Pitfalls of statistics-limited X-ray polarization analysis

V. Mikhalev^{1,2}

¹ KTH Royal Institute of Technology, Department of Physics, 106 91 Stockholm, Sweden
e-mail: mikhalev@kth.se

² The Oskar Klein Centre for Cosmoparticle Physics, AlbaNova University Centre, 106 91 Stockholm, Sweden

Received 19 September 2017 / Accepted 23 March 2018

ABSTRACT

Context. One of the difficulties with performing polarization analysis is that the mean polarization fraction of sub-divided data sets is larger than the polarization fraction for the integrated measurement. The resulting bias is one of the properties of the generating distribution discussed in this work. The limitations of Gaussian approximations in standard analysis based on Stokes parameters for estimating polarization parameters and their uncertainties are explored by comparing with a Bayesian analysis. The effect of uncertainty on the modulation factor is also shown, since it can have a large impact on the performance of gamma-ray burst polarimeters. Results are related to the minimum detectable polarization (MDP), a common figure of merit, making them easily applicable to any X-ray polarimeter.

Aims. The aim of this work is to quantify the systematic errors induced on polarization parameters and their uncertainties when using Gaussian approximations and to show when such effects are non-negligible.

Methods. The probability density function is used to deduce the properties of reconstructed polarization parameters. The reconstructed polarization parameters are used as sufficient statistics for finding a simple form of the likelihood. Bayes theorem is used to derive the posterior and to include nuisance parameters.

Results. The systematic errors originating from Gaussian approximations as a function of instrument sensitivity are quantified here. Different signal-to-background scenarios are considered making the analysis relevant for a large variety of observations. Additionally, the change of posterior shape and instrument performance MDP due to uncertainties on the polarimetric response of the instrument is shown.

Key words. polarization – methods: data analysis – methods: statistical

1. Introduction

The first observations of astrophysical X-ray polarization were made more than forty years ago (Novick et al. 1972; Weisskopf et al. 1976). The field has been reinvigorated in the past decade by a series of measurements by satellite, including INTEGRAL/SPI (Dean et al. 2008; Chauvin et al. 2013); INTEGRAL/IBIS (Forot et al. 2008; Moran et al. 2016), AstroSat/CZTI (Vadawale et al. 2018); and IKAROS/GAP (Yonetoku et al. 2011) as well as balloon-borne instruments PoGOLite (Chauvin et al. 2016) and PoGO+ (Chauvin et al. 2017). Results from several on-going missions are expected in the near future (POLAR, Produit et al. 2005; X-Calibur, Beilicke et al. 2014), and a dedicated satellite mission is in development (IXPE, Weisskopf et al. 2016). Although some instruments, for example, IXPE, will be able to make polarization measurements of astrophysical sources with high precision, where the statistical analysis becomes relatively trivial, it is always desirable to observe weaker sources or to sub-divide the data. Fine splitting of data may be necessary for understanding physical phenomena, for example, gamma-ray bursts (GRBs) and pulsars require temporal binning, nebulae require spatial binning, and spectral binning is interesting for all objects. It is therefore important to know when the number of photons is sufficient for making an accurate analysis using simple methods and when a more rigorous approach is necessary.

The parameters describing linear polarization are the polarization fraction and the polarization angle. In the frequentist approach, the major challenge in estimating the polarization

fraction is that it is a positive definite quantity, meaning that a non-zero fraction is measured even for an unpolarized source, thus introducing a bias. Ways of correcting for this bias have been studied by Simmons & Stewart (1985) and more recently by Maier et al. (2014). A Bayesian approach was first introduced by Vaillancourt (2006) and extensively expanded upon by Quinn (2012), where the shapes of the resulting parameter distributions are described. These preceding works focus on optical measurements of polarization for which some formulae differ from the X-ray counterpart (Kislat et al. 2015).

This work quantifies, as a function of measurement sensitivity, the error incurred when using the conventional Stokes parameter analysis. This is shown for both the polarization fraction and angle as well as for their uncertainties. Henceforth the expression “low statistics” is used to indicate poor data quality having a low $S/\sqrt{N}$, where S is the number of signal photons and N is the total number of photons (signal and background). Since a point-source polarimeter working in the low-statistics regime is likely to have a low signal-to-background ratio R , whereas a GRB polarimeter is expected to have a higher R (especially if the GRB is bright and short in duration), the study is conducted for different values of R where meaningful. Additionally, the effects of uncertainty on the modulation factor μ_0 , defined as the polarization fraction measured for a 100% polarized beam, on the performance of the polarimeter are studied. The uncertainty is typically large for GRB polarimeters which are not optimized for localization of GRBs since μ_0 varies with the photon incidence angle.

2. Parameter estimation

For a Compton scattering (Lei et al. 1997) or photo-electric polarimeters (Bellazzini & Muleri 2010) with a uniform response, the conditional distribution for a measured polarization angle ψ_i given a polarization fraction p_0 , a polarization angle ψ_0 and a modulation factor μ_0 follows

$$f(\psi_i|p_0, \psi_0, \mu_0) = \frac{1}{2\pi(S+B)} \times (S \times (1 + \mu_0 p_0 \cos(2(\psi - \psi_0))) + B), \quad (1)$$

where $\psi = \phi - 90^\circ$ for the Compton process and $\psi = \phi$ for the photo-electric effect, ϕ is the measured scattering angle, S is the number of signal photons and B is the number of background photons. Treating signal and background as latent variables is beyond the scope of this paper so it is assumed that the background is unpolarized and is known with much higher precision than the polarization parameters. The modulation factor μ_0 is a property of the polarimeter and is derived from calibration. In what follows, the notation with subscript “0” means a physical parameter generating the data, in this case a set of measured scattering angles transformed to the polarization frame $\{\psi_i\}$, while subscript “r” denotes reconstructed parameters from this set. The reconstructed parameters are actually sufficient statistics for the data set allowing it to be represented by 2 scalars rather than a set of N angles.

The two most common ways of computing p_r and ψ_r are by performing a χ^2 -fit of Eq. (1) to a histogram of scattering angles or by computing the Stokes parameters. This work only considers the latter as it avoids complicating the analytical form of the likelihood due to binning effects.

2.1. Stokes parameters

Polarimeters operating in the radio, infra-red, or optical domain measure photon intensities rather than individual photons as opposed to X-ray polarimeters. Clarke et al. (1983) discuss how the Stokes parameter distributions vary depending on the instrumental technique used for measuring these intensities. Since Eq. (1) is not used in the low-energy domain, not all results presented in that publication can be extrapolated to the X-ray energy band. In particular, the standard deviations and the correlation coefficient of Stokes parameters are affected.

In the X-ray energy band, the Stokes parameters are derived from individual photon events comprising two quantities

$$\begin{aligned} q_i &= \cos(2\psi_i) \\ u_i &= \sin(2\psi_i). \end{aligned} \quad (2)$$

For a total number of photons $N = S + B$ the normalized Stokes parameters are written as

$$\begin{aligned} Q_r &= \frac{1}{S} \sum_{i=1}^N q_i \\ U_r &= \frac{1}{S} \sum_{i=1}^N u_i. \end{aligned} \quad (3)$$

The normalization is only proportional to the signal because the background is assumed to be unpolarized. Therefore the background contributes only to the variance of (Q_r, U_r) .

The Central Limit Theorem (CLT) makes Q_r and U_r Gaussian distributed as long as N is sufficiently large. This is referred

to here as the CLT approximation. Thus Q_r and U_r follow the bivariate Gaussian distribution

$$\begin{aligned} B(Q_r, U_r|Q_0, U_0) &= \frac{1}{2\pi\sigma_Q\sigma_U\sqrt{1-\rho^2}} \\ &\times \exp\left[-\frac{1}{2(1-\rho^2)}\left(\frac{(Q_r-Q_0)^2}{\sigma_Q^2} + \frac{(U_r-U_0)^2}{\sigma_U^2} - \frac{2\rho(Q_r-Q_0)(U_r-U_0)}{\sigma_Q\sigma_U}\right)\right]. \end{aligned} \quad (4)$$

Since the mean, the standard deviation, and the correlation coefficient are sufficient statistics for Gaussian distributed data, so are Q_r and U_r together with the second moments of the data derived in Appendix B. Although not written explicitly in the conditioning, Q_0 , U_0 , σ_U , σ_Q and ρ are functions of p_0 and ψ_0 .

2.2. Polar coordinates

Stokes parameters (Q, U) are in Cartesian coordinates but can be transformed to polar coordinates to allow the polarization parameters to be expressed in terms of (p, ψ) . The sufficient statistics become

$$p_r = 2\sqrt{Q_r^2 + U_r^2}/\mu_r, \quad (5)$$

$$\psi_r = \frac{1}{2} \arctan \frac{U_r}{Q_r}, \quad (6)$$

as per the derivations in Appendix A.

Now Eq. (4) can be transformed to polar coordinates by using the general coordinate transformation for a probability density function (p.d.f.) yielding the likelihood

$$L(p_r, \psi_r|p_0, \psi_0) = B(Q_r, U_r|Q_0, U_0) \times |\det(J)|, \quad (7)$$

where $\det(J)$ is the determinant of the Jacobian given by

$$\begin{vmatrix} \frac{\partial Q_r}{\partial p_r} & \frac{\partial Q_r}{\partial \psi_r} \\ \frac{\partial U_r}{\partial p_r} & \frac{\partial U_r}{\partial \psi_r} \end{vmatrix} = \begin{vmatrix} \frac{\mu_r}{2} \cos(2\psi_r) & -\mu_r p_r \sin(2\psi_r) \\ \frac{\mu_r}{2} \sin(2\psi_r) & \mu_r p_r \cos(2\psi_r) \end{vmatrix} = \frac{p_r \mu_r^2}{2}. \quad (8)$$

Although the polarization parameters have a very similar form to that of the sufficient statistics (simply replacing the subscript “r” by “0”)

$$p_0 = 2\sqrt{Q_0^2 + U_0^2}/\mu_0, \quad (9)$$

$$\psi_0 = \frac{1}{2} \arctan \frac{U_0}{Q_0}, \quad (10)$$

p_r does not correspond to the most probable estimate of p_0 . This occurs because $(\sigma_Q, \sigma_U, \rho)$ depend on (p_0, ψ_0) and the fact that $0 \leq p_0 \leq 1$ (since a polarization greater than 100% is unphysical) but there is no upper limit on p_r , that is, $0 \leq p_r$. Therefore, using the sufficient statistic p_r as the estimator of the polarization fraction p_0 incurs an error.

The error is non-negligible for the one-dimensional likelihood of p_0 (obtained after marginalizing over ψ_0) if the statistical significance of the measurement is low. It results in $\arg\max_{p_0} L(p_r, \psi_r|p_0) \neq p_r$ and $\langle p_r \rangle \neq p_0$, where $\langle p_r \rangle$ is the expected value of p_r . Hence p_r is neither the maximum likelihood nor an unbiased estimator of p_0 . This is the case even if $\mu_r = \mu_0$ which occurs when there is no uncertainty on μ as is assumed throughout this section.

As the statistical precision of a measurement improves $\lim_{S \rightarrow \infty} \langle p_r \rangle = p_0$ and $\lim_{S \rightarrow \infty} \arg\max_{p_0} L(p_r, \psi_r|p_0) = p_r$. Conversely, due to symmetry, there is no bias on ψ so $\langle \psi_r \rangle = \psi_0$ and $\arg\max_{\psi_0} L(p_r, \psi_r|\psi_0) = \psi_r$.

Fig. 1. The comparison of the CLT approximation to the full likelihood, as given by Eqs. (7) and (11), respectively. For $p_0 = 0.5$, both pure signal and mixed scenarios are well described by the CLT approximation. For $p_0 = 1.0$, the pure signal scenario deviates farther from the approximation than the scenario with background because the pure signal scenario has fewer photons. All scenarios use $\mu_0 = 0.5$ and $\text{MDP} = 2$.

2.3. Central Limit Theorem approximation of the likelihood

The CLT approximation in $B(Q_r, U_r|Q_0, U_0)$ provides an always easily computable analytical form of the likelihood given by Eq. (7). However, if the polarization fraction is high and the number of photons is low, the likelihood is not well described by such an approximation. The full form of the likelihood is given by a product over Eq. (1).

$$L(\{\psi_i\}|p_0, \psi_0, \mu_0) = \prod_{i=1}^N \frac{1}{2\pi(S+B)} \times (S \times (1 + \mu_0 p_0 \cos(2(\psi_i - \psi_0))) + B). \quad (11)$$

It has N factors and is therefore cumbersome to evaluate when the number of photons is large. This is unlike optical polarimetry where photon intensities of (Q, U) are directly measured instead of individual scattering angles.

Figure 1 shows the difference between the CLT approximation and $L(p_r|p_0, \psi_0, \mu_0)$ as calculated by generating a data set $\{\psi_i\}$ from Eq. (11), computing p_r using the Stokes parameters, repeating for many iterations and making a normalized histogram of p_r .

In what follows, the concept of minimum detectable polarization (MDP) at 99% confidence level is important. The MDP Weisskopf et al. (2010) is given by

$$\text{MDP} = \frac{4.29}{\mu_0 S} \sqrt{S+B}. \quad (12)$$

Its statistical meaning is that, given an unpolarized source ($p_0 = 0$), the probability of measuring $p_r > \text{MDP}$ is 1%. This quantity is a standard figure of merit for polarimeter performance and can easily be calculated for any measurement. The number of signal photons in Fig. 1 is chosen such that $\text{MDP} = 2$, since this is the highest (“worst”) MDP considered in later sections. As seen from Fig. 1, the CLT approximation becomes more accurate as the number of photons increases and the polarization fraction p_0 decreases. In particular, only measurements with high signal-to-background ratio and high polarization fraction require the full likelihood given by Eq. (11). Although the

error incurred using the CLT approximation is negligible in most cases, it becomes larger when using Eq. (7) to derive the posterior, as is discussed below. Figure 1 is for qualitative purposes only; the equivalent figure for ψ_r has been omitted since it leads to the same conclusions.

2.4. Magnitude and importance of bias

Bias is a frequentist concept which relies on fixing (p_0, ψ_0) and investigating $\langle p_r \rangle$. This approach provides an intuitive understanding for how an unpolarized source can produce a polarized signal ($p_r > 0$). Several previous measurements use the sufficient statistic p_r as an estimate of the polarization fraction p_0 (e.g., Weisskopf et al. 1978; Slowikowska et al. 2009) and it is therefore necessary to understand when the bias is negligible and when a more sophisticated approach is necessary.

In the frequentist approach, the p.d.f. $B(Q_r, U_r|Q_0, U_0)$ is not a function of ψ_0 (it is a fixed parameter) but of ψ_r . It is assumed, without loss of generality due to the angular symmetry of the problem, that $\psi_0 = 0$ so that $\rho = 0$. This can be thought of as rotating the angular coordinate system which does not have any special reference point. The definition of the “zero angle” of such a system has no influence on the polarization fraction. The p.d.f. now simplifies to the product of two normal distributions

$$B(Q_r, U_r|Q_0, U_0)_{\psi_0=0} = \frac{1}{2\pi\sigma_Q\sigma_U} \times \exp\left[-\frac{1}{2}\left(\frac{(Q - \langle Q \rangle)^2}{\sigma_Q^2} + \frac{(U - \langle U \rangle)^2}{\sigma_U^2}\right)\right]. \quad (13)$$

It can now be transformed to polar coordinates similarly to Eq. (7) yielding

$$f(p_r, \psi_r|p_0)_{\psi_0=0} = B(Q, U)_{\psi_0=0} |\det(J)|. \quad (14)$$

The relative mean bias β is now given by

$$\beta = \frac{\langle p_r \rangle - p_0}{p_0} = \frac{\int_0^\infty \int_0^\pi p_r f(p_r, \Delta\psi|p_0)_{\psi_0=0} d\Delta\psi dp_r - p_0}{p_0}, \quad (15)$$

where $\Delta\psi = \psi_r - \psi_0 = \psi_r$.

To understand which parameters have a significant impact on β an approximate analytical expression can be derived by introducing

$$\sigma \equiv \frac{\sqrt{2N}}{\mu_0 S} \approx \frac{2\sigma_Q}{\mu_0} \approx \frac{2\sigma_U}{\mu_0}. \quad (16)$$

It is now possible to write Eq. (14) as

$$f(p_r, \Delta\psi|p_0, 0) = \frac{p_r}{\pi\sigma^2} \times \exp\left(-\frac{p_r^2 + p_0^2 - 2p_r p_0 \cos(2\Delta\psi)}{2\sigma^2}\right). \quad (17)$$

Integrating over $\Delta\psi$ results in

$$f(p_r|p_0) = \frac{p_r}{\sigma^2} \exp\left(-\frac{p_r^2 + p_0^2}{2\sigma^2}\right) \times I_0\left(\frac{p_r p_0}{\sigma^2}\right), \quad (18)$$

Fig. 2. Relative mean bias β as given by Eq. (15) (solid colored lines) and the approximation of Eq. (20) (black dashed line). Here $p_0 = \mu_0 = 1$ has been used to show the maximum difference between the different signal-to-background ratios R . A log-log plot is shown in the inset.

which is the Rice distribution where I_0 is the modified Bessel function of zeroth order. In the limit of high statistics, the relative mean bias is given by

$$\lim_{p_0/\sigma \rightarrow \infty} \beta \approx \frac{\sigma^2}{2p_0^2} = \frac{N}{S^2 \mu_0^2 p_0^2} = \left(\frac{\text{MDP}}{4.29 p_0} \right)^2, \quad (19)$$

as shown in Appendix C.

Since p_0 is not known a priori, Eq. (19) needs to be expressed as a function of p_r by recursion. After some simplification the result is

$$\beta \approx \frac{1 - 2x^2 - \sqrt{1 - 4x^2}}{2x^2}, \quad (20)$$

where $x \equiv \text{MDP}/4.29 p_r$. This shows that MDP/p_r is a good choice of independent variable. Equation (20) is shown in Fig. 2 where it is seen that $\beta > 0$ and increases monotonously, i.e., on average the reconstructed polarization p_r will always be greater than the fixed polarization p_0 . Exact numerical integration of Eq. (15) is also provided for different signal-to-background ratios R yielding similar results to the approximation in Eq. (20). Here $R = 0$ is the limit of large N , low S and yet finite MDP. To avoid inaccurately computing β for small S (a problem under the CLT approximation) Eq. (11) is used for calculating $L(p_r|p_0, \psi_0, \mu_0)$ and ultimately its mean when $S < 200$.

To understand when the bias is significant with respect to the statistical uncertainty on p_r , the bias fraction $(\langle p_r \rangle - p_0)/\sigma_{p_r} = \beta \times p_0/\sigma_{p_r}$ is shown in Fig. 3. Here σ_{p_r} is derived (see Appendix D) using standard error propagation yielding

$$\sigma_{p_r} = \frac{2}{\mu_r} \sqrt{\frac{1}{S} \left(\frac{N}{2S} - \frac{\mu_0^2 p_0^2}{4} \right)}. \quad (21)$$

It becomes clear that this bias cannot be ignored in the low statistics regime. For a measured polarization below MDP, i.e., when $\text{MDP}/p_r > 1$, the bias is more than 18% of the statistical uncertainty. Hence, one should not use Eq. (5) for estimating p_0 in this regime.

A conclusion of this analysis is that significant errors are incurred when dividing the data into several data sets as is done

Fig. 3. Ratio of the absolute bias to the statistical uncertainty $(\langle p_r \rangle - p_0)/\sigma_{p_r} = \beta \times p_0/\sigma_{p_r}$. The β in the approximation (black dashed line) is given by Eq. (20). Here $p_0 = \mu_0 = 1$ has been used to show the maximum difference between the different signal-to-background ratios R .

for example by Dean et al. (2008) in order to estimate the statistical uncertainty. The smaller the data set, the bigger the bias, and thus, on average, the result of an integrated measurement will be lower than the mean of its constituent data sets. This is also relevant when fitting models to polarization fraction sub-divided with respect to energy or time. Binning the data will result in a higher reconstructed polarization fraction, thus biasing the fit and therefore physical conclusions should not be drawn from p_r , as is done for example in Vadawale et al. (2018) where the evolution of polarization parameters throughout the Crab pulsar pulse phase is investigated.

3. Maximum a posteriori estimate

The previous section showed the bias arising when using Eq. (5) for estimating p_0 independently of the polarization angle. It is not obvious if such a bias exists when using (p_r, ψ_r) for the joint point estimate of (p_0, ψ_0) . However, in this case it does not make sense to find a $\langle p_r, \psi_r \rangle$ for a fixed (p_0, ψ_0) so here the word “bias” is interpreted as the difference between the Maximum A Posteriori (MAP) estimate, p_{MAP} , and p_r . Such an analysis requires a Bayesian approach.

3.1. Central Limit Theorem approximation of the posterior

The posterior $P(p_0, \Delta\psi_0|p_r)$ is derived by applying the Bayes theorem to the likelihood given by Eq. (7)

$$P(p_0, \Delta\psi_0|p_r) = \mathcal{N} \times P(p_0, \Delta\psi_0) \times L(p_r, 0|p_0, \Delta\psi_0), \quad (22)$$

where $\mathcal{N}$ is the normalization factor, $P(p_0, \psi_0)$ is the prior and $\Delta\psi_0 = \psi_0 - \psi_r$. Since in a Bayesian approach the parameters (p_0, ψ_0) vary and the data (p_r, ψ_r) are fixed, $L(p_r, 0|p_0, \Delta\psi_0)$ is a function of ψ_0 . It follows that $\sigma_Q \neq \sigma_U$ and $\rho \neq 0$. This differs from the typical assumption of $\sigma_Q = \sigma_U$ and no correlation between Q and U as done in other works (Simmons & Stewart 1985; Vaillancourt 2006; Quinn 2012; Maier et al. 2014).

The Jeffreys prior has the desirable property of being invariant under re-parametrization (Jeffreys 1939). In this case it is a uniform prior in the Cartesian coordinates (Q_0, U_0) . It will result in a preference of high polarization fractions after transforming to polar coordinates, as discussed in Quinn (2012).

Fig. 4. Highest posterior density credibility regions for the posterior given by Eq. (22) for $\text{MDP}/p_r = 1$, $R = 0$ and $\mu_0 = 1$. The contours correspond to 1σ , 2σ and 3σ probability content. The posterior is truncated by the prior at $p_0 = 1$. The unbound prior scenario (black line) corresponds to the unnormalized prior $p_0 \geq 0$.

This is unphysical because it is more difficult to make a highly polarized photon beam in nature since any disorder in the emission region will lower the polarization fraction. A more realistic approach is to instead take a uniform prior $0 \leq p_0 \leq 1$ in polar coordinates. The prior for ψ_0 is also uniform due to symmetry.

Figure 4 shows the posterior as given by Eq. (22) for a low-statistics measurement $\text{MDP}/p_r = 1$. The asymmetry is apparent as the posterior broadens for low values of p_0 . The effect of the prior is illustrated by including the unbound prior scenario where $p_0 \geq 0$.

Since the posterior is derived using a likelihood which utilizes the CLT approximation, the posterior may, for certain combinations of parameters, be inaccurate. To check the magnitude of the effect, pairs of (p_0, ψ_0) are sampled from the prior and then used to generate a data set. Data sets falling within a narrow window of a chosen p_r and ψ_r ($p_r \pm 0.002$ and $\psi_r \pm 0.36^\circ$) are selected and the posterior for each data set is found using Eq. (11) as the likelihood. Finally, the intervals containing 90% of the posteriors marginalized over the polarization angle (for clarity) are shown in Fig. 5 along with the CLT approximation. The situation is similar for the posterior of the polarization angle marginalized over the polarization fraction.

The most important feature of Fig. 5 is that the posterior is not the same for fixed (p_r, ψ_r) when the number of signal photons is low. Hence (p_r, ψ_r) are not always sufficient statistics and the full data set $\{\psi_i\}$ is required for deriving the posterior. However, the posterior converges quickly towards the CLT approximation when $R < 0.5$ since any meaningful measurement with such a low R requires a large number of photons. An effect not shown in the figure is that the difference between the full posterior and the CLT approximation increases as p_r and μ_r increase requiring more photons for good convergence. However, few X-ray polarimeters have μ_0 higher than 0.5 (Krawczynski et al. 2011; Kaaret 2014) and synchrotron emission, a proposed mechanism for many astrophysical sources, is not expected to produce a polarization fraction above ~ 0.6 (Lyutikov et al. 2003).

In conclusion, the full likelihood in Eq. (11) rather than the CLT approximation in Eq. (7) is required for sources where the

Fig. 5. Intervals containing 90% of the posteriors for $(p_r = 0.5, \psi_r = 22.5^\circ)$ and the CLT approximation. As the number of signal photons increases, the posteriors converge towards the approximation. Here $\mu_0 = \mu_r = 0.5$ and $\text{MDP}/p_r = 1$.

Table 1. Relative MAP bias β_{MAP} for $\text{MDP}/p_r = 2$, different signal-to-background ratios R , sufficient statistic p_r and modulation factor μ_0 .

R	p_r	μ_0	β_{MAP}
0	0.5	0.5	-8×10^{-6}
0	0.8	1.0	-3×10^{-4}
0.5	0.5	0.25	-7×10^{-5}

signal-to-background ratio is high but the total number of photons is low (e.g., short-duration GRBs). A guideline is to use the CLT approximation only for data sets where $S > 10^3$. As mentioned before, polarimeters foreseen to be active in the near future (Astrosat/CZTI, POLAR, IXPE, X-Calibur) have $\mu_0 < 0.5$ so this guideline translates to the CLT approximation being generally valid for $\text{MDP} < 0.27$.

3.2. Relative MAP bias

The MAP estimate is at $p_{\text{MAP}} > p_r$ and $\psi_0 = \psi_r$. The inequality occurs because σ_Q and σ_U depend on p_0 but only measurements with high R and $p_0^2 \mu_0^2 \gg 0$ are significantly affected. In all other cases $\sigma_Q \approx \sigma_U \approx \sqrt{2N}/\mu_0 S$ and $p_{\text{MAP}} \approx p_r$.

The relative MAP bias is defined as

$$\beta_{\text{MAP}} \equiv \frac{p_r - p_{\text{MAP}}}{p_{\text{MAP}}}. \quad (23)$$

Table 1 shows β_{MAP} for the extreme case $\text{MDP}/p_r = 2$ (a monotonously decreasing function) for different polarization parameters. Only scenarios with a sufficient number ($S > 884$ based on the results shown in Fig. 5) of signal photons are considered. The results show that β_{MAP} increases with p_r and μ_0 but its value is negligible for all cases where the CLT approximation is valid. Therefore, when (p_r, ψ_r) are sufficient statistics of the data, they are a good estimate for the MAP.

The broadening of the joint posterior at low values of p_0 , as shown in Fig. 4, results in p_{MAP} being a poor estimate for p_0 when marginalizing over ψ_0 . More generally, this occurs for any two-dimensional distribution when marginalizing over the nuisance parameter if the estimated parameter has an asymmetric distribution.

Fig. 6. Skewness of the marginalized posterior for the polarization fraction. The unbound prior scenario (black line) uses $p_0 \geq 0$ and $p_r = 1$ (here the result is independent of p_r). A positive skewness implies an asymmetric distribution with a longer tail on the right, while a negative skewness results in a tail on the left. In all cases $R = 0$ is assumed.

4. Uncertainties on polarization parameters

There are two reasons why it is not always appropriate to use naive Stokes estimates for uncertainties of polarization parameters: non-Gaussianity and the prior. Equation (21) gives the correct uncertainty on the polarization fraction when statistics are high and the reconstructed fraction p_r is well below 1. If these conditions are not fulfilled, the marginalized posterior becomes highly asymmetric either due to low statistics or because a part of the likelihood is truncated by the prior at $p_0 = 1$ and is thus poorly approximated by a Gaussian.

The asymmetry of the marginalized posterior for the polarization fraction is shown in terms of the skewness in Fig. 6. The unbound prior scenario employs the unnormalized uniform prior $p_0 \geq 0$ and $p_r = 1$ (here all values of p_r produce the same result). It shows that when statistics are low, the marginalized posterior gets a longer tail on the right (positive skewness). This is because the peak moves left, which is equivalent to measuring a higher p_r for a fixed p_0 , that is, the same interpretation as for the relative mean bias shown in Fig. 2.

Although the marginalized posterior for the polarization angle is always symmetric, it deviates from Gaussianity when statistics are low, for example, approaching the uniform distribution since all angles are equally likely. It is also affected by the prior when p_r is sufficiently high, resulting in a distribution with longer tails because, as shown in Fig. 4, the prior only truncates the part of the likelihood contributing to the central part of the marginalized distribution. Therefore, the uncertainty given by

$$\sigma_{\psi_r} = \frac{1}{\sqrt{2}\mu_r p_r} \frac{\sqrt{N}}{S} \quad (24)$$

(derived in Appendix D) is not always valid.

In a Bayesian approach, the posterior is used instead of a parameter estimate and its uncertainty since the posterior contains the full information about the parameter. The posterior cannot easily be described in text so a simplified description is necessary such as its peak and the region of Highest Posterior Density (HPD) containing the probability content corresponding to one Gaussian standard deviation.

It is possible to quantify when the HPD region must be derived from the posterior and when Eqs. (21) and (24)

Fig. 7. Ratio $\sigma_{p_{\text{HPD}}} / \sigma_{p_r}$ of the uncertainty derived from the HPD region of the marginalized posterior and the Gaussian uncertainty given by Eq. (21) for the polarization fraction. Since the posterior is asymmetric, an effective $\sigma_{p_{\text{HPD}}}$ is defined as half of the region containing 1σ Gaussian probability. The unbound prior scenario corresponds to the unnormalized prior $0 \leq p_0$ and $p_r = 1$ (here the result is independent of p_r). In all cases $R = 0$ is assumed.

are good approximations by considering the ratios $\sigma_{p_{\text{HPD}}} / \sigma_{p_r}$ and $\sigma_{\psi_{\text{HPD}}} / \sigma_{\psi_r}$ as functions of MDP/p_r , which are shown in Figs. 7 and 8, respectively. In both figures, the unbound prior scenario uses the unnormalized uniform prior $p_0 \geq 0$ and $p_r = 1$. Only scenarios with $R = 0$ are considered, meaning that the CLT approximation is valid in the entire range of MDP/p_r for any p_r . Increasing R does not change the results shown in this section but may invalidate the CLT approximation. Such cases cannot be represented because p_r is not a sufficient statistic. However, since the posterior shape does not change drastically (as shown in Fig. 5), the following discussion is a good qualitative description of its behavior for any R .

For ease of comparison, $\sigma_{p_{\text{HPD}}}$ is defined as half of the region containing 1σ Gaussian probability. Additionally, the Gaussian interval is limited to account for the prior, for example, $p_0 = 0.8 \pm 0.3$ is the interval $[0.5, 1.0]$ which gives an uncertainty of $\sigma_{p_r} = 0.25$, and therefore it does not contain 68.3% (1σ) probability content.

The unbound prior scenario in Fig. 7 shows that for $\text{MDP}/p_r \sim 1$ the marginalized posterior has longer tails than a Gaussian. Additionally, the prior truncates the posterior for high p_r or for low statistics. This can result in either a higher or a lower $\sigma_{p_{\text{HPD}}} / \sigma_{p_r}$ depending on where exactly the distribution is truncated. Unless p_r is high, σ_{p_r} is a good approximation for the magnitude of the uncertainty at $\text{MDP}/p_r \sim 1$ but it is important to remember that the posterior is asymmetric for such low statistics.

For the polarization angle, Fig. 8 shows that $\sigma_{\psi_{\text{HPD}}} > \sigma_{\psi_r}$. The situation is simpler than for p_0 because the angle has a symmetric posterior. The unbound prior shows the deviation from a Gaussian manifesting in longer tails. Figure 4 shows that when adding the prior $0 \leq p_0 \leq 1$, the part of the likelihood which extends past $p_0 = 1$ is truncated. This part contributes to the density at the center of the marginalized posterior. Removing it further extends the tails. It follows that Eq. (24) underestimates the uncertainty on the polarization angle by relative 10–20% at $\text{MDP}/p_r \sim 1$.

This analysis shows that in the limit of low statistics, $\text{MDP}/p_r \sim 1$, the uncertainty on the polarization angle is not well-described by Eq. (24) and a Bayesian treatment is necessary.

Fig. 8. Ratio $\sigma_{\psi_{\text{HPD}}}/\sigma_{\psi_r}$ of the uncertainty derived from the HPD region of the marginalized posterior and the Gaussian uncertainty given by Eq. (24) for the polarization angle. The unbound prior scenario corresponds to the unnormalized prior $0 \leq p_0$ and $p_r = 1$ (here the result is independent of p_r). In all cases $R = 0$ is assumed.

Table 2. Examples of the difference between a Bayesian approach and a Gaussian approximation.

MDP/ p_r	p_r (%)	p_0 (%)	ψ_r	ψ_0
1.3	19.5 ± 8.3	$17.4^{+8.6}_{-9.3}$	$137 \pm 12^\circ$	$137 \pm 15^\circ$
1	80 ± 26	76^{+21}_{-22}	$90 \pm 9^\circ$	$90 \pm 11^\circ$

Notes. The first row is a recent measurement of the Crab nebula by PoGO+ (Chauvin et al. 2017). The second row is a hypothetical measurement of a high reconstructed polarization fraction.

For high p_r , such a treatment is required even for high-statistics measurements, $\text{MDP}/p_r \sim 0.5$, because of the asymmetry in the uncertainty on the polarization fraction.

To illustrate the effects described above, two examples are presented in Table 2. The first is a recent measurement of the Crab nebula by PoGO+ (Chauvin et al. 2017) and the second is a hypothetical measurement of a high reconstructed polarization fraction highlighting the importance of the prior $0 \leq p_0 \leq 1$. The largest differences are in the polarization fraction and the uncertainty on the polarization angle.

5. Uncertainties on the modulation factor

The uncertainty σ_μ on the modulation factor μ_0 can be minimized for most polarimeters designed for measuring point sources by increasing the statistics during calibration tests. However, for polarimeters measuring GRBs it is often impossible to make σ_μ arbitrarily small. The problem is that μ varies depending on the location of the GRB, typically having the highest values for on-axis GRBs but significantly lower values for GRBs located at a large angular separation from the detector axis. Determining the location of the GRB by using a polarimeter involves large uncertainties since it is often not optimized for the task. The uncertainty on the location propagates to a non-negligible σ_μ . Additionally, the primary spectrum of GRBs may not be reconstructed with sufficient precision by the polarimeter, introducing further uncertainties in the simulation required for deducing the μ_0 for a particular GRB. If the GRB is simultaneously observed by a dedicated GRB monitor, the uncertainty on

its location and spectrum will be smaller, but no GRB monitor has complete sky coverage or 100% duty-cycle, so it is inevitable that some GRBs will only be seen by the polarimeter.

An example of a GRB polarimeter is POLAR (Produit et al. 2005). For a typical bright GRB, POLAR is expected to have a σ_μ/μ of between 5 and 15% assuming that the burst occurs on-axis (Suarez-Garcia et al. 2010). The simplest way to account for this additional uncertainty is to propagate σ_μ through

$$p_r = \frac{M}{\mu_r}, \quad (25)$$

where M is the measured modulation. The total symmetric uncertainty is then given by

$$\sigma_{\text{tot}} = p_r \sqrt{\left(\frac{\sigma_\mu}{\mu_r}\right)^2 + \left(\frac{\mu_r \sigma_{p_r}}{M}\right)^2}. \quad (26)$$

To check for which parameters the symmetry is a good approximation, an additional prior can be added to the posterior Eq. (22). For simplicity, this prior is assumed to be Gaussian but it can vary depending on localization sensitivity. The nuisance parameter μ_0 can then be marginalized over, yielding

$$P(p_0, \Delta\psi_0 | p_r, \mu_r) = \mathcal{N} \int_0^1 P(p_0, \Delta\psi_0) \times \exp\left[-\frac{(\mu_r - \mu_0)^2}{2\sigma_\mu^2}\right] \times L(p_r, 0, \mu_r | p_0, \Delta\psi_0) d\mu_0, \quad (27)$$

where the integral is taken between 0 and 1 because $\mu_0 > 1$ is unphysical. As an example, the posterior (marginalized over the polarization angle) for a signal-only measurement at $\text{MDP} = 10\%$ (CLT approximation is valid because S is large), $p_r = 0.4$ and $\mu = 0.4$ is shown in Fig. 9 for different σ_μ/μ . Due to the high statistics, $\text{MDP}/p_r = 0.25$, the distribution is symmetric for low σ_μ/μ , however, the tail on the right grows rapidly as σ_μ/μ increases. The behavior is governed by the reciprocal Gaussian distribution which the posterior approaches in the limit of high photon statistics

$$g(p_0 | p_r, \mu_r) = \frac{1}{\sqrt{2\pi}\sigma_\mu} \frac{p_r \mu_r}{p_0^2} \exp\left[-\frac{(p_r \mu_r / p_0 - \mu_r)^2}{2\sigma_\mu^2}\right], \quad (28)$$

shown as an approximation in Fig. 9. If not for the prior $0 \leq p_0 \leq 1$, the moments of this distribution would be undefined.

Figure 9 demonstrates that for $\sigma_\mu/\mu > 10\%$ the shape of the posterior changes significantly and Eq. (27) is required. In the extreme case of $\sigma_\mu/\mu = 20\%$, the half-width of the HPD region is 0.092, whereas Eq. (26) yields 0.086, implying not only that the uncertainty is asymmetric but also that it is significantly larger.

These results can be related to the instrument performance by studying the effect σ_μ/μ has on the MDP. A frequentist approach must be followed since MDP is a frequentist concept. MDP is derived by finding the 99% upper limit for the Rayleigh distribution – a special case of the Rice distribution where $p_0 = 0$. However, since there is an uncertainty on μ , the Rayleigh distribution must be multiplied by the likelihood for μ (assumed to be a Gaussian for simplicity) and integrated to yield

$$f(p_r | p_0 = 0, \mu_0) = \int_0^\infty \frac{p_r \mu_r}{\sqrt{2\pi}\sigma_\mu^2} \times \exp\left[-\frac{\mu_r^2 p_r^2}{2\sigma^2} - \frac{(\mu_r - \mu_0)^2}{2\sigma_\mu^2}\right] d\mu_r. \quad (29)$$

Finally, $f(p_r | p_0 = 0, \mu_0)$ can be integrated to find the 99% upper limit so that $\int_0^{\text{MDP}_r} f(p_r | p_0 = 0, \mu_0) dp_r = 0.99$. Figure 10

Fig. 9. Polarization fraction posterior density for $MDP/p_r = 0.25$, $p_r = 0.4$ and $\mu_r = 0.4$ for different σ_μ/μ . As σ_μ/μ increases, the distribution becomes wider and more asymmetric. The approximation is the reciprocal Gaussian distribution in Eq. (28).

Fig. 10. Relative increase in MDP, defined as the ratio MDP_σ/MDP where MDP_σ is the solution to $\int_0^{MDP_\sigma} f(p_r|p_0 = 0, \mu_0) dp_r = 0.99$ as a function of the relative uncertainty on μ . The results are independent of the initial MDP.

shows the relative increase in the MDP, defined as the ratio MDP_σ/MDP where MDP is given by Eq. (12). Although the effect is small for low σ_μ/μ (1% at $\sigma_\mu/\mu = 5\%$), it becomes significant at larger σ_μ/μ (14% at $\sigma_\mu/\mu = 15\%$) and deteriorates the instrument performance for extreme values (80% at $\sigma_\mu/\mu = 25\%$). The result is independent of the initial MDP.

6. Conclusions

The results presented here provide a means of quantifying the errors incurred when using the simple estimators for polarization parameters as well as for their uncertainties. These errors are related to the well-established figure of merit MDP and the reconstructed polarization fraction, making the results easily applicable to any X-ray polarimeter. Unlike some previous works, this analysis does not ignore the correlation between the Stokes parameters Q and U .

Additionally, the extent to which the reconstructed fraction p_r and polarization angle ψ_r can be used as sufficient statistics is explored. In certain situations, such as for high polarization fraction and high signal-to-background ratio, Q and U are not Gaussian distributed and their likelihood is non-Gaussian, resulting in a significantly different posterior for the polarization

parameters when the number of photons is low. In such cases, the Stokes parameters cannot be used at all because (p_r, ψ_r) are not sufficient statistics.

The fact that the reconstructed polarization fraction is biased towards higher values implies that binning a high-statistics data set into smaller subsets to estimate the evolution of the polarization fraction will yield incorrect results when using the simple Gaussian estimator p_r . This work provides a way to decide how coarsely the data must be binned in order to justify the use of Gaussian estimators.

The uncertainties on the estimated parameters are as important as the parameter estimates themselves. When statistics are high ($MDP/p_r < 0.5$) the errors are small and can usually be neglected (justifying Gaussian assumptions made when using simple estimators) unless the reconstructed polarization fraction is high, for example, > 0.8 . However, in the statistics-limited regime ($MDP/p_r > 1$) the systematic error made from using such simple estimators is non-negligible compared to the statistical uncertainty. Additionally, it is shown how quickly the posterior becomes asymmetric as the statistical power decreases. This effect is strongly dependent on the reconstructed polarization fraction because of the prior $0 \leq p_0 \leq 1$.

Lastly, the effect of uncertainties on the modulation factor μ is studied. It is shown to be important once the relative uncertainty exceeds 10% and to dominate the performance of an instrument when it is above 20%. This is relevant for the optimization of future GRB polarimeters, since they tend to have large uncertainties on μ due to the difficulty of localizing bursts and measuring the primary GRB spectrum. It also shows that simple Gaussian estimators cannot be used in the high-photon-counts regime for GRB polarimeters when localization uncertainties are high.

Acknowledgements. This research was supported by the Swedish National Space Board. M. Kiss, M. Pearce, F. Xie and the Referee are thanked for providing constructive feedback on the manuscript.

References

- Beilicke, M., Kislak, F., Zajczyk, A., et al., 2014, *JAI*, **3**, 1440008
- Bellazzini, R., & Muleri, F., 2010, *NIM A*, **623**, 766
- Chauvin, M., Roques, J. P., Clark, D. J., et al., 2013, *ApJ*, **769**, 137
- Chauvin, M., Florén, H.-G., Jackson, M., et al., 2016, *MNRAS*, **456**, L86
- Chauvin, M., Florén, H.-G., Friis, M., et al., 2017, *Sci. Rep.*, **7**, 7816
- Clarke, D., Stewart, B. G., Schwarz, H. E., et al., 1983, *A&A*, **126**, 260
- Dean, A. J., Clark, D. J., Stephen, J. B., et al., 2008, *Science*, **321**, 1183
- Forot, M., Laurent, P., Grenier, I. A., et al., 2008, *ApJ*, **688**, L29
- Jeffreys, H., 1939, *Theory of Probability* (Oxford: Oxford University Press)
- Kaaret, P., 2014, ArXiv e-print [[arXiv:1408.5899](https://arxiv.org/abs/1408.5899)]
- Kislak, F., Clark, B., Beilicke, M., et al., 2015, *Astropart. Phys.*, **68**, 45
- Krawczynski, H., Garson, A., Guo, Q., et al., 2011, *Astropart. Phys.*, **34**, 550-67
- Lei, F., Dean, A. J., Hills, G. L., 1997, *Space Sci. Rev.*, **82**, 309-88
- Lyutikov, M., Pariev, V. I., & Blandford, R. D., 2003, *ApJ*, **597**, 998
- Maier, D., Tenzer, C., & Santangelo, A., 2014, *PASP*, **126**, 459
- Moran, P., Kyne, G., Gouiffès, C., et al., 2016, *MNRAS*, **456**, 2974
- Novick, R., Weisskopf, M. C., Berthelsdorf, R., et al., 1972, *ApJ*, **174**, L1
- Produit, N., Barao, F., Deluit, S., et al., 2005, *NIM A*, **550**, 616
- Quinn, J. L., 2012, *A&A*, **538**, A65
- Simmons, J. F. L., & Stewart, B. G., 1985, *A&A*, **142**, 100
- Słowikowska, A., Kanbac, G., Kramer, M., et al., 2009, *MNRAS*, **397**, 103
- Suarez-Garcia, E., Haas D., Hajdas W., et al., 2010, *NIM A*, **624**, 624
- Vadawale S. V., Chattopadhyay, T., Mithun N. P. S., et al., 2018, *Nat. Astron.*, **2**, 50
- Vaillancourt, J. E., 2006, *PASP*, **118**, 1340
- Weisskopf, M. C., Cohen, G. G., Kestenbaum, H. L., Long, K. S., et al., 1976, *ApJ*, **208**, L125
- Weisskopf, M. C., Silver, E. H., Kestenbaum, H. L., et al., 1978, *ApJ*, **220**, L117
- Weisskopf, M. C., Elsner R. F., & O'Dell S. L., 2010, *Proc. SPIE*, **7732**, 77320E
- Weisskopf, M. C., Ramsey, B., O'Dell, S., et al., 2016, *Proc. SPIE*, **9905**, 990517
- Yonetoku, D., Murakami, T., Gunji, S., et al., 2011, *ApJ*, **743**, L30

Appendix A: Stokes parameters

The Stokes parameters can be constructed from

$$\begin{aligned} q_i &= \cos(2\psi_i) \\ u_i &= \sin(2\psi_i). \end{aligned} \quad (\text{A.1})$$

To find an expression for p_0 , the means $\langle q \rangle$ and $\langle u \rangle$ must be computed.

$$\langle q \rangle = \int_0^{2\pi} \cos(2\psi) f(\psi) d\psi, \quad (\text{A.2})$$

$$\langle q \rangle = \frac{\mu_0 p_0 S}{2(B+S)} \cos(2\psi_0), \quad (\text{A.3})$$

where $f(\psi)$ is given by Eq. (1). Similarly

$$\langle u \rangle = \frac{\mu_0 p_0 S}{2(B+S)} \sin(2\psi_0). \quad (\text{A.4})$$

The polarization fraction p_0 is then given by

$$\begin{aligned} p_0 &= 2 \sqrt{Q_0^2 + U_0^2} / \mu_0 \\ \psi_0 &= \frac{1}{2} \arctan \frac{U_0}{Q_0}, \end{aligned} \quad (\text{A.5})$$

where Q_0 and U_0 are the normalized Stokes parameters defined as

$$\begin{aligned} Q_0 &= \frac{B+S}{S} \langle q \rangle = \frac{\mu_0 p_0}{2} \cos(2\psi_0) \\ U_0 &= \frac{B+S}{S} \langle u \rangle = \frac{\mu_0 p_0}{2} \sin(2\psi_0). \end{aligned} \quad (\text{A.6})$$

The sufficient statistics for data generated from (Q_0, U_0) are (Q_r, U_r) defined as

$$\begin{aligned} Q_r &= \frac{1}{S} \sum_{i=1}^N q_i = \frac{\mu_r p_r}{2} \cos(2\psi_r) \\ U_r &= \frac{1}{S} \sum_{i=1}^N u_i = \frac{\mu_r p_r}{2} \sin(2\psi_r), \end{aligned} \quad (\text{A.7})$$

where $N = B + S$. In polar coordinates this becomes

$$\begin{aligned} p_r &= 2 \sqrt{Q_r^2 + U_r^2} / \mu_r \\ \psi_r &= \frac{1}{2} \arctan \frac{U_r}{Q_r}. \end{aligned} \quad (\text{A.8})$$

Appendix B: Uncertainties on Stokes parameters

The uncertainties on Q and U as well as their covariance can be derived by computing their second moments.

$$\langle q^2 \rangle = \int_0^{2\pi} \cos^2(2\psi) f(\psi) d\psi = \frac{1}{2}. \quad (\text{B.1})$$

Similarly

$$\langle u^2 \rangle = \frac{1}{2}, \quad (\text{B.2})$$

and the cross-term is

$$\langle qu \rangle = \frac{1}{2\pi(S+B)} \int_0^{2\pi} \cos(2\psi) \sin(2\psi) \times f(\psi) d\psi = 0. \quad (\text{B.3})$$

Combining Eqs. (A.2) and (B.1) yields the standard deviation

$$\sigma_q = \sqrt{\frac{1}{2} - \left(\frac{\mu_0 p_0 S \cos(2\psi_0)}{2(B+S)} \right)^2}. \quad (\text{B.4})$$

Here σ_q describes the dispersion of q but it is σ_Q which is of interest. It is given by

$$\sigma_Q^2 = \sigma_q^2 \times \frac{N}{S^2}, \quad (\text{B.5})$$

where the division by S^2 comes from the definition of normalized Stokes parameters.

$$\sigma_Q = \sqrt{\frac{1}{S} \left(\frac{N}{2S} - \frac{\mu_0^2 p_0^2 \cos^2(2\psi_0)}{4} \right)}. \quad (\text{B.6})$$

Similarly

$$\sigma_U = \sqrt{\frac{1}{S} \left(\frac{N}{2S} - \frac{\mu_0^2 p_0^2 \sin^2(2\psi_0)}{4} \right)}. \quad (\text{B.7})$$

It is finally possible to compute the covariance and the correlation coefficient ρ

$$\text{Cov}(Q, U) = \text{Cov}(q, u) \times \frac{N^2}{S^3} = -\frac{\mu_0^2 p_0^2}{8S} \sin(4\psi_0), \quad (\text{B.8})$$

$$\rho = -\frac{S \mu_0^2 p_0^2 \sin(4\psi_0)}{\sqrt{16N^2 - 8NS \mu_0^2 p_0^2 + S^2 \mu_0^4 p_0^4 \sin^2(4\psi_0)}}. \quad (\text{B.9})$$

Appendix C: Relative mean bias

An approximate expression for the relative mean bias can be derived from the Rice distribution

$$f(p_r | p_0) = \frac{p_r}{\sigma^2} \exp\left(-\frac{p_r^2 + p_0^2}{2\sigma^2}\right) \times I_0\left(\frac{p_r p_0}{\sigma^2}\right), \quad (\text{C.1})$$

where $\sigma \equiv \sqrt{2N}/\mu_0 S$. The mean $\langle p_r \rangle$ is given by

$$\begin{aligned} \langle p_r \rangle &= \sqrt{\frac{\pi}{2}} \sigma \times \exp\left(-\frac{p_0^2}{4\sigma^2}\right) \\ &\times \left(\left(1 + \frac{p_0^2}{2\sigma^2}\right) I_0\left(\frac{p_0^2}{4\sigma^2}\right) + \frac{p_0^2}{2\sigma^2} I_1\left(\frac{p_0^2}{4\sigma^2}\right) \right), \end{aligned} \quad (\text{C.2})$$

where I_0 and I_1 are modified Bessel functions of the zeroth and first order, respectively. For high-statistics measurements (where $p_0 \gg 2\sigma$) the Bessel functions can be approximated by expanding them to second order.

$$\lim_{x \rightarrow \infty} I_0(x) = \frac{e^x}{\sqrt{2\pi x}} \left(1 + \frac{1}{8x}\right), \quad (\text{C.3})$$

$$\lim_{x \rightarrow \infty} I_1(x) = \frac{e^x}{\sqrt{2\pi x}} \left(1 - \frac{3}{8x}\right). \quad (\text{C.4})$$

Finally

$$\lim_{p_0/\sigma \rightarrow \infty} \langle p_r \rangle = \frac{\sigma^2}{2p_0} + p_0. \quad (\text{C.5})$$

So in the limit of high statistics, the relative mean bias β is given by

$$\lim_{p_0/\sigma \rightarrow \infty} \beta \approx \frac{\sigma^2}{2p_0^2} = \frac{N}{S^2 \mu_0^2 p_0^2} = \left(\frac{\text{MDP}}{4.29 p_0} \right)^2. \quad (\text{C.6})$$

By using recursion, β can be expressed as a function of p_r instead of p_0 which is not known a priori. This simplifies to

$$\begin{aligned} \beta &\approx \left(\frac{\text{MDP}}{4.29 p_0} \right)^2 = \sum_{n=1}^{\infty} C_n x^{2n} \\ &= \sum_{n=1}^{\infty} \frac{1}{n+1} \binom{2n}{n} x^{2n} = \frac{1 - 2x^2 - \sqrt{1 - 4x^2}}{2x^2}, \end{aligned} \quad (\text{C.7})$$

where $x = \text{MDP}/4.29 p_r$ and C_n is the n th Catalan number.

Appendix D: Gaussian uncertainties on polarization parameters

The uncertainty on p_r can be derived by using the standard uncertainty propagation formula

$$\sigma_{p_r}^2 = \left| \frac{dp_r}{dQ_r} \right|^2 \sigma_Q^2 + \left| \frac{dp_r}{dU_r} \right|^2 \sigma_U^2 + 2 \left| \frac{dp_r}{dQ_r} \right| \left| \frac{dp_r}{dU_r} \right| \text{Cov}(Q, U), \quad (\text{D.1})$$

yielding

$$\sigma_{p_r} = \frac{2}{\mu_r} \sqrt{\frac{1}{S} \left(\frac{N}{2S} - \frac{\mu_0^2 p_0^2}{4} \right)}, \quad (\text{D.2})$$

and similarly for ψ_r

$$\sigma_{\psi_r} = \frac{1}{\sqrt{2} \mu_r p_r} \frac{\sqrt{N}}{S}. \quad (\text{D.3})$$