, where the posterior distribution (13) converges to (4) as M tends to infinity. Note, however, that while the question of uniform convergence in x was made a point in our example, this issue is not considered by Bioche and Druilhet (2016).

At first glance it seems that the above cited result on convergence of posteriors justifies the common excuse for using improper priors, namely that they are limits of proper priors and hence that the posteriors are limits of posteriors based on proper priors. However, we have already seen problems connected to such a view. Next we shall see another type of misinterpretation of improper limits of proper distributions, which in turn may give completely misleading results regarding posterior distributions.

5.2. The Jeffreys-Lindley paradox (Bioche and Druilhet)

Let $X|\theta \sim N(\theta, 1)$, and consider testing of the null hypothesis $H_0 : \theta = 0$ versus the alternative $H_1 : \theta \neq 0$. Suppose we have a prior distribution for θ given by

$$\pi(\theta) = \frac{1}{2}\delta_0 + \frac{1}{2}I(\theta \neq 0),$$

where δ_0 is a point mass at $\theta = 0$ and $I(\cdot)$ is the indicator function. This means that we have a prior belief of $1/2$ in H_0 , while the remaining probability $1/2$ is distributed according to Lebesgue measure on H_1 . A straightforward calculation gives

$$\pi(0|x) = \left(1 + \sqrt{2\pi}e^{x^2/2}\right)^{-1}$$

implying $\pi(0|x) \leq (1 + \sqrt{2\pi})^{-1} \approx 0.285$ whatever be the data x .

Using instead the proper prior measure

$$\pi_n(\theta) = \frac{1}{2}\delta_0 + \frac{1}{2}N(0, n^2)$$

we get

$$\pi_n(0|x) = \left(1 + \sqrt{\frac{1}{1+n^2}}e^{\frac{n^2 x^2}{2(1+n^2)}}\right)^{-1}.$$

But this converges to 1 as $n \rightarrow \infty$, in conflict with the above calculation which was based on an apparently equivalent argument using the limiting prior. The result has therefore been considered as a paradox.

The clue, as presented by Bioche and Druilhet (2016), is that while $N(0, n^2)$ converges q-vaguely to Lebesgue measure on the real line, the measure $\frac{1}{2}\delta_0 + \frac{1}{2}N(0, n^2)$ converges to $\frac{1}{2}\delta_0 \sim \delta_0$ and not to $\frac{1}{2}\delta_0 + \text{Lebesgue measure}$ which one might believe. This explains the paradox, noting that by the convergence result for posteriors, the limiting posterior is a point mass at 0 as well.

6. Intrinsic Gaussian Markov random fields (IGMRF)

Intrinsic conditional autoregressions (ICAR) are widely used in spatial statistics and dynamic linear models (Besag et al., 1991). These models are improper versions of conditional autoregressive models (CAR) as introduced as spatial models by Besag (1974). Important special cases of CAR and ICAR models are Gaussian Markov Random Fields (GMRF) and the intrinsic (improper) versions denoted IGMRF, see Rue and Held (2005) for a thorough treatment including applications.

As discussed by Lavine and Hodges (2012), the fact that the intrinsic models correspond to improper distributions, implies that care should be taken in their use and interpretation.

6.1. The first order random walk

Following Rue and Held (2005) we use this simple special case of an IGMRF to illustrate some of the main issues regarding IGMRF models.

Let $\mathbf{x} = (x_1, x_2, \dots, x_n)$ be the successive observations of a random walk, assuming independent increments

$$\Delta x_i = x_{i+1} - x_i \sim_{iid} N(0, \kappa^{-1}), \quad i = 1, 2, \dots, n-1.$$

The IGMRF model specifies the density of \mathbf{x} to be the density obtained from these increments (only), giving

$$\begin{aligned} f(\mathbf{x}|\kappa) &\propto \kappa^{(n-1)/2} \exp\left(-\frac{\kappa}{2} \sum_{i=1}^{n-1} (\Delta x_i)^2\right) \\ &= \kappa^{(n-1)/2} \exp\left(-\frac{\kappa}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x}\right). \end{aligned} \quad (17)$$

Here the *structure matrix* \mathbf{Q} (displayed in Rue and Held (2005), p. 96) is positive semi-definite, with exactly one eigenvalue equal to 0, implying that $f(\mathbf{x}|\kappa)$ is an improper density.

Statistical inference in models involving IGMRFs may involve making inference about the precision parameter κ . In a Bayesian analysis, one must typically assign to κ a hyperprior and work with (17) as a likelihood function. In this connection, Lavine and Hodges (2012) question the use of the constant $\kappa^{(n-1)/2}$ appearing in (17). In the following discussion let us replace (17) by

$$f(\mathbf{x}|\kappa) \propto c(\kappa) \exp\left(-\frac{\kappa}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x}\right), \quad (18)$$

thus making the appropriate choice of $c(\kappa)$ the main issue. As reported by Lavine and Hodges (2012), this choice has been discussed in several papers during the last two decades. Besag et al. (1991) in fact used $c(\kappa) = \kappa^{n/2}$, which was used by WinBUGS (Lunn et al. (2000)) until it was changed to $c(\kappa) = \kappa^{(n-1)/2}$ following derivations appearing in, e.g., Knorr-Held (2003) and Hodges et al. (2003).

Rue and Held (2005) justify the density (17) as follows: Consider first the 1-1 transformation $(x_1, x_2, \dots, x_n) \leftrightarrow (\Delta x_1, \dots, \Delta x_{n-1}, \bar{x})$ where \bar{x} is the average of the x_i . Assuming that \mathbf{x} is multivariate normal, $(\Delta x_1, \dots, \Delta x_{n-1})$ and \bar{x} are stochastically independent. We may hence write down the following proper density for \mathbf{x} , indexed by k for the purpose of later taking limits,

$$\tilde{f}_k(\mathbf{x}|\kappa) = f(\mathbf{x}|\kappa) \cdot \check{f}_k(\bar{x}) \quad (19)$$

Here $f(\mathbf{x}|\kappa)$ is the density (17) while $\check{f}_k(\bar{x})$ is normal with zero expectation and precision $\gamma_k > 0$. Suppose $\gamma_k \rightarrow 0$ as $k \rightarrow \infty$. We may then invoke Proposition 2.15 of Bioche and Druilhet (2016) to show that

$$\tilde{f}_k(\mathbf{x}|\kappa) \rightarrow f(\mathbf{x}|\kappa) \text{ as } k \rightarrow \infty \quad (20)$$

The interpretation of this is that the improper density (17) is the limit of a sequence of proper densities for \mathbf{x} . This derivation can also be interpreted as adding to the model for the Δx_i a prior specification for \bar{x} given in the form of a constant prior.

Lavine and Hodges (2012) point, however, to a problem with this conclusion, having to do with the non-uniqueness of marginals in cases involving improper distributions and related to our discussion in Section 2. To illustrate, essentially following Lavine and Hodges, we consider the modified 1-1 transformation of \mathbf{x} given as

$$(\Delta x_1, \dots, \Delta x_{n-1}, \bar{x} \Delta x_1).$$

It follows by the ordinary transformation formula (involving a Jacobi-determinant), starting from (19), that we have

$$\tilde{f}_k(\mathbf{x}|\kappa) = f(\mathbf{x}|\kappa) \cdot \check{f}_k(\bar{x}/\Delta x_1) \cdot \frac{1}{|\Delta x_1|}.$$

Again, letting $\gamma_k \rightarrow 0$, we get

$$\tilde{f}_k(\mathbf{x}|\kappa) \rightarrow f(\mathbf{x}|\kappa) \cdot \frac{1}{|\Delta x_1|} \text{ as } k \rightarrow \infty,$$

thus giving a limit different from (20).

Lavine and Hodges (2012) conclude that essentially all the arguments given in the literature for the value of the constant $c(\kappa)$ in some way are flawed. Their conclusion is therefore that any value of this constant may do. This is of course also in accordance with the previous section where the quotient topology for distributions was used, and where improper (as well as proper) distributions were identified with equivalence classes only.

Having said this, there seem to be good reasons to use the form (17). It follows from Rue and Held (2005), p. 90-91, who considered a more general case, that (17) when restricted to \mathbf{x} such that $\bar{x} = \mu$, is the conditional density of \mathbf{x} given $\bar{x} = \mu$. Here μ can be any real number, but it seems that $\mu = 0$ is commonly used. Furthermore, the specification of μ enables one to simulate from the distribution (17) (see Rue and Held (2005), p. 92).

6.2. Bayesian analysis with IGMRFs

In a Bayesian inference with κ as a parameter we consider $f(\mathbf{x}|\kappa)$ in (18) as a likelihood function. It should then be noted that $f(\mathbf{x}|\kappa)$ is improper and hence not proportional to a proper distribution, which is the case for commonly considered likelihood functions.

Let $\pi(\kappa)$ be the prior density of κ , possibly improper. The natural definition of the joint distribution of (\mathbf{x}, κ) is then

$$f(\mathbf{x}, \kappa) = f(\mathbf{x}|\kappa)\pi(\kappa). \quad (21)$$

Thus the marginal density of κ is

$$\int f(\mathbf{x}, \kappa) d\mathbf{x} = \int f(\mathbf{x}|\kappa)\pi(\kappa) d\mathbf{x} = \infty,$$

so $\pi(\kappa)$ is in fact not the marginal distribution of κ . But still, by the theory of Section 2, the posterior density $\pi(\kappa|\mathbf{x})$ is well defined provided \mathbf{x} is σ -finite. This holds if the integral over κ of (21) is finite for (almost) all \mathbf{x} , i.e., if

$$\int \pi(\kappa)c(\kappa) \exp\left(-\frac{\kappa}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x}\right) d\kappa < \infty.$$

A sufficient condition for this is clearly that $\int \pi(\kappa)c(\kappa)d\kappa < \infty$. The conclusion of the above is that Bayesian inference for κ is well-behaved under reasonable restrictions, as soon as the constant $c(\kappa)$ has been determined.

7. Concluding remarks

In this paper we have presented, and discussed in view of several examples, a simple theoretical approach which enables the inclusion of improper priors in Bayesian analyses. A special feature of the approach is that both parameters and observations are represented as random quantities defined on a common underlying space Ω . The clue has been to allow the probability \Pr in Kolmogorov's axioms to be a σ -finite law with $\Pr(\Omega) = \infty$. In fact it was shown in Section 2 that $\Pr(\Omega) = \infty$ is *necessary* if improper priors are to be included.

What makes this a sensible theory is the fact that all conditional distributions, given σ -finite random quantities, are proper distributions. In particular this property leads to a consistent treatment of statistical models and a theoretically based condition for posterior propriety.

The relation to Renyi's theory of conditional probability spaces has been mentioned earlier. In this connection we would also like to quote from Lindley (1965). In the Preface to his classical test on probabilities he writes:

The axiomatic structure used here is not the usual one associated with the name of Kolmogorov. Instead one based on the ideas of Renyi has been used. The essential difference between the two approaches is that Renyi's is stated in terms of conditional probabilities, whereas Kolmogorov's is in terms of absolute probabilities, and conditional probabilities are defined in terms of them. Our treatment always refers to the probability of A , given B , and not simply to the probability of A . In my experience students benefit from having to think of probability as a function of two arguments, A and B , right from the beginning. The conditioning event, B , is then not easily forgotten and misunderstandings are avoided. These ideas are particularly important in Bayesian inference where one's views are influenced by the changes in the conditioning event.

References

- Akaike, H., 1980. The interpretation of improper prior distributions as limits of data dependent proper prior distributions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 46–52.
- Besag, J., 1974. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 192–236.
- Besag, J., York, J., Mollié, A., 1991. Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics* 43 (1), 1–20.
- Billingsley, P., 2008. *Probability and Measure*. John Wiley & Sons, Hoboken, New Jersey.
- Bioche, C., Druilhet, P., 2016. Approximation of improper priors. *Bernoulli* 22 (3), 1709–1728.

- Casella, G., Berger, R. L., 2002. *Statistical Inference*, 2nd Ed. Duxbury, Pacific Grove, CA.
- Chang, J. T., Pollard, D., 1997. Conditioning as disintegration. *Statistica Neerlandica* 51 (3), 287–317.
- Gelfand, A. E., Sahu, S. K., 1996. Identifiability, propriety, and parametrization with regard to simulation-based fitting of generalized linear mixed models. Tech. rep., 96-36, Department of Statistics, University of Connecticut.
- Gelfand, A. E., Sahu, S. K., 1999. Identifiability, improper priors, and Gibbs sampling for generalized linear models. *Journal of the American Statistical Association* 94 (445), 247–253.
- Hartigan, J. A., 1983. *Bayes Theory*. Springer Science, New York.
- Hobert, J. P., Casella, G., 1996. The effect of improper priors on Gibbs sampling in hierarchical linear mixed models. *Journal of the American Statistical Association* 91 (436), 1461–1473.
- Hodges, J. S., Carlin, B. P., Fan, Q., 2003. On the precision of the conditionally autoregressive prior in spatial models. *Biometrics* 59 (2), 317–322.
- Irony, T. Z., Singpurwalla, N. D., 1997. Non-informative priors do not exist. A dialogue with José M. Bernardo. *Journal of Statistical Planning and Inference* 65 (1), 159–177.
- Knorr-Held, L., 2003. Some remarks on Gaussian Markov random field models for disease mapping. In: Green, P., Hjort, N., Richardson, S. (Eds.), *Highly Structured Stochastic Systems*. Oxford University Press, Oxford.
- Lavine, M. L., Hodges, J. S., 2012. On rigorous specification of ICAR models. *The American Statistician* 66 (1), 42–49.
- Lindley, D. V., 1965. *Introduction to Probability and Statistics from Bayesian Viewpoint*. Vol. 1-2. Cambridge University Press, Cambridge.
- Lindqvist, B. H., Taraldsen, G., 2005. Monte Carlo conditioning on a sufficient statistic. *Biometrika* 92 (2), 451–464.

- Lunn, D. J., Thomas, A., Best, N., Spiegelhalter, D., 2000. WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing* 10 (4), 325–337.
- McCullagh, P., Han, H., et al., 2011. On Bayes’ theorem for improper mixtures. *The Annals of Statistics* 39 (4), 2007–2020.
- Renyi, A., 1962. *Probability Theory*. North-Holland, Amsterdam.
- Renyi, A., 1970. *Foundations of Probability*. North-Holland, Amsterdam.
- Royden, H., 1968. *Real Analysis: 2nd Ed.* Macmillan, London.
- Rue, H., Held, L., 2005. *Gaussian Markov Random Fields: Theory and Applications*. CRC Press, London.
- Stone, M., Dawid, A., 1972. Un-Bayesian implications of improper Bayes inference in routine statistical problems. *Biometrika* 59 (2), 369–375.
- Taraldsen, G., Lindqvist, B. H., 2010. Improper priors are not improper. *The American Statistician* 64 (2), 154–158.
- Taraldsen, G., Lindqvist, B. H., 2013. Fiducial theory and optimal inference. *The Annals of Statistics* 41 (1), 323–341.
- Taraldsen, G., Lindqvist, B. H., 2015. Fiducial and posterior sampling. *Communications in Statistics—Theory and Methods* 44 (17), 3754–3767.
- Taraldsen, G., Lindqvist, B. H., 2016. Conditional probability and improper priors. *Communications in Statistics—Theory and Methods* 45 (17), 5007–5016.