Archeologia e Calcolatori
21, 2010, 185-200

STATISTICAL ANALYSIS OF SOME LOOMWEIGHTS FROM POMPEII: A POSTSCRIPT

1. Introduction

In «Archaeologia e Calcolatori» no. 19 we presented an analysis based on the dimensions of 95 complete loomweights found during excavations in Insula VI.1, Pompeii (Baxter, Cool 2008, henceforth BC08). This was published in the belief that these were the complete set of complete loomweights from the site, but during post-excavation work in the 2009 season a further 42 complete loomweights from other areas in the insula were discovered in boxes that had been placed in another part of the store. These have allowed us to address questions of sample size and to update and extend the analyses given in BC08. Some recently published results relating to experiments with warp-weighted looms (Mårtensson, Nosch, Strand *et al.* 2009) have also made it possible to explore what the patterns uncovered in the data may be implying.

An interesting, and recurring, question in the statistical analysis of archaeological data is how large sample sizes should be for conclusions based on an analysis to be sound. This begs a lot of questions that do not admit easy answers. To go beyond a purely descriptive analysis the "sample" available for analysis has to be conceived of as in some sense "representative" of a larger population (or if the sample is biased in any way this needs to be understood). In the present case the sample is initially regarded as representative of loomweights used in Insula VI.1. Whether or not it is typical of other areas of Pompeii is something that could, in principle, be checked if sufficient loomweights were recovered and recorded elsewhere.

A general problem is that a "sample" available for analysis, particularly from a completely excavated site, will often be the only material that will ever be available. In BC08 a tentative attempt was made to draw archaeological conclusions from the statistical patterns observed in the data, raising the question of whether the sample size is large enough (or the patterns stable enough) to justify this. The additional loomweights can be regarded, initially, as sampled from the same population as those discussed in BC08. This provides a perhaps rare opportunity to investigate the effect on conclusions of increasing the sample size. The major finding of BC08 was that the distribution of the weights of loomweights was bimodal, inviting a functional interpretation if it is judged that this reflects a population characteristic. The new data helps strengthen this finding, justifying the detailed attempt at interpretation in the third part of the present paper.

As shorthand, the data that was used in BC08 will be called the “old” data; the additional loomweights will be referred to as the “new” data; and the total set as the “combined” data. The terms refer to the date of cataloguing and have no chronological significance. The analyses to follow both combine and contrast the “old” and “new” data sets to see what, if any, changes result. The most important findings of BC08 are unaffected, but there are details raised by the discovery of the “new” weights that merit comment.

Loomweights are mostly made of fired clay. Three of the new loomweights were made of gritty sandstone, and exhibited characteristics untypical of the fired clay loomweights. These are excluded from further analysis.

BC08 may be referred to for details of the protocols followed.

2. Statistical Analysis

2.1 *Dimensional analysis*

Fig. 1 shows a kernel density estimate (KDE) of the weights using all the data. There are two main modes in the data with a shallow trail of unusually large weights over 400 g indicated by the small “bumps” to the right. Fig. 2 shows a plot of height against weight using all the data. Both figures look similar to comparable figures in BC08.

There is one very small example weighing 15 g. It is unique in this assemblage, but weights of similar size and shape are on display in the Museo Archeologico Georges Vallet at Piano di Sorrento. They come from the Punta della Campanella at the tip of the Sorrentine Peninsula which is the site of a sanctuary devoted to the goddess Minerva (Caputo 2004, 84-87). It is very likely that these were votive miniatures rather than practical loomweights. They would have been appropriate for her, given she was associated with handicrafts and weaving. As will be discussed later, such light loomweights would be totally impractical for serious weaving.

In BC08 a “modified” data set was defined and used for some analyses. The same procedure is followed here, omitting the small loomweight and those above 400 g. Additionally Fig. 2 shows a “new” loomweight at about 300 g that plots below the main body of the data and this will also be treated as unusual and omitted from the modified data set, which thus consists of 125 loomweights.

Figs. 3 and 4, based on the modified data set, contrast KDEs for the “old” and “new” data, for weight and height respectively. For weight, the left modes for both data sets are very similar; the right mode is greater for the “new” data than for the “old”. The modes are similar for the contrasted height data, except that there is the suggestion of a third smaller mode to the right. If the data sets are combined (the figure is not shown) the evidence

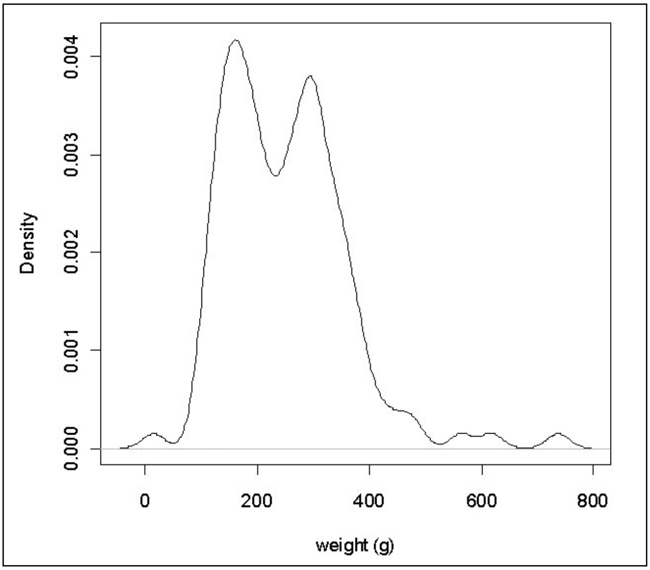

Fig. 1 – A KDE using all the weight data; the main modes are at 162 g and 293 g. Compare with Fig. 2 in BC08.

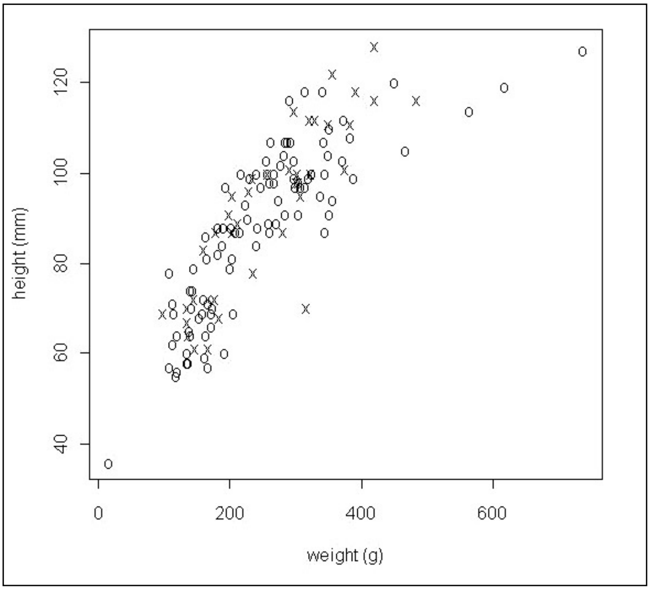

Fig. 2 – Height against weight labelled by “old” (o) or “new” (x).

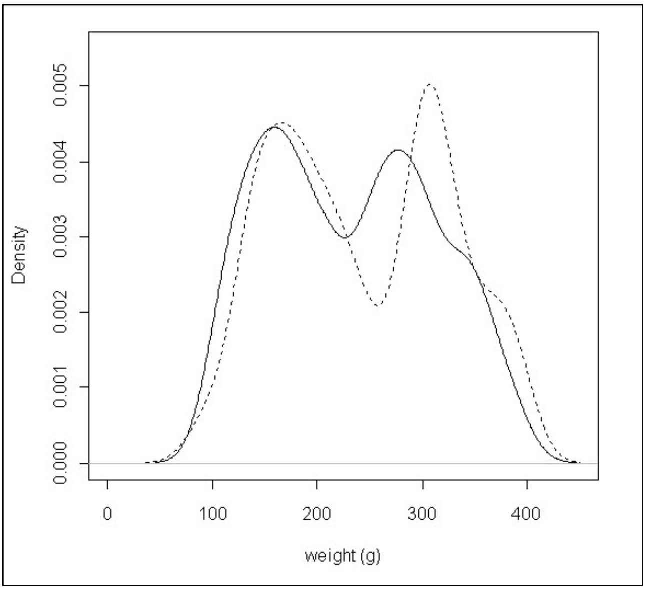

Fig. 3 – KDEs for the "old" (solid line) and "new" data (dashed line) for weight, from the modified data set. Modes are at 159 and 274 for the old data; and 162 and 301 for the new data. Sample sizes are 89 for the "old" data and 36 for the "new" data.

for a third mode disappears. A two-dimensional KDE, for the weight and height data in the modified data set, was visually difficult to distinguish from the similar Fig. 8 in BC08 and leads to similar conclusions about bivariate bimodality.

In BC08 loomweights were categorised as "large" or "small" according to whether they weighed 239 g or more, or 230 g or less. There are no loomweights in the "old" data set between 230 g and 239 g, and the anti-mode in the KDE was about 235 g. For the "new" data there are two weights of 233 g and 235 g respectively, the next largest after these two weighing 256 g. This suggests that the distinction between small and large weights, suggested by the bimodality of the KDEs, remains valid, with the minor redefinition of "small" as 235 g or less.

Although the distribution of "old" and "new" weights between the categories "small" and "large" is virtually identical, the larger "new" weights tend to be slightly heavier accounting for the shift in the KDE for "new" weights to the right in Fig. 3. For the modified "old" data set the mean weights for the small and large loomweights are 175 g and 302g; for the "new" data the means are 175 g and 322 g.

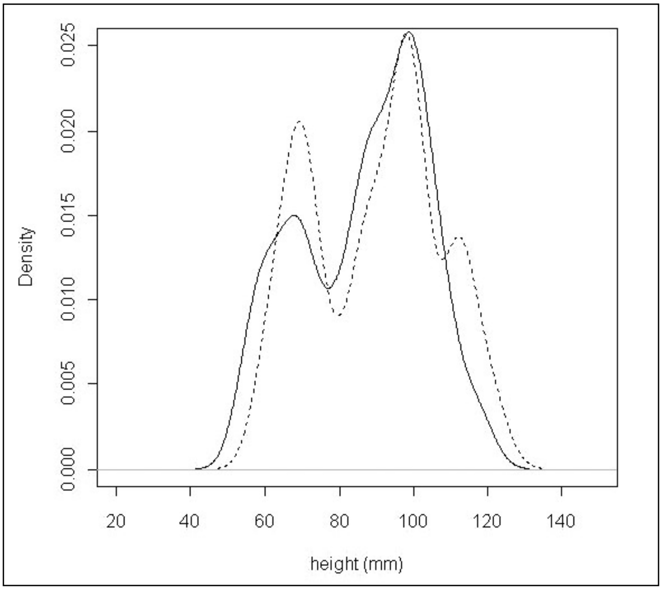

Fig. 4 – As Fig. 3, but for height. The modes are at 68 and 98 for the "old" and 69, 99 and 112 for the "new" data.

2.2 *Volume analysis*

The volume, v, of an idealised truncated pyramidal loomweight with rectangular base and top is

$$v = h(b_1t_2 + t_1b_2 + 2b_1t2 + 2\ t_1t2)/6$$

where (b_1, b_2) and (t_1, t_2) are the dimensions of the sides of the base and top and h is the height. It is to be expected that volume is related to weight (w) with some variation as there are departures from the ideal, and the density of the loomweights may vary.

Fig. 5 shows a KDE for the volume data for the modified data set with a subjectively chosen bandwidth that perhaps undersmooths, but suggests the possibility of a third mode, as was the case with height in Fig. 4 for the "new" data.

Fig. 6 shows a bivariate contour plot, based on a two-dimensional KDE, for the volume and weight data from the modified data set. One feature of this is the suggestion of tri-modality. The dotted line shows a division between "small" and "large" loomweights, suggested by Fig. 7, which is slightly different from that based on weight alone.

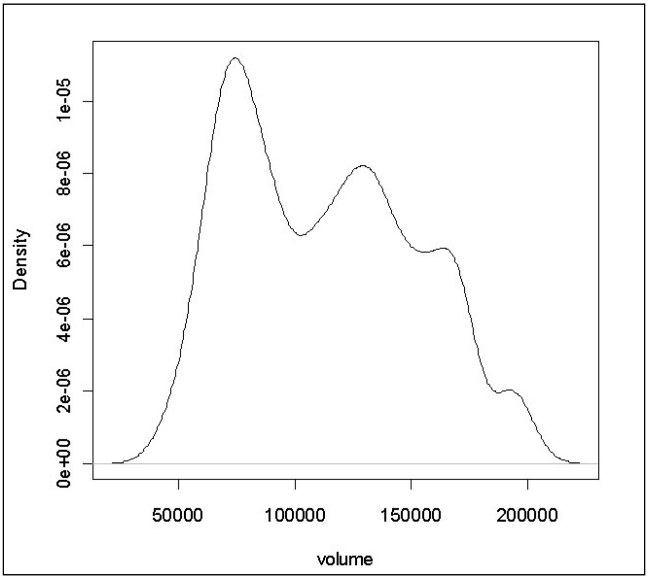

Fig. 5 – KDE for the volume data, using a subjectively determined bandwidth of 8000.

The (Pearson) correlation coefficient for v and w of 0.89 is, as expected, strong. There is some interest in predicting weight from volume. The regression equation for the modified data is

$$\hat{w} = 29.688 + 0.001821v$$

with an R^2 value of 78%. The regression line is shown in the plot of weight against volume in Fig. 7, where the horizontal dashed line shows the classification – "large" above the line, and "small" below – suggested by an analysis of weight alone. The dotted line proposes an alternative categorisation suggested by visual inspection of the plot. Eight or nine loomweights would be differently classified using this rule, but most of these are very close to the boundaries suggested, so that the two rules give practically similar results.

The interest in predicting weight from volume arises from the possibility of estimating weights for incomplete loomweights (or those with accretions), where the dimensions necessary to calculate volume survive. This, in turn, would increase the number of securely dated loomweights with known or estimated weights. Previous investigation has shown that other variables available for weight prediction, including base area, do not perform well. The requirement to measure the loomweights "in the field" (they are not al-

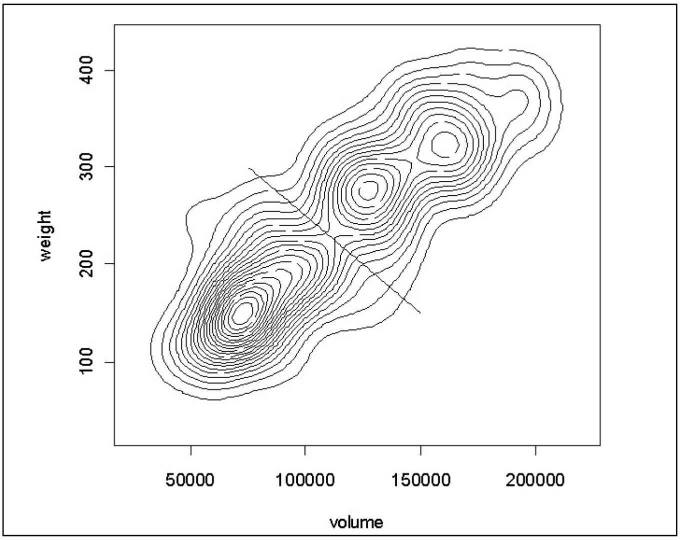

Fig. 6 – A bivariate contour plot based on a two-dimensional kernel density estimate for the volume and weight data. The kde2 function from the MASS library in R was used, with bandwidths of 40000 and 100 suggested by the Sheather-Jones estimates for each variable separately. The dotted line shows the division into large and small suggested by Fig. 7.

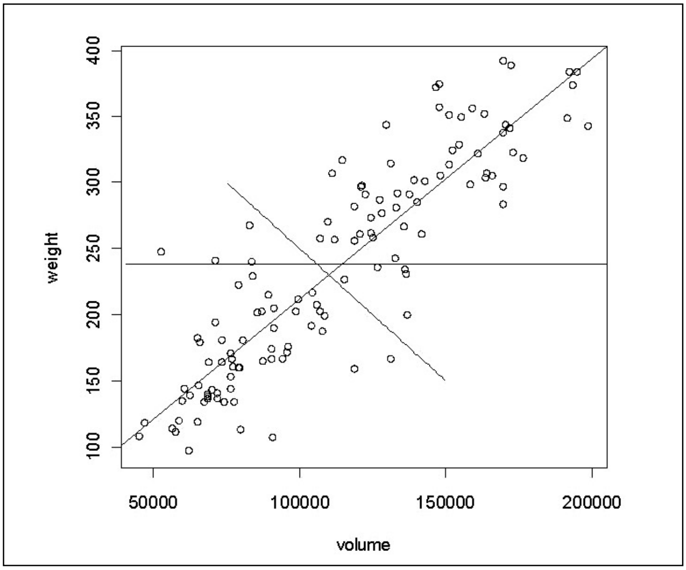

Fig. 7 – Plot of estimated volume against height. The solid line is the regression line that can be used to predict weight from height; the horizontal dashed line shows the division into "large" and "small" weights suggested by an analysis of weight only; the dotted line is subjectively determined from this figure to suggest an alternative division into "large" and "small".

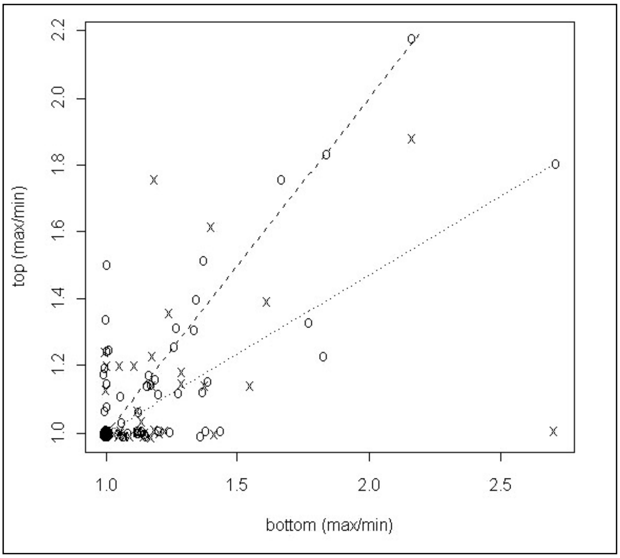

Fig. 8 – A plot of the maximum to minimum ratio of the tops of the loomweights against a similar ratio for the bottom of the loomweights; "old" data (o) and "new" data (x). See the text for further explanation. Compare BC08 Fig. 10.

lowed off site) over a "season" of four weeks, with many competing demands has meant that incomplete weights have not yet been subjected to the same measurement protocols as complete weights, but the analyses here suggest that it may be worthwhile to do so.

2.3 *Shape analysis*

Fig. 8 is similar to Fig. 10 in BC08 and shows a plot of the maximum to minimum ratios for the tops against those for the bottoms. The plot has been jittered – that is each point is displaced by a small random amount – so that the "blob" in the (1,1) position corresponds to loomweights with square tops and bottoms.

Cases vertically above the (1,1) co-ordinate have square bottoms and rectangular tops, with the reverse the situation for cases horizontally along from the (1,1) co-ordinate. The dashed line corresponds to the ideal where loomweights with rectangular tops and bottoms would plot if the ratios were the same; the dotted line corresponds to loomweights where the ratio for the bottom is 1.5 times that for the top.

In BC08 (p. 248) it was judged that most of the "old" data fell into or close to the categories implied above, prompting the development of a shape-based typology with five categories. Several "new" loomweights do not fit readily into the categories defined (except as "other" which is category 0 in BC08). The typology will be retained for present purposes but must be treated as tentative.

BC08, using the "old" data, presented several analyses suggesting among other things that Type 4 loomweights (square top, rectangular bottom) were almost all large, and Type 1 (square top, square bottom) and Type 4 loomweights were disproportionately represented among the heavier weights from about 300 g, the proportion increasing as the definition of "heavy" increases.

Using all the data there are 85 weights with square tops, 63% of the total. Defining "heavy" as greater than 250 g gives 65 "heavy" weights of which 50 have a square top, or 70%. As the definition of "heavy" is increased in steps of 50 g, for 300 g the corresponding results are 41, 34 and 83%; for 350 g the results are 19, 17 and 89%; and for 400 g the results are 7, 8 and 88% (the sample size being rather small). This reinforces the observation in BC08 that the larger weights tend to have square tops.

3. Comparison to the experimental work

Since BC08 was published a new avenue of investigation has been made possible by the recent publication of quantified experimental work which explored how warp-weighted looms worked (Mårtensson *et al.* 2009, henceforth M09). The experiments were based on weaving a simple tabby cloth and examined what the optimal loom set-up would be for a given loomweight. They showed how the properties of the cloth altered depending on the tension the threads were being placed under, and that the best result was obtained when the total thickness of the loomweights used was equivalent to the width of the piece of cloth being woven. Practicalities of weaving relating to the desire not to have too many or too few threads attached to each loomweight also played a part.

The results can be used to investigate the VI.1 loomweights. First the possible and probable tensions they were used at will be explored. From this it is then possible to calculate the range of the warp threads per centimetre implied by the size of the loomweight and the tension it was most likely used at. Warp threads per centimetre are an important element in describing a cloth, and these calculations allow some considerations about the cloth the VI.1 loomweights were being used to make. Only the loomweights of the modified data set, excluding the outliers, are used. In Tables 1 and 2 figures have been rounded to the nearest integer with measurements of exactly 0.5 rounded up.

Tension is calculated by dividing the weight of the loomweight by the desired tension for each warp thread. Thus to keep threads at 10 g warp tension the weight is divided by 10, for 20 g by 20, for 30 g by 30 and so

Weight (g)	10 g warp tension	20 g warp tension	30 g warp tension	40 g warp tension	50 g warp tension
15	2	1	0	0	0
96	**10**	**5**	3	2	2
175	**18**	**9**	**6**	4	4
308	31	**15**	**10**	**8**	**6**
737	74	37	**25**	**18**	**15**

Tab. 1 – Number of warp threads per loomweight at different tensions.

on. The same loomweight can, of course, be used for different loom setups at different warp tensions. As an example, Table 1 shows the number of threads at different warp tension for a selection of the loomweights from VI.1. The examples chosen are the largest, the two smallest and two with the typical mean values of the large and small (Tab. 1).

In M09 it was suggested that an optimal setup consists of 5 to 30 warp threads per loomweight, the number depending on the thickness of the thread. Inspection of Table 1 with this in mind immediately suggests that some of the combinations of numbers of threads and particular tensions can be rejected. The numbers in bold indicate combinations that appear plausible within the experimental set-up. This table is useful as it confirms that the very small loomweight recovered would be useless as a functional item, so a votive purpose does seem very likely.

It is thus possible to decide which combinations of numbers of warp threads at a given tension each loomweight could have been used at in an optimal setting. Any examples that fall above 30 or below 5 warp threads per loomweight at a given tension have been disregarded. Table 2 summarises the remainder showing how many loomweights there are in each combination of tension and warp thread numbers according to whether the weights fall into the large or small groupings already established. An individual loomweight can occur in the table more than once at different tensions. Structure in the table is immediately apparent. For the small loomweights the median values of number of warp threads at 10, 20, 30 and 40 g tension are 17, 8, 6 and 5 respectively. In the case of 30 and 40 g the median values are close to the lower boundary of what was considered optimal, and indeed in two of the experimental set-ups that number of warp threads at those tensions was only considered possible rather than optimal (M09, Tables 2 and 3). Many of the small loomweights were too light to be used at the 30 and 40 g tension and so it seems most likely they were used when weaving at 10 to 20 g tensions. For the large loomweights the equivalent median number of warp threads was 28, 15, 10 and 8. Approximately half of the larger weights were too heavy to be used effectively at a 10 g tension, and the median value of those that could be so used is very close to

Warp threads	Small 10 g WT	Small 20 g WT	Small 30 g WT	Small 40g WT	Large 10 g WT	Large 20 g WT	Large 30 g WT	Large 40 g WT
5	-	3	18	18	-	-	-	-
6	-	6	17	6	-	-	-	8
7	-	14	12	-	-	-	-	24
8	-	8	5	-	-	-	4	17
9	-	10	-	-	-	-	18	13
10	1	10	-	-	-	-	19	3
11	5	5	-	-	-	-	12	-
12	3	3	-	-	-	4	9	-
13	4	-	-	-	-	11	3	-
14	9	-	-	-	-	10	-	-
15	2	-	-	-	-	14	-	-
16	7	-	-	-	-	9	-	-
17	8	-	-	-	-	7	-	-
18	5	-	-	-	-	4	-	-
19	4	-	-	-	-	5	-	-
20	7	-	-	-	-	1	-	-
21	3	-	-	-	-	-	-	-
22	2	-	-	-	-	-	-	-
23	4	-	-	-	-	-	-	-
24	1	-	-	-	3	-	-	-
25	-	-	-	-	1	-	-	-
26	-	-	-	-	7	-	-	-
27	-	-	-	-	5	-	-	-
28	-	-	-	-	6	-	-	-
29	-	-	-	-	6	-	-	-
30	-	-	-		9	-	-	-

Tab. 2 – Summary of the distribution of warp threads per loomweight for the modified data set.

	Small				Large			
WT per cm.	10g	20g	30g	40g	10g	20g	30g	40g
2	-	-	6	9	-	-	-	-
3	-	7	33	15	-	-	-	45
4	-	22	12	-	-	-	3	11
5	1	24	1	-	-	3	29	1
6	1	10	-	-	-	18	25	-
7	11	2	-	-	-	26	7	-
8	10	-	-	-	-	13	-	-
9	14	-	-	-	-	3	-	-
10	13	-	-	-	2	1	-	-
11	4	-	-	-	1	-	-	-
12	6	-	-	-	8	-	-	-
13	4	-	-	-	11	-	-	-
14	1	-	-	-	5	-	-	-
15	-	-	-	-	2	-	-	-
16	-	-	-	-	2	-	-	-
17	-	-	-	-	2	-	-	-

Tab. 3 – Summary of the warp thread per centimetre range of the loomweights in the modified data set, according to the different tensions they could be used at in an optimal loom set-up.

the upper optimal boundary. It thus seems unlikely that the loomweights of the larger group were intended to be used at the 10 g tension, and one of c. 20 g or more would be appropriate (Tab. 2).

The experimental work showed that if the thickness of the loomweight and the number of threads attached to it is known, the number of warp threads per centimetre in the finished cloth can be calculated by adding together the total number of warp threads for a pair of loomweights (front and back) and dividing by the thickness of the loomweight in centimetres (full details are given in M09, 393). In our data set the measurement described as bottom minimum would be equivalent to thickness as defined in that paper (M09, fig. 7). Table 3 summarises the calculated warp threads per centimetre that can be calculated for fabric woven employing those loomweights which can be used in an optimal loom setup at a given tension (Tab. 3).

The first question to be asked of this table is whether the number of warp threads per centimetre indicated would fall within the expected range for contemporary cloth. Problems of preservation mean that more is known about Roman cloth from the north-western provinces and the eastern ones than from Italy itself so directly comparable material is lacking. In his survey of the textiles from the north-western provinces Wild usefully included warp threads per centimetre for the textiles then known (Wild 1970, tables A and B). The date range belongs to the Imperial era and so will be later than most of our loomweights, but the survey does give some indication of what might be expected. For wool plain tabby the range was 4 to 20 warp threads with the inter-quartile range between 7 and 12 and the median at 10. For linen the range was 7 to 30 with the interquartile range between 10 and 16 and the median at 12. Wild (1988, fig. 113) has also published a useful graph covering both eastern and north-western finds showing that the warp thread per centimetre for 2/1 twill mostly lie between about 6 and 14. In the latter publication he also notes that for Roman woollen fabric a warp thread count of 10 per centimetre would be regarded as a medium weight fabric whilst one of 20 would be regarded as fine.

If these figures can be taken as a guideline, then it can be seen from the figures in Table 3 that though some of the small weights could have been used at the higher tensions, they provide a more natural spread of warp threads per centimetre at 10 g. Equally the large loomweights look as though they might have been intended to function mainly at 20 to 30 g. This supports the conclusions that can be drawn from Table 2. Clearly much more detailed analysis and collaboration with textile specialists will be required to understand the patterning, but it does appear that the bimodality that has been revealed by the statistical analysis may well be reflecting the fact that they were being used for different types of cloth.

	Small at 10 g warp thread tension					Large at 20 g warp thread tension				
WT/cm	Ph. 2	Ph. 3	Ph. 4	Ph. 5	Ph. 6	Ph. 2	Ph. 3	Ph. 4	Ph. 5	Ph. 6
5	-	-	-	-	1	-	-	-	-	-
6	-	-	-	-	-	-	-	1	-	3
7	1	(1) 3	-	-	1	-	1	-	4	1
8	-	-	-	-	-	-	-	-	-	-
9	-	1	1	1	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-	-	-
11	-	-	-	-	1	-	-	-	-	-
12	-	-	-	(2) 1	1	-	-	-	-	-
13	-	(3) 1	-	-	-	-	-	-	-	-
Total 1	1	5	1	2	4	-	1	1	4	4
Total 2	1	3	1	3	5		1	1	4	4

Tab. 4 – The distribution of complete loomweights from the Casa del Chirugo by phase and number of warp threads per centimetre. Phase 2 is the activity prior to the building of the Casa del Chirugo in Phase 3 (c. 200 BC). Phase 4 is the primary definition of property boundary wall and some internal changes to the house after about 100 BC. Phase 5 is a period of widespread restructuring within the Casa del Chirurgo during the late 1st c. BC to early 1st c. AD. Phase 6 includes the construction of final decorative floors and the creation of upper stories within the house in the mid 1st c. AD. (Abbreviations Ph. = Phase; WT/cm = warp threads per centimetre. Notes: (1) might be Phase 5 but unlikely; (2) might be Phase 6; (3) might be Phase 5).

4. Discussion

This paper arose through the opportunity to extend our analysis of the loomweights in Insula VI.1 by the discovery of additional examples in the stores. These have confirmed the bimodal nature of the data. This was clearly intentional and so an explanation needs to be sought. The opportunity to explore the dataset by using the experimental results of work of Mårtensson and her colleagues has suggested that a major explanation may lie in the desire of the weavers to have loomweights suitable for particular types of loom setups. Clearly other factors may also contribute to the observed variations such as chronology and spatial distribution. In the earlier paper it was suggested that the bimodality might have been partially connected with chronology with the smaller and the less standard shapes being earlier. It was stressed that this was a very tentative suggestion given that the number from dated contexts was very small because the stratigraphic analysis of the insula as a whole was not far advanced. The phasing of the Casa del Chirurgo is now complete and Table 4 summarises the complete loomweights according to warp threads per centimetre and site phasing. There are a few problems over assigning three small loomweights to a phase. Where this happens it is noted in the table. Total 1 gives the phase total if they are assigned to the earliest possible phase, and total 2 if they are assigned to the latest possible.

Following the discussion in the previous section the small weights are considered at 10 g and the larger at 20 g tension. The numbers are still small but some interesting features do arise (Tab. 4).

The tendency of the weights in the earlier contexts to be the smaller ones can still be seen, even when the examples from contexts whose phasing is uncertain are assigned to the latest possible phase (Total 2). Of particular interest is the spread of the numbers of warp threads seen in the small ones compared with the much more concentrated spread seen in the large ones. A comparison with Table 3 shows that the small loomweights have a spread that is covering the whole range of the complete dataset, whilst the large ones are only in the lower part of it. Of course the higher the tension desired, the fewer the number of warp threads that can be attached to a particular loomweight; and that, as has been shown, can be directly related to the warp thread per centimetre measure. This does not appear to be the explanation here, as even if the warp thread per centimetre figures are examined for the large loomweights at 10 g tension the same restriction in range compared to the 10 g figures in Table 3 is seen. Given that at the Casa del Chirurgo the large loomweights are predominantly late it will be interesting to see if this is a general chronological trend in the data over the insula as a whole once more phasing information is available. If replicated it would open up the possibilities of investigating the types of cloth being woven at different times.

Even with a large data set such as this, the numbers in each date category are likely to be fairly small once the need for them to be complete has been factored in. The number presented from the Casa del Chirurgo in Table 4, for example, are less than half of the total known from the property. This is why the volume analysis reported on in section 2 is important as it should allow us to make some use of the numerous incomplete, but dated, examples.

Our previous finding, that the weights of the loomweights appeared to be bimodal, continues to be supported with the increased sample size. Work, as yet unpublished, using formal tests of the null hypothesis that the data are sampled from a population with a single mode allows the null hypothesis to be emphatically rejected (this is also true for the smaller sample size used in BC08). Thus we can be reassured that the sample bimodality reflects a population reality, encouraging the search for a functional interpretation described in the third section of this paper.

As we have discussed elsewhere (Cool, Baxter 1999), the important element of analysing any archaeological assemblage is to identify what the dominant factor causing any variation is. Once that has been identified, it can be set to one side so that underlying variations can be explored. What we have attempted to do here is to show how an observed pattern in the data can be explored in such a way as to retrieve an underlying explanation. The bimodality in the data has been confirmed, and exploring it via the practicalities of warp thread tension seems to be a very promising avenue. With this established, the data can be interrogated in other ways such as chronology. What is clear is that this very common, but rather neglected, class of artefact has much to tell us about the past providing they are correctly recorded.

Michael J. Baxter
Department of Physics and Mathematical Sciences
Nottingham Trent University

Hilary E.M. Cool
Barbican Research Associates
Nottingham

Michael A. Anderson
Department of Classics
San Francisco State University

Acknowledgements

We are most grateful to P.G. Guzzo, A. D'Ambrosio and all their colleagues in the Soprintendenza Archeologica di Pompei for their continued support and encouragement of the Anglo-American Project, and for the permission to work on the artefacts from the insula. The Anglo-American Project is directed by Rick Jones and Damian Robinson and we thank them for inviting Hilary E.M. Cool to analyse the finds. Especial thanks goes to Victoria Keitel and David Griffiths for finding the additional weights during their work on the pottery in 2009.

REFERENCES

Allason-Jones L., Bishop M.C. 1988, *Excavations at Roman Corbridge: the Hoard*, HBMCE Archaeological Report 7, London.

Baxter M.J., Cool H.E.M. 2008, *Notes on the statistical analysis of some loomweights from Pompeii*, «Archeologia e Calcolatori», 19, 239-258.

Cool H.E.M., Baxter M.J. 1999, *Peeling the onion: an approach to comparing vessel glass assemblages*, «Journal of Roman Archaeology», 12, 72-100.

Caputo T. 2004, *Il promontorio di Minerva: nuove acquisizioni topografiche*, in Senatore 2004, 51-101.

Mårtensson L., Nosch M-L., Strand E.A. 2009, *Shape of things: understanding a loom weight*, «Oxford Journal of Archaeology», 28, 4, 373-398.

Senatore E. 2004, *Pompei, Capri e la penisola sorrentina. Atti del Quinto ciclo di Conferenze di geologia, storia e archeologia (Pompei, Anacapri, Scafati, Castellamare di Stabia 2002-2003)*, Capri, Oebalus.

Wild J.P. 1970, *Textile Manufacture in the Northern Roman Provinces*, Cambridge, Cambridge University Press.

Wild J.P. 1988, *The textiles*, in Allason-Jones, Bishop 1988, 107-108.

ABSTRACT

In 2008 the Authors undertook a detailed statistical analysis of the dimensions of a large sample of loomweights from Insula VI.1, Pompeii. A major finding was that the weights of the loomweights appeared to have a bimodal distribution. Further analysis using loomweights that have come to light since the original work strengthens this observation. An analysis of loomweight volume has been undertaken with a view to predicting weight from volume for incomplete loomweights where sufficient information is retained to allow a volume calculation. Recently published experimental work allows an interpretation of the bimodality of the weights in terms of the loom set-up and the cloth being woven, and this is explored in some detail, along with further, tentative, observations on the chronology of the weights.