Bayesian analysis of polarization measurements

J. L. Quinn

Pontificia Universidad Católica de Chile, Departamento de Astronomía y Astrofísica, Vicuña Mackenna 4860, Macul 690-4411, Santiago, Chile
 e-mail: jquinn@astro.puc.cl

Received 19 September 2010 / Accepted 25 October 2011

ABSTRACT

Context. A detailed and formal account of polarization measurements using Bayesian analysis is given based on the assumption of gaussian error for the Stokes parameters. This analysis is crucial for the measurement of the polarization degree and angle at low and high signal-to-noise. The treatment serves as a framework for customized analysis of data based on a particular prior suited to the experiment.

Aims. The aim is to provide a rigorous and self-consistent Bayesian treatment of polarization measurements and their statistical error focused on the case of a single measurement.

Methods. Bayes Theorem is used to derive a variety of posterior distributions for polarization measurements.

Results. A framework that may be used to construct accurate polarization point estimates and confidence intervals based on Bayesian ideas is given. The results may be customized for a prior and loss function chosen for a particular experiment.

Key words. polarization – methods: data analysis

1. Introduction

Measuring polarization presents several challenges. The most important are that a naive calculation of small polarization seems to be biased towards higher values (Serkowski 1958), the error bars of the measured polarization are non-symmetric below a threshold signal-to-noise (Simmons & Stewart 1985), and the angular distribution is non-gaussian (Naghizadeh-Khouei & Clarke 1993). Given measured values of the Stokes parameters, what are the best estimates for the true values of the polarization degree and angle and how should error bars be assigned?

The polarization state of quasi-monochromatic light may be stated in terms of I_0 , Q_0 , U_0 , and V_0 Stokes parameters which are defined as certain averages of the electric field along a pair of orthogonal axes perpendicular to the direction of the light's propagation (see, e.g., del Toro Iniesta (2003) or Landi Degl'Innocenti (2002) for more information). I_0 is the intensity, Q_0 and U_0 are related to linear polarization, and V_0 is circular polarization. The Stokes parameters are subject to two conditions:

$$I_0^2 \geq Q_0^2 + U_0^2 + V_0^2 \quad (1)$$

and

$$I_0 \geq 0. \quad (2)$$

If incoming light is totally polarized, then the equality holds in Eq. (1), otherwise it is partially polarized or unpolarized. To keep the equations of manageable size, the rest of this paper makes the assumption that circular polarization is negligible (i.e., $V_0 = 0$). It should be straight-forward to generalize the results to include a non-zero V_0 . This restriction is not very severe because circular polarization is small compared to linear polarization for many astrophysical sources. It will also be taken that the Stokes parameters are constant during the course of a single observation.

Stokes parameters are calculated from intensities that follow a Poisson distribution governed by photon counts. Clarke et al. (1983) and Maronna et al. (1992) have investigated this to make

optimal estimates of them from repeated measurements. In general, however, a full Poisson treatment is computationally difficult when large intensities are involved. To avoid concerns about Poisson statistics, it is assumed that all intensities used to calculate Q and U (and also I) are large enough to be treated as having gaussian distributions.

The specifics regarding the calculation of I , Q , U and their errors (Σ_I , Σ_Q , Σ_U) depend on the experimental design (Keller 2002). For example, if one is using a dual-beam polarimeter with a half-wave-plate, Patat & Romaniello (2006) present an optimal set of equations for calculating the Stokes parameters from the intensities of the ordinary and extraordinary beams at different waveplate angles. Since the assumption that these intensities are large enough to be approximated by a Gaussian is being made, this implies that one may use the usual Gaussian error propagation formula for Σ_I , Σ_Q , and Σ_U based on the errors for the individual intensities (square root of the counts) that went into their calculation. Regardless how I , Q , and U are actually calculated, they may be treated as logically separate (e.g., one could imagine one instrument measuring intensity with error Σ_I in conjunction with completely separate instruments measuring Q and U with errors Σ_Q and Σ_U , respectively). Hereon, the error on the measured total intensity, Σ_I , is treated as being independent of Σ_Q and Σ_U .

Most of the polarization error literature makes the further assumption that the measured Stokes parameters are normally distributed about their true values, Q_0 and U_0 , with standard deviation Σ_Q and Σ_U . It is also usually assumed that the Q and U distributions are characterized by equal dispersion. Unfortunately, $\Sigma_Q = \Sigma_U$ is not generally satisfied for arbitrary datasets. This assumption is a pragmatic one to avoid ellipsoidal distributions which would complicate the analysis. Careful design of the experiment can make it such that $\Sigma_Q \approx \Sigma_U$. Extension of the results to allow for unequal variance is reserved for future work. It is now assumed that Σ_Q and Σ_U are equal and called Σ . The quantity Σ_I only tangentially enters the analysis, leaving Σ

as the lone important error quantity. In practice, Σ could be set equal to the average of Σ_Q and Σ_U or the maximum value if one is more conservative.

The units on I , Q , U and Σ_I , Σ_Q , Σ_U are energy per time. Shortly, new variables will be defined that are unitless. They will be best interpreted as percentages or ratios and the issue of units vanishes.

2. Theoretical background

In this section, the various distributions that will be needed are derived. The distributions will be presented in several different coordinate systems because this is intended to serve as a reference and because it is exceedingly easy to miss factors introduced by the Jacobian.

2.1. The sampling distribution

2.1.1. Variable introductions and large intensity limit

The fundamental assumption of this paper is that the measurement of the (unnormalized) Stokes parameters, Q and U , which are assumed to be uncorrelated, is described by a two-dimensional gaussian with equal standard deviation, Σ , in both directions. This distribution, F_C , is

$$F_C(Q, U|Q_0, U_0, \Sigma) = \frac{1}{2\pi\Sigma^2} \exp\left(-\frac{(Q - Q_0)^2 + (U - U_0)^2}{2\Sigma^2}\right). \quad (3)$$

(Most of the work for this paper will be done in polar coordinates. A C -subscript will be used for distributions in Cartesian coordinates.) The value of Σ is a positive constant assumed to be known precisely and, in practice, is estimated from the data. The parameters, Q_0 and U_0 , must be elements of a disk of radius I_0 centered on the origin in the Q_0 - U_0 plane. The range of both Q and U is $(-\infty, \infty)$ due to measurement error although the probability of measuring a value outside a disk of radius I_0 diminishes rapidly even for values of Q_0 and U_0 near the rim. Equation (3) is normalized ($\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F_C dQ dU = 1$).

It is often helpful to work in normalized Stokes parameters. Define new variables $q \equiv Q/I_0$, $u \equiv U/I_0$, $q_0 \equiv Q_0/I_0$, $u_0 \equiv U_0/I_0$, and $\sigma \equiv \Sigma/I_0$. The scaled error, σ , is an admixture between the intensity and the mean Q - U error. It or its inverse, $1/\sigma = I_0/\Sigma$, may be regarded as a measure of data quality (cf. I_0/Σ_I , Q/Σ_Q , and U/Σ_U). Small σ (large $1/\sigma$) implies good data. The values for q_0 and u_0 are restricted to a unit disk centered on the origin. The new probability density, f_C , after the change of coordinates is

$$f_C(q, u|q_0, u_0, \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(q - q_0)^2 + (u - u_0)^2}{2\sigma^2}\right). \quad (4)$$

This equation is also normalized ($\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_C dq du = 1$).

The normalized Stokes parameters q and u present a minor problem. How can they be calculated if I_0 is an unknown quantity? It might be thought that the definition of q and u should use the measured intensity, I , instead of the true intensity, I_0 . If that alternative definition is used, however, the differentials of q and u are more complicated than just $dq = dQ/I_0$ and $du = dU/I_0$. This makes a change of variables for the distribution difficult. It is better to keep the first definition and require that $I_0 \approx I$. This occurs when the total number of counts that went into the measurement of I is large, which also implies $\Sigma_I \ll I$.

Another helpful set of variables are “signal-to-noise” ratios. They are defined by $\bar{q} \equiv q/\sigma$, $\bar{u} \equiv u/\sigma$, $\bar{q}_0 \equiv q_0/\sigma$, and $\bar{u}_0 \equiv u_0/\sigma$. In these barred variables, the new distribution, $\bar{f}_C$, is

$$\bar{f}_C(\bar{q}, \bar{u}|\bar{q}_0, \bar{u}_0) = \frac{1}{2\pi} \exp\left(-\frac{(\bar{q} - \bar{q}_0)^2 + (\bar{u} - \bar{u}_0)^2}{2}\right). \quad (5)$$

This time, the Jacobian causes the $1/\sigma^2$ leftover after the change of variables to disappear. This is normalized as well ($\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \bar{f}_C d\bar{q} d\bar{u} = 1$); but, in this case, $\bar{q}_0$ and $\bar{u}_0$ are restricted to a disk of radius $1/\sigma$.

2.1.2. Transformation to polar coordinates

The true polarization degree, p_0 , and true angle on the sky, ϕ_0 , may be defined in terms of the true (i.e., perfectly known) q_0 and u_0 Stokes parameters by

$$p_0 \equiv \sqrt{q_0^2 + u_0^2} \quad \text{and} \quad \phi_0 \equiv \frac{1}{2} \text{atan}\left(\frac{u_0}{q_0}\right). \quad (6)$$

The true angle on the sky is related to the true angle in the q_0 - u_0 plane by the relation $\theta_0 \equiv 2\phi_0$. Similarly,

$$p \equiv \sqrt{q^2 + u^2} \quad \text{and} \quad \phi \equiv \frac{1}{2} \text{atan}\left(\frac{u}{q}\right) \quad (7)$$

and $\theta \equiv 2\phi$. The range of p is $[0, \infty)$ and the range of p_0 is $[0, 1]$. Define $\bar{p} \equiv p/\sigma$ and $\bar{p}_0 \equiv p_0/\sigma$. It is important to notice that the range of $\bar{p}_0$ is now $[0, 1/\sigma]$ or $[0, I_0/\Sigma]$ because in the Bayesian analysis, this will be the range of integration. The polarization angle defines the plane of vibration of the electric field on the sky. By convention, it is usually defined such that 0° corresponds to North and increases eastward such that East is $\phi = 90^\circ$.

Equation (4) may be converted to polar coordinates by inverting Eqs. (6) and (7). Using θ instead of ϕ , this yields $q = p \cos \theta$ and $u = p \sin \theta$ and corresponding equations for the zero-subscripted true values. The normalized probability density distribution in polar coordinates (p, θ) is

$$f'(p, \theta|p_0, \theta_0, \sigma) = \frac{p}{2\pi\sigma^2} \exp\left(-\frac{p^2 + p_0^2 - 2pp_0 \cos(\theta - \theta_0)}{2\sigma^2}\right). \quad (8)$$

A prime is used as part of the function's name to serve as a warning that the angle in the q - u plane, θ , is being used as the angular variable. The factor of p in front enters because $dq du = p dp d\theta$ (i.e., due to the Jacobian of the transformation). Some authors choose to leave the factor of p off of the p, θ -distribution and insert it manually when needed. Probability distributions transform as a scalar density under coordinate transformations so they should “pick-up” the Jacobian factor. Caution is needed when reading the literature.

The distribution in terms of the sky angle is

$$f(p, \phi|p_0, \phi_0, \sigma) = \frac{p}{\pi\sigma^2} \exp\left(-\frac{p^2 + p_0^2 - 2pp_0 \cos(2(\phi - \phi_0))}{2\sigma^2}\right) \quad (9)$$

and in barred variables,

$$\bar{f}(\bar{p}, \phi|\bar{p}_0, \phi_0) = \frac{\bar{p}}{\pi} \exp\left(-\frac{\bar{p}^2 + \bar{p}_0^2 - 2\bar{p}\bar{p}_0 \cos(2(\phi - \phi_0))}{2}\right). \quad (10)$$

The $\bar{f}(\bar{p}, \phi | \bar{p}_0, \phi_0)$ distribution is normalized under the support $[0, \infty) \times (-\pi/2, \pi/2]$ (i.e., $\int_{-\pi/2}^{\pi/2} \int_0^\infty \bar{f} d\bar{p} d\phi = 1$). The origin has infinitesimal measure so usually it does not contribute finitely to an integral. In later sections, integration over delta functions (aka delta distributions) centered at the origin will occur. This requires some formality in the polar coordinate definition. The parameter space technically has $(p = 0, \phi)$ identified for all ϕ as the origin and (p, ϕ) is identified with $(p, \phi = \phi + n\pi)$ for all $p > 0$ and any integer n . The resulting quotient space has $[0, \infty) \times (-\frac{\pi}{2}, \frac{\pi}{2}]$ as a fundamental cell. Functions on this space are periodic in ϕ and they must satisfy $f(p, \phi) = f(p, ((\phi + \frac{\pi}{2}) \bmod \pi) - \frac{\pi}{2})$. Here, “mod” is the modulus operator. The $\pi/2$ operations perform shifts needed because the angles are defined as $(-\frac{\pi}{2}, \frac{\pi}{2}]$ instead of $[0, \pi)$. This formality is sometimes important in calculations, like median estimates of ϕ_0 , where it is natural to have the range of integration extend outside of the $(-\frac{\pi}{2}, \frac{\pi}{2}]$ domain.

Some intuition is gained by examining the distribution in Eq. (10) for differing values of p_0/σ . Figures 1a–c show the distributions for $\phi_0 = 0$ and $p_0/\sigma = 0.0, 0.5$ and 2.0 , respectively. Choosing a different value of ϕ_0 , for a given value of p_0/σ , only translates the distribution in the ϕ -direction and does not change its shape; therefore, all plots simply use $\phi_0 = 0$. When $p_0/\sigma = 0$ the angular distribution is flat and much of the probability is concentrated in a band around $p/\sigma = 1$. As p_0/σ increases, an oval-shaped “probability bubble” forms. This shape persists even at large values. The probability under the distribution to the left and right of the maximum is asymmetric, with more to the left. As p_0/σ continues to increase, the probability on each side of the maximum approaches 0.5. This is true even though the oval shape stays present.

The sampling distribution for N measurements is simply a product of those for individual measurements when the measurements are independent of one another,

$$\mathcal{F}(\{p, \phi\}_N | p_0, \phi_0, \sigma) = \prod_{n=1}^N f((p, \phi)_n | p_0, \phi_0, \sigma). \quad (11)$$

In this notation, $(p, \phi)_n$ denotes the n -th measurement and $\{p, \phi\}_N$ is the whole set of measurements. This paper focuses on the case of a single measurement, such that $\mathcal{F} = f$.

2.2. The “most probable” and maximum likelihood estimators

Two classical analytic estimators for the unknown parameters are easily found from the sampling distribution. These are the maximum likelihood and “most probable” estimators. Later new estimators will be introduced using the posterior distribution.

If $f(p, \phi | p_0, \phi_0, \sigma)$ is viewed as a function of p_0 and ϕ_0 for a fixed p and ϕ , it is called the likelihood function. Solving the system $\frac{\partial f}{\partial p_0} = 0$ and $\frac{\partial f}{\partial \phi_0} = 0$ yields the maximum likelihood estimate (ML),

$$\hat{p}_{0,ML} = p \quad \text{and} \quad \hat{\phi}_{0,ML} = \phi. \quad (12)$$

This solution is not particularly useful as it does not correct for bias. There is potential for confusion here: if the same construction is applied to the sampling distribution marginalized over the angle (i.e., Rice distribution), the maximum likelihood estimator *does* correct for some bias. The Rice distribution will be covered in Sect. 2.3.

Wang et al. (1997) suggest simultaneously maximizing Eq. (8) with respect to p and θ to estimate p_0 and θ_0 . This yields

Fig. 1. Contour plots of $\bar{f}(\bar{p}, \phi | \bar{p}_0, \phi_0)$ (Eq. (10)) for a few different values of p_0/σ with $\phi_0 = 0$. The dots indicate the maximums of the distributions. Contour spacings are at intervals of 0.02.

another estimator sometimes called the “most probable” estimator (MP) in the literature (or Wardle and Kronberg estimator when applied to the Rice distribution). The system $\frac{\partial f}{\partial p} = 0$ and $\frac{\partial f}{\partial \phi} = 0$ has the following solution for a maximum:

$$\hat{p}_{0,MP} = \left(p - \frac{\sigma^2}{p} \right) h(p - \sigma) \quad \text{and} \quad \hat{\phi}_{0,MP} = \phi, \quad (13)$$

where $h(x)$ is the Heaviside step function. The Heaviside function can be motivated by contemplation of Fig. 1a. Equation (13) may be used to find the maximums in Fig. 1. It is worth noting that this formula is also expressible as $\hat{p}_{0,MP} = \left(\left(p - \frac{\sigma^2}{p} \right) + \left| p - \frac{\sigma^2}{p} \right| \right) / 2$, which may be more practical in data

reduction pipelines, and $\hat{p}_{0,MP} = \left(\bar{p} - \frac{1}{\bar{p}}\right)h(\bar{p} - 1)$ in terms of the barred variables.

The goal is to replace these estimators with new ones based on Bayesian ideas and Decision Theory. The ML estimate continues to be relevant in that context because it is equal to some Bayesian estimators when a uniform prior is used.

2.3. The Rice distribution

The Rice distribution (Rice 1945), R , is the marginal distribution of p over f ,

$$R(p|p_0, \sigma) = \int_{-\pi/2}^{\pi/2} f(p, \phi|p_0, \phi_0, \sigma) d\phi. \quad (14)$$

This integration may be accomplished using the integral $\int_{-\pi/2}^{\pi/2} e^{a \cos(2x)} dx = \pi I_0(a)$ so that

$$R(p|p_0, \sigma) = \frac{p}{\sigma^2} \exp\left(-\frac{p^2 + p_0^2}{2\sigma^2}\right) I_0\left(\frac{pp_0}{\sigma^2}\right). \quad (15)$$

Here, $I_0(x)$ denotes the zeroth-order modified Bessel function of the first kind¹. It is an even *real* function when its argument is real and $I_0(0) = 1$. The Rice distribution does not depend on ϕ_0 or ϕ after the integration over ϕ so they have been dropped from the notation.

In barred variables, the Rice distribution is

$$\bar{R}(\bar{p}|\bar{p}_0) = \bar{p} \exp\left(-\frac{\bar{p}^2 + \bar{p}_0^2}{2}\right) I_0(\bar{p}\bar{p}_0). \quad (16)$$

This is shown as a contour plot in Fig. 2a. In this plot, $\bar{R}$ is plotted against $\bar{p}_0$ and $\bar{p}$. The contours are at 0.1 intervals. There are two special points in the plot. A global maximum of $1/\sqrt{e} \approx 0.6065$ is reached at (1, 0) (black dot) and a critical point at $(\sqrt{2}, 0)$ (gray dot) is found if maximums along vertical slices are examined. Recall $\bar{p} = p/\sigma$ and $\bar{p}_0 = p_0/\sigma$. The maximum along a vertical slice in the interval $0 \leq p/\sigma \leq \sqrt{2}$ lies on the p/σ -axis while for values greater than $\sqrt{2}$ it lies above the axis. Tracing the maximum of the vertical slices for $\frac{p}{\sigma} > \sqrt{2}$ implicitly defines a curve which is given by

$$\bar{p} I_1(\bar{p}\bar{p}_0) - \bar{p}_0 I_0(\bar{p}\bar{p}_0) = 0. \quad (17)$$

This is the ML estimator of Simmons & Stewart (1985) (with a factor of i corrected). $I_1(x)$ is the first-order modified Bessel function of the first kind. It is an odd real function when its argument is real. Similarly, horizontal slices define a curve via

$$(\bar{p}^2 - 1)I_0(\bar{p}\bar{p}_0) - \bar{p}\bar{p}_0 I_1(\bar{p}\bar{p}_0) = 0. \quad (18)$$

This is the Wardle and Kronberg estimator (Wardle & Kronberg 1974). These are the two dashed curves in the plot. The curve for the horizontal slice maximums is the one that terminates at the global maximum. The red dotted line is the Wang estimator from Eq. (13). There is a significant difference between the Rice distribution-based estimators and the two-dimensional estimator. For large values of $\bar{p}$, it estimates a smaller value of $\bar{p}_0$ lower than the two other curves.

¹ Some papers in the literature use Bessel functions, $\mathcal{J}$, instead of modified Bessel functions, I . This introduces distracting complex factors of i into the equations. It is worth noting that $\mathcal{J}_0(ix) = I_0(x)$ and $\mathcal{J}_1(ix) = iI_1(x)$ to aid in reading those papers.

(a) A view of the Rice distribution near the origin. The value of σ is mostly irrelevant so long as $1/\sigma$ is at least 3 such that the domain plotted is valid. For instance, it may be taken to be 100 for comparison with later figures. The meaning of the dashed and dotted curves and dots is discussed in the text.

(b) A view of the Rice distribution for the full domain of $\bar{p}_0$ when $\sigma = 1/8$. Notice that the linear-like band that develops persists to the maximum value of $\bar{p}_0$, which is 8 in this case. This is different than the behavior of the posterior distribution as will be seen.

Fig. 2. Two contour plots of the Rice distribution (Eq. (16)). The contours are spaced by 0.1 intervals. The value of σ is mostly unimportant for $\bar{R}$. It does not affect the shape of the distribution. The only purpose of mentioning it is to insure $\bar{p}$ exists for values above $1/\sigma$.

Figure 2b plots the Rice distribution with the full range of $\bar{p}_0$ when $\sigma = 1/8$. Except near the origin, much of the probability is distributed in a diagonal linear band that continues uninterrupted until the maximum value of $\bar{p}_0$. Later it will be seen that the posterior distribution becomes non-linear for large values of polarization.

2.4. The posterior distribution

The distribution given by Eq. (9) is the distribution of the measured values given true values as input parameters. In practice, one usually wishes to estimate the true values from the measured

values. This is accomplished through the posterior distribution, B , given by Bayes Theorem provided one accepts some prior distribution, $\kappa(p_0, \phi_0)$, for the model parameters. The posterior distribution (for a single measurement) is,

$$B(p_0, \phi_0 | p, \phi, \sigma) = \frac{f(p, \phi | p_0, \phi_0, \sigma) \kappa(p_0, \phi_0)}{\int_{-\pi/2}^{\pi/2} \int_0^1 f(p, \phi | p'_0, \phi'_0, \sigma) \kappa(p'_0, \phi'_0) dp'_0 d\phi'_0} \quad (19)$$

Care must be taken with the limits of the p_0 integration and it must always be remembered if one is using the unnormalized, normalized, or barred variables. The barred version is

$$\bar{B}(\bar{p}_0, \phi_0 | \bar{p}, \phi, \sigma) = \frac{\bar{f}(\bar{p}, \phi | \bar{p}_0, \phi_0) \bar{\kappa}(\bar{p}_0, \phi_0)}{\int_{-\pi/2}^{\pi/2} \int_0^{1/\sigma} \bar{f}(\bar{p}, \phi | \bar{p}'_0, \phi'_0) \bar{\kappa}(\bar{p}'_0, \phi'_0) d\bar{p}'_0 d\phi'_0} \quad (20)$$

with $\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma \kappa(\sigma \bar{p}_0, \phi_0)$. It is critical to notice the upper limit of the integration of $\bar{p}_0$ is $1/\sigma$, which equals I_0/Σ . It is not infinity as is sometimes used. In practice, however, many of the integrals over $\bar{p}_0$ that occur in Bayesian polarization equations, tend to be extremely insensitive to the value of $1/\sigma$ so long as it is above a value of about 4. Thus, using infinity as the upper limit often produces a very reasonable approximation. This has the added advantage that integrals sometimes have explicit closed forms when otherwise they might not.

There is great freedom in defining a Bayesian prior. The function is ostensibly required to be a probability distribution. This necessitates that $\kappa(p_0, \phi_0)$ be non-negative everywhere; however, in most situations, the additional requirement that the distribution be normalizable may be relaxed as the normalization constant cancels. The choice of prior will be further discussed in Sect. 3.

Equation (20) is shown in Figs. 3 and 4 for $I_0/\Sigma = 100$ and two different choices of prior. Figure 3 uses $\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma^2 2\bar{p}_0/\pi$ and Fig. 4 uses $\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma/\pi$. These priors are discussed in more detail in Sect. 3. Each plot shows several different values of p/σ . As before, changing the value of ϕ merely shifts the distribution along the ϕ_0 axis so $\phi = 0$ has been used. In Figs. 2.5 and 4a, $p/\sigma = 0$ and there is no preference for any ϕ_0 over another. In Figs. 2.5 and 4b, $p/\sigma = 0.5$. Here a probability bubble is beginning to form. In Figs. 2.5 and 4c, $p/\sigma = 2.0$ and the probability bubble is now fairly mature. The panels of Fig. 4 are similar. Notice that Fig. 3 is nearly identical to the distribution plotted in Fig. 1. There is a very weak dependence of $\bar{B}$ on σ that becomes larger for larger values of p_0 . If p_0 were plotted near the maximum, the plots would differ radically. When $1/\sigma \rightarrow \infty$, the plots do become identical. The similarity (and differences) will be discussed later in conjunction with the one-dimensional, marginalized version of this plot.

2.5. Simple Bayesian estimators

Bayesian estimators based on the mean, median, and mode can be defined for the posterior distribution.

The mode estimator (usually called the MAP or maximum a posteriori estimate) is the maximum of the posterior distribution. If any delta function is used in the prior, this causes delta functions to appear in the posterior distribution, which trumps any finite maximum of the posterior distribution and renders the mode estimate somewhat useless without further modification. If no delta functions and a uniform prior is used, then the mode estimate is the same as the ML estimate (Eq. (12)).

The median estimator, $\hat{p}_{0,\text{MED}}$, is the value such that

$$\int_{-\pi/2}^{\pi/2} \int_0^{\hat{p}_{0,\text{MED}}} B dp_0 d\phi_0 = \frac{1}{2}. \quad (21)$$

This estimator will usually have to be found numerically. A median estimator $\hat{\phi}_{0,\text{MED}}$ is usually easy to find from symmetry considerations but for our parameter space must be defined as

$$\int_0^1 \int_{\hat{\phi}_{0,\text{MED}}}^{\hat{\phi}_{0,\text{MED}} + \frac{\pi}{2}} B d\phi_0 dp_0 = \frac{1}{2} \text{ and } \int_0^1 \int_{\hat{\phi}_{0,\text{MED}} - \frac{\pi}{2}}^{\hat{\phi}_{0,\text{MED}}} B d\phi_0 dp_0 = \frac{1}{2}. \quad (22)$$

Under special circumstances, $\hat{\phi}_{0,\text{MED}}$ may not be unique but after a non-zero measurement of p , it will generally be the case that it is.

The mean estimators are the values

$$\hat{p}_{0,\text{MEAN}} = \int_{-\pi/2}^{\pi/2} \int_0^1 p_0 B dp_0 d\phi_0 \quad (23)$$

and

$$\hat{\phi}_{0,\text{MEAN}} = \int_{-\pi/2}^{\pi/2} \int_0^1 \phi_0 B dp_0 d\phi_0. \quad (24)$$

These three estimators will be different in general.

2.6. Bayesian decision theory

The posterior distribution contains *all* information about the relative likelihood of the model parameters. It represents what has been learned from the observation. For many problems it is usually desired to summarize the posterior distribution by a few statistics such as an estimate of the “best” value and some confidence interval (often called “credible sets” in this context). In Bayesian Decision and Estimation Theory determining the “best” estimate of the model parameters requires defining a loss function and a decision rule (there are many good standard texts offering much more detail such as [Berger 1985](#); and [Robert 1994](#)). The loss function, L , assigns a weight to deviations from the true value such that measured values that are far away from it incur a greater penalty than those that are closer. The decision rule assigns an estimate of the true value based upon the data, which will usually be the estimate that minimizes the expected posterior probable loss.

Deciding upon a loss function is one of the toughest parts of a Bayesian analysis. The three most common are squared-error loss, absolute deviation loss, and the “0–1” loss. These loss functions have as their solutions for each variable the mean, median, and mode values of the posterior, respectively ([Robert 1994](#)).

The squared-deviation loss function for the polarization problem may be defined as the squared distance between the true value (q_0, u_0) and the best estimate (q_a, u_a) under a Euclidean metric (scaled by σ^2),

$$L(q_0, u_0, q_a, u_a) = ((q_0 - q_a)^2 + (u_0 - u_a)^2)/\sigma^2. \quad (25)$$

Transforming Eq. (25) to polar coordinates gives,

$$L(p_0, \phi_0, p_a, \phi_a) = (p_0^2 + p_a^2 - 2p_0 p_a \cos(2(\phi_0 - \phi_a)))/\sigma^2. \quad (26)$$

The posterior expected loss, Z , under B is

$$Z(p_a, \phi_a) = \int_{-\pi/2}^{\pi/2} \int_0^1 L(p_0, \phi_0, p_a, \phi_a) B(p_0, \phi_0 | p, \phi, \sigma) dp_0 d\phi_0. \quad (27)$$

(a) $p/\sigma = 0.0$ and $\phi = 0$. (a) $p/\sigma = 0.0$ and $\phi = 0$. (b) $p/\sigma = 0.5$ and $\phi = 0$. (b) $p/\sigma = 0.5$ and $\phi = 0$. (c) $p/\sigma = 2.0$ and $\phi = 0$. (c) $p/\sigma = 2.0$ and $\phi = 0$.

Fig. 3. Contour plots of $\bar{B}(\bar{p}, \phi | \bar{p}_0, \phi_0)$ (Eq. (20)) for a few different values of p/σ with $\phi = 0$ assuming a Jeffreys prior ($\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma^2 2\bar{p}_0/\pi$) and $I_0/\Sigma = 100$. The dots indicate the maximums of the distributions. Contour spacings are at intervals of 0.02. This figure is nearly the same as Fig. 1 with the roles of p and p_0 reversed. While for small values of p/σ and p_0/σ the plots appear identical, they are not. Their similarity diverges for larger values of signal-to-noise, as they must since p_0/σ has a maximum value while p/σ does not.

Fig. 4. Contour plots of $\bar{B}(\bar{p}_0, \phi_0 | \bar{p}, \phi, \sigma)$ (Eq. (20)) for a few different values of p/σ with $\phi = 0$ assuming a uniform polar prior ($\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma/\pi$) and $I_0/\Sigma = 100$. The dots indicate the maximums of the distributions. Contour spacings are at intervals of 0.02.

are a solution and $\frac{\partial^2 Z}{\partial p_a^2} + \frac{\partial^2 Z}{\partial \phi_a^2} - \left(\frac{\partial^2 Z}{\partial p_a \partial \phi_a}\right)^2 > 0$ when $p_a, p_0 > 0$ so the solutions are a minimum.

The absolute loss function is

$$L(q_0, u_0, q_a, u_a) = \sqrt{(q_0 - q_a)^2 + (u_0 - u_a)^2}/\sigma. \quad (30)$$

It will be assumed that the usual median estimators in polar coordinates are in fact a solution to this loss function.

The “0–1” loss function is

$$L(q_0, u_0, q_a, u_a) = \begin{cases} 1 & \text{if } q_a = q_0 \text{ and } u_a = u_0 \\ 0 & \text{if } q_a \neq q_0 \text{ and } u_a \neq u_0. \end{cases} \quad (31)$$

Loss can be minimized by solving $\frac{\partial Z}{\partial p_a} = 0$ and $\frac{\partial Z}{\partial \phi_a} = 0$ and checking that the solution is a minimum via the second derivative test. These two equations are difficult to solve completely but it can be shown that if $\kappa(p_0, \phi_0) = \kappa(p_0)$ then

$$\hat{\phi}_a = \phi_0 = \phi \quad (28)$$

and

$$\hat{p}_a = \int_{-\pi/2}^{\pi/2} \int_0^1 p_0 B(p_0, \phi_0 | p, \phi, \sigma) dp_0 d\phi_0 \quad (29)$$

It will also be assumed that the mode is the solution of this loss function. When the prior contains a delta function, this estimator is not very useful.

2.7. Posterior odds

An analog of hypothesis testing in Bayesian analysis is the comparison of posterior odds. In the simplest case, suppose one wishes to test the hypothesis that the model parameters lie in some subset Θ_1 of the parameter space. This is accomplished by integrating the posterior distribution over the subset. If the probability is greater than 50%, the hypothesis is accepted. If it is less than 50%, it is rejected. The simplicity of these tests is one of the benefits of the Bayesian approach. The subset need not consist of more than a single point but in that case the associated probability will usually be infinitesimal unless distributions are allowed for priors. In the next section, priors will be introduced that use a delta function at the origin to represent the probability that a source has negligible polarization. After the measurement, one may compute the posterior odds for the origin (a single point) to test if the measurement is consistent with zero. One of the benefits of $\hat{p}_{0,\text{MED}}$ is that it naturally performs this test for a delta function at the origin.

Two probabilities are of interest, the probability that p_0 is zero and the probability that p_0 is greater than zero. In mathematical form,

$$\text{Prob}(p_0 = 0) = \lim_{\alpha \rightarrow 0^+} \int_{-\pi/2}^{\pi/2} \int_0^\alpha B(p_0, \phi_0 | p, \phi, \sigma) dp_0 d\phi_0 \quad (32)$$

and

$$\text{Prob}(p_0 > 0) = \lim_{\alpha \rightarrow 0^+} \int_{-\pi/2}^{\pi/2} \int_\alpha^1 B(p_0, \phi_0 | p, \phi, \sigma) dp_0 d\phi_0. \quad (33)$$

Since the distribution is normalized, only Eq. (32) need actually be calculated because $\text{Prob}(p_0 = 0) + \text{Prob}(p_0 > 0) = 1$. If $\text{Prob}(p_0 = 0) > 1/2$, then it is more likely the source has negligible than appreciable polarization. If $\text{Prob}(p_0 = 0) < 1/2$, then it is more likely the source has appreciable than negligible polarization.

It is now time to discuss choosing a prior so that these equations may be used.

3. Elicitation of priors

In Bayesian statistics, the choice of the prior distribution is non-trivial and can be very controversial because it is frequently subjective and based on previous experience for the problem at hand. The purpose of the prior density is to represent knowledge of the model parameters *before* an experiment. One common situation is where one has no prior knowledge of their values and wishes to construct a prior to represent this ignorance. Such priors are called “objective”. Constructing objective priors can be extremely subtle. The Bertrand Paradox illustrates that the method of producing or defining a “random” value can have important consequences for the resulting probability distributions (Bertrand 1888). In general, the solution to the Bertrand Paradox is that “real randomness” requires certain translation and scaling invariance properties (Jaynes 1973; Tissier 1984; Di Porto et al. 2010). In practice, such objective priors may contain too little information to satisfactorily model the phenomena being studied.

In general, the polarization state of an unresolved source at a particular wavelength is described by a point of the Poincaré

sphere, which is the three-dimensional set of the physically possible values for the q , u , and v Stokes parameters given an intensity. An assumption was made in the introduction that circular polarization is negligible. This is a *physical* assumption which influences the meaning of the polarization state of a “random” source. The resulting q – u distribution of random sources is not to be confused with the expected q – u distribution of a random source whose v component was ignored. For a uniform distribution of points within the Poincaré sphere, this latter distribution would be the density distribution obtained by projection onto the q – u plane, which would have a larger density near the origin. This distribution is important and interesting but its three-dimensional motivation is somewhat outside two-dimensional scope of this work. A brief discussion of it is given in Appendix B.

Investigation is still required into the meaning of “randomness” even under just the $V = 0$ point-of-view. It is necessary to understand what role the choice of coordinates plays and how “uniform” does not imply “random”.

3.1. Uniform prior in Cartesian coordinates (the Jeffreys prior)

The prior most likely to coincide with a person’s intuitive notion of a “random point” in the unit disk is a uniform distribution in q_0 – u_0 coordinates is

$$\kappa_C(q_0, u_0) = \frac{1}{\pi}. \quad (34)$$

In polar coordinates, this is equivalent to

$$\kappa(p_0, \phi_0) = \frac{2p_0}{\pi} \quad \text{or} \quad \bar{\kappa}(\bar{p}_0, \phi_0) = \sigma^2 \frac{2\bar{p}_0}{\pi}. \quad (35)$$

The barred versions are always $\bar{\kappa}(\bar{p}_0, \phi_0) = \sigma \kappa(\sigma \bar{p}_0, \phi_0)$. Both distributions are normalized. Random points generated using this distribution have an expected density that is constant across the unit disk (see Fig. 5a). One aspect of this prior is that a randomly generated point will generally be more likely to have a strong polarization than a weak one because an annulus associated with a large radius has more area than for a smaller radius.

It can be shown that the prior given by Eq. (34) is the normalized Jeffreys prior for the polarization problem. The Jeffreys prior is an objective prior constructed from the square root of the determinate of the Fisher information matrix (Jeffreys 1939, 1946; Robert et al. 2009a,b). One of the features of the Jeffreys prior is that it is invariant under a re-parametrization. As an objective prior, the Jeffreys prior is in some sense the prior desired when absolutely no previous knowledge of the model parameters is known.

While the objectivity of this prior is admirable, in practice, most astronomical sources are not expected to have very large polarizations (masers can be a notable exception); so, in many cases, it is undesirable that the polarization be more likely to be large than small. In this respect, this prior may not be a good choice for some astrophysical observations despite its objectivity: it simply contains too little information to be suitable and overestimates our lack of knowledge and previous experience.

3.2. Uniform prior in polar coordinates

The uniform prior in polar coordinates captures some features that are usually desired for astronomical sources than the

(a) A set of 2000 points generated using the Jeffreys prior.

(b) A second set of 2000 points in p_0 - ϕ_0 space generated for the uniform polar prior.

(c) The same set as in Fig. 5b in q_0 - u_0 space.

Fig. 5. A comparison of “random” points generated with the Jeffreys prior and the uniform polar prior.

previous case. Its (normalized) formula is

$$\kappa(p_0, \phi_0) = \frac{1}{\pi} \quad \text{or} \quad \bar{\kappa}(\bar{p}_0, \phi_0) = \frac{\sigma}{\pi}. \quad (36)$$

While this prior may be “uniform” in p_0 - ϕ_0 space (see Fig. 5b), it is definitely not uniform in q_0 - u_0 space. In fact, in that space, the density of points produced randomly by this prior increases near the origin, as is shown in Fig. 5c. This is a subjective prior in the sense that it prefers points closer to the origin.

3.3. The addition of a delta term at the origin

One subtle feature of the previous priors is that they do not allow for a finite probability that an object’s polarization is exactly zero. Instead there is only a finite chance that the polarization will be in some interval that includes zero. A Dirac delta function at the origin may be used to allow for the possibility that an object truly has zero polarization. “Exactly zero” polarization may be interpreted as equivalent to the statement that there is a large probability that the polarization is negligible. In this case, the delta function then becomes just a mathematical tool to replace a probability density that rises extremely sharply at the origin in such a way to bound an appreciable probability.

When one is investigating if some class of objects are polarized, it seems reasonable to assign a 50% chance to the probability that the object is unpolarized and a 50% chance to it being polarized. More generally, the percentage can be parametrized, such that the probability of zero polarization is A and greater than zero polarization is $(1 - A)$. Let the prior consists of two components: a normalized delta function term at the origin, $\kappa_1(p_0, \phi_0)$, and a normalized non-delta function term, $\kappa_2(p_0, \phi_0)$. The general form for the prior can now be written as

$$\kappa(p_0, \phi_0) = A\kappa_1(p_0, \phi_0) + (1 - A)\kappa_2(p_0, \phi_0) \quad (37)$$

or

$$\bar{\kappa}(\bar{p}_0, \phi_0) = A\bar{\kappa}_1(\bar{p}_0, \phi_0) + (1 - A)\bar{\kappa}_2(\bar{p}_0, \phi_0). \quad (38)$$

The delta term must be

$$\kappa_1(p_0, \phi_0) = \frac{2\delta(p_0)}{\pi} \quad \text{or} \quad \bar{\kappa}_1(\bar{p}_0, \phi_0) = \frac{2\delta(\bar{p}_0)}{\pi}, \quad (39)$$

respectively. The two versions in Eq. (39) have the same form because the delta function transforms in a way that cancels the σ ’s that typically will appear in the barred version. Notice too that the p_0 -delta function occurs at the lower limit of the p_0 integration. This is a somewhat quirky situation but by remembering that delta functions are constructed as a limit of a sequences of functions it is easily seen that $\int_0^\alpha f(x)\delta(x)dx = f(0)/2$ for $\alpha > 0$. Taking into account the factors of one-half that arise in this way, it can be shown that Eqs. (38) and (39) are both normalized properly². A delta function could also be placed at any other point in the disk of physically possible values but the origin is unique as it is the only case that preserves rotational symmetry.

The previous priors with the addition of a delta term become

$$\bar{\kappa}(\bar{p}_0, \phi_0) = \frac{2A\delta(\bar{p}_0)}{\pi} + (1 - A)\sigma^2 \frac{2\bar{p}_0}{\pi} \quad (\text{Jeffreys}) \quad (40)$$

and

$$\bar{\kappa}(\bar{p}_0, \phi_0) = \frac{2A\delta(\bar{p}_0)}{\pi} + \sigma \frac{1 - A}{\pi} \quad (\text{uniform polar}). \quad (41)$$

These priors simplify to the more typical cases when $A = 0$. The delta term however is an interesting addition. It is only with the addition of the delta term and the use of the median estimator that one is allowed to state the source is consistent with no polarization from the posterior distribution. It seems most reasonable to assign A the value 0.5 when it is unknown if the source

² This idea of delta functions on boundaries (or even boundaries of boundaries) helps motivate and generally requires the concept of Lebesgue density. Extensions of this paper to include circular polarization, V , would have to be careful to get the Lebesgue density factors correct.

is polarized. Another interesting choice is $A = \frac{1}{1+\pi}$. This value causes the $\text{Prob}(p_0 = 0)$ for the prior to be equal to the probability density of the non-delta function component. There is a temptation to say that it recaptures some translation invariance properties that are otherwise lost upon the introduction of the delta term. It is unclear to the author if this has any mathematical significance or indeed how to justify a particular choice of A in general. Using $A = 0$ is the most conservative choice and is probably the most appropriate in many experiments unless a non-zero A is explicitly required. A non-zero value of A , as will be seen, can cause the median estimator of $\bar{p}_0$ to be zero below some critical value of $\bar{p}$. One way of interpreting this is that the median estimator determines what value of polarization must be measured to convince a person who is skeptical of polarization with degree A that a source is actually polarized. Ultimately, it is up to the researcher to decide the best prior for the experiment.

Under Eq. (38), the posterior distribution reduces to

$$\bar{B}(\bar{p}_0, \phi_0 | \bar{p}, \phi, \sigma) = \frac{A \bar{f} \bar{\kappa}_1 + (1-A) \bar{f} \bar{\kappa}_2}{A \frac{\bar{p}}{\pi} \exp\left(-\frac{\bar{p}^2}{2}\right) + (1-A) \int_{-\pi/2}^{\pi/2} \int_0^{1/\sigma} \bar{f} \bar{\kappa}_2 d\bar{p}'_0 d\phi'_0} \quad (42)$$

Figures 3 and 4 were special cases of this formula.

3.4. Other constructions

Bayesian statisticians have invented a number of other techniques for constructing priors such as using conjugate families and building hierarchical models. These are not covered.

4. Marginal distributions of polarization magnitude

The marginal posterior distribution, S , of p_0 over B is given by

$$S(p_0 | p, \phi, \sigma) = \int_{-\pi/2}^{\pi/2} B(p_0, \phi_0 | p, \phi, \sigma) d\phi_0. \quad (43)$$

The denominator of B is a function of p and ϕ only and may be pulled through the integral. The numerator becomes $\int_{-\pi/2}^{\pi/2} f(p, \phi | p_0, \phi_0, \sigma) \kappa(p_0, \phi_0) d\phi_0$. No further simplifications can be made unless a prior is chosen. While not necessary, it will frequently be the case that priors for a first measurement are independent of angle, that is, $\kappa(p_0, \phi_0) = \kappa(p_0)$. If so, then it is easy to show that

$$S(p_0 | p, \sigma) = \frac{R(p | p_0, \sigma) \kappa(p_0)}{\int_0^1 R(p | p'_0, \sigma) \kappa(p'_0) dp'_0}. \quad (44)$$

(An integration of the sampling function under consideration, $f(p, \phi | p_0, \phi_0, \sigma)$, arrives at the same result, Eq. (15), regardless if the integration is over ϕ_0 or ϕ due to the symmetry in those variables). The resulting new posterior, which may be treated as a prior for the next measurement, will not, in general, be independent of angle. The condition $\kappa(p_0, \phi_0) = \kappa(p_0)$ therefore does not extend to multiple measurements. Those are not the focus of this investigation and none of the κ -priors examined so far vary with angle. Equation (44) turns out to be exactly the same form as the posterior Rice distribution, which is now examined.

4.1. Posterior Rice distributions

Vaillancourt (2006) uses a Bayesian approach to argue for the use of the posterior Rice distribution to calculate the maximum likelihood and error bars on measured polarization. This special case was the motivation for the more general theory of this paper.

Introduce new one-dimensional prior $\tau(p_0)$. The posterior Rice distribution, ρ , is

$$\rho(p_0 | p, \sigma) = \frac{R(p | p_0, \sigma) \tau(p_0)}{\int_0^1 R(p | p'_0, \sigma) \tau(p'_0) dp'_0}. \quad (45)$$

The conversion to barred priors is given by $\bar{\tau}(\bar{p}_0) = \sigma \tau(\sigma \bar{p}_0)$. The barred version of the formula is

$$\bar{\rho}(\bar{p}_0 | \bar{p}, \sigma) = \frac{\bar{R}(\bar{p} | \bar{p}_0) \bar{\tau}(\bar{p}_0)}{\int_0^{1/\sigma} \bar{R}(\bar{p} | \bar{p}'_0) \bar{\tau}(\bar{p}'_0) d\bar{p}'_0}. \quad (46)$$

The objective prior in one dimension is $\tau(p_0) = 1$ so technically the one-dimensional Jeffreys prior and the uniform polar prior are the same. As can be seen from Eqs. (44) and (45), when $\kappa(p_0)$ is independent of angle, S and ρ have the same form with $\kappa(p_0)$ playing the role of $\tau(p_0)$. The “one-dimensional” version of the two-dimensional Jeffreys prior is $\tau(p_0) = 2p_0$ and the one-dimensional version of the two-dimensional uniform polar prior is $\tau(p_0) = 1$. There is only an immaterial factor of $1/\pi$ difference between $\kappa(p_0)$ and $\tau(p_0)$. In a slight abuse of terminology, $\tau(p_0) = 2p_0$ will still be called the “Jeffreys prior” despite switching to the one-dimensional notation. The two-dimensional theory is still within context because of Eq. (44). In barred variables, the priors are $\bar{\tau}(\bar{p}_0) = \sigma^2 2\bar{p}_0$ and $\bar{\tau}(\bar{p}_0) = \sigma$, respectively.

4.1.1. Mean estimator for the posterior Rice distribution

The mean estimator, $\hat{\bar{p}}_{0, \text{MEAN}}$, is given by

$$\hat{\bar{p}}_{0, \text{MEAN}} = \int_0^{1/\sigma} \bar{p}_0 \bar{\rho}(\bar{p}_0 | \bar{p}, \sigma) d\bar{p}_0. \quad (47)$$

Its value is always greater than zero for any given $\bar{p}$ and will usually have to be found numerically.

4.1.2. Median estimator for the posterior Rice distribution

The median estimator, $\hat{\bar{p}}_{0, \text{MED}}$, for the posterior Rice distribution is the value of the upper integrand such that

$$\int_0^{\hat{\bar{p}}_{0, \text{MED}}} \bar{\rho}(\bar{p}_0 | \bar{p}, \sigma) d\bar{p}_0 = \frac{1}{2}. \quad (48)$$

This expression will also have to be solved numerically in most cases.

4.2. Selected plots of the posterior Rice distribution

Figure 6 shows two versions of the posterior Rice distribution using the Jeffreys prior, $\bar{\tau}(\bar{p}_0) = \sigma^2 2\bar{p}_0$. Figure 6a shows the distribution near the origin and assuming $1/\sigma = 100$. Figure 6b shows the full range of $\bar{p}_0$ assuming $1/\sigma = 8$. This figure exhibits several key results. For the Jeffreys prior, the posterior Rice distribution is similar to just the transpose of the Rice distribution. This similarity breaks down at values of $\bar{p}$ near or exceeding the maximum value of $\bar{p}_0$. The mean (black) and median (blue) estimator curves, are very similar. They are also similar to the

(a) This plot uses $I_0/\Sigma = 100$. The blue dotted line is the median estimator. The black dashed line is the mean estimator.

(a) This plot uses $I_0/\Sigma = 100$. The blue dotted line is the median estimator. The black dashed line is the mean estimator.

(b) This plot uses $I_0/\Sigma = 8$ (a very small value) and shows the behavior for $\bar{p} > I_0/\Sigma$. It is clearly seen that for large $\bar{p}$, there is a large probability that the value of $\bar{p}_0$ is near the maximum (8 in this case). Figure 6a would exhibit similar behavior in its large polarization regime near and past $\bar{p} = 100$.

(b) This plot uses $I_0/\Sigma = 8$ (a very small value) showing the behavior for $\bar{p} > I_0/\Sigma$. It is clearly seen that for large $\bar{p}$, there is a large probability that the value of $\bar{p}_0$ is near the maximum (8 in this case). Fig. 7a would exhibit similar behavior in its large polarization regime near and past $\bar{p} = 100$.

Fig. 6. Two contour plots of the posterior Rice distribution (Eq. (46)) with $\bar{\tau}(\bar{p}_0) = \sigma^2 2\bar{p}_0$ and two values of σ . The contours are spaced by 0.1 intervals.

Fig. 7. Two contour plots of the posterior Rice distribution (Eq. (46)) with $\bar{\tau}(\bar{p}_0) = \sigma$ and two values of σ . The contours are spaced by 0.1 intervals.

(transposed) one-dimensional Rice distribution estimator curves but that similarity also breaks down at near the maximum value of $\bar{p}_0$.

Figure 7 uses $\bar{\tau}(\bar{p}_0) = \sigma$ and two different values of σ . Figure 7a is nearly the same plot as presented in Vaillancourt (2006), which used $\tau(p_0) = 1$ and the unstated assumption that $I_0/\Sigma \rightarrow \infty$, which allows direct integration of the denominator of Eq. (46) at the expense of having an unnormalizable prior. While this approximation will be acceptable most of the time when the signal-to-noise is large, it is unphysical. The true value of $\bar{p}_0$ cannot exceed $1/\sigma (=I_0/\Sigma)$, which is finite. This is best seen by plotting the full range of $\bar{p}_0$ as is done in Fig. 7b for $I_0/\Sigma = 8$ (a small value). There it is seen that for values of $\bar{p}$ that are near or greater than the maximum possible value of $\bar{p}_0$ (8 in this case), the probability density function “crowds up” near the

maximum value. This time the mean (black) and median (blue) curves are rather different. The median curve prefers value closer to $\bar{p}_0 = 0$. The median estimator seems to be a better statistic in this case as the mean estimator curve is affected too much by outlier possibility.

A delta term may also be added to the prior as was previously motivated. Let $\tau(p_0) = A\tau_1(p_0) + (1-A)\tau_2(p_0)$, where $\tau_1(p_0)$ is normalized a delta term at the origin and $\tau_2(p_0)$ is a normalized component without any delta terms. The normalized barred prior is $\bar{\tau}(\bar{p}_0) = A\bar{\tau}_1(\bar{p}_0) + (1-A)\bar{\tau}_2(\bar{p}_0)$. This time $\tau_1(p_0) = 2\delta(p_0)$ and $\bar{\tau}_1(\bar{p}_0) = 2\delta(\bar{p}_0)$. Under this prior, Eq. (46) reduces to

$$\bar{p}(\bar{p}_0|\bar{p}, \sigma) = \frac{A\bar{R}\bar{\tau}_1 + (1-A)\bar{R}\bar{\tau}_2}{A\bar{p} \exp\left(-\frac{\bar{p}^2}{2}\right) + (1-A) \int_0^{1/\sigma} \bar{R}\bar{\tau}_2 d\bar{p}'_0}. \quad (49)$$

Fig. 8. These four panels show the posterior distribution $\bar{\rho}(\bar{p}_0|\bar{p}, \sigma)$ with prior $\bar{\tau}(\bar{p}_0) = 2A\delta(\bar{p}_0) + (1-A)\sigma^2 2\bar{p}_0$ for four combinations of value of A and σ . The blue dotted line is the median estimator. The black dashed line is the mean estimator. The red, dashed line at $\bar{p}_0 = 0$ is a reminder of the finite amount of probability that exists along the $\bar{p}$ -axis due to the delta function term. No such line is needed in Figs. 6a and 6b because $A = 0$ there and the delta term does not contribute. This probability along the $\bar{p}$ -axis is plotted underneath each contour plot.

Here, the function arguments have been dropped for space. For the two-component prior, the definition of the median estimator (Eq. (48)) leads to the following expression

$$\int_0^{\bar{p}_{0,\text{MED}}} \bar{R} \bar{\tau}_2 d\bar{p}_0 = \frac{1}{2} \int_0^{1/\sigma} \bar{R} \bar{\tau}_2 d\bar{p}_0 - \frac{A}{2(1-A)} \bar{p} e^{-\bar{p}^2/2}. \quad (50)$$

Figures 8 and 9 show Eq. (49) for the Jeffreys and uniform priors, respectively, and four different combination of values of A and σ each. The red, dashed lines at $\bar{p}_0 = 0$ are a reminder that Eq. (49) contains a delta term that cannot be plotted in the usual fashion. Finite probability is attached to each point along the red line and is graphed beneath each contour plot. This is critical to fully understand these plots. Several new features are seen for non-zero A . In Fig. 8 for the Jeffreys prior, the probability band seems to bend towards the $\bar{p}$ -axis (Fig. 8a) but in actuality abruptly turns and heads towards the $\bar{p}_0$ -axis (Fig. 8b). This behavior is however strongly damped by residual probability that

Fig. 9. These four panels show the posterior distribution $\bar{\rho}(\bar{p}_0|\bar{p}, \sigma)$ with prior $\bar{\tau}(\bar{p}_0) = 2A\delta(\bar{p}_0) + (1-A)\sigma$ for four combinations of value of A and σ . The blue dotted line is the median estimator. The black dashed line is the mean estimator. The red, dashed line at $\bar{p}_0 = 0$ is a reminder of the finite amount of probability that exists along the $\bar{p}$ -axis due to the delta function term. No such line is needed in Figs. 7a and 7b because $A = 0$ there and the delta term does not contribute. This probability along the $\bar{p}$ -axis is plotted underneath each contour plot.

the polarization is actually zero (as in Figs. 8c and d). Figure 9 for the uniform prior is similar except this time the probability band, after heading towards the $\bar{p}$ -axis as well, turns towards the origin if the behavior is not damped by the probability that the polarization is zero (as in Figs. 9c and d).

The mean and median estimator curves in Figs. 8 and 9 have changed with the introduction of a non-zero A . The mean curve is “pulled” more towards the $\bar{p}$ -axis at small values of $\bar{p}$ for larger values of A or smaller values of σ . The median curve has a totally new feature: it intercepts the $\bar{p}$ -axis at a critical value, $\bar{p}_{\text{crit}}$, and is zero for smaller values. This occurs when the right-hand side of Eq. (50) equals zero. Tables of some critical values, including those corresponding to the figures, are presented in Appendix A.

When $A = 0$ and $\sigma = 0.01$, a contour plot of Eq. (49) for the Jeffrey prior reproduces Fig. 6a. This plot resembles the Rice distribution in Fig. 2a, only transposed! Figure 8b is a good example of the transitional form that occurs in the contour shape between the standard $A = 0$ form and the form seen for larger

values of A and/or smaller values of σ . This is an important result. It means that the Rice distribution results can effectively be transposed to find best estimates of $\bar{p}_0$ for small values of polarization when $1/\sigma \rightarrow \infty$.

When $A = 0$ and $\sigma = 0.01$, a contour plot of Eq. (49) for the uniform prior reproduces Fig. 7a. There is no delta function and the global maximum of this distribution occurs at the origin. The wrench-like shape characteristic of this plot arises from small values of A for a given value of σ . Figure 9b is a good example of the transitional form.

All the usual mathematical machinery may now be applied to $\bar{p}(\bar{p}_0|\bar{p}, \sigma)$ to produce estimators for $\bar{p}_0$ given a measured value of $\bar{p}$.

Previously “bias corrections” were performed on measurements at low signal-to-noise to make them closer to zero. As is seen in the figures with $A = 0$, the mean and median estimators seem to act for very small values of p like a “bias correction” that goes the wrong way! (The median estimator curves with p -axis intercepts are an exception.) Bias corrections should not be used. A non-zero value for the estimate of p_0 even when $p = 0$ is not a deficiency with the estimators. It is a result that can be understood intuitively when all possible values of p_0 and ϕ_0 that could have produced a given p and ϕ are considered.

5. Marginal distributions of polarization angle

The marginal posterior distribution, T , of ϕ_0 is given by

$$T(\phi_0|p, \phi, \sigma) = \int_0^1 B(p_0, \phi_0|p, \phi, \sigma) dp_0. \quad (51)$$

As before, the denominator of B is a function of p and ϕ only and may be pulled through the integral. The numerator becomes $\int_0^1 f(p, \phi|p_0, \phi_0, \sigma) \kappa(p_0, \phi_0) dp_0$. If the prior is a function of p_0 to some non-negative integer power n (i.e., $\kappa(p_0) \propto p_0^n$), there appears to exist a family of solutions via repeated integration by parts. Explicit solutions were found up to $n = 10$ using the computer algebra system Mathematica. Beyond $n = 1$, these solutions quickly become impractically large.

The marginal distribution, G , of ϕ over f has been investigated by several authors (Vinokur 1965; Clarke & Stewart 1986; Naghizadeh-Khouei & Clarke 1993). Its defining formula in terms of the sky angle is

$$G(\phi|p_0, \phi_0, \sigma) = \int_0^\infty f(p, \phi|p_0, \phi_0, \sigma) dp.$$

Unfortunately, the solution to this equation in Clarke & Stewart (1986) has some sign mistakes and the derivation in Appendix B of Naghizadeh-Khouei & Clarke (1993) confuses the sky angle, ϕ , with the angle in the q - u plane, θ . It is worth explicitly stating both versions to avoid confusion:

$$G(\phi|p_0, \phi_0, \sigma) = \frac{1}{\sqrt{\pi}} \left(\frac{1}{\sqrt{\pi}} + \eta_0 e^{\eta_0^2} [1 + \text{erf}(\eta_0)] \right) e^{-\frac{p_0^2}{2\sigma^2}} \quad (52)$$

and

$$G'(\theta|p_0, \theta_0, \sigma) = \frac{1}{2\sqrt{\pi}} \left(\frac{1}{\sqrt{\pi}} + \eta'_0 e^{\eta'^0_2} [1 + \text{erf}(\eta'_0)] \right) e^{-\frac{p_0^2}{2\sigma^2}}, \quad (53)$$

where $\eta_0 = \frac{p_0}{\sqrt{2}\sigma} \cos(2(\phi - \phi_0))$ and $\eta'_0 = \frac{p_0}{\sqrt{2}\sigma} \cos(\theta - \theta_0)$. It is useful to notice that the distribution is symmetric about ϕ_0 and using a value of ϕ_0 different than zero only translates the distribution along the angular axis. Therefore when calculating

confidence intervals, only $\phi_0 = 0$ need be used. Unlike the Rice distribution, which is independent of the angular variable, this distribution depends on p_0 . G is normalized (i.e., $\int_{-\pi/2}^{\pi/2} G d\phi = 1$).

The distribution T cannot in general be related to G as S was to R through ρ , even if the prior is independent of angle, because f lacks the symmetry in p and p_0 that it has for ϕ and ϕ_0 . T and G may have different form. If however $\kappa(p_0, \phi_0) \propto p_0$, as with the Jeffreys prior, it is easy to show that $T \rightarrow G$ as $1/\sigma \rightarrow \infty$.

Figure 10 compares T to G for the Jeffreys prior, i.e., $\kappa(p_0, \phi_0) = 2p_0/\pi$, and $I_0/\Sigma = 100$ and three different signal-to-noise values. As expected, the T distribution and the G distribution are almost identical because $1/\sigma$ is fairly large. The two sets of curves overlap in the figure but the T curves are not precisely identical to the G curves. The difference between the T and G curves is larger if $1/\sigma$ is small or for polarization measurements near $1/\sigma$. In practice, the classical error bars on the angle based on G are perfectly acceptable under this prior.

Figure 11 compares T to G for the case of a uniform polar prior, i.e., $\kappa(p_0, \phi_0) = 1/\pi$, and $I_0/\Sigma = 100$. In the figure, T (blue) is plotted for $\phi = 0$ and $p/\sigma = 0.5$ (dotted), 1.0 (solid), 2.0 (dashed) and G (red) is plotted for $\phi_0 = 0$ and $p_0/\sigma = 0.5$ (dotted), 1.0 (solid), 2.0 (dashed). The figure clearly demonstrates that using G , the non-Bayesian formula, to compute errors bars results in errors bars that are too small. This is especially important when the polarization degree to σ ratio is about one.

If a delta term is added to the priors as was done with the polarization magnitude plots, the result can roughly be described by saying that the distribution is scaled by $(1 - A)$ and it is then shifted upwards by a constant of A/π . This constant shift occurs because the origin has an indeterminate angle. The equation in barred variables is

$$\bar{T}(\phi_0|\bar{p}, \phi, \sigma) = \frac{\frac{A}{\pi} \bar{p} \exp\left(-\frac{\bar{p}^2}{2}\right) + (1-A) \int_0^{1/\sigma} \bar{f} \bar{\kappa}_2 d\bar{p}_0}{A \frac{A}{\pi} \exp\left(-\frac{\bar{p}^2}{2}\right) + (1-A) \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_0^{1/\sigma} \bar{f} \bar{\kappa}_2 d\bar{p}_0 d\phi_0}. \quad (54)$$

When $A = 1.0$, the distribution is flat with value $1/\pi$, as it must be to be normalized.

6. Summary of derivations

Figure 12 is a flow chart to help remember the relationships among the many distributions presented. All the main functions result from a change of variables, Bayes Theorem, or marginalization (denoted by “*Marg.*” in the figure). Attempts were made to remain consistent with the function names used in previous published papers when possible.

7. Confidence intervals (aka credible sets)

All the pieces needed to construct confidence intervals (known as “credible sets” in Bayesian statistics) for p_0 and ϕ_0 have been given. In general, these sets will be two-dimensional. Usually they will be constructed so as to contain the “best” estimate ($\hat{p}_a$ and $\hat{\phi}_a$) that was found. In practice, especially with polarimetric spectra, using two-dimensional sets as “error bars” is impossible. In these cases, the marginal distributions, T and S , may be used to construct error bars.

High-quality confidence interval tables, suitable for use in data reduction pipelines, will be the subject of a future paper.

Fig. 10. This plot compares the classical angle distribution, $G(\phi|p_0, \phi_0, \sigma)$, to the Bayesian distribution, $T(\phi_0|p, \phi, \sigma)$, with $\kappa(p_0, \phi_0) = 2p_0/\pi$ and $I_0/\Sigma = 100$. The Bayesian curves are virtually identical to the classical curves and overlap in the figure. Both plots simply use zero for the angle since the choice is immaterial. Three different values for the signal-to-noise, i.e., p_0/σ (G) or p/σ (T), are used: 0.5 (dashed lines), 1.0 (solid lines), and 2.0 (dotted lines).

Fig. 11. This plot compares the classical angle distribution, $G(\phi|p_0, \phi_0, \sigma)$ (red), to the Bayesian distribution, $T(\phi_0|p, \phi, \sigma)$ (blue), with $\kappa(p_0, \phi_0) = 1/\pi$ and $I_0/\Sigma = 100$. The Bayesian approach now suggests the classical approach underestimates the error bars. Both plots simply use zero for the angle since the choice is immaterial. Three different values for the signal-to-noise, i.e., p_0/σ (G) or p/σ (T), are used: 0.5 (dashed lines), 1.0 (solid lines), and 2.0 (dotted lines).

8. Conclusion

A foundation for a Bayesian treatment of polarization measurements has been presented based when $V = 0$. The treatment

Fig. 12. This figure is a flow chart to help remember the relationships between the distributions presented. “Marg.” stands for marginalization over the accompanying variable.

covers the regime where the measured polarization has large or small Stokes parameters but all the intensities that were used in the calculation of the Stokes parameters are large enough that Poisson statistics is unimportant. This justifies the use of a gaussian approach for the Stokes parameters’ error.

A Bayesian analysis of polarization allows the observer some freedom in choosing a prior and loss function. These should be stated and justified when presenting one’s results. In particular, the importance of choosing an appropriate prior has been stressed.

The transposed Rice distribution is intimately connected with the posterior distribution with a two-dimensional Jeffreys prior. As $1/\sigma \rightarrow \infty$, the distributions are the same. The classical angular distribution is also the same as the posterior angular distribution in this limit. A uniform polar prior may be better for sources with small expected polarization. In this case, the posterior distributions are rather different than the classical ones.

Further advances may be made by generalizing the analysis to include ellipsoidal, perhaps even correlated, distributions for Q and U , Poisson statistics, circular polarization, or repeated measurements. Monte Carlo methods will likely be necessary to handle the integrations for more general approaches. Polarimetric spectra and time series data that involve measurements that are not independent and may be highly correlated; modeling the wavelength-dependence or time-evolution could further improve results in those situations.

Table A.1. Median estimator critical values of p such that $\text{Prob}(p_0 = 0) = 1/2$ for the Jeffreys prior.

A	$I_0/\Sigma = 25$	$I_0/\Sigma = 100$	$I_0/\Sigma = 500$	$I_0/\Sigma = 1000$
0.01	0.060649	0.028007	0.0075581	0.0041300
0.02	0.076999	0.030415	0.0079219	0.0042963
0.1	0.10654	0.035553	0.0087356	0.0046744
0.2	0.11810	0.037765	0.0090994	0.0048450
0.3	0.12519	0.039166	0.0093331	0.0049550
0.4	0.13071	0.040278	0.0095206	0.0050431
0.5	0.13558	0.041273	0.0096897	0.0051231
0.6	0.14029	0.042244	0.0098556	0.0052013
0.7	0.14524	0.043277	0.010033	0.0052856
0.8	0.15106	0.044505	0.010246	0.0053869
0.9	0.15942	0.046291	0.010558	0.0055350
0.99	0.18190	0.051210	0.011430	0.0059525

The results presented have focused on the interpretation of the quantities p_0 and ϕ_0 as polarization degree and angle but the formulas are applicable to any application that uses a bounded radial function of the form $\sqrt{x^2 + y^2}$, where x and y are measured with equal gaussian error.

Appendix A: Critical values for the median estimator curves with the uniform polar prior

Tables A.1 and A.2 show some critical values of the median estimator curve for the Jeffreys and uniform polar prior and an array of choices for A and $1/\sigma$ ($= I_0/\Sigma$). Entries with “...” never intercept the $\bar{p}$ -axis. As is expected, for a given value of A , the critical value decreases as the data quality increases (i.e., $1/\sigma$ gets larger). Also as expected, the critical values increase as A increases for a given value of data quality. These tables are created using Eq. (50). The critical values occur when the right-hand side of that equation is equal to zero.

Unbarred variables are used in the tables. The behavior in barred variables is somewhat counter-intuitive and should be avoided to prevent confusion. As can be seen by comparing the median estimator intercepts for $\bar{p}_0$ in Figs. 8c, d, 9c and d, the intercept is larger for $\sigma = 0.01$ than $\sigma = 0.04$ even though the data quality is better for $\sigma = 0.01$. This does not mean that, as the data gets better, a larger measured value of polarization is required to prove polarization is non-zero! There is a human tendency to interpret the barred plots as if $\sigma = 1$, but in practice σ is always less than 1. When converted to unbarred variables, the polarization behaves as expected.

More tables such as this are intended to be included in a future work.

Appendix B: “Projected” prior

A third interesting prior was discussed in Sect. 3. This prior, called here the “projected” prior, is motivated by considering an instrument that only measures linear polarization but ignores circular polarization. While a distribution for q and u can be derived for this case and the machinery given in the main text can be applied to it, this is somewhat inappropriate since it violates the leading assumption that $V = 0$ upon which the subsequent equations were derived. Nevertheless, it is worth recording here some of the salient formulas.

The starting point is a normalized uniform Cartesian prior on the Poincaré sphere

$$\kappa_C(q_0, u_0, v_0) = \frac{3}{4\pi}. \quad (\text{B.1})$$

Table A.2. Median estimator critical values of p such that $\text{Prob}(p_0 = 0) = 1/2$ for the uniform polar prior.

A	$I_0/\Sigma = 25$	$I_0/\Sigma = 100$	$I_0/\Sigma = 500$	$I_0/\Sigma = 1000$
0.01	...	...	0.0042099	0.0024939
0.02	...	0.013263	0.0049980	0.0028143
0.1	0.066096	0.025408	0.0063962	0.0034318
0.2	0.089414	0.028977	0.0069392	0.0036822
0.3	0.10092	0.031047	0.0072724	0.0038378
0.4	0.10910	0.032617	0.0075324	0.0039599
0.5	0.11591	0.033976	0.0077618	0.0040683
0.6	0.12221	0.035270	0.0079835	0.0041734
0.7	0.12861	0.036617	0.0082172	0.0042846
0.8	0.13590	0.038183	0.0084923	0.0044159
0.9	0.14603	0.040404	0.0088880	0.0046056
0.99	0.17197	0.046273	0.0099570	0.0051216

Notes. Entries marked “...” mean that $\text{Prob}(p_0 = 0) < 1/2$ for all values of p so no solution exists for p_{crit} .

Fig. B.1. The posterior Rice distribution using the “projected” prior of Eq. (B.5) with $I_0/\Sigma = 8$ (a very small value). The distribution of the measured values near the physically-maximum value is notably different than the previous corresponding figures for the Jeffreys and uniform polar prior but the small-value behavior is nearly exact to the Jeffreys case.

Upon projection onto the q - u plane, the probability density function becomes

$$\kappa_{\text{proj}}(q_0, u_0) = \frac{3}{2\pi} \sqrt{1 - q_0^2 - u_0^2}. \quad (\text{B.2})$$

Equivalently in polar coordinates, this is

$$\kappa_{\text{proj}}(p_0, \phi_0) = \frac{3}{\pi} p_0 \sqrt{1 - p_0^2} \quad (\text{B.3})$$

or

$$\bar{\kappa}_{\text{proj}}(\bar{p}_0, \phi_0) = \frac{3\sigma^2}{\pi} \bar{p}_0 \sqrt{1 - \sigma^2 \bar{p}_0^2}. \quad (\text{B.4})$$

Notice that for small values of $\bar{p}_0$, Eq. (B.5) has the same form as Eq. (35) up to an irrelevant scale factor. Thus for small values of $\bar{p}_0$, the posterior distribution should be similar to that of the Jeffreys posterior. The large value behavior will be slightly different. Marginalization of $\bar{\kappa}_{\text{proj}}(\bar{p}_0, \phi_0)$ over ϕ produces

$$\bar{\tau}_{\text{proj}}(\bar{p}_0) = 3\sigma^2 \bar{p}_0 \sqrt{1 - \sigma^2 \bar{p}_0^2} \quad (\text{B.5})$$

which is plotted in Fig. B.1 and may be contrasted with the previous similar figures.

Acknowledgements. The authors acknowledge support from the Millennium Center for Supernova Science (MCSS) through grant P10-064-F funded by “Iniciativa Científica Milenio”. The contour plots were generated using computer algebra system *Mathematica* (version 7.01.0). The authors wish to thank Alejandro Clocchiatti and Lifan Wang for helpful discussion and the anonymous referee for valuable comments.

References

- Berger, J. O. 1985, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. (Springer)
- Bertrand, J. 1888, *Calcul des probabilités* (Paris: Gauthier-Villars)
- Clarke, D., & Stewart, B. G. 1986, *Vist. Astron.*, 29, 27
- Clarke, D., Stewart, B. G., Schwarz, H. E., & Brooks, A. 1983, *A&A*, 126, 260
- del Toro Iniesta, J. C. 2003, *Introduction to Spectropolarimetry* (Cambridge University Press)
- Di Porto, P., Crosignani, B., Ciattoni, A., & Liu, H. C. 2010 [arXiv:1008.1878]
- Jaynes, E. T. 1973, *Foundations of Physics*, 3, 477
- Jeffreys, H. 1939, *Theory of probability* (Oxford University Press)
- Jeffreys, H. 1946, *Proc. Royal Society of London Series A, Mathematical and Physical Sciences*, 186, 453
- Keller, C. U. 2002, in *Astrophysical Spectropolarimetry*, ed. J. Trujillo-Bueno, F. Moreno-Insertis, & F. Sánchez, 303
- Landi Degl’Innocenti, E. 2002, in *Astrophysical Spectropolarimetry*, ed. J. Trujillo-Bueno, F. Moreno-Insertis, & F. Sánchez, 1
- Maronna, R., Feinstein, C., & Clocchiatti, A. 1992, *A&A*, 260, 525
- Naghizadeh-Khouei, J., & Clarke, D. 1993, *A&A*, 274, 968
- Patat, F., & Romaniello, M. 2006, *PASP*, 118, 146
- Rice, S. O. 1945, *Bell Systems Tech. J.*, 24, 46
- Robert, C. P. 1994, *The Bayesian Choice: A Decision-Theoretic Motivation* (Springer)
- Robert, C. P., Chopin, N., & Rousseau, J. 2009a, *Statist. Sci.*, 24, 141
- Robert, C. P., Chopin, N., & Rousseau, J. 2009b, *Statist. Sci.*, 24, 191
- Serkowski, K. 1958, *Acta Astron.*, 8, 135
- Simmons, J. F. L., & Stewart, B. G. 1985, *A&A*, 142, 100
- Tissier, P. E. 1984, *The Mathematical Gazette*, 68, pp. 15
- Vaillancourt, J. E. 2006, *PASP*, 118, 1340
- Vinokur, M. 1965, *Ann. d’Astrophys.*, 28, 412
- Wang, L., Wheeler, J. C., & Hoefflich, P. 1997, *ApJ*, 476, L27
- Wardle, J. F. C., & Kronberg, P. P. 1974, *ApJ*, 194, 249