

IMPROPER PRIORS & FIDUCIAL INFERENCE

Gunnar Taraldsen & Bo Henry Lindqvist
Norwegian University of Science and Technology



The 4th Bayesian, Fiducial and Frequentist Workshop
Harvard University
May 2017

ABSTRACT

The use of improper priors flourish in applications and is as such a central part of contemporary statistics. Unfortunately, this is most often presented without a theoretical basis:

“Improper priors are just limits of proper priors ... ”

We present ingredients in a mathematical theory for statistics which generalize the axioms of Kolmogorov so that improper priors are included. A particular by-product is an elimination of the famous marginalization paradoxes in Bayesian and structural inference. Secondly, we demonstrate that structural and fiducial inference can be formulated naturally in this theory of conditional probability spaces. A particular by-product is then a proof of conditions which ensure coincidence between a Bayesian posterior and the fiducial distribution. The concept of a conditional fiducial model is introduced, and the interpretation of the fiducial distribution is discussed. It is in particular explained that the information given by the prior distribution in Bayesian analysis is replaced by the information given by the fiducial relation in fiducial inference.

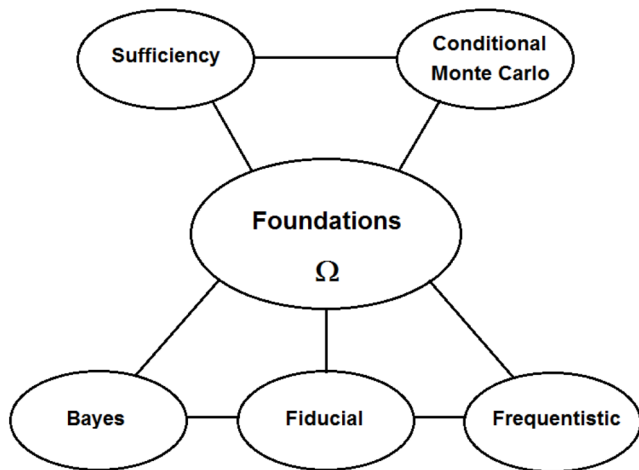
TABLE OF CONTENTS

INTRODUCTION

STATISTICS WITH IMPROPER PRIORS

FIDUCIAL INFERENCE

THE LARGE PICTURE



A MOTIVATING PROBLEM GIVES IT ALL


- ▶ Initial problem: Generate data $X = \chi(U, \theta)$ conditionally given a sufficient statistic $T = \tau(U, \theta) = t$.
- ▶ Tentative solution: Adjust parameter value θ for simulated data so that the sufficient statistic is kept fixed equal to t (Trotter-Tukey, 1956; Engen-Lillegård, *Biometrika* 1997).
- ▶ Corrected solution: The simulated data must be weighted, and the weight depends on an arbitrarily chosen improper distribution for the parameter (*Biometrika* 2003 & 2005).
- ▶ Realization after many years: Oh ... the resulting unweighted adjusted parameters are fiducial and the weighted are Bayes posterior ... (*Comm.Stat.* 2015)
- ▶ Ingredients: Optimal inference (*Ann.Stat.* 2013), Improper distributions (*Am.Stat.*, 2010), Sufficiency, Fiducial versus Bayesian distribution, Borel paradox, Marginalization paradox, Axioms for statistics (*Comm.Stat.* 2016).

TABLE OF CONTENTS

INTRODUCTION

STATISTICS WITH IMPROPER PRIORS

FIDUCIAL INFERENCE



AXIOMS FOR A STATISTICAL MODEL (Ω, X, Θ)

- ▶ The **basic space** Ω is a conditional probability space

$$(1) \quad (\Omega, \mathcal{E}, P)$$

given by axioms that generalize the axioms of Kolmogorov.

- ▶ The **model observation** X is a measurable function

$$(2) \quad X : \Omega \rightarrow \Omega_X$$

- ▶ The **model parameter** Θ is a measurable function

$$(3) \quad \Theta : \Omega \rightarrow \Omega_\Theta$$

CONDITIONAL PROBABILITY AND MORE

- ▶ Conditional probability and conditional expectation are defined by conditioning on a σ -finite σ -field \mathcal{F} of events.

- ▶ If Y is σ -finite, then:
 - (A) $(\Omega_Y, \mathcal{E}_Y, P_Y)$ is a conditional probability space.
 - (B) $P^y(A) = P(A | Y = y)$ is well defined.

- ▶ If Θ and X are σ -finite, then **model** $\{P_X^\theta(A) = P^\theta(X \in A)\}$ and **posterior** $\{P_\Theta^x(B) = P^x(\Theta \in B)\}$ are well defined.

- ▶ Convergence $Y_n \rightarrow Y$ in distribution can be defined by q -vague convergence (Bioche & Druilhet, *Bernoulli* 2016).

THE UNIFORM LAW P_{Θ} ON THE REAL LINE $\mathbb{R} \dots$

- ▶ A common, but sadly imprecise statement:

“It’s just a limit of the uniform law $U[-n, n]$ ”

- ▶ What kind of limit?
- ▶ Limit in what space?
- ▶ Compare with:

“Real numbers are just limits of rational numbers”

THE UNIFORM LAW P_{Θ} ON THE REAL LINE \mathbb{R}

- ▶ Renyi: The uniform law P_{Θ} for Θ is defined by

$$(4) \quad \forall n \in \mathbb{N}, \quad (\Theta \mid -n \leq \Theta \leq n) \sim U[-n, n]$$

- ▶ The family of sets $\mathcal{B} = \{[-n, n] \mid n \in \mathbb{N}\}$ is a *bunch* of sets, and the family of conditional probabilities

$$(5) \quad \{P_{\Theta}(\cdot \mid B) \mid B \in \mathcal{B}\}$$

defines a conditional probability space.

- ▶ The uniform law is not defined by a limit. Consider instead all the conditional uniform laws together as a (new) concept:

A conditional probability space

$$(6) \quad (\Omega_{\Theta}, \mathcal{E}_{\Theta}, P_{\Theta})$$

IMPROPER PRIORS AND SUFFICIENCY

- ▶ T is sufficient relatively to X for the parameter θ if $X|(T = t, \Theta = \theta)$ does not depend on θ .
- ▶ Equivalently (\forall priors): X, Θ are conditionally independent given T , so $(\Theta | X, T) \sim (\Theta | T)$. If, additionally, $T = \hat{\tau}(X)$, then $(\Theta | X) \sim (\Theta | T) \sim$ Bayes posterior.
- ▶ Theorem: Sufficiency implies that $[X|(T = t, \Theta = \theta)] \sim [X|T = t]$.
- ▶ Proof in discrete case: $E^\theta(\phi(X) | T = t) =$

$$\frac{E^\theta(\phi(X)(T = t))}{E^\theta(T = t)} = \frac{\int \pi(\theta) E^\theta(\phi(X)(T = t)) d\theta}{\int \pi(\theta) E^\theta(T = t) d\theta}$$

SUFFICIENCY AND OPTIMAL INFERENCE

- ▶ Sufficiency principle: Valid inference must be based on T for all sufficient T . It can be disputed. Halmos argument relies on accepting randomized procedures.
- ▶ If $\phi(X)$ is an estimator, then $E^\theta(\phi(X) | T)$ is an estimator with smaller (or equal) convex risk. If T is complete (and minimal), then it is the unique optimal estimator.
- ▶ An exact test for $(H_0 : \alpha = \alpha_0, \theta \text{ arbitrary})$ is obtainable if T is sufficient for the nuisance parameter θ . This gives exact confidence distributions. Lehmann: Optimality follows from this with additional assumptions.
- ▶ Accepting randomized procedures is equivalent to accepting construction of an improved alternative experiment. It can be disputed.

STATISTICAL INFERENCE (BFF)

- ▶ **Bayes:** Uncertainty P_{Θ}^x directly for the particular experiment at hand based on observation x , model P_X^{θ} , and prior P_{Θ} .

Challenge: Calculate characteristics of the posterior P_{Θ}^x .

- ▶ **Fiducial:** To be discussed on the next slides!
- ▶ **Frequentist:** Uncertainty indirectly from properties of the instrument in use based on observation x and model P_X^{θ} .

Challenge: Construct a suitable instrument ϕ .

- ▶ Bayesian and Fiducial arguments can sometimes be used to obtain excellent instruments $y = \phi(x)$ beyond the case of $y = \phi(x)$ equal to a confidence distribution (*Ann.Stat.*, 2013).

TABLE OF CONTENTS

INTRODUCTION

STATISTICS WITH IMPROPER PRIORS

FIDUCIAL INFERENCE

A FIDUCIAL MODEL

- ▶ Based on $x = \chi(u, \theta)$ for simulation of data x .
- ▶ A fiducial model (U, χ) for the observation X

$$(7) \quad \begin{array}{ccccc} \Omega & \xrightarrow{U, \Theta} & \Omega_U \times \Omega_\Theta & \xrightarrow{\chi} & \Omega_X \\ & \searrow & & \nearrow & \\ & & & & X \end{array}$$

- ▶ The law P_U^θ and the fiducial relation χ give the law P_X^θ of the statistical model.
- ▶ There exist many possible fiducial models for a given statistical model. This is an advantage!

SIMPLE FIDUCIAL INFERENCE

- ▶ Let $X = \Theta + U$ with $(U | \Theta = \theta) \sim N(0, I)$.
- ▶ Assume that $x = X(\omega) = (x_1, x_2)$ has been observed.
- ▶ Can **You** give a probability judgement regarding the unknown parameter $\theta = \Theta(\omega)$ when ω is unknown?

The simple fiducial argument:

1. **Known:** $x = X(\omega) = \Theta(\omega) + U(\omega)$
2. **Epistemic probability:** $x = \Theta^x + U^x$ with $U^x \sim N(0, I)$
3. **Fiducial distribution:** $\Theta^x = x - U^x \sim N(x, I)$

The judgement $U^x \sim (U | \Theta = \theta)$ based on the fiducial model and x replaces the use of a prior judgement P_{Θ} .

CONDITIONAL FIDUCIAL INFERENCE

- ▶ Let $X = \Theta + U$ with $(U | \Theta = \theta) \sim N(0, I)$.
- ▶ Assume that $x = (x_1, x_2)$ has been observed.
- ▶ Assume that it is known that $C(\theta) = \theta_1 - \theta_2 = 0$.
- ▶ Can **You** give a probability judgement regarding the unknown parameter θ ?

The conditional fiducial argument:


1. Unconditional fiducial distribution:

$$\Theta_u^x = x - U_u^x \sim N(x, I)$$

2. Fiducial distribution: $\Theta^x \sim (\Theta_u^x | C(\Theta_u^x) = 0)$

The judgement $U^x \sim (U_u^x | C(U_u^x) = x_1 - x_2)$ with $U_u^x \sim N(0, I)$ and $x = \Theta^x + U^x$ give the same result.

CONDITIONAL FIDUCIAL MODELS



Level curves for a bivariate fiducial together with three possible curves for restriction on the parameter space (Fisher, 1973 ed, p.138).

- ▶ A conditional fiducial model (U, χ, C) is given by a fiducial model (U, χ) for the observation X and a condition $C(\Theta) = c$.
- ▶ Fiducial inference is then based on the observation x , the law P_U^θ , the fiducial relation χ , and the condition $C(\Theta) = c$.
- ▶ $C_s(\theta) = \theta_1/\theta_2 = 1$ and $C_l(\theta) = \theta_1 - \theta_2 = 0$ both give the level curve $\theta_1 = \theta_2$, but the conditions are not equal!

FIDUCIAL INFERENCE FOR QUASI-GROUP MODELS

- ▶ Fiducial model: $X = \Theta U$
- ▶ Fiducial: $\Theta^x = x[U^x]^{-1}$ with $U^x \sim (U | \Theta = \theta)$
- ▶ Theorem: The fiducial gives a confidence distribution
- ▶ Theorem: If $\Theta u \sim \Theta$, then $\Theta^x \sim (\Theta | X = x)$
- ▶ Theorem: A right-invariant measure does not always exist
- ▶ Theorem: The fiducial determines optimal inference

OPTIMAL INFERENCE FROM FIDUCIAL ARGUMENTS

A simple calculation demonstrates that a fiducial model can be used in frequentistic decision theory to determine possible optimal instruments ϕ (*Ann.Stat.* 2013):

$r = E^\theta \ell [\theta, \phi(X)]$	Risk = expected loss
$= E^\theta \ell [\theta, \phi(\theta U)]$	Fiducial model $X = \Theta U$
$= E^\theta \ell [\theta, \theta \phi(U)]$	Equivariance of instrument ϕ
$= E^\theta \ell [e, \phi(U)]$	Invariance of loss ℓ
$= E^\theta \ell [\Theta^x, \Theta^x \phi(U)]$	Invariance of loss ℓ
$= E^\theta \ell [\Theta^x, \phi(\Theta^x U)]$	Equivariance of ϕ
$= E^\theta \ell [\Theta^x, \phi(x)]$	Fiducial equation

Conclusion: The risk is determined by the fiducial distribution, and an optimal instrument ϕ - if it exists - is determined by the fiducial distribution. There is no need for a Bayes prior in this argument!

FINAL COMMENTS ON THE INVOLVED IDEAS

- ▶ **Intuition: The information in the prior is replaced by the (weaker) information in the fiducial relation.**
- ▶ Instead of deciding a prior P_{Θ} : Decide on a distribution for U^x for a given observation x and the given fiducial relation.
- ▶ Fiducial distributions can give more than confidence intervals: Good, possibly optimal, instruments more generally.
- ▶ A theory with improper priors has been used repeatedly above. This is useful also more generally. It gives for instance precise limit statements involving priors, and resolves marginalization type of paradoxes.
- ▶ In the above arguments the law of $(U | \Theta = \theta)$ does not depend on θ . Fraser considers interesting models where $(U | \Theta = \theta)$ has a distribution that depends on θ through a shape parameter.