Photon Counting - The Statistics of Light

Dylan Nelson
University of California Berkeley
dnelson@berkeley.edu

September 12, 2006

Group Members - Donald VanNess, Trey Jalbert

Abstract

The detection of quantized light in the form of photons offers an accessible method for exploring many of the fundamental limitations of scientific experimentation. When counting photons, we run into physical limitations, such as the ability to detect discrete light quanta. There are also theoretical limitations, such as the way in which the brightness or intensity of a light source can be specified, and to what precision. We first use a photomultiplier tube to count incident photons, and then analyze this data in IDL to compare the observed statistical distributions with the theoretical parent populations. Both the Poisson and Gaussian distributions successfully model our experimental data, though in different regimes. Preliminary results also indicate a strong correspondence between the observed and theoretical distributions, where the error involved can be influenced by tailoring the experimental parameters.

Contents

1 Introduction

The quantum mechanical process responsible for the emission of photons is probabilistic. That is, if an atom is in an excited state with one or more atoms above their ground level energies, we can never know for certain when or if a photon will be emitted, we can only know the probability that it may be emitted. When using a photomultiplier tube (PMT) to count photons, we are essentially taking a finite number of samples of this probability function. Thus, in order to understand the data that will be collected, we must first recognize that it will be governed by the statistical mechanisms underlying photon emission at the quantum level. If we treat the probability function for photon emission as the parent population, our samples will obey the statistical laws of the regimes which they fall in.

Several important statistical properties require explicit definition so that they can be applied uniformly throughout the remainder of this paper. The experimental mean $\bar{x}$ is calculated as in Equation 1:[1]

$$\bar{x} = \frac{1}{N}\sum_{i}^{N} x_i \tag{1}$$

The actual mean μ of the parent population is given, in the theoretical limit as $N \to \infty$, by Equation 2:[2]

$$\mu = \lim_{N\to\infty}\left(\frac{1}{N}\sum_{i}^{N} x_i\right) \tag{2}$$

The compensated variance s^2 (where s is the standard deviation) is calculated from a sample as in Equation 3:[3]

$$s^2 = \frac{1}{N-1}\sum(x_i - \bar{x})^2 \tag{3}$$

Finally, the corresponding expression for the variance σ^2 of the theoretical parent population is expressed, like μ, in the limit of $N \to \infty$. The full expression is given in Equation 4:[4]

$$\sigma^2 = \lim_{N\to\infty}\left[\frac{1}{N}\sum(x_i - \mu)^2\right] = \lim_{N\to\infty}\left(\frac{1}{N}\sum x_i^2\right) - \mu^2 \tag{4}$$

Where in all four cases $\bar{x}$ - the mean of the sample distribution - is the estimate of the true mean μ, and s^2 - the compensated variance - is the estimate of the variance σ^2. These expressions are used extensively throughout the rest of this paper in order to provide a metric in comparing our collected data with theoretical prediction. We will proceed by first describing the collection apparatus, verifying its functionality, and computing the above statistical properties for an acquired data set. Then we

[1] Taylor: "Error Analysis" p98 eq4.5
[2] Graham: "Errors and Statistics" p16 eq2.5
[3] Taylor: "Error Analysis" p100 eq4.9
[4] Graham: "Errors and Statistics" p18 eq2.10

will explore the parameter space in both rate and the number of samples, comparing gathered data with theoretical expectations. Finally, we end with comments on the relationship between statistical measures and experimental parameters, and the consequences of this on applications of photon counting, including those involved in astrophysical observations.

2 Photon Counting

2.1 Approach

The principal aparatus used was aapparatusltiplier tube (PMT), which has three main components. In our apparatus light from a variable intensity LED enters a photocathode, which then emits electrons. These electrons are amplified by an electron multiplier, and then collected by an anode. The strength of the PMT is that it can handle individual photon detections as discrete pulses, which are digitized before output. If the light intensity is sufficiently low such that the seperation between separationgreater than the pulse width, output pulses from the anode are discrete. At this point there exists a direct proportionality between the number of output pulses and the intensity of incident light.

A request for the PMT to conduct a single experiment proceeds as follows. We pass three variables to the sendphot command - the number of requested samples, the rate in Hz at which to collect, and the filename to store the results in. Since the counting board accepts only integer divisions of the maximum count rate (10kHz), allowable values for the rate are 5000Hz, 3333Hz, 2500Hz, and so on. For example, to collect one hundred samples at a rate of 1000Hz we would execute the command:

```
echo counter nsamples=100 rate=1000 fname=100_1000 | sendphot
```

The duration for which the PMT counter is active during such a request is given by Equation 5:

$$t = \frac{N}{\nu} \tag{5}$$

Where N is the requested number of samples, and ν is the rate. In order to streamline this data collection process, a csh script was made which (a) calculates the allowable rates; (b) for a given number of samples, gathers data for all allowable rates up to a requested amount; and (c) generates and executes an IDL batch script which runs the primary analysis on the data. In order to investigate the standard deviation of the mean, the script can also hold the rate constant while varying the number of samples, as well as take any number of multiple samples for a fixed configuration.

2.2 Observations

In order to visualize a single experiment, the corresponding data file was read into IDL and a time sequence plot made, where the x-axis runs in time through the duration of the experiment, and the y-axis represents the number of photons counted for that particular time slice. Following the previous example, the time sequence plot for an experiment of 100 samples at 1000Hz is given in Figure 1.

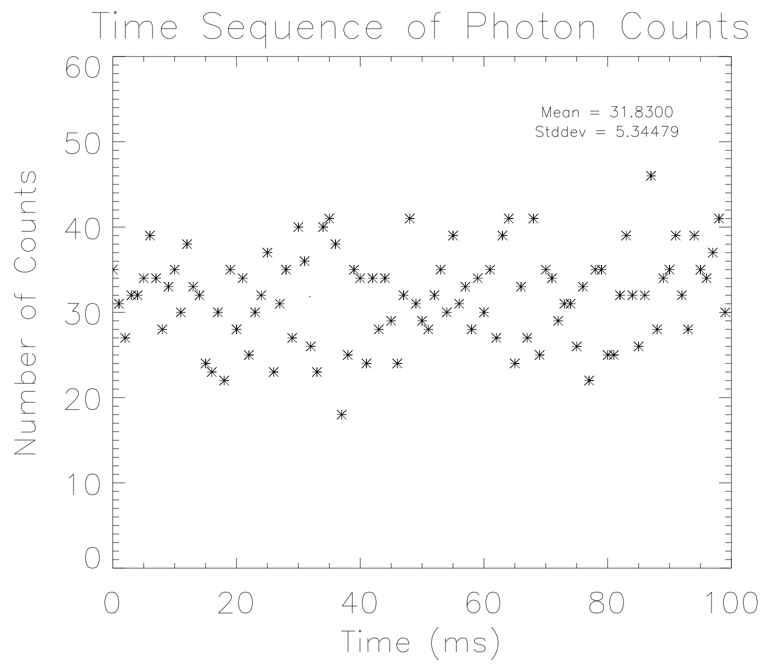

Figure 1. A time sequence plot of the data from a sample experiment. The results consist of 100 samples each taken with a collection rate of 1000Hz. Each sample has an effective collecting duration of 1/10 of a second.

To determine variability between experiments using the same configuration parameters, two tests were run. In each case six unique sets of data were collected with a sample rate of 1000Hz. The first test consisted of 100 samples per data set, the second of 1000 samples per data set. The mean of such a data set is given by Equation 1, while the standard deviation is the square root of the variance given in Equation 3. The results of these two tests are given in Table 1.

It follows immediately that the mean and standard deviation of these sequences varies for each data set. We would expect that, if measuring a phenomenon with infinite precision, repeated measurements should yield exactly the same value. The differing measurements in this case are a direct consequence of the incomplete sampling of the theoretical parent population. Thus, since in a finite time interval we can only gather a finite number of samples, repeated experiments are never guaranteed to give the same result. Effectively, we are measuring $\bar{x}$ and s, not μ and σ.

From Table 1 the differences between the measurement groups of 100 versus 1000 samples is also apparent. The spread of both the mean and standard deviation is much smaller with the greater number of samples. This holds theoretically, as a larger number of collected samples represents a better approximation to the parent population values, and consequently less variation with repeated measurements.

N (Histogram)	Mean ($\bar{x}$)	Stddev (s)
1 (a)	14.2	3.74
2 (b)	15.9	4.05
3 (c)	17.9	4.10
4 (d)	15.5	4.04
5 (e)	15.8	3.69
6 (f)	16.3	3.73

(a)

N	Mean ($\bar{x}$)	Stddev (s)
1	14.1	3.82
2	13.7	3.62
3	14.1	3.82
4	13.1	3.63
5	13.5	3.58
6	13.5	3.70

(b)

Table 1. (a) Results of the two tests, using a rate of 1000Hz for each of the six data sets, using (a) 100 samples per experiment vs. (b) 1000 samples per experiment.

Empirically, our increase of the number of samples by a factor of ten resulted in a decrease in the spread of $\bar{x}$ from 3.7 to 1.0, and a decrease in the spread of s from 0.41 to 0.24. In fact, in the regime where Poisson statistics hold, if we wanted to improve the accuracy of a measurement of the mean by a factor of λ, the number of samples would need to be increased by a factor of λ^2. Equivalently, increasing the number of samples by a factor ρ is equivalent to decreasing the error by $\sqrt{\rho}$.

3 Data Reduction and Analysis

For visualization we can bin the data into intervals and make histograms of the counts. Figure 2 represents the histograms of the previously described test, with the number of samples equal to 100.

We observe that the six histograms, despite representing the same experiment parameters, all have different shapes. This is a consequence of the same argument presented for the mean and standard deviation, and is a fundamental characteristic of sampling the parent population far from the $N \to \infty$ limit. Additionally, there is the effect of random background variation, whereby sources other than our variable intensity LED contribute erroneously to the photon counts. Such a "background rate" could be determined with the LED turned off, and then subtracted out of subsequent measurements as a form of error control, although this was not done explicity for the purposes of this paper.

In order to further test the level to which our measurements met theoretical expectations, we attempted to determine a relation between the number of counts and the standard deviation. This was accomplished by taking a sequential set of data using decreasing sample rates (i.e. increasingly long sample times). In total, 40 measurements were run, each consisting of 100 samples, ranging in rate from 5000Hz to 250Hz. We then plot the variance σ (the standard deviation squared) versus the mean μ, as shown in Figure 3. Here the overplotted line represents $x = y$, or the mean versus itself.

In the regime of $\mu \ll n$ (that is, where the average number of successes is much

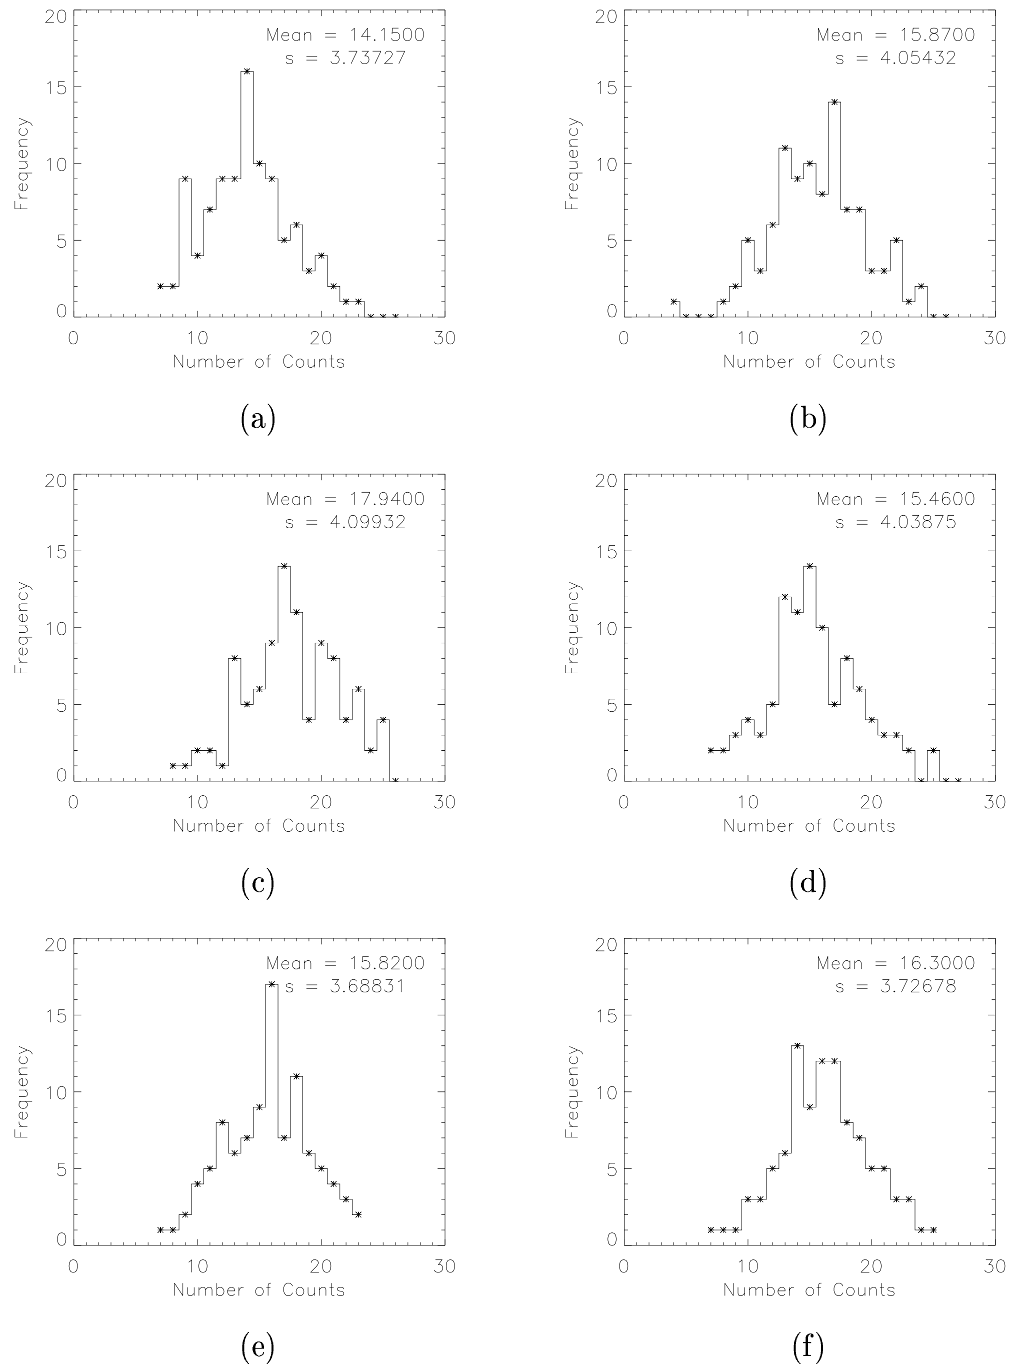

Figure 2. Histograms of six different data sets each consisting of 100 samples at 1000Hz. (a) through (f) correspond to 1 through 6 in Table 1(a).

smaller than the total possible number) the variance appears equal to the mean. To show this follows from the Poisson distribution, the expectation value of the variance

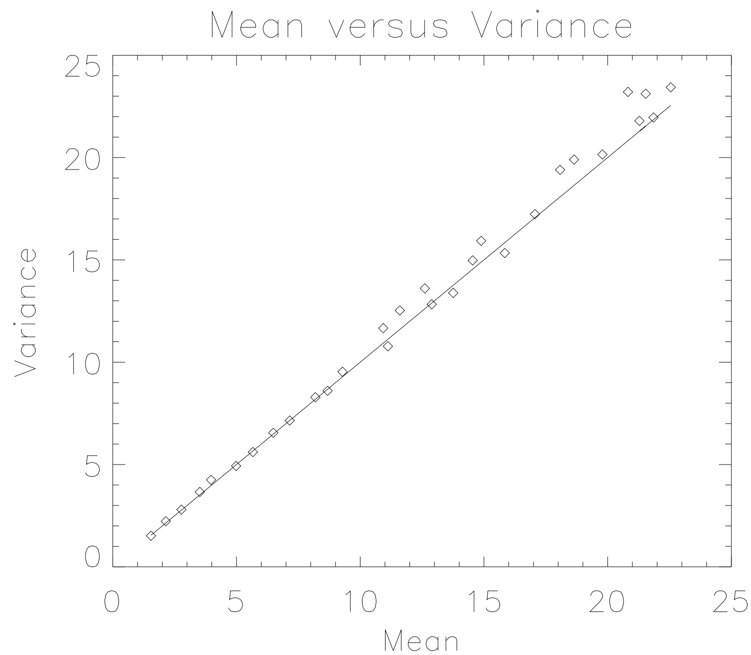

Figure 3. The variance plotted against the mean over a range of sample rate space.

can be evaluated as follows:[5]

$$\sigma^2 = \langle (x-\mu)^2 \rangle = \sum_{x=0}^{\infty} \left[(x-\mu)^2 \frac{\mu^x}{x!} e^{-\mu} \right] = \mu \tag{6}$$

Our results agree with theoretical prediction in the regime for which the Poisson distribution provides more accuracy than the binomial distribution. We do, however, begin to see increasing deviation of the sample data from the prediction with increasing mean. This could be a result of the experiments leaving the $\mu \ll n$ limit for which the Poisson distribution holds, but is more likely a result of experimental errors such as the occurrence of zero frequency hits for a given number of counts.

4 Comparison to Theoretical Models

In order to directly compare these results with theoretical expectations, we overplot the histograms with both the Poisson and Gaussian distribution curves. The Poisson distribution is given by:[6]

$$P(x,\mu) = \frac{\mu^x}{x!} e^{-\mu} \tag{7}$$

[5] Graham: "Errors and Statistics" p31 eq3.23
[6] Taylor: "Error Analysis" p246 eq11.2

And the Gaussian probability distribution is given by:[7]

$$P(x,\mu,\sigma) = \frac{\exp(\frac{1}{2}(\frac{x-\mu}{\sigma})^2)}{\sigma\sqrt{2\pi}} \tag{8}$$

Where μ and σ are given in Equations 2 and 4, respectively. We note a fundamental difference between the two distributions, that the Gaussian is symmetric about the mean, whereas the Poisson is not. We see the consequences of these symmetries in the comparison to our actual data. In order to compare a histogram to either of these theoretical curves, however, a normalization factor is needed. Since both the Poisson and Gaussian distributions give a probability, from $0-1$, multiplying by the number of samples should yield the frequency of a certain number of counts. This is valid when using a binsize of one within the histogram procedure, else the normalization factor must be additionally modified by the binsize. Alternatively, the histogram height could be normalized to one, effectively making it a probability histogram.

The four histograms in Figure 4 were normalized by the number of samples, and were obtained with the light intensity inside the PMT at a relatively low setting. This insured that we could evaluate $x!$ over the range of the histogram without it becoming unmanageably large.

Each of these four histograms represents discrete sampling of the parent distribution, while the overplotted curves represent an idealized sampling of that distribution in the limit that the number of experiments $N \to \infty$. The solid curve represents the Poisson distribution, and the dotted curve represents the Gaussian distribution. It should be noted that the Poisson distribution curve was calculated at non-integer values by using the Γ-function generalization of the factorial function.

In order to compare the effectiveness of the Poisson distribution versus the Gaussian distribution in describing the data, we also took a set of data covering the identical parameter space, but with the light source set to a higher intensity, such that the average number of counts was in the hundreds. Figure 5 again contains four histograms at different sample rates, with the Gaussian curve overplotted as a dotted line.

From both sets of histograms we conclude that the Poisson distribution provides an excellent description of the data. That is, other than the finite number of samples resulting in the piece-wise shape of the histogram, our data closely agrees with the theoretical prediction. The effectiveness of this fit appears independent of the variable sample rate. The Gaussian distribution, on the other hand, is initially a poor fit to the data when the count rate is high ($\sim\geq 2500Hz$). However, it quickly becomes a good representation of the data as the sample rate decreases (more precisely, when the mean count rate increases). Additionally, the Gaussian curve approximates the Poisson distribution precisely when there are a smaller number n of possible events, making exact evaluation of the Gaussian distribution possible.

It should be noted, however, that there are several physicall unrealistic aspects of the Gaussian. Firstly, it evaluates to a finite, though small, value for all x from negative ∞ to positive ∞. This includes the regime of negative x. In other words, the

[7] Taylor: "Error Analysis" p133 eq5.25

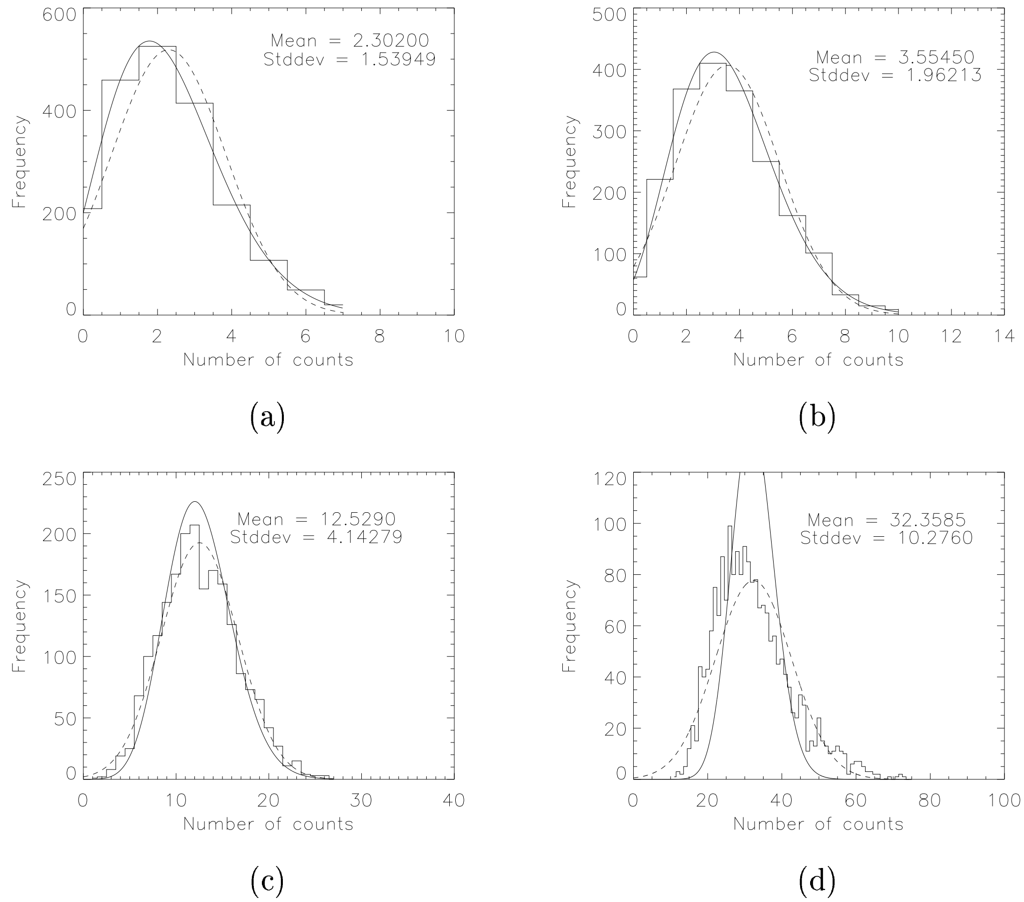

Figure 4. Four histograms of photon count data obtained with 2000 samples per set, and a rate of (a) 2500Hz, (b) 1666Hz, (c) 833Hz, and (d) 322Hz. In each case the Poisson distribution is overplotted with a solid line, and the Gaussian distribution is overplotted with a dotted line. The means of these experiments were designed to be small, between 5-25.

Gaussian distribution predicts a finite probability of counting a negative number of photons - a clearly impossible occurrence. Thus when using the Gaussian distribution as a model its specific limitations must always be considered.

5 Accuracy and Limitations

In order to quantitatively evaluate the relation between the number of samples and the accuracy with which we can measure the number of counts, the following procedure was conducted. The sample rate was fixed at 1000Hz, while the number of samples taken was varied as 2, 4, 8, 16, etc. (that is, 2^1, 2^2, 2^3, 2^4, etc.). For each "number of samples" parameter ten unique data sets were taken. Due to statistical variations, the means among each of the data sets will differ, even when all the experiment

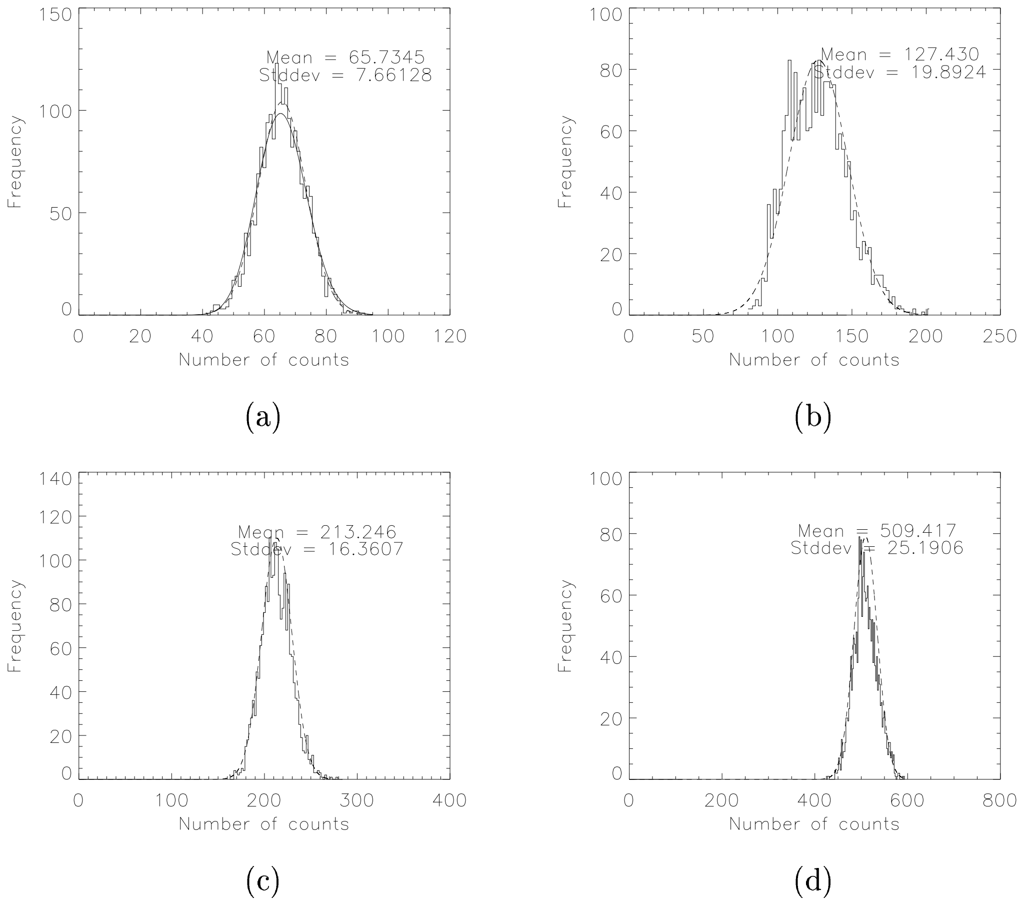

Figure 5. Four histograms of photon count data obtained with 2000 samples per set, and a rate of (a) 2500Hz, (b) 1666Hz, (c) 833Hz, and (d) 322Hz. In each case the Poisson distribution is overplotted with a solid line, and the Gaussian distribution is overplotted with a dotted line. The means of these experiments were designed to be large, in the hundreds.

parameters are the same. We thus calculate the mean of the means (MOM) for each "number of samples" using Equation 1, as well as the standard deviation of the means (SDOM), using Equation 3. Figure 6 is a plot of the MOM versus the number of samples, and Figure 7 is a plot of the SDOM versus the number of samples.

The mean of the mean (MOM) appears to asymptotically approach a constant value as the number of samples varies, while the standard deviation of the mean (SDOM) appears to decrease as $x^{\frac{1}{2}}$. Indeed, from Poisson statistics we predict that the SDOM will vary as $s \propto 1/\sqrt{N}$. This theoretical curve is plotted as the dotted line in Figure 7, and shows excellent agreement with the measured points.

Finally, given this oversampling of sorts, we can give a best estimate of the count rate, as well as the accuracy of that value. Specifically, the best estimate comes from our highest number of samples - 4096. The mean of those 10 samples was 6.63, and the standard deviation 0.32, so our best estimate is 6.63 ± 0.32.

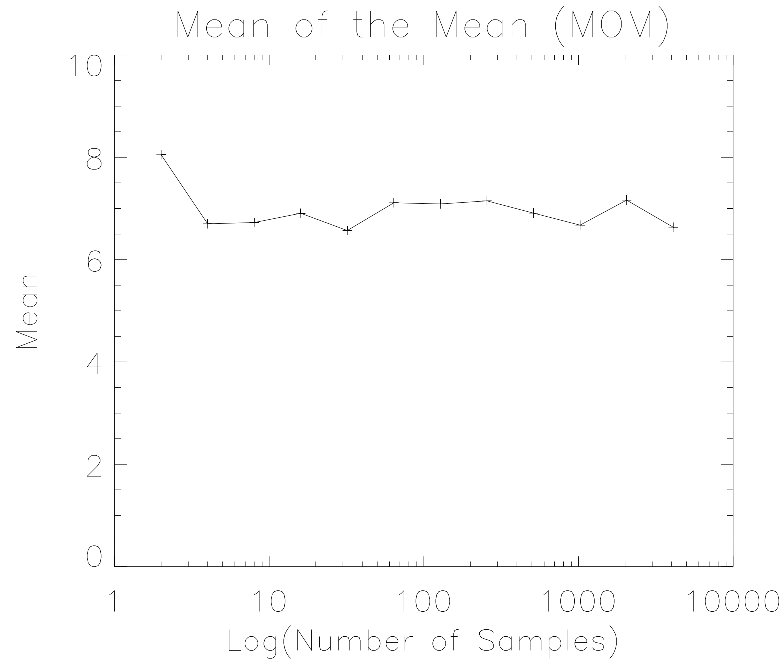

Figure 6. The mean of the means (MOM) versus the number of samples. The number of samples ranges from 2 to 4096, with a logarithmic x-axis.

6 Conclusion

Throughout this paper we have noticed deviation between data sets resulting from approximation of a parent population obtained through finite sampling. However, we cannot disregard non-statistical sources of error. These include sources of interference from outside the PMT affecting the photon counts. This was perhaps the largest source of error, noticed several times during the experiment, whereby the movement or location of external bodies (i.e. people) around the PMT caused extreme variation on the frequency of photon counts. Additionally, errors within the equipment were present, such as a lack of time precision in the counting board, as well as time-variability of the light source.

This "time-variability" of the light source, as we have seen, is not a flaw in the mechanical design, but rather an intrinsic property of the quantum mechanism at the heart of photon emission. Clearly it would be impossible to construct a light source which would show zero variations in the count rate. The number of atoms in the light source are so numerous that the regime in which the Poisson distribution holds can never realistically be escaped. That is, there are such a large number of possible events n, each of which has such a minute possibility p of occurring, that it is impossible to predict exactly the number of photon detections which will occur. Rather, we can only quote a probability of counting a certain number of photons - the actual number could only be obtained within the theoretical regime of the parent population. As a result, the statistics underlying photon emission are a critical factor whenever

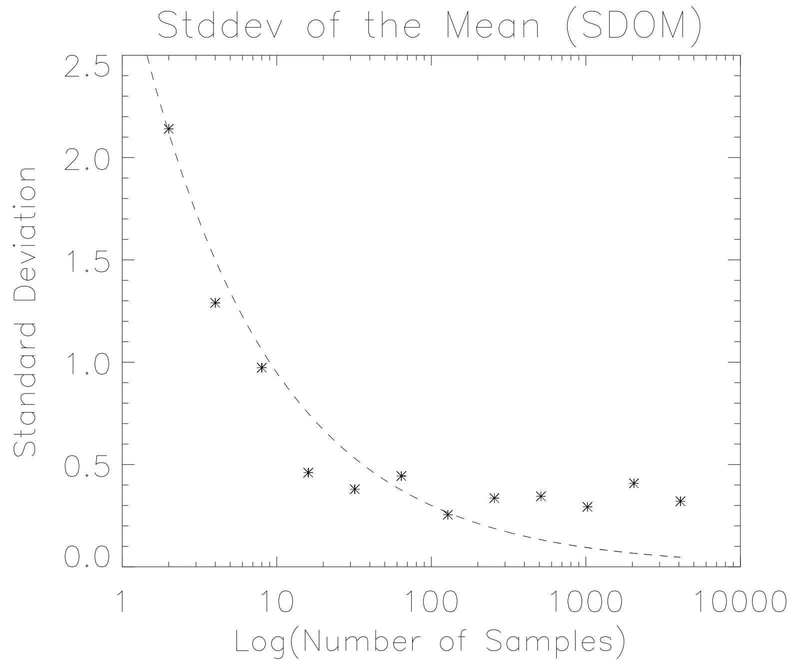

Figure 7. The standard deviation of the means (SDOM) versus the number of samples. The number of samples is as in Figure 5. Overplotted is a curve of $1/\sqrt{N}$, the predicted falloff of the SDOM.

conducting experiments based on the detection of light. Since astrophysical observations are almost entirely based on the detection of electromagnetic waves, including the measurement of photons with CCDs, the statistical consequences explored herein have undeniable importance for the accuracy and variation of such experiments.