Notes on bioburden distribution metrics: The log-normal distribution

Mark Bailey, March 2010

Introduction

The shape of distributions of bioburden measurements on devices is usually treated in a very simple manner. The common use however of the mean and standard deviation has the inbuilt assumption that the measurements must follow normal or Gaussian statistics, and this forces the results to be treated in a manner which often causes problems for the microbiologists:

- How are warning and action levels to be set if the statistics do not follow the normal distribution? (I.e. if the mean and standard deviation are actually meaningless.)
- What is the justification for the values set for warning and action levels?

'Spikes' are the major cause of concern as a high 'spike' can lead to a high estimate of the mean, which is not representative of the actual population. In setting warning and action limits in bioburden monitoring, for example, the use of a scientifically-based method of bioburden determination including an understanding of the origin of spikes and of the correct distribution followed, should lead to better justified warning and action levels. This in turn would avoid unnecessary work (for eg. a repeat of a dose audit) arising from false non-conformances.

In sterilization dose establishment, the underlying assumptions in the measurement of bioburden are not fully explained in BS EN ISO 11137-1, BS EN ISO 11137-2 or BS EN ISO 11737-1 although we should note that dose establishment methods, based on the ISO 11137-2 Method 1 microbial population of standard distribution of radiation resistance (SDR), are intended as a severe and conservative test of the bioburden in terms of microbial numbers and radiation resistance and actually include distributions containing "spikes". These methods were mathematically modelled and validated using actual distributions, so no assumptions were made about the type of statistical model that should be applied. This means that methods in ISO 11137-2 involving the SDR should be robust against the methods used for establishing the bioburden. It may however be possible that, in future revisions of the dose establishment methods of ISO 11137, a simplification of the procedures may follow from a better knowledge of the statistical distributions followed by the bioburden.

Bioburden examples

Examples of real data have been reviewed to identify the statistical distribution followed by the bioburden values. The most interesting examples were those where (a) there was a high bioburden (expressed as colony forming units, CFU) so that the statistical distribution could be identified with reasonable certainty, and (b) the bioburden did not change over the time. Analysis of measurements made on devices with low bioburden would not enable a conclusion to be drawn with great confidence.

An example, comprising a bioburden which changed with time, did not lend itself to a simple analysis; for example, in *Figure 1* below bioburden control has clearly improved with time from high levels, with some apparent spikes, to lower values later on as control improved. The time-averaged bioburden in this type of example cannot be expressed as one population, and the improvement in control implies that a change in the warning and action levels might be necessary.

A high-bioburden example (*Figure 2*) with roughly unchanging characteristics within the time of the measurements has been selected for statistical analysis. This example was also selected so that artefacts arising from application of a large recovery factor on low bioburden counts, or from the use of a portion (SIP) of the product, were avoided.

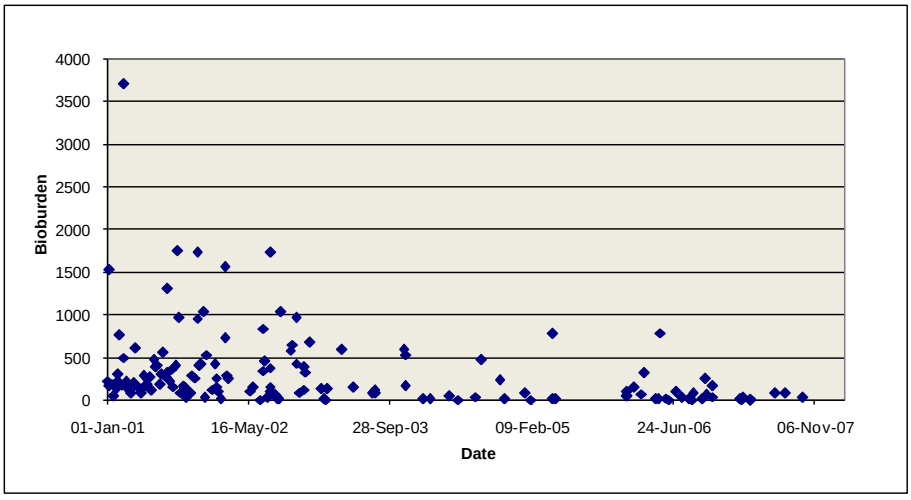

Figure 1: *Example of product with high bioburden. Note a significant improvement during Autumn 2002, a gap in production through much of 2005, and further improvement in control from mid-2006.*

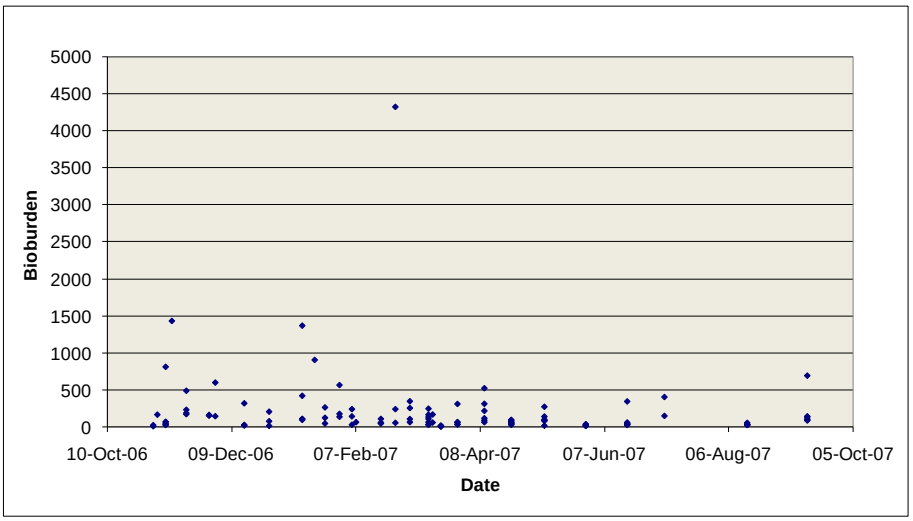

Figure 2: *Example of product where bioburden was high but unchanged over the time. Note the 'spike' of 4300 cfu during March 2007.*

Taking a simple average will clearly result in a distortion upward, if it includes events like the March 2007 spike in *Figure 2*. However, spikes are a normal, expected feature of bioburden measurements. Clearly therefore, the assumption that normal statistics may be applied (where measurement may be characterised by a simple mean and standard deviation) is false.

The mean of the example in *Figure 2* is about 190, with a standard deviation of about 450 when the spike is included. These mean and standard deviation values are clearly nonsense and demonstrate that setting warning and action limits using the assumption of normal statistics will not work. The standard deviation here is meaningless; if this distribution followed Poisson statistics, the standard deviation would be about 14.

So, how might sensible warning and action limits be derived?

There are now many papers and references published which describe statistical distributions appropriate for use in measurements of natural phenomena including microbiological populations (for example; Limpert,Stahel & Abbt, 2001 and Sutton, 2006). The most useful distribution, which also has some physical justification, is the log-normal distribution (also known as multiplicative normal

distribution). Biological populations predominantly show exponential growth, so a distribution reflecting the statistics of exponential growth is what would be expected: The log-normal distribution exactly fits this situation, as the standard deviation that arises is multiplicative, not additive. Instead of adding or subtracting a quantity from a mean value, the width of the distribution is expressed by taking the median and multiplying or dividing by some quantity.

The Log-Normal Distribution

A variable such as bioburden, is log-normally distributed if the natural logarithm of the bioburden is itself normally distributed. This turns out to be true in this case. *Figures 3* and *4* show the same data as in *Figures 1* and *2* respectively, but each with a logarithmic bioburden axis. It is clear that the scatter is much more uniform than that seen for a linear bioburden axis; the data are not clumped near the bottom of the graph.

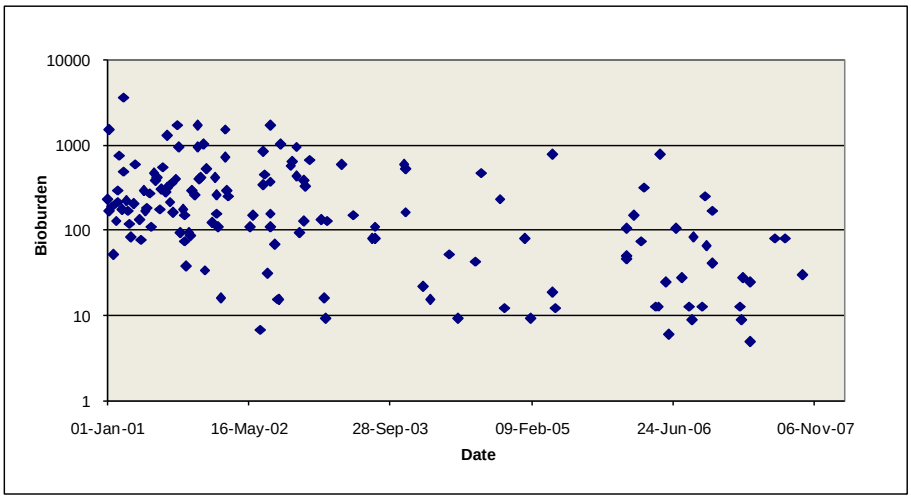

Figure 3: *As figure 1, with a logarithmic bioburden axis. The data are more evenly spread, and the reduction in bioburden after 2003 is clear.*

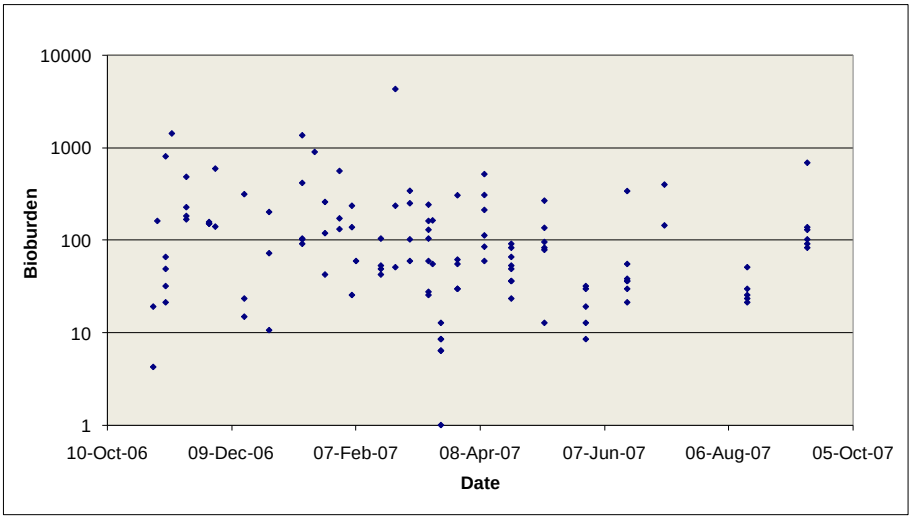

Figure 4: *As figure 2, with a logarithmic bioburden axis. This again shows that the data are very evenly spread, strongly supporting the suggestion that ln(bioburden) is the correct metric to use.*

The data in Figure 4 may be recast as a histogram of bioburden frequency and these are plotted in *Figure 5*. Bioburden frequency is the frequency of occurrence of bioburden values between given limits N and $N+\Delta N$, where ΔN is the width of a histogram 'bin' – here, 20 CFU. This represents the 'traditional'

view of the data, strongly peaked towards zero and with occasional 'spike' values at high bioburden. Analysing this form of distribution is clearly difficult.

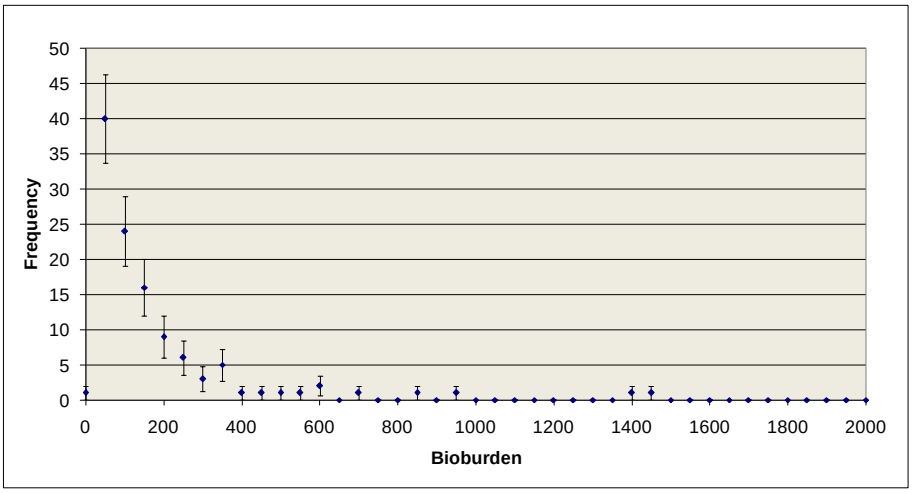

Figure 5: *Histogram of bioburden frequency using the same data as used in Figures 2 and 4.*

If, instead, we plot the histogram for the natural logarithm of the bioburden, expressed as ln(bioburden) (see *Figure 6*), we can now easily fit a Gaussian curve, with a traditional mean and standard deviation. Here, the mean derived from the best-fit Gaussian curve to this histogram is 4.64 and sigma is 1.24.

Mean and Standard Deviation Derived From Log-Normal Distribution

Most product batches will only have a limited number of measurements of the bioburden. Fitting a normal distribution curve directly to a histogram will therefore be inefficient, and the parameters of the resultant curve will be strongly dependent on details of the size and width of the histogram bins.

It is simpler, and indeed more accurate, to calculate the mean and standard deviation of the values of ln(bioburden) directly from tabulated data (see Appendix). In this case, the mean is 4.39, and the standard deviation is 1.26. *Figure 6* illustrates curves produced by both methods of calculation, and the close agreement between the fitted curve through the histogram and the curve derived from unbiased estimates of the mean and standard deviation strongly supports the hypothesis that the bioburden is distributed log-normally.

The mean of the Gaussian distribution here, corresponds to the *median* of the log-normal distribution. The median is the value for which there is an equal probability of data being higher or lower, and for the Gaussian distribution is identical to the mean. The standard deviation corresponds to a value by which the median is *multiplied*, to give values above or below the median corresponding to the probability of the bioburden falling to $1/\sqrt{e}$ of the value at the median. Here, e takes its usual meaning of the base of natural logarithms, $e \approx 2.718281828$. It is this multiplication which ties the distribution to the exponential nature of the microbial growth: for the normal or Gaussian distribution, the standard deviation is additive. For the log-normal distribution, it is multiplicative.

For example, consider a normal distribution with a mean located at $x = 4$ and a standard deviation of 2. The distribution would fall to $1/\sqrt{e}$ of the peak value (at $x = 4$), at $x = 2$ and at $x = 6$. If this distribution were describing a log-normal distribution so that $x = \ln(y)$, the median for the distribution would be located at $y = e^4$ (which is approximately 55), and the curve would fall to $1/\sqrt{e}$ of that peak value at $y = e^2$ (about 7) and $y = e^6$ (about 400).

It now becomes clear that the 'spikes' of high bioburden are simply part of the distribution. The value at ln(bioburden) ~ 8.5, corresponds to the highest measured bioburden of 4300. Thus by using this distribution we can remove the necessity for special treatment of high bioburden spikes. Also, non-

conformances arising from artificially low warning and action limits on bioburdens, derived from a false assumption of normal statistics, can be avoided.

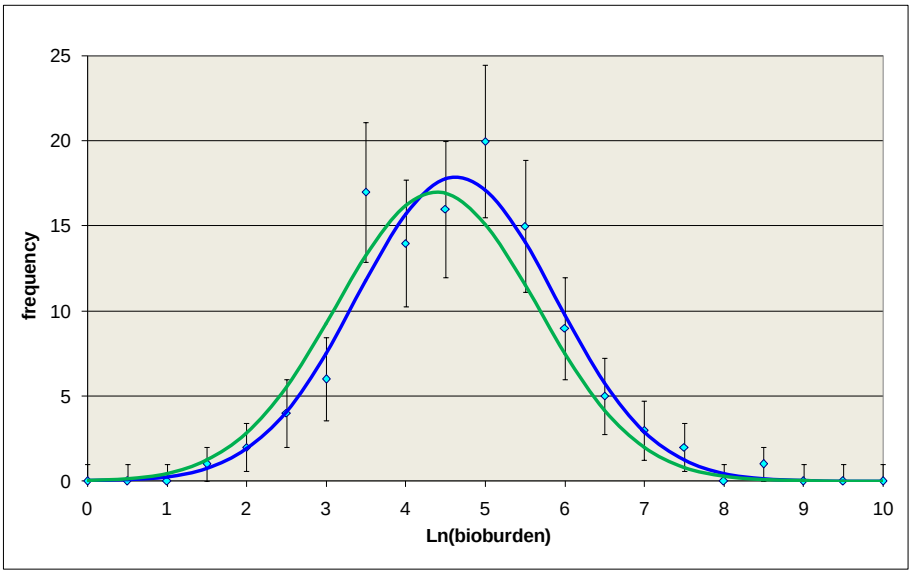

Figure 6: *Bioburden plotted as ln(bioburden) with superimposed best-fit Gaussian curve (blue), and Gaussian based on unbiased estimates of the mean and standard deviation (green), using the same data as used in Figures 2, 4 and 5.*

Note that the value in any given histogram bin, is however normally distributed. The values here have uncertainties, which vary as the square root of the value. Clearly, if more data are used to plot the histogram, then a better-determined fitted Gaussian will be obtained. This will lead then, to better values (with lower uncertainties) for the median μ^* (as distinct from the mean, usually expressed as μ) and for the "multiplicative standard deviation" σ^* (again, as distinct from the true standard deviation, usually expressed as σ). Here then, μ^* is approximately 80 CFU (from the unbiased estimate of the mean, 4.39) and σ^* is about 3.53 (corresponding to $\sigma = 1.26$). Here, $\mu^* = \exp(\mu)$ and $\sigma^* = \exp(\sigma)$.

Similarly, for product with low bioburden it is possible to demonstrate that the logarithms of the data may be approximated by a Gaussian distribution, although care should be taken if the number of CFU on a particular device is zero (as ln(0) is not defined). *Figure 7* shows similar plots for two products with different bioburdens.

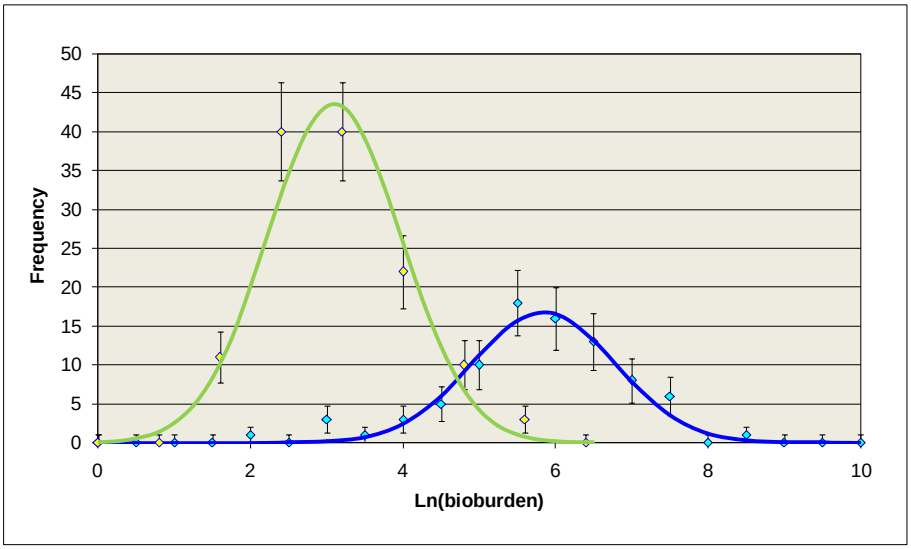

Figure 7: *Different products, one relatively high in bioburden (blue), the other very low (yellow/green). Both demonstrate that Gaussian fits to the ln(bioburden) gives excellent agreement.*

Care should be taken in the analysis of data derived from low bioburden values where in order to avoid histogram bins with zero counts (as may happen, for example, when high multipliers on the recovery or SIP are used) the bins should be broadened. In this case it may be that the lowest non-zero bioburden will not be 1, but the reciprocal of the SIP fraction and this may also distort the analysis.

Conclusions and recommendations

Use of the log-normal distribution for the analysis of bioburden data more closely matches the distributions found in nature and therefore will give statistically robust results. In fact, on page 20 of ISO 11737-1:2006, we find

> *For product with an overall average bioburden exceeding 10 CFU per unit, the distribution could again tend to be skewed with a long tail to the right. In practice, such data can be best approximated by a lognormal distribution. In practice, taking logarithms of the individual counts will make the data approximate to a normal distribution.*

The suggestion here is that *all* levels of bioburden can in fact be expressed using the log-normal distribution: In the limit, for measurements of quantities with a high mean and low standard deviation, the log-normal distribution is in fact indistinguishable from the normal distribution. However for bioburdens, the log-normal distribution is much more appropriate. Generating warning and action levels based on the statistics of the ln(bioburden) will avoid special treatment of 'spikes', which have been demonstrated to be simply an expected feature of the behaviour of bioburden statistical distributions. So, warning and actions levels on the natural logarithm of the bioburden can be set.

In the main example (Figure 6), with mean 4.39 and standard deviation 1.26, bioburden limits expressed as ln(bioburden) might be as follows:

warning level = (2.5 x 1.26) + 4.39 = 7.5
action level = (3.5 x 1.26) + 4.39 = 8.8

corresponding to 2.5σ and 3.5σ respectively. 2.5σ and 3.5σ are commonly used in process control and reflect the stochastic behaviour (i.e. reflecting the statistical nature) of the bioburden measurements, heavily dependent on the uncertainties in the derived values of the mean and standard deviation on a limited set of measurements. Interestingly, in the example above, the highest measurement of 4300 CFU is actually below the suggested action level as ln(4300) is just 8.4.

The point here is that these occasional 'spikes' will be seen; they are an expected part of the statistics of microbial populations. It may still be considered important to carry out an investigation to ensure that there is no other contributing factor to the high count. If however such an investigation finds nothing, then the spike can be safely and validly considered to be part of the statistical distribution.

In addition, a trending analysis will show whether the microbiological population is in fact decreasing or increasing with time. *Figure 3* for example clearly shows a decreasing trend, reflecting improvements in the bioburden control.

Ultimately, use of this approach may also lead to improvements in the methods of dose establishment given in ISO 11137-2, or in their successors. However as noted in the Introduction, the current methods are conservative; product safety is assured using the Methods described in ISO 11137-2 and the statistical approach here suggested for bioburden measurement will not change this.

Microbiologists working in the sterile medical device industry, may see a significant advantage of this approach in the analysis of data over time in terms of maximising on trend analysis.

Suggested useful references

Limpert E Stahel W A and Abbt M: *Log-normal Distributions across the Sciences: Keys and Clues* Bioscience **51** (2001) 341-352

Sutton S: *Counting colonies* PMF Newsletter September 2006

BS EN ISO 11137-1:2006: *Sterilization of health care products — Radiation — Part 1: Requirements for development, validation and routine control of a sterilization process for medical devices*

BS EN ISO 11137-2:2007: *Sterilization of health care products — Radiation — Part 2: Establishing the sterilization dose*

BS EN ISO 11737-1:2006: *Sterilization of medical devices — Microbiological methods — Part 1: Determination of a population of microorganisms on products*

Appendix: How to use the log-normal distribution

In order to carry out an analysis of bioburden data using the log-normal distribution, the logarithms of the bioburden data are tabulated. *Table 1* shows bioburden data together with natural logarithms of the bioburden for three different products. These are real data representing the final values after use of recovery factors etc. but the details of their application are ignored.

30 is a sufficient number of determinations to obtain a good estimate of the median, μ^*, and of the multiplicative standard deviation, σ^*. ISO 11137 recommends 30 for the initial determinations for Method 1, with subsequent determinations for dose audit on 10 units, and there is no need to change these recommendations. In fact, since the mean and standard deviation for the ln(bioburden) is much better defined than for bioburden itself, any conclusions drawn from the statistics are much safer.

	Product 1		Product 2		Product 3	
	No. CFU	Ln(No. CFU)	No. CFU	Ln(No. CFU)	No. CFU	Ln(No. CFU)
1	12.6	2.53	50	3.91	268.1	5.59
2	50	3.91	10	2.30	83.0	4.42
3	47	3.85	120	4.79	136.2	4.91
4	107	4.67	40	3.69	78.7	4.37
5	151	5.02	10	2.30	95.7	4.56
6	76	4.33	30	3.40	31.9	3.46
7	326	5.79	20	3.00	8.5	2.14
8	13	2.56	0		12.8	2.55
9	13	2.56	10	2.30	29.8	3.39
10	793	6.68	50	3.91	19.1	2.95
11	25	3.22	30	3.40	38.3	3.65
12	6	1.79	140	4.94	29.8	3.39
13	107	4.67	60	4.09	36.2	3.59
14	28	3.33	50	3.91	55.3	4.01
15	13	2.56	60	4.09	36.2	3.59
16	9	2.20	50	3.91	340.4	5.83
17	84	4.43	60	4.09	21.3	3.06
18	66	4.19	50	3.91	400.0	5.99
19	258	5.55	30	3.40	144.7	4.97
20	13	2.56	30	3.40	23.4	3.15
21	41	3.71	20	3.00	21.3	3.06
22	173	5.15	20	3.00	29.8	3.39
23	9	2.20	10	2.30	51.1	3.93
24	13	2.56	10	2.30	25.5	3.24
25	28	3.33	10	2.30	689.4	6.54
26	5	1.61	10	2.30	129.8	4.87
27	25	3.22	10	2.30	91.5	4.52
28	83	4.42	10	2.30	102.1	4.63
29	83	4.42	10	2.30	83.0	4.42
30	30	3.40	10	2.30	138.3	4.93

Table 1: *Numbers of colony-forming units on devices, with ln(no. CFU), for three devices. The bioburdens are quoted to the level of precision supplied, and reflects the use of SIPs for some devices and recovery factors in others. Note that as ln(0) is not defined, in this case the zero count for Product 2 is not included in the analysis. (In cases where zero counts occur frequently due to a very low bioburden, some averaging should be carried out in order that the problems at zero are avoided.)*

The means and standard deviations of the ln(bioburden) data may then be calculated. Confirmation that the data follow the log-normal distributions can be obtained by plotting the histograms of the data in the same way as *Figures 6* and *7* and then demonstrating that these may be approximated by the normal distribution . As mentioned above, the details of the fit may be strongly affected by histogram bin width details for low bioburden products.

Here, for example, the median of product 1 is $e^{3.68} = 40$, and the multiplicative standard deviation (or multiplicative sigma) is $e^{1.26} = 3.5$.

Product	ln(bioburden) Mean	SD	Bioburden median	mult. sigma
1	3.68	1.26	40	3.5
2	3.21	0.84	25	2.3
3	4.10	1.06	61	2.9

Table 2: *Mean and standard deviation for the ln(bioburden) data, and median and multiplicative standard deviations (multiplicative sigma) for the bioburdens themselves*

The warning and action levels are then established, based on these results.

Product	Warning level ln(bioburden)	Bioburden	Action level ln(bioburden)	bioburden
1	6.8	920	8.1	3230
2	5.3	200	6.2	470
3	6.7	850	7.8	2450

Table 3: *Warning and action levels based on 2.5σ and 3.5σ. Note that the warning and action levels are rounded to the nearest ten; with higher bioburden figures, rounding to the nearest 100 or even 1000 may be appropriate.*