Statistics for X-Ray Polarization
Technical Report

C. G. Montgomery
J. H. Swank

March 27, 2015

1. ABSTRACT

This Report presents methods for simulating and for analyzing the azimuthal distribution of scatterings produced by a partially or completely linearly polarized beam of X-rays. The results should be useful both in the design of missions to measure the polarization of X-rays from astronomical sources and in the drawing of conclusions from data produced by such missions. Besides confirming previous results, in particular those of Strohmayer & Kallman (2013, ApJ, 773, 103), for low polarization amplitudes, they provide an analytical treatment for the full range of possible amplitudes. This work is also described in (Montgomery & Swank 2015, ApJ, 801, 21; arXiv:1501.02222), which also discusses, and provides references to, its relation to other work on this topic.

2. GENERATION AND ANALYSIS OF SIMULATED DATA

A beam of X-rays that is partially linearly polarized, when passing through an appropriate medium, can produce scattered photons, by Compton scattering, or scattered (ejected) electrons, by the photoelectric effect. The distribution of scattering into different directions in the plane perpendicular to the beam is influenced by the polarization of the X-rays. Observations of this distribution, as indicated by detections of photons or electrons scattered in different azimuthal directions, enable estimates of the amount of polarization in the beam and its direction. The goal of this work is to find how the observed distribution can be used to obtain estimates of the extent and direction of polarization in the incoming beam, along with quantitative evaluations of the uncertainties in those estimates.

Since the probabilities of scattering either photons or electrons has the same dependence on scattering angle, most of what follows could be applied to either physical mechanism, *mutatis mutandis;* for simplicity we will speak only of measurements using the photoelectric effect.

The partially polarized incident beam of X-rays can scatter electrons into various azimuthal angles around the axis of the incident beam. We want to estimate the degree of polarization in the incident beam, and its direction, from the observed numbers of electrons scattered into different azimuthal angles, and also to estimate the uncertainties in those estimates.

The polarized component of the beam scatters electrons into an azimuthal angle ϕ, measured from the direction of the incident polarization, with a probability proportional to $\cos^2(\phi)$, while the unpolarized component of the beam scatters electrons isotropically. In general, if the true polarization is at an angle ϕ_0 in a coordinate system, ϕ should simply be replaced by $(\phi - \phi_0)$.

So we suppose that the number of scattered electrons reaching detectors, per unit time, at an azimuthal angle between ϕ and $(\phi + d\phi)$, can be expressed as proportional to $[I_0 + U_0 \cos(2\phi)]d\phi$.

This form ranges from $I_0 - U_0$ to $I_0 + U_0$ and has an amplitude of variation $a_0 = U_0/I_0$.
a_0 is one of the physically significant quantities that we would like to estimate.

The detectors are taken to divide the interval from $-\pi$ to $+\pi$ into M equal angular bins, each of size $2\pi/M$ radians. M is an odd integer large compared with unity. The bins can be labelled with an index j which runs from $-(M-1)/2$ to $+(M-1)/2$, with the center of the jth bin at $\phi_j = 2\pi j/M$.

The expected number of counts in the jth bin in a time T is

$$< n_j >= \frac{\kappa}{M}[I_0 + U_0 cos(2\phi_j)] = \frac{\kappa I_0}{M}[1 + a_0 \cos(2\phi_j)]. \tag{1}$$

κ includes geometric and efficiency factors; it is proportional to T.

It is important to recognize that while the flux of electrons in different directions is correctly described by equation (1) the values of n_j are independent Poisson-distributed random variables, with these mean values (and variances). The total number of counts

$$N = \sum_j n_j \tag{2}$$

is then Poisson-distributed with a mean (and variance) equal to

$$< N >= \kappa I_0. \tag{3}$$

So a set of samples generated with the same $< N >$ will have different total count numbers.

The mathematical properties of the n_j values enable a number of useful facts about the statistical properties of various quantities related to them to be derived.

To estimate the incident polarization amount and direction from an observed set of n_j values, we seek to represent their angular distribution by a function describing the counts as a function of angle:

$$f(\phi) = I_f + U_f \cos(2\phi) + Q_f \sin(2\phi). \tag{4}$$

Then the amount and direction of the incoming beam can be estimated in the usual way. For the amplitude of polarization the estimate is

$$a_e = \left(U_f^2 + Q_f^2\right)^{1/2} / I_f \tag{5}$$

and the polarization angle estimate is

$$\phi_e = \frac{1}{2}\arctan(Q_f/U_f). \tag{6}$$

The arctangent function in this equation is actually the two-argument arctangent function [represented in many programming languages as $atan2(Q_{f,}, U_f)$]. Technically, it is the principal value of the argument of the complex number $U_f + iQ_f$. Its values range from from $-\pi$ to $+\pi$ radians, so ϕ_e ranges from $-90°$ to $+90°$.

Multiplying the function $f(\phi)$ by any constant produces no change in the amplitude and angle estimates.

There is a useful graphical representation of the results. To simplify the notation somewhat, let U' stand for U_f/I_f and Q' stand for Q_f/I_f. Then the amplitude and angle estimates are just given by

$$a_e = \sqrt{U'^2 + Q'^2} \tag{7}$$

and

$$\phi_e = \frac{1}{2}\arctan(Q'/U'). \tag{8}$$

If we think of U' and Q' as the rectangular coordinates of a point, these are the polar coordinates of that point: a_e is the distance of the point from the origin and $2\phi_e$ is the angle between the line from the origin to the point and the U' axis. It will be seen that it can be convenient to discuss the geometry of U' and Q' in the $U'Q'$ plane and then, if desired, represent the results in plots of ϕ_e and a_e.

The angular distribution of the set of n_j values for a sample will of course include deviations from the simple form of equation (4) because of the statistical nature of the counting process. In order to select a set of three values, I_f, U_f, and Q_f, to describe the behavior of an observed distribution of counts, we view f as part of a discrete trigonometric interpolating polynomial, where the terms in the polynomial series are the ones that reflect the properties of the distribution that are of physical interest. They reflect the polarization of the incoming radiation and enable estimates of its direction and amplitude. This approach, based on the basic theory of discrete Fourier transforms, provides results with useful mathematical properties.

We choose to express the angular distribution as

$$f(\phi_j) = \frac{M}{2\pi} n_j \tag{9}$$

The coefficients I_f , U_f , and Q_f are then sums over the angles ϕ_j:

$$I_f = \frac{1}{2\pi}\sum_j n_j \tag{10}$$

$$U_f = \frac{1}{\pi}\sum_j n_j \cos(2\phi_j)$$

$$Q_f = \frac{1}{\pi} \sum_j n_j \sin(2\phi_j)$$

The coefficients I_f, U_f, and Q_f are random variables, the sums of linear combinations of the random n_j values. By the Lyapunov central limit theorem their distributions become normal as M becomes large enough, with means and variances that can be calculated using the sums over j of polynomials in $\cos(2\phi_j)$ and $\sin(2\phi_j)$.

The mean value of I_f is

$$< I_f >= \frac{1}{2\pi} \sum_j < n_j >= \kappa I_0/(2\pi) =< N > /(2\pi) \quad (11)$$

The other mean values are

$$< U_f >= \kappa U_0/(2\pi) =< N > a_0/(2\pi) \quad (12)$$

$$< Q_f >= 0$$

The variance of I_f is

$$\mathrm{Var(I_f)} = \frac{1}{4\pi^2} \sum_{\mathrm{j}} \mathrm{Var(n_j)} = \frac{1}{4\pi^2} \sum_{\mathrm{j}} < \mathrm{n_{j>}} =< \mathrm{N} > /(4\pi^2) \quad (13)$$

and by similar calculations the variances of U_f and Q_f are each equal to

$$\sigma_f^2 = \kappa I_0/(2\pi^2) =< N > /(2\pi^2). \quad (14)$$

The covariance of Q_f with either I_f or U_f is zero, but the covariance of U_f with I_f is $< N > a_0/4\pi^2$ and the correlation of U_f and I_f is $a_0/\sqrt{2}$. (For $\phi_0 \neq 0$, the correlation of Q_f with I_f is also nonzero.)

The coefficients I_f, U_f, and Q_f are not statistically independent, since they are calculated from the same set of n_j values. In fact I_f, U_f, and Q_f are trivariate normal. This is helpful in derivations of the distributions of the physically important quantities U' and Q'. Correct distributions for U' and Q' and thus for a_e and ϕ_e can be generated by simulations in which the M values of n_j are randomly generated for each sample, and the resulting values of I_f, U_f, Q_f and thence U' and Q', are used to get a_e and ϕ_e. The simulated samples are found to agree well with the theoretical distributions described in what follows.

The points for a collection of samples will be centered around the U', Q' point with coordinates a_0, 0. Let us call that point Z. The probability density for points, that is, the probability per unit area in this plane, describes what fraction of the points from a set of many samples will have those locations. A line of constant probability, a closed curve enclosing the point Z, identifies the region within which a specific fraction of samples will lie, and thus the likelihood of that set of amplitude and angle estimates. Figure 1 provides an example for a low true amplitude., showing four sets of 100 samples each, together with the theoretical contour corresponding to 1-sigma.

The joint distribution of I_f, U_f, and Q_f obtained in this way as the parameters describing the angular distribution of counts allows a calculation of the marginal distribution of U' and Q'. For large $< N >$, U' and Q' are independent and normally distributed. Details are given in Appendix A, which also provides the joint distribution of U' and Q' before any approximation of large $< N >$.

The calculation also provides the means and variances of U' and Q':

$$< U' >= a_0 \, ; \quad < Q' >= 0 \, ; \tag{15}$$

$$\mathrm{Var(U')} = \frac{2}{< \mathrm{N} >}(1 - \mathrm{a_0^2}/2) \, ; \; \mathrm{Var(Q')} = \frac{2}{< \mathrm{N} >}. \tag{16}$$

These results are confirmed by the results of simulations for the full range of amplitude values from 0 to 1.

The probability density for a system point would be the product of the probability density for U' and the probability density for Q' . But these are simply Gaussians with the derived variances and means. Figure 2 shows the results of simulations for a true polarization $a_0 = 3/4$, along with the Gaussian distributions with the calculated means and variances.

With the definition of λ as the ratio of the standard deviations:

$$\lambda^2 = \frac{\sigma^2_{Q'}}{\sigma^2_{U'}} = \frac{1}{1 - a_0^2/2} \tag{17}$$

and denoting $\sigma_{Q'}$ simply as σ', the probability per unit area in the $U'Q'$ plane is also a Gaussian:

$$\begin{aligned} P_{U'Q'}(U', Q') &= \frac{\lambda}{2\pi\sigma'^2} \exp[-(U' - a_0)^2\lambda^2/2\sigma'^2] \cdot \exp(-Q'^2/2\sigma'^2) \\ &= \frac{\lambda}{2\pi\sigma'^2} \exp(-\frac{D^2}{2\sigma'^2}), \end{aligned} \tag{18}$$

with

$$D^2 = (U' - a_0)^2\lambda^2 + Q'^2. \tag{19}$$

It is worth noting that while these results were derived with the point Z on the U' axis (i.e., with $\phi_0 = 0$) they can be expressed in a way that is more general. Let us define 2η as the angular difference between the location of Z and the location of a data point. When Z is on the U' axis, η is just ϕ_e. But in general U' in the above equations is $a_e \cos(2\eta)$ while Q' is $a_e \sin(2\eta)$. Then Equation (19) becomes

$$D^2 = \frac{[a_e \cos(2\eta) - a_0]^2}{1 - a_0^2/2} + a_e^2 \sin^2(2\eta). \tag{20}$$

In this form it involves only the distances of Z and the data point from the origin and the angle between them, and is correct for any choice of a reference direction.

Lines of constant probability in the $U'Q'$ plane are ellipses, centered on the point Z, with minor axes along the line from the origin to the point Z, and the ratio of major to minor axes equal to

λ. An example is shown in Figure 3. D is the semi major axis of the ellipse. The area within an ellipse with a given D is $\pi D^2/\lambda$. The total probability of a point lying outside the ellipse with a given D is just $\exp(-D^2/2\sigma'^2)$ (as long as $D^2 << 1$.) See Figure 4 for an example of agreement between the probability distribution and a simulation with many samples.

This provides an extension of previous work, covering the full range of incident amplitudes, even when the correlation of I_f with U_f changes the variance of U' enough to be important. The results for U' and Q' also give

$$\mathrm{Var(a_e)} = \sigma'^2(1 - a_0^2/2), \tag{21}$$

$$\mathrm{Var}(\eta) = \frac{1}{4}\frac{\sigma'^2}{a_0^2}. \tag{22}$$

Since λ is quite close to unity unless the incident polarization amplitude is quite large, however, a simpler approximation should often be applicable. Setting $\lambda = 1$ reduces D^2 to

$$R^2 = (U' - a_0)^2 + Q'^2 \tag{23}$$

This is simply the square of the distance from the point with rectangular coordinates U' and Q' to the point Z. The probability density is

$$P_{U'Q'}(U', Q') = \frac{1}{2\pi\sigma'^2}\exp(-R^2/2\sigma'^2). \tag{24}$$

The lines of constant probability are simply circles centered at Z. The probability that a point is farther from Z than R is $\exp(-R^2/2\sigma'^2)$.

It is straightforward to reexpress this in terms of the polar coordinates a_e and $2\phi_e$, noting that the element of area in these coordinates is $a_e \cdot d(a_e) \cdot d(2\phi_e)$. Also, the variance $\sigma'^2 = 2/ < N >$ can be replaced by $2/N$ if the number of counts is large. The result is

$$P(a_e, \phi_e) = \frac{N a_e}{4\pi}\exp[-\frac{N}{4}(a_e^2 + a_0^2 - 2a_e a_0\cos(2\phi_e))] \tag{25}$$

which is the formula commonly used in discussions of this topic (with $\phi_0 = 0$).

Some details of our simulations and calculations are given in Appendix B, including the Fortran program used in the simulations.

3. DRAWING CONCLUSIONS FROM AN OBSERVED DATASET

We turn now to the question of how we can draw conclusions about the magnitude and direction of incoming X-rays from a single data set. We have the number of counts at each azimuthal angle with respect to some angle on the sky. We calculate I_f, U_f, and Q_f for our angular distribution of counts, and from them we compute two quantities: U' and Q'. We can think of these as the rectangular components of a point in a coordinate system with the U' axis in our reference direction. The origin of this coordinate system is the point corresponding to a completely unpolarized beam, for which both components would vanish.

What we do not know but would like to draw conclusions about is the location of the point Z in this coordinate system. By definition, Z is the point whose distance from the origin is the true polarization amplitude a_0 and which is in the same direction as the true polarization direction.

A contour line enclosing the observed data point can be defined by the fact that the probability of the observed U' and Q' values for any Z_C (the C stands for "candidate") located at a point on that contour has the same value. Any point inside the contour is a location for Z_C that gives a probability of the data that is higher, and any point outside the contour is a location for Z_C such that the data has a lower probability. This defines a confidence region for Z_C.

We want to find the coordinates of the points Z_C for which the probability of the data point, with coordinates U' and Q', has a specified probability. The geometry of this situation is very similar to the case previously studied, where one wanted to find which data points have a specified probability when the true amplitude is known. We define the angle $2\eta_C$ as the angle between lines from the origin to the data point and to Z_C, the difference between the angular position of the data point and the angular position of the point Z_C.

In order to have a specified probability, the data point must lie somewhere on an ellipse whose center is at the point Z_C, whose major axis is perpendicular to the line from the origin to Z_C and has the length $2D$, and whose minor axis is $2D/\lambda_C$, with $\lambda_C = 1/\sqrt{1 - a_C^2/2}$. D is related to the probability in the same way as in the previous discussions; it depends only on the total number of counts in the data set and the specified probability.

From the earlier section, the requirement for the desired probability is given by Equation (20), which can be rewritten for the present purpose as

$$D^2 = \frac{[a_e \cos(2\eta_C) - a_C]^2}{1 - a_C^2/2} + a_e^2 \sin^2(2\eta_C). \tag{26}$$

We define a pair of new variables

$$u = a_C \cos(2\eta_C) \ ; \ \ v = a_C \sin(2\eta_C) \tag{27}$$

which specify the location of Z_C relative to the data point. Using these, the condition that the data point have the correct position relative to Z_C can be written as

$$u^2(1 + D^2/2) - 2a_e u + v^2[1 + (D^2 - a_e^2)/2] = D^2 - a_e^2 \tag{28}$$

This is a quadratic in u and v and thus defines an ellipse. It is even in v, so the axes lie along and perpendicular to the line from the origin to the data point. The quadratic equation for u when $v = 0$ has two roots equidistant from $u = a_e/(1 + D^2/2)$, so the center of the ellipse is located at $u = a_e/(1 + D^2/2)$; $v = 0$. The v semi-axis is $D/\sqrt{1 + D^2/2}$ and the u semi-axis is $[D/(1 + D^2/2)]\sqrt{1 + D^2/2 - a_e^2/2}$. (A more thorough and detailed derivation is provided in Appendix C.)

Since the derivation of probabilities in the earlier section included neglecting D^2 compared with unity, it is not inconsistent to do the same here. The result is simpler than might have been expected. The ellipse is centered on the data point and has axes of $2D$ and $2D/\lambda_e$. The data point and the candidate have changed places. Contour lines around the data point (in the $U'Q'$ plane) with a specified value of D are given by:

$$D^2 = (u - a_e)^2\lambda_e^2 + v^2 \tag{29}$$

where $u = a_C \cos(2\eta)$ and $v = a_C \sin(2\eta)$ as defined in Equation (31), and 2η is the angle between the data point and the point on the contour.

Since the minor axis of the ellipse is on the line that passes through the origin and the data point, the points which are closest to and farthest from the origin, and correspond to the smallest and largest polarization amplitude, are the end points of the minor axis. In an amplitude-angle plot, the maximum and minimum amplitudes occur at the same angle, and are equally spaced from the center of the constant-probability contour. The reflection symmetry of the ellipse about its minor axis guarantees that for any amplitude the two angles are above and below the central angle by the same amount, so there is also a reflection symmetry in an amplitude-angle plot of the contour.

One can derive a description of probabilities similar to the situation of a known incident amplitude. The area of an ellipse with a given value of D would be just $\pi D^2/\lambda$, and the incremental area between ellipses for D and $D + dD$ would be proportional to DdD. An integral from D to infinity is then proportional to $\exp(-D^2/2\sigma'^2)$. Normalization makes the probability outside the contour equal to $\exp(-D^2/2\sigma'^2)$. It seems likely that in any practical situation there will be large enough additional uncertainties in data acquisition not included in this statistical analysis that the difference in the contours for approximate or full inclusion of D^2dependence would not be important.

Figure 5 shows the implication of the results for a particular estimated polarization from a measurement sample.

The marginal probability of U' and Q' derived in the previous section is in terms of $< N >$ for many samples. For a single sample, one has only the N for that sample. Keeping in mind that it already the case that $< N >$ is large enough, and D small enough, to result in a Gaussian distribution of probabilities in the $U'Q'$ plane, it is reasonable to obtain the probability contour for a desired probability level by using the value for D that would be the correct one if $< N >$ were equal to N. That contour is still a line of constant probability, but for a different probability if $< N >$ differs from the known N. The dependence of the probability on $< N >$ is not a strong one though, it is unlikely that $< N >$ and N differ greatly, and the corresponding uncertainty may well not be important. As an example, suppose we have a sample with N equal to 10000, and we want a “1 sigma” contour. Our contour is actually the “1.01 sigma” contour if $< N >$is actually 10133 instead the 10000 value we used. If $< N >$ is actually 9889, our contour is actually correct for

"0.99 sigma". So a difference of one standard deviation in how much N differs from $< N >$ results in a difference of about 1 percent in the "sigma value". For the "3 sigma" contour, it is actually the "3.02 sigma" contour if $< N >$ is actually 10111, or the "2.98 sigma" if $< N >$ is actually 9889. It seems unlikely that this level of uncertainty in the true confidence levels for different probabilities is important for indications of which physical situations and processes may be present at the source object. It is not difficult to calculate the corresponding conclusions for other values of N if they are wanted. (Here we are using the conventional description of probabilities with "sigma values": "x sigma" corresponds to the probability that a random value from a normal distribution differs from the mean by x standard deviations or more.)

It has been mentioned that statistical variations in the recorded counts result in a positive probability for any pair of U' and Q' values, even improbable ones that might lead to amplitude estimates that are larger than 1, while the true amplitude can only be zero or positive and less than or equal to 1. This could complicate the correct conclusions that can be drawn from a single data set if the total number of counts is too low or the confidence level chosen leads to a formal contour that extends into regions where $U'^2 + Q'^2$ is greater than 1. The simplest way to avoid such problems is just to get more data, or choose a different confidence level, or both, so that the contour is completely within the acceptable range. But if this is not feasible or desirable, one could impose restrictions and provide a truncated confidence region.

4. SUMMARY

The methods presented here for generating simulated counts of scattered electrons at different azimuthal angles, and for analyzing the resulting angular distribution, have the special advantage of making it possible to derive rigorous results about the estimation of polarization amplitudes and directions and their uncertainties. Alternative methods for generating simulations, and for analyzing the angular distribution, should give results that are essentially similar to these, but some of the predicted properties may only be observed rather than derived.

We confirm previous work on amplitude and angle estimates for cases in which the incident amplitude is not too large. With our approach we are able to provide analytical treatment for larger incident amplitudes, so that the entire physical range of possibilities is covered.

A. THE DISTRIBUTIONS OF U_f, Q_f AND OF U', Q'

It is useful to introduce another variable, linear in U_f and I_f. Because $\langle U_f \rangle = a_0 \langle I_f \rangle$, the mean value of $(U_f - a_0 I_f)$ is zero and because the covariance of I_f with U_f is just a_0 times the variance of I_f, the covariance of $(U_f - a_0 I_f)$ with I_f is zero. The variance of $(U_f - a_0 I_f)$ is $\sigma_f^2(1 - a_0^2/2)$. We introduce the quantity

$$\lambda = 1/\sqrt{1 - a_0^2/2} \tag{A.1}$$

and define the new variable

$$V_f = \lambda(U_f - a_0 I_f). \tag{A.2}$$

V_f has a mean value of zero and the same variance as Q_f and is uncorrelated with I_f and Q_f.

The three variables, I_f, V_f, and Q_f are trivariate normal, and mutually uncorrelated, so they are independent. Then the joint probability distribution function is given by

$$P_{I,V,Q}(I_f, V_f, Q_f) = \frac{\sqrt{2}}{(2\pi\sigma_f^2)^{3/2}} \exp[-\frac{1}{2\sigma_f^2}(2(I_f - \langle I_f \rangle)^2 + V_f^2 + Q_f^2)]. \tag{A.3}$$

What we would like to have is the joint probability of $Q' = Q_f/I_f$ and $V' = V_f/I_f = \lambda(U' - a_0)$. To obtain this, we transform to V' and Q' as our variables, and integrate the probability of I_f, V', Q' over I_f. For the transformation we have $dV'dQ' = I_f^2 dV_f dQ_f$. Since V' and Q' only appear in the combination

$$D^2 = V'^2 + Q'^2 \tag{A.4}$$

and $\langle I_f \rangle = \langle N \rangle/(2\pi) = \pi\sigma_f^2$, the integral to be evaluated is

$$P_{V'Q'}(V', Q') = \frac{\sqrt{2}}{(2\pi\sigma_f^2)^{3/2}} \int_{-\infty}^{\infty} dI_f I_f^2 \exp[-\frac{I_f^2}{\sigma_f^2}(1 + D^2/2) + 2\pi I_f - \pi^2{\sigma_f}^2], \tag{A.5}$$

which can be evaluated exactly. It is convenient to introduce

$$\sigma'^2 = \frac{1}{\pi^2\sigma_f^2} = \frac{2}{\langle N \rangle} \tag{A.6}$$

. Then

$$P_{V'Q'}(V', Q') = \frac{1 + (1 + D^2/2)/\langle N \rangle}{(1 + D^2/2)^{5/2}} \frac{1}{2\pi\sigma'^2} \exp(-\frac{D^2}{2\sigma'^2}). \tag{A.7}$$

In any practical application, the value of $\langle N \rangle$ will be so large that the term proportional to $1/\langle N \rangle$ can be neglected in comparison with unity. Moreover, the values of D that might be of interest will be small enough to justify replacing $(1 + D^2/2)$ with unity. Otherwise, the probabilities involved are so small that they would be of no value. Accordingly the initial fraction in Equation(A.7) can be dropped, leaving

$$P_{V'Q'}(V', Q') = \frac{1}{2\pi\sigma'^2} \exp(-\frac{D^2}{2\sigma'^2}). \tag{A.8}$$

The distribution of D^2 is just a chi-squared distribution with two degrees of freedom.

Thus it is clear that

$$P_{U'Q'}(U',Q') = \frac{\lambda}{2\pi\sigma'^2}\exp(-\frac{D^2}{2\sigma'^2}). \tag{A.9}$$

This is the form in which U' and Q' are approximately independent normal variables with the variance in U' lower than that of Q' by $1-a_0^2/2$. These results for the variance of U' and Q' agree with the estimates given by considering $\langle\delta^2 U'\rangle$ and $\langle\delta^2 Q'\rangle$, where expanding around the mean values gives $\delta U' = \delta U_f/\langle I_f\rangle - \langle U_f\rangle\delta I_f/\langle I_f\rangle^2$ and similarly for $\delta Q'$.

If $\phi_0 \neq 0$, the following means and covariances have a ϕ_0 dependence:

$$\begin{aligned}
\langle U_f\rangle &= \frac{\langle N\rangle}{2\pi}a_0\cos(2\phi_0), \\
\langle Q_f\rangle &= \frac{\langle N\rangle}{2\pi}a_0\sin(2\phi_0), \\
\langle\delta U_f\delta I_f\rangle &= \frac{\langle N\rangle}{4\pi^2}a_0\cos(2\phi_0), \\
\langle\delta Q_f\delta I_f\rangle &= \frac{\langle N\rangle}{4\pi^2}a_0\sin(2\phi_0).
\end{aligned} \tag{A.10}$$

It remains true that

$$\langle\delta U_f\delta Q_f\rangle = 0.$$

With

$$\begin{aligned}
X &= U\cos(2\phi_0) + Q\sin(2\phi_0), \\
Y &= -U\sin(2\phi_0) + Q\cos(2\phi_0), \\
Z &= \lambda(X - a_0 I_f),
\end{aligned} \tag{A.11}$$

corresponding to Equation (A.3), we have

$$P_{IZY}(I_f, Z_f, Y_f) = \frac{\sqrt{2}}{(2\pi\sigma_f^2)^{3/2}}\exp[-\frac{1}{2\sigma_f^2}(2(I_f - \langle I_f\rangle)^2 + Z_f^2 + Y_f^2)]. \tag{A.12}$$

X, Y are just a set of axes rotated from U, Q by $2\phi_0$. Since

$$\begin{aligned}
U &= a_e\cos(2\phi_e) \\
Q &= a_e\sin(2\phi_e),
\end{aligned}$$

we have

$$\begin{aligned}
X &= a_e\cos(2(\phi_e - \phi_0)) \\
Y &= a_e\sin(2(\phi_e - \phi_0)).
\end{aligned}$$

The probability only depends on the angle between $\boldsymbol{a_e}$ and the true polarization $\boldsymbol{a_0}$, which is $2\eta = 2(\phi_e - \phi_0)$.

With $X' = X/I'$, $Y' = Y/I'$, and $D^2 = \lambda^2(X' - a_0)^2 + Y'^2$ the general result is derived as before, and similarly, for large $\langle N\rangle$,

$$P_{X'Y'}(X', Y') \approx \frac{\lambda}{2\pi\sigma'^2} \exp(-\frac{D^2}{2\sigma'^2}). \tag{A.13}$$

B. COMPUTATIONAL REMARKS

This technical note describes the simulations done to explore the result of least squares calculations of the Stokes parameters. These simulations were first done with Fortran routines. They are discussed below and the Fortran program is given. Later the results of generating the data and analyzing it in this way were compared, as discussed in Montgomery & Swank (2015), to the results of generating it and analyzing it in the way described by Strohmayer & Kallman (2013). Much of the comparison work used IDL with version 8.2 RANDOMU to generate Poisson variates. Fits were carried out using IDL and the fitting routines based on MINPACK-1 (Markwardt 2009 ASP Conf. Ser. 411, Astronomical Data Analysis Software and Systems XVIII, p. 251). As discussed in the paper, differences found were ascribable to the differences between weighted and unweighted least squares fitting.

While the discussion in this paper and in this technical note provides rather simple, closed form expressions for amplitude and angle estimate uncertainties which should suffice for most applications, the generation of samples is easy and not very time-consuming. In particular, for a given number of angular bins, M, all the angular quantities required for evaluating the polynomial coefficients will be the same for every sample, and can be computed once, before specific n_j values have been generated. Then each sum used in evaluating the values of I_f, U_f, and Q_f is reduced to a simple scalar product.

The following short and simple Fortran program can process several thousand samples, of several thousand points each, in a few seconds on a fairly modest desktop or laptop computer. Besides providing checks on the equations, it could be useful in situations where some of the approximations used may not be valid. For example, if the number of bins is too small for the asymptotic normality provided by the central limit theorem, it is easy to redistribute counts into the number of bins available and proceed to find the amplitude and angle estimates for a large number of samples.

The program itself is written in Standard Fortran77, and so is also Standard F90 and F95 and F2003 and F2008, and should work with no problems on any contemporary fortran system, but should also be easy to rewrite in another language if desired: e.g. IDL, Matlab, Python, Haskell.

A couple of remarks about this specific program:

The random number generator used is the RAND function, which is not standard but almost universally provided by Fortran compilers. It could, and probably should, be replaced by the standard F95 subroutines RANDOM_SEED and RANDOM_NUMBER, which are also available on any fortran system less than 20 years old. Replacement with any other random number generator would be simple. This application does not need a particularly good generator; any default is probably adequate.

The function IPOISSON is a minimal one, adequate for this use. There are many more sophisticated ones available. If the number of counts per bin gets up into the hundreds an improvement might be desirable (or the number of bins increased). [The limitation comes from underflow in the exponential of a large negative number, and is therefore quite system-dependent.]

```
c   polsim1.f    simulation of counts and analysis for x-ray polarization
c   cgm  2013
c   Standard Fortran77 except for "implicit none" and "rand" c
c   Variable names intended to resemble those in text
c   Nbar is intended value for number of counts
c   Ntot is actual value in a particular sample
c   Phiest is angle estimate in degrees
c   Uz, Qz are displacements from point Z
c   R2 is square of distance from Z
c   D2 is square of ellipse semi-major axis
c
c   Usage: choose values for Nsamp,Nbar,Azero
c          adjust write statement if desired
c          compile, link, execute.

      implicit none
      integer Mbins, Nsamp, Ntot, ipoisson
      parameter (Mbins = 499)
      integer i, k, ks
      integer nj(Mbins), seed(3)
      real Nbar, Azero, OneoverPi, p, f, g, R
      real Vcos(Mbins), Vsin(Mbins)
      real Ic, Uc, Qc, Aest, Phiest
      real Uz, Qz, R2, D2, lam2

      Nsamp = 49000
      Nbar = 4000.
      Azero = 0.50

      OneoverPi = 0.25/atan(1.0)
      ks = (Mbins+1)/2
      f = 16*atan(1.0)/Mbins
      g = Nbar/Mbins
      lam2 = 1.0/(1.0-0.5*Azero**2)

      DO 10 k=1,Mbins
         p = f*(k-ks)
         Vcos(k) = cos(p)
         Vsin(k) = sin(p)
 10   CONTINUE

      CALL itime(seed)
      R=rand(seed(1)*seed(2)+seed(3))
```

```
      OPEN (13, file='psimout')

      DO 50 i=1,Nsamp
         ntot = 0
         DO 20 k=1,Mbins
            R = g*( 1.0+Azero*Vcos(k) )
            nj(k) = ipoisson(R)
            Ntot = Ntot + nj(k)
 20      CONTINUE

         Uc = 0.0
         Qc = 0.0
         DO 30 k=1,Mbins
            Uc=Uc + nj(k)*Vcos(k)
            Qc=Qc + nj(k)*Vsin(k)
 30      CONTINUE

         Uc=OneoverPi * Uc
         Qc=OneoverPi * Qc
         Ic = 0.5*OneoverPi*Ntot
         Uz=Uc/Ic-Azero
         Qz=Qc/Ic
         R2=Uz**2 + Qz**2
         D2=Qz**2 + lam2*Uz**2
         Aest = sqrt(Uc**2 + Qc**2)/Ic
         Phiest = 90*OneoverPi*atan2(Qc,Uc)
         WRITE(13,*) Uz, Qz, Aest, Phiest, Ic, Uc, Qc

 50   CONTINUE
      CLOSE(13)
      END
C******************************
        function ipoisson(a)
C-----------returns integer from Poisson distribution with mean a
C-------------uses intrinsic rand(flag), which is probably C srand
C-------------can be slow for arguments much larger than 10
C-------------not reliable for arguments over 400 or so
       real a
       double precision b, c
       integer i
       b=exp(-1.d0*a)
       i=-1
       c=1
```

```
10  i=i+1
    c=c*rand(0)
    if (c .gt. b) goto 10
    ipoisson=i
    return
```

C. DERIVATION OF CANDIDATE CONTOURS

The probability of a particular measured polarization, given a known true polarization, depends only on D^2 (Equation (20) in the text), where

$$D^2 = \frac{[a_e \cos(2\eta) - a_0]^2}{1 - a_0^2/2} + a_e^2 \sin^2(2\eta). \tag{C.1}$$

We wish to determine the candidate true polarizations a_C at angle 2η from a measured polarization amplitude and direction that would have a particular probability. The values must obey the same equation with $a_0 \to a_C$. The candidate point Z_C has the projections $u = a_C \cos 2\eta$ and $v = a_C \sin 2\eta$ parallel and perpendicular to $\mathbf{a_e}$. Then, using these substitutions and $a_C^2 = u^2 + v^2$, Equation (28) in the text is obtained, which can be written as

$$D^2 - a_e^2 - (u^2 + v^2)(1 + D^2/2) + a_e^2 \frac{v^2}{2} + 2a_e u = 0. \tag{C.2}$$

The following sequence of reorganizations:

$$\frac{(D^2 - a_e^2)}{1 + D^2/2} - u^2 - v^2 + \frac{a_e^2 v^2/2}{1 + D^2/2} + \frac{2a_e u}{1 + D^2/2} = 0, \tag{C.3}$$

$$\frac{(D^2 - a_e^2)}{1 + D^2/2} - (u - \frac{a_e}{1 + D^2/2})^2 + \frac{a_e^2}{(1 + D^2/2)^2} - v^2(1 - \frac{a_e^2/2}{1 + D^2/2}) = 0, \tag{C.4}$$

$$\frac{D^2}{1 + D^2/2} - \frac{a_e^2}{1 + D^2/2}(1 - \frac{1}{1 + D^2/2}) - (u - \frac{a_e}{1 + D^2/2})^2 - v^2(1 - \frac{a_e^2/2}{1 + D^2/2}) = 0, \tag{C.5}$$

$$\frac{D^2}{1 + D^2/2}(1 - \frac{a_e^2/2}{1 + D^2/2}) - (u - \frac{a_e}{1 + D^2/2})^2 - v^2(1 - \frac{a_e^2/2}{1 + D^2/2}) = 0. \tag{C.6}$$

leads to

$$\frac{D^2}{1 + D^2/2} = (u - \frac{a_e}{1 + D^2/2})^2/(1 - \frac{a_e^2/2}{1 + D^2/2}) + v^2. \tag{C.7}$$

This is the ellipse centered on $a_e/(1 + D^2/2), 0$ with semi-axes $(D/(1 + D^2/2))\sqrt{1 - a^2/2 + D^2/2}$ and $D/\sqrt{1 + D^2/2}$ for u and v, respectively. Now neglecting $D^2/2$ compared to 1,

$$D^2 = \frac{(u - a_e)^2}{1 - a_e^2/2} + v^2. \tag{C.8}$$

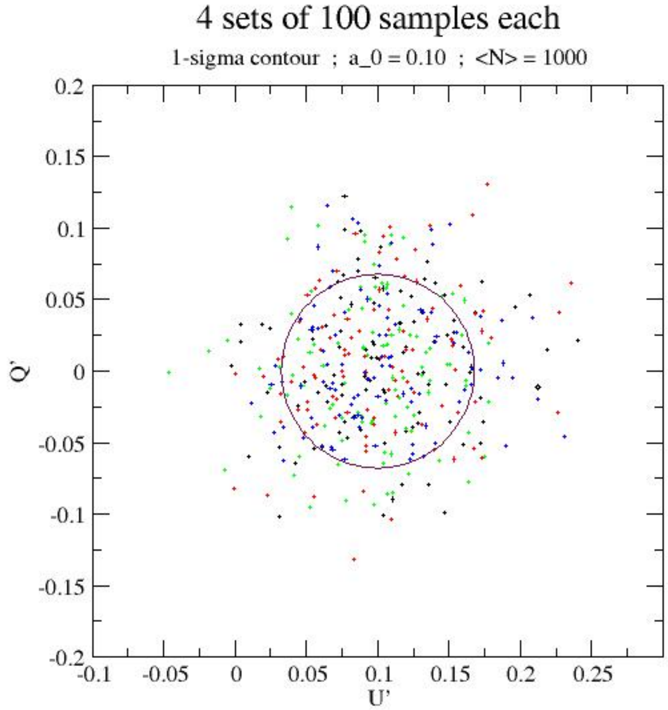

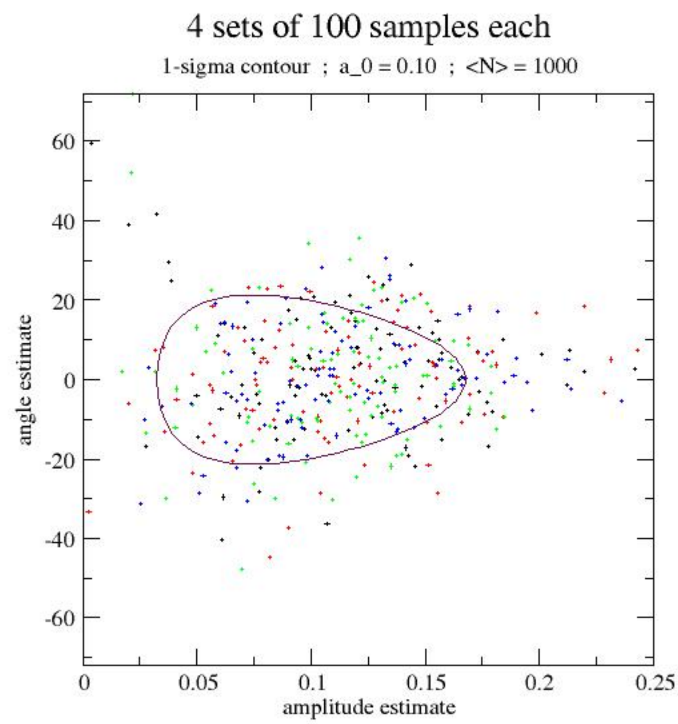

Fig. 1.— (top) U' and Q' for 4 sets of 100 simulations each for $< N >= 1000$, with the true polarization along the U_f axis. Points from the different sets are indicated by triangles of different orientation. The 68.3 % probability contour is centered on the point Z at $a_0 = 0.1, 0$. (bottom) The polarization amplitude a_e and angle ϕ_e corresponding to the U' and Q' results. ϕ_e is half the angle between the U' axis and the direction to the point U', Q'.

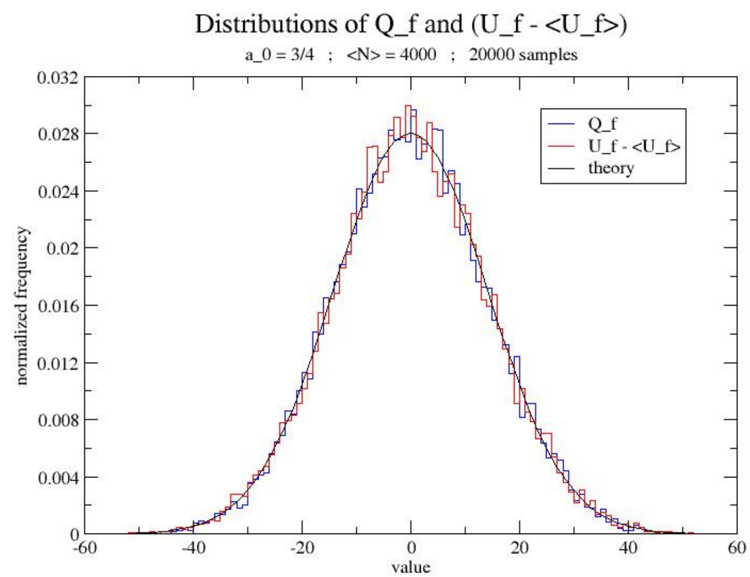

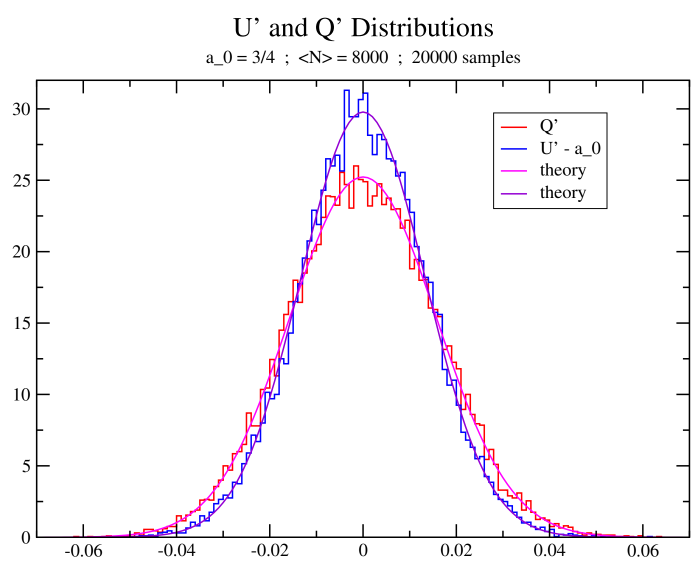

Fig. 2.— (top) Marginal distributions of U_f and Q_f for 20,000 simulations with $< N >= 4000$ for a polarization $a_0 = 3/4$ along the U_f axis. The theoretical distributions are the same. (bottom) Marginal distributions of U' and Q' for 20,000 simulations with $< N >= 8000$ for a polarization $a_0 = 3/4$ along the U_f axis. Here the ordinate is the number of simulations in a bin. The curves are the predictions for independent normal distributions.

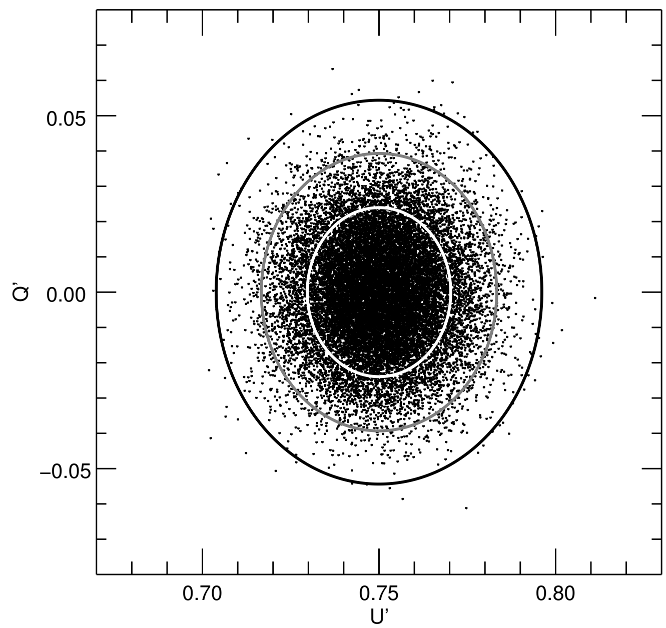

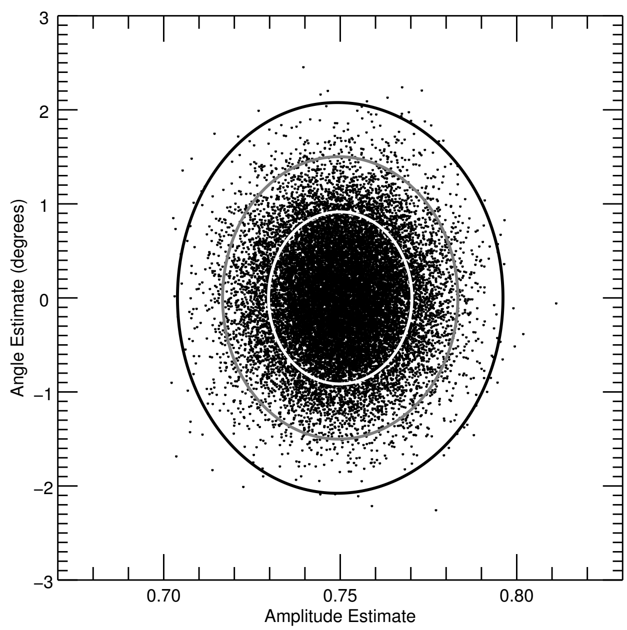

Fig. 3.— (top)Distribution of U' and Q' for $20,000$ simulations (the same as for Fig 2, bottom) with $< N >= 8000, a_0 = 3/4$, and $\phi_0 = 0$. Contours of the predicted 1, 2, and 3 sigma levels (68.27, 95.45, and 99.73 %) are superposed. 13597, 19010, and 19935 simulations fell within those contours, in comparison to 13654, 19090, and 19946 expected. The semi-major axes are along the Q' axis and the semi-minor axes along the U' axis, centered on $a_0, 0$. (bottom) The corresponding a_e and ϕ_e with their theoretical contours.

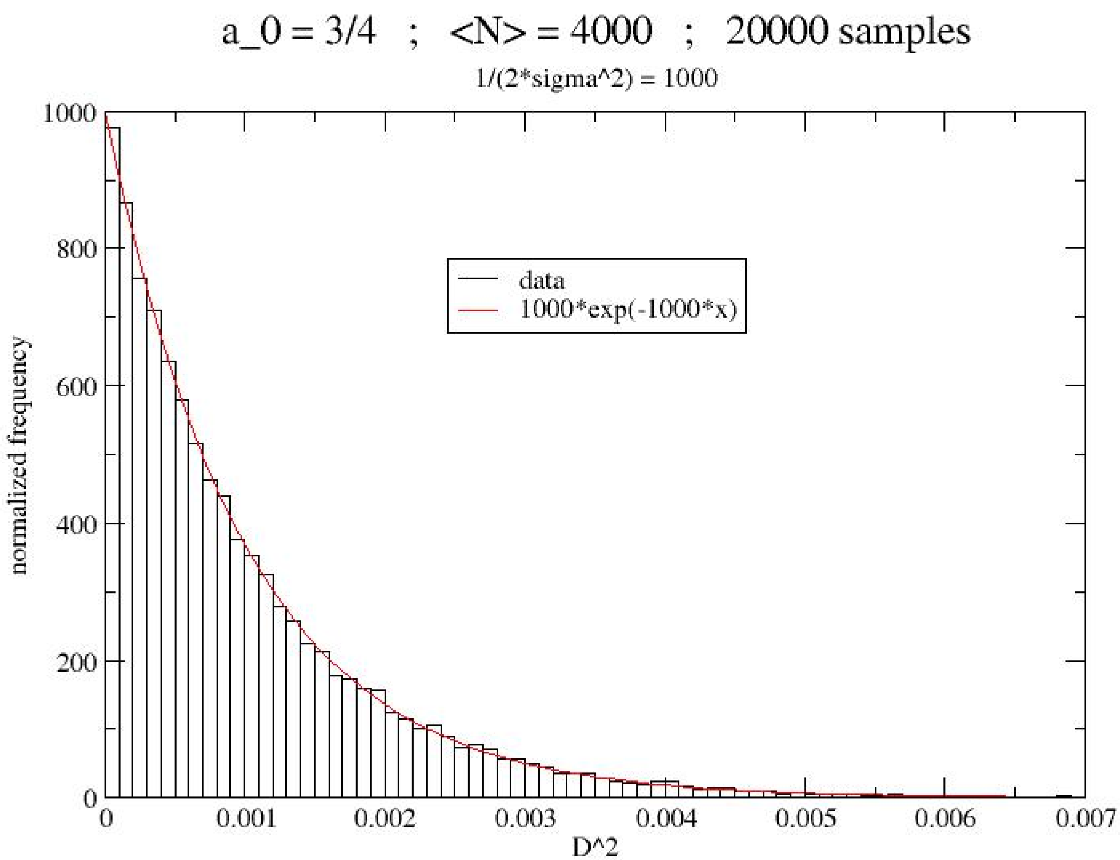

Fig. 4.— Distribution in D^2 for a simulation with $a_0 = 3/4$ and $< N >= 4000$, for $N_{sample} = 20,000$. The number of samples in a D^2 increment $\Delta D^2 = 10^{-4}$, is plotted, normalized by $N_{sample}\Delta D^2(= 2)$, together with the expected values. For this case $1/2\sigma'^2 = 1000$. The probability for the measured polarization to lie outside of D^2 matches the predicted exponential.

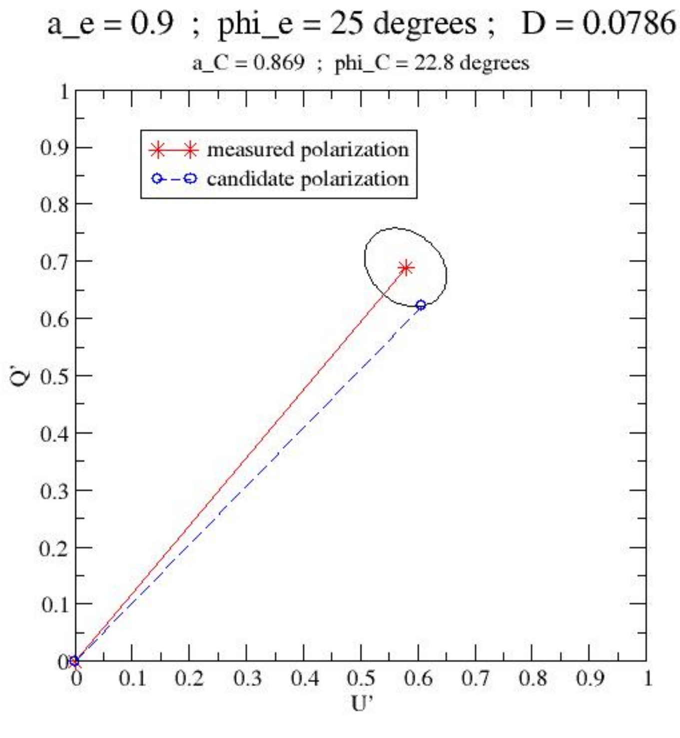

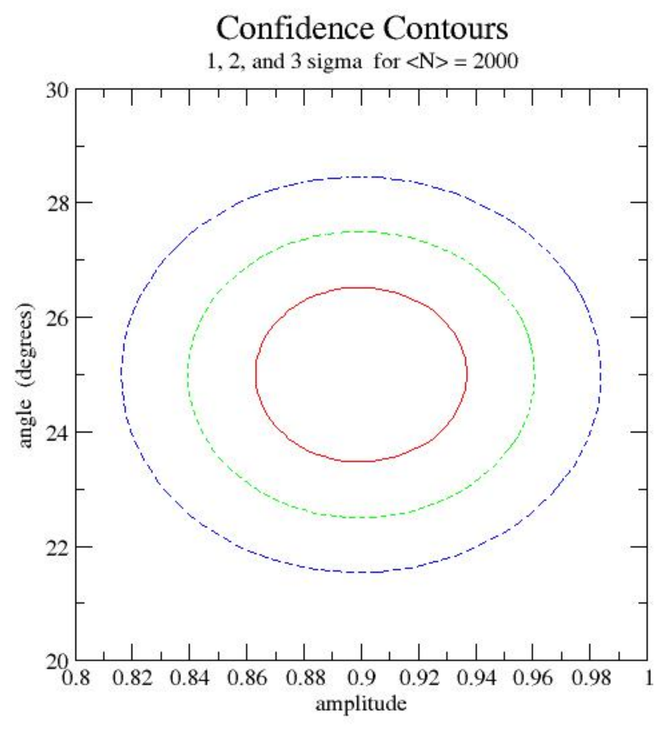

Fig. 5.— Examples of contours for the candidate true polarization quantities for a measurement of $a_e = 0.9$ and $\phi_e = 25$ degrees. (top) Contour for 95.0 % (theoretical) confidence in $U'Q'$ space for the true polarization, for $D = 0.0786$. The dashed line indicates a candidate true polarization at distance a_C from the origin and angle 2η from the measured polarization. (bottom) Contours (solid and dashed lines, respectively) of 1 sigma (68.3 %, red), 2 sigma (95.0 %, green) and 3 sigma (99.7 %, violet) for the true amplitude and angle, for $N = 2000$ events.