Effects of Poisson-Distributed Photon Noise on Spitzer Photometry

John W. Fowler

8 June 2004

Several erroneous statements have been made recently regarding the role of Poisson-distributed photon noise on Spitzer photometry. The specific case involved MIPS-24 BCDs and mosaics, but the fundamental aspects of these remarks are more general and may be applied to any photometric data dominated by photon noise.

Claim No. 1

It was claimed that flux was not "conserved" because photon noise caused higher flux measurements to have higher uncertainties, penalizing them when inverse-variance weighting was applied in flux averaging and thereby introducing a "bias". The inverse-variance weighting is a feature of the refinement of measurements involving Gaussian noise, but the photon noise, as described (quite accurately) by Poisson statistics was well into its Gaussian limit (e.g., at 10^4 electrons, the Poisson skewness, which is the inverse of the square root of the mean, is 1%, and for most purposes the Gaussian approximation is more than acceptable; for bright but not saturated fluxes, the electron count is typically more like 10^5 electrons, even more safely into the Gaussian limit).

It should be noted that any claim that flux is not conserved presumes a knowledge of the correct answer with respect to which the alleged violation is deficient (or excessive). In this case, however, circular reasoning had been used in the claim of flux loss: the standard to which inverse-variance weighted and unweighted averaging was compared was an external result based on unweighted averaging with no given justification. It is not surprising that the unweighted result agreed with the external standard. A similar fallacy was presented involving averaging of BCD fluxes.

It can be easily shown that placing the blame for the observed anomalies on the use of Poisson noise in its Gaussian limit is without foundation. First, it should be noted that when a Poisson process is in its Gaussian limit, all that this means is that we have a Gaussian distribution whose variance is equal to its mean. Clearly this can be applied only to dimensionless quantities that can be counted, in this case electrons. The Gaussian refinement formula for two measurements is

$$\bar{x} = \frac{\dfrac{\bar{x}_1}{\sigma_1^2} + \dfrac{\bar{x}_2}{\sigma_2^2}}{\dfrac{1}{\sigma_1^2} + \dfrac{1}{\sigma_2^2}} = \frac{\dfrac{\bar{x}_1}{\bar{x}_1} + \dfrac{\bar{x}_2}{\bar{x}_2}}{\dfrac{1}{\bar{x}_1} + \dfrac{1}{\bar{x}_2}} = \frac{2}{\dfrac{\bar{x}_1 + \bar{x}_2}{\bar{x}_1 \bar{x}_2}} = \frac{\bar{x}_1 \bar{x}_2}{\bar{x}_u}$$

where x_u is the unweighted mean given by

$$\bar{x}_u \equiv \frac{\bar{x}_1 + \bar{x}_2}{2}$$

and

$$\bar{x}_i \equiv \textit{observed result \#}i$$
$$\sigma_i \equiv \textit{one-sigma error for } \bar{x}_i$$

Let us define the relative difference between the weighted and unweighted average as

$$R \equiv \frac{\bar{x}_u - \bar{x}}{\bar{x}_u}$$

Given discrepant measurements, without loss of generality we can define the observed values such that the second is the larger, and the ratio of the larger to the smaller is

$$\alpha \equiv \frac{\bar{x}_2}{\bar{x}_1}, \alpha > 1, \bar{x}_2 > \bar{x}_1$$

Then

$$\bar{x} = \frac{\bar{x}_1 \bar{x}_2}{\bar{x}_u} = \frac{\alpha \, \bar{x}_1^2}{\frac{\bar{x}_1 + \alpha \, \bar{x}_1}{2}} = \frac{2\alpha \, \bar{x}_1^2}{(\alpha + 1)\, \bar{x}_1} = \frac{2\alpha}{\alpha + 1} \bar{x}_1$$

$$R = \frac{\left(\frac{\alpha + 1}{2} - \frac{2\alpha}{\alpha + 1} \right) \bar{x}_1}{\left(\frac{\alpha + 1}{2} \right) \bar{x}_1} = 1 - \frac{4\alpha}{(\alpha + 1)^2}$$

For example, a 20% difference in quantities to be averaged would involve

$$\alpha = 1.2$$
$$R = 0.826\%$$

Thus a 20% difference in quantities averaged would yield less than 1% difference between the weighted mean and the unweighted mean; this would hardly draw attention to any “bias”. On the other hand, for a bright but unsaturated source, a typical pixel value would be about 10^5 electrons in the well; the Poisson sigma would be the square root of this, or 316.23. So then a 20% difference, i.e., 20,000 electrons, would correspond to a chi-square value of

$$\chi^2 = \frac{20{,}000^2}{316.23^2 + 346.41^2} = 1818.2$$

or in other words, we would have averaged two measurements that differed by 42.6 sigma. Note that the chances of getting a 42.6 sigma case are one in 10^{396}. At face value, even unweighted averaging is a very poor option. But even this would yield less than 1% difference between weighted and unweighted means. The claim was that the difference between means was about 5%, which would require a 57.6% difference in the quantities averaged, which would be a 113.5-sigma discrepancy in what was averaged (one chance in 10^{2800}). Clearly the property of Poisson-distributed photon noise having larger uncertainties for larger fluxes is not capable of producing the discrepancies reported. Something else is seriously wrong, and blind use of unweighted averaging without further probing into what is wrong would be extremely foolhardy.

Claim No. 2

It was said that photon (Poisson) noise is a feature of "the process" and not of the "measurement", and as such it should be ignored in any averaging of separate measurements. This would be true if and only if (a.) only photon noise were involved; (b.) each measurement was drawn from the exact same Poisson process. Neither of these conditions is satisfied in our case, but for simplicity, assume that (a.) is satisfied. Condition (b.) is still significantly violated in the real world because of the variation in quantum efficiency over the pixels, hence the need to perform flat fielding.

For example, consider two pixels, one with a responsivity (as reflected in the flat field model) of 0.75 and one with 1.25 (these are not at all extreme values). With an input photon stream averaging 10^5 electrons in a standard integration time on a pixel with responsivity = 1.0, these two detectors would have expected values of 75,000 and 125,000 respectively, with one-sigma uncertainties of 274 and 354 electrons, respectively. After applying the flat field, these become 100,000 electrons each in expected value, but with one-sigma values of 365 and 283 electrons, respectively.

The flat-field destroys the pure Poisson character of the distribution, i.e., the means are no longer equal to the variances. But both distributions are in their Gaussian limits. One happens to have a significantly smaller sigma than the other. In practice, these detectors will exhibit random draws based on their means of 100,000 each and sigmas of 365 and 283, respectively, with the more responsive pixel almost always having the smaller uncertainty, even when it has the higher flux. Inverse-variance weighting of any averages of these detectors' outputs is therefore highly desirable and will always yield slightly smaller uncertainties than unweighted averaging (as shown below). The following table lists some aspects to compare between the two detectors.

	Detector A	Detector B
Responsivity	0.75	1.25
<x> for 10^5 electron input	75,000	125,000
sigma	274	354
S/N	274	354
mean - one-sigma response	74,726	124,646
mean + one-sigma response	75,274	125,354
after flat-fielding:		
mean - one-sigma response	99,635	99,717
mean + one-sigma response	100,365	100,283
uncertainty	365	283

For illustrative purposes, consider all four combinations of observed flux averaging both with inverse-variance weighting and unweighted. The fluxes after flat-fielding will be denoted A- for detector A's mean - one-sigma, A+ for detector A's mean + one-sigma, and similarly B- and B+ for detector B's mean minus and plus one sigma, respectively. The options for averaging are as listed below.

	Unweighted Mean /Sigma/Error	Inverse-Variance-Weighted Mean/Sigma/Error
A-B-	99,676 / 223 / -324	99,686 / 223 / -314
A-B+	99,959 / 224 / -41	100,040 / 224 / 40
A+B-	100,041 / 224 / 41	99,960 / 224 / -40
A+B+	100,324 / 224 / 324	100,314 / 224 / 314

In all combinations, the inverse-variance-weighted mean is closer to the correct answer. For both forms of averaging, the uncertainty variance is given by the final mean divided by two. This can be seen as follows. For the inverse-weighted mean, we have the usual Gaussian expression plus the Poisson relation between mean and variance:

$$\frac{1}{\sigma^2} = \frac{1}{\sigma_1^{\,2}} + \frac{1}{\sigma_2^{\,2}} = \frac{1}{\bar{x}_1} + \frac{1}{\bar{x}_2} = \frac{\bar{x}_1 + \bar{x}_2}{\bar{x}_1 \bar{x}_2}$$

But above we showed that

$$\bar{x} = \frac{\bar{x}_1 \bar{x}_2}{\bar{x}_u}$$

where

$$\bar{x}_u = \frac{\bar{x}_1 + \bar{x}_2}{2}$$

so that the inverse-variance-weighted uncertainty is

$$\sigma^2 = \frac{\bar{x}}{2}$$

For the unweighted mean, we have from the theory of functions of a random variable:

$$x_u = \frac{x_1 + x_2}{2} \Rightarrow \sigma_u^2 = \frac{\sigma_1^2 + \sigma_2^2}{4} = \frac{\bar{x}_1 + \bar{x}_2}{4} = \frac{\bar{x}_u}{2}$$

where we used the fact that the measurement errors are independent. Since for Poisson noise, assuming unequal measurements, higher fluxes have higher error variances, and we always have

$$\bar{x} < \bar{x}_u \Rightarrow \sigma^2 < \sigma_u^2$$

although when we take the square roots of the variances and round off, the standard deviations may be equal, as in the table above. We also note that since different detectors have different quantum efficiencies, their effective photon noise *is* peculiar to them, it *is* a part of the measurement, and it cannot be considered something to be ignored in averaging processes. There are also other effects that decouple the pixels into separate Poisson distributions, e.g., different portions of ramps lost to radiation hits in MIPS-24 RAW mode.