Medieval Ceramics 19 (1995) 89-98.

Notes on some statistical aspects of pottery quantification

M. J. BAXTER* and H. E. M. COOL**

SUMMARY

The quantification of finds assemblages from excavations, as an aid to the comparative study of assemblages and sites, is a developing and increasingly important aspect of post-excavation analysis. In the area of pottery studies a major influence has been the work ofClive Orton and Paul Tyers, culminating in the recent release of the 'pie-slice' package for computer analysis.

Much of the published literature is either highly technical, or of an expository nature which needs a great deal of the technical material and underlying assumptions to be taken on trust. The present paper is intended to be intermediate between these two levels.

We address some of the more complex or less obvious issues involved in application of the pie-slice 'philosophy'. A worked example is given to highlight aspects of some of the assumptions and calculations involved. Some attention is given to what can be done outside the pie-slice package. One concern is the analysis of assemblages that have been quantified using estimated vessel equivalents (eves), but not in a manner that allows conversion to pottery information equivalents (pies) in the pie-slice package.

i. irs 1 *ISXJl-/* u ^ 1 *1KJ1M*

This paper has arisen from research, supported by English Heritage, into the possibility of quantifying glass assemblages in a manner inspired by the work of Clive Orton, and developed for pottery assemblages, over a period of years. Orton's work has culminated in the publication of the pie-slice package (Orton and Tyers 1993) which implements many of the ideas previously published in Orton (1992), Orton and Tyers (1989, 1990, 1991, 1992), and Tyers and Orton (1991). These publications could be seen to have their roots in Orton (1975); the literature is scattered and, for the potentially interested archaeologist, not always accessible in terms of either location or content. The more general accounts of the underlying ideas (Orton 1993; Orton *et al* 1993; and Orton and Tyers 1993), are less technical but require the reader to accept some of the methodology on trust.

We have reviewed the literature, and the application of the pie-slice package, in some detail and in this paper attempt an expository and critical review of selected aspects of the methodology embodied in the package. This review is selective in the sense that it focusses on what we believe to be important aspects of the methodology, at a level intermediate among those published to date.

The pie-slice package requires that the pottery has been quantified by *estimated vessel equivalents* or *eves.* An eve is a measure of the amount of the vessel that survives. In pottery studies it is usually based on the proportion of the rim or base present, calculated by using calibrated rim charts (Orton 1980, pi. 14). Additionally, proper implementation of the package requires that in any context all of the rim or base fragments from a particular vessel are measured as a single eve (one *sherd family* per record). While recording in sherd families is desirable, it is the case that data have sometimes been collected in other ways, for example, where all of the rim fragments of a particular type from a context are recorded as a single eve irrespective of the number of vessels the fragments came from. This method of recording militates against the use of the pie-slice package and other solutions are needed. We will examine the extent to which some of the ideas in the package can be applied to such data.

In Section 2 the main ideas underlying the philosophy of the pie-slice package are briefly reviewed. This is followed by Sections 3 and 4 on the eve and its conversion to a *pottery information equivalent* or *pie,* which lies at the heart of pie-slice. A pie is a somewhat more abstract concept than an eve and less readily defined in intuitive terms. One pie contains as much statistical information as one whole vessel. It inherits the (desirable) properties of the eve but, unlike the eve, is constructed so that it behaves as a number which is comparable between assemblages and can be processed using statistical methods not applicable to eves. Section 4 focusses on this and some of the more difficult and less accessible ideas involved. In Section 5 a worked example is given to illustrate some of the calculations involved. In Section 6 we highlight what might very easily be done outside that package, using any reasonable spreadsheet or statistical software package. We also suggest what we think can be done with data that are in eves, but not suitable for use with pie-slice. We do not attempt to describe the statistical methods in pie-slice in any detail. We try to explain important ideas in words where this is possible; a technical appendix includes the formulae used in the worked example of Section 5.

2. BASIC IDEAS

Quantification of pottery is essential to effect comparison between assemblages in order to answer important archaeological questions. These involve the comparison of *relative proportions* of types within *target populations,* of which excavated assemblages are samples. This requires *unbiased* estimates of proportions to be derived from the sample, and of various measures that have been proposed only the *estimated vessel equivalent (eve),* satisfies the requirement that it produce unbiased estimates both within and between assemblages.

Using eves, proportions of types within assemblages can be descriptively compared. Assessing the statistical significance of any differences is more problematical. Statistical techniques for data recorded as counts (i.e. as whole numbers, such as the number of individuals in a sample belonging to a particular age group and social class) cannot be applied to eve totals. These are context-dependent (see section 3) so that, for example, the eve total for a type in an assemblage is not directly comparable to the total in a second assemblage.The main thrust of the work leading to the pie-slice package has been to develop a measure, the *pottery information equivalent* or *pie,* derived from the eve, that behaves as if it is a count that can be compared across assemblages. Statistically this allows the use of standard methodology, not otherwise available, to compare assemblages; archaeologically it opens up the possibility of integrating pottery studies with other kinds of finds. The methodology has been adapted to the study of glass by Cool (1994) and to bone finds by Moreno-Garcia *et al.* (1996).

The pie-slice package embodies quite complex statistical analyses but is written for users who will not usually be professionally trained to cope with such complexity. Inevitably, therefore, users have to accept its prescriptions and treat it as something of a 'black-box'. This paper illustrates some of the calculations that take place — with a small but realistic set of data — and shows how other analytical packages might be used for some aspects of the analysis.

3. ALL ABOUT EVE

The case for using eves, rather than other possible measures such as sherd count or estimated vessels represented, has been made in several of the publications cited and will not be rehearsed here. We confine ourselves to noting the assumptions under which eves were shown to be superior to other measures, using standard statistical sampling theory (Orton 1975).

The number of sherds into which a vessel breaks may depend on both the type of vessel and its context; and the probability that a sherd is found, from the target population of which the context provides a sample, may depend on the context. Both the number, and size, of sherds found in a context may thus be 'context-dependent', and this militates against the use of raw eves (and other measures) as the unit of comparison between contexts.

What is shown is that, given the assumptions, within an assemblage, eves lead to an estimate of the relative proportion of two types that is *unbiased.* This means, loosely, that the sample proportion is expected to be a reasonable reflection of the proportion in the target population. If the relative proportion of types is compared for two assemblages then their ratio is also unbiased. The eve is the only measure that has both these properties.

4. PIE-ING THE EVE

4.1 The importance of recording by sherd families

The main source for this discussion is Orton and Tyers (1990, 104-6). Using eves gives unbiased estimates of the proportions of particular vessel types within the target population. In statistical parlance these are examples of *point estimates.* Any point estimate will differ from the 'true' population or target value by an unknown amount and the uncertainty of the estimate can be measured in various ways, including the *variance.* Other things being equal, the larger a variance is the less reliable is the point estimate.

The ability to estimate a variance (using the theory of ratio estimation) for the proportion of a different type within an assemblage is crucial in moving from eves to pies. This theory assumes that the pottery is catalogued as records, and that these records are independent; this implies that the *sherd family* is the smallest permissible recording unit which can be transformed from eve to pie (Orton andTyers 1990, 104-5). If, within a context, over-detailed records are kept, in the sense that sherds from the same pot are split over more than one record, the records would be correlated and the results developed by Orton and Tyers cannot be used. Conflated records, containing measures of more than one vessel, can be used but it is inefficient to do so in the sense that variances are inflated and results become less reliable.

4.2 How many pies — a paradox?

In this sub-section we shall first attempt to describe what is involved *mathematically* to convert eves to pies. The *archaeological* assumption needed to justify the procedure is then considered.

For a single type the 'trick' is to estimate in two different ways the variance of the proportion of that type within an assemblage. The first estimate (Orton andTyers 1990, 89), based on the theory of ratio estimation, operates directly on the observed data and gives — in principle — a numerical value for the variance. The second estimate is obtained by proceeding *as if* the observed proportion had been obtained from a sample of n complete vessels. Using statistical theory associated with the binomial distribution gives a formula for the variance that involves the unknown n. Equating the two estimates and solving for n gives a pie estimate *for the assemblage.*

In other words, to obtain a pie for the assemblage we 'pretend' that the proportion has been found by sampling a notional number, n, of complete vessels (of all types) from the assemblage. Pies for each type are found simply by multiplying each eve proportion by the pie total.

The estimated *proportion* of a type within an assemblage is the same whether eves or pies are used, but the total pie will differ from assemblage to assemblage so that the multiplying factor from eve to pie depends on the assemblage. Grasping what n means can be rather difficult since it is something of an abstract quantity (essentially it generates numbers that behave as if they are samples from multinomial distributions that can be treated using statistical methods not applicable to eves).

In practice one need not determine the two separate variance estimates; the formulae for the two estimates can be rearranged to get a single estimate of n. There is a further catch in all the above. The value of n depends on the type used, so that if there are p types there will be p separate values of n. This problem can be overcome if it is legitimate to 'average' or pool results, and mathematical details are given in, for example, Orton and Tyers (1992), with the final result also given in the appendix in this paper and illustrated later (section 5).

The fundamental archaeological assumption that justifies the averaging procedure is that the actual vessel equivalents within a context (as opposed to their estimates) have the same distribution for each type. In particular the means and variances of the actual vessel equivalents for all types should be the same. Ideally the *estimates* of the vessel equivalents, the eves, should reflect this property, and this can be checked. For each type the variance and mean of the eves may be calculated, and from these a separate n for each type. If, for example, the values of n are markedly different, then *either* the assumption about the distribution of the actual vessel equivalents is wrong *or* the particular eve used does not truly reflect the actual vessel equivalent. The example in section 5 illustrates the ideas.

If the latter view is adopted then some corrective action is required, and this is discussed in the next sub-section. Note that from this standpoint the problem may arise because a particular estimate is unsuited to a particular type. If this is judged to be the case it is perfectly legitimate to use an estimate that is different from that used for other types; for example, it might be more appropriate to use base eves, rather than rim eves, for some types.

4.3 The 'chunkiness' problem

The problem noted in the last sub-section is discussed in Orton andTyers (1990, 96-8) and Orton and Tyers (1992, 170-71). It is argued that low average eve values for a type may be indicative of *residuality,* a function of post-depositional history, while large average- values reflect *chunkiness,* a function of fabric type.

An archaeological justification for the procedure to be described that resolves the problem is as follows; it is assumed rim eves are being used. Non-chunky type vessels fragment into 'small' sherds, including the rim, and rim sherds representing such a vessel are likely to be recovered in an excavated assemblage. Chunky-type vessels fragment into large fragments and, since not all sherds comprising the original vessel are recovered, there is more chance that rim sherds (from pottery represented in the assemblage) will be absent.

If within a context the eve distribution for one type differs markedly from that of other types then — if the assumption about the distribution of actual vessel equivalents is true — the eve cannot be correctly reflecting the actual vessel equivalent, which is the real quantity of interest. The proposed solution to the problem is to retain the eve total for the chunky types but to 'redistribute' the total over an imaginary number of records in such a way that the distribution of values mimics that across the

Table 1.

Record	Type	eve	Squared eve
1	1	40	1600
2	1	40	1600
3	1	40	1600
4	1	40	1600
5	2	80	6400
6	2	20	400
7	2	40	1600
.	.	.	.
.	.	.	.
.	.	.	.

non-chunky types (Orton andTyers, 1990, 98).The procedure described is designed to reflect more closely the actual vessel equivalent distribution inferred from other types in the context and is known as correction.

It can be shown that the correction procedure gives rise to a larger number of records than originally used, and that on recalculating the pies a larger number will, in general, be obtained compared with the 'uncorrected' analysis. [There appears to be a typographical error in equation (8) of Orton and Tyers (1990). Some details are given in the appendix to this paper.]

5. A WORKED EXAMPLE

Given the pie calculations, which as we shall see are readily calculated outside pie-slice, what is special to the pie-slice package is the statistical procedures that may be used for analysis of the pies. Before looking at these a simple worked example is given to illustrate how readily pies may be calculated using the theory developed by Orton and Tyers.

The data used are actually glass eves, obtained using methods developed in Cool (1994; see also Cool and Baxter, to appear), for material from a single context from excavations in Roman Chester. Glass is usually less abundant than pottery and to get the results shown some amalgamation of the originally defined types and contexts was needed. Even so, it will be noted that some of the types generate rather small numbers of records and eves, and investigation continues into how serious a problem this is. Type 5, bottles, dominates the analysis — as it does that of other contexts. The results we give below are thus illustrative rather than definitive.

Table 1 shows the first 7 of 84 records and indicates the type, eve value (multiplied by 100) and squared eve value (needed for later calculation). All the records for type 1 are given. This is in slightly simpler form than might be needed for pottery, where a fabric/form combination would often correspond to our single type of glass.

Using the notation given in the appendix, *m.* is the number of records for type *j* and *W.* is the eve total for type *j;* for example, for type 1 the number of records is $m_t = 4$, the eve total is $W_t = 160$ and the sum of squared eves is $S^2 = 6400$. For all 84 records these quantities have been calculated for each type separately and the information summarised in the first four columns in Table 2.

If pooling the results (see above section 4.2) is legitimate only the first three totals, $m = 84$, $W = 1694$ and $S^2 = 44924$ are needed. From equation (1) in the appendix the estimated pie total is

$$n = (83/84) \times (1694^2/44924) = 63.17$$

The estimated pie for type 1, for example, is then

$$n \times \{\text{eve for type } j\}/(\text{eve total for context}) = n \times W_j/W = 63.17 \times 160/1694 = 5.97$$

The fact that we have a small number of records for some types has already been noted (e.g. types 1, 3, and 4 with 4, 5 and 5 records), and in general the results for such types must be treated with caution, since the associated variances may be large. For illustration we will ignore this problem. The contents of the 5th column of the table should be similar for each type for the pooling leading to *n* to be justified. Informally, they look rather different. Furthermore, the mean eve values (column 6) ought to be similar and inspection suggests they are not (Orton andTyers 1990; 1992, discuss statistical tests of this). The final column of the table shows

Table 2.

		Column				
1 Type	2 Count (m_j)	3 eve total (W_j)	4 S	5 $\frac{S_j^2}{W_j}$	6 $\bar{W}_j$	7 n_j
1	4	160	6400	40.00	40.00	43.37
2	15	422	15780	37.39	28.13	49.54
3	5	98	2156	22.00	19.60	75.14
4	5	104	2360	22.70	20.80	73.00
5	55	910	18228	20.00	16.60	60.70
Total	m = 84	W = 1694	$S^2 = 44924$	20.20		

the results of calculating the pie total separately for each type, using equation (3) in the appendix. There are some rather large differences (admittedly associated with types with a small number of records). Suppose, for illustration, we decide that type 1 is 'chunky'. Applying the results of the appendix, and equation (2) in particular, results in a modified n of 66.67.This is a modest increase (not surprisingly, since both calculations are still dominated by the bottles).

There can be considerable merit in examining the data in such a table, or graphically (e.g. the distribution of eve values within type), since this may highlight problems with the data. For example Table 2 might suggest that type 2 is also 'chunky'. However, we note from Table 1 that there is an eve value of 80 for record 5, and in fact no other record has a value in excess of 40. If the 80 record is removed from the calculations then the mean drops to 24.4 which, along with the other statistics that can be calculated, is much more in line with types 3-5. The problem here is not that the type is chunky, but that a single vessel is much more complete than others (in terms of the eve calculations). With a smallish number of records per type this kind of problem may need attention before analysis.

It is straightforward to do the calculations outlined here. We used MINITAB and a calculator for the above analysis, but any reasonable spreadsheet or statistics package ought to be suitable. Identifying and dealing with unusual data is not, in general, straightforward and will be more complicated for pottery data than the glass data used here. For the glass data two aspects only were used — type and context — and the eves can take on only a limited number of values so that unusual data stand out. For pottery, three aspects - fabric, form and context — may be recorded and the eve measure can have any value between 0 and 100 (i.e. a complete rim), so that genuinely unusual data may be harder to identify. A sensible approach, when some observations look unusual, is to do calculations with and without them to see if the substantive conclusions are affected. To avoid any suspicion that data are being manipulated to attain a 'desirable' result, it is important in such circumstances to report all analyses. Either conclusions are the same in all cases, reinforcing one's confidence in them, or they are not and must therefore be treated with circumspection.

6. STATISTICAL ANALYSIS OF PIES

The pie-slice package assumes that data are presented 'three-dimensionally' as context by fabric by form. From such data it is possible to generate three tables — context-by-form, context-by-fabric, and fabric-by-form. These *two-way tables,* which are 'collapsed' versions of the full three-dimensional table and ignore one of the variables, are also called *marginal tables.* The statistical procedures in pie-slice operate in various ways on the full three-dimensional table and on the two-way tables. We will note most of these procedures only briefly, but comment in more detail on the graphical presentation of the two-way tables, since there is some hope here of using eve values that are not ideally suited to analysis with pie-slice.

A major practical problem is that with real data, the tables of context-by-form-by-fabric pies generated will often be sparse; that is, many of the pies will be zero or 'small'. The same may be equally true of entries in the two-way tables, and indeed for individual context, form or fabric totals. Such sparseness can cause problems for the statistical procedures used and must be dealt with. The merging or deleting of contexts, forms and/or fabrics necessary to obtain a useable reduced table should preferably be done on an archaeological rather than statistical basis. Where, for whatever reason, this is not viable, pie-slice allows the automatic merging or deleting of contexts, forms or fabrics, using a form of cluster analysis. Any merges will take place for a (statistical) reason, and can themselves provide additional archaeological insight into the data. If a merge does not, in retrospect, make archaeological sense, it can always be reversed.

Given a suitably reduced three-dimensional table, a main feature of pie-slice is the subsequent application of a technique (known as quasi- log-linear analysis) to investigate the relationships between context, fabric and form. The most complex case is that fabric and form are related, and that the nature of the relationship depends on context. Often the relationships are simpler than this, making it easier to understand the data. The simplest case is that context, fabric and form are mutually unrelated; archaeologically this seems unrealistic in most circumstances, so that the main hope in modelling is that the important two-way relationships can be determined (e.g. fabric and form are related but not in a way that depends on context).

If relationships are to be investigated formally in a statistical manner then pies must be used, but *correspondence analysis,* used in pie-slice to visualise the two-way marginal tables and to aid in interpretation of the results does not require pies. The following paragraphs, which are a little more technical than what has gone before, discuss how correspondence analysis might be applied to eves, and be expected to give similar results to the analysis of pies.

Correspondence analysis is a graphical method for the exploratory analysis of two-way tables of

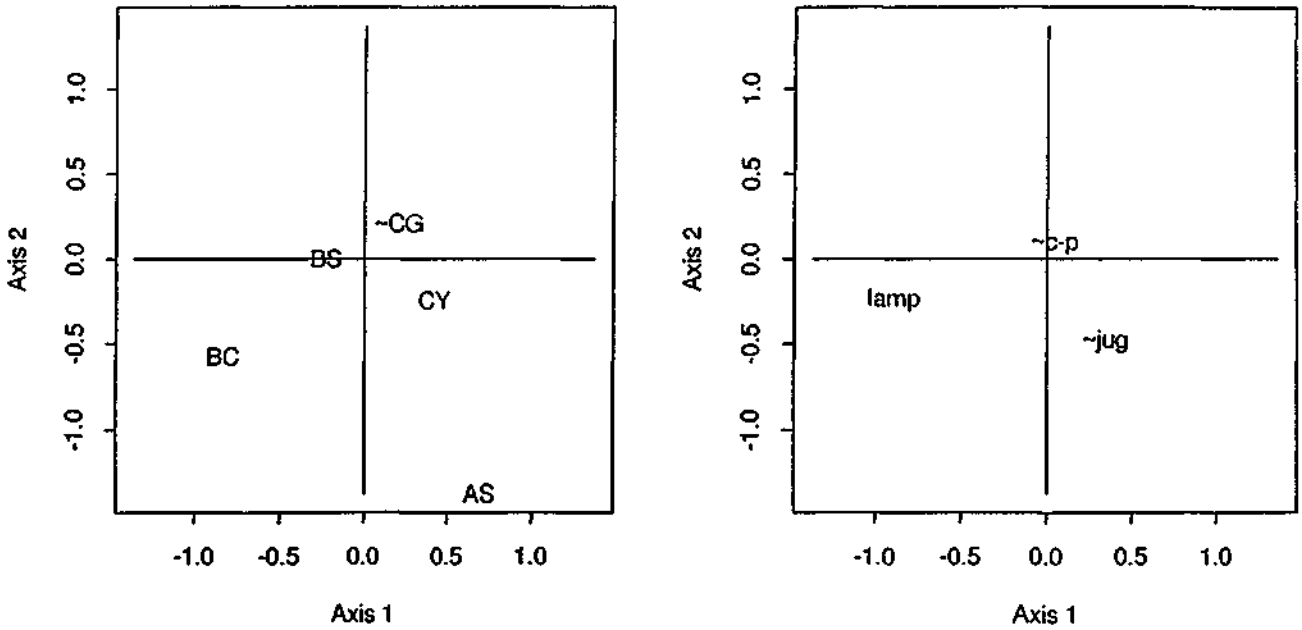

Fig. 1. Correspondence analysts of pies based on the 5 by 3 context by form data given in Orton and Tyers (1993, 20). The left plot is for contexts and the right for forms; labels identify contexts and forms with a ~ indicating that amalgamation of contexts or forms has occurred.

counts. It results in two 'maps' or plots; the scaling should be the same on both axes, and distances on the plot can therefore be read as real distances, as in a map (this fact has often been ignored).The first plot represents each row as a point; points close to each other on the plot represent rows with a similar profile. Thus if in a table, rows corresponded to contexts and columns to types (however defined) then similar rows correspond to contexts with a similar distribution of types. The second plot represents the columns (i.e. types). If the plots are 'overlaid' then the relative positions of the row (context) and column (type) points shows which types tend to characterise which contexts. More details are given in the appendix.

Correspondence analysis can be formally applied to any ordinary table (strictly, to a two-way table of non-negative numbers), and in particular may be applied directly to a table based on pies. In the method by which correspondence analysis is carried out the first step is to convert rows to 'profiles', or proportions, of the columns (types). In the subsequent analysis which leads to the row plot, rows with similar profiles should be close together, but more weight is given to rows based on larger pie totals while the representation of rows based on small pie totals may be unreliable (Pack and Jolliffe 1992). Similar comments apply to columns.

Correspondence analysis cannot be applied to a table of eves (e.g. context by form) directly because eve values are context-dependent and therefore not directly comparable (see Section 3). If, however, eves are converted to row proportions, where rows correspond to contexts, then the proportions are comparable and correspondence analysis can be formally applied. These row proportions are the same as those that would be calculated for the pie totals in the first stage of correspondence analysis, so that we might expect the row plot to be similar. The main difference is that in using eves we give each row (context) equal weight and ignore the fact that some may be based on far less data than others. For contexts with small eve totals more caution must thus be exercised in interpreting the positions on the plot (though this is also true of pies).

There is less reason to expect the column plots for eves converted to proportions to be similar to the column plot for pies. In each case correspondence analysis effectively rescales column values such that the column total is one, and — unlike the row rescaling — there is no reason why this should result in similar sets of figures. The similarity between the two analyses will be greatest when the pie totals for rows are similar.

To illustrate ideas an example based on that given in Orton and Tyers (1993, 15-23) will be used. Eve values were included for 1227 records from Winchester, classified by context, form and fabric.The eves are converted to pies, thus allowing formal statistical analysis. This takes place after some of the contexts, forms and/or fabrics are merged, a step that is necessary to get numbers large enough to permit the desired analysis (see above p. 93). Their analysis then showed that, for

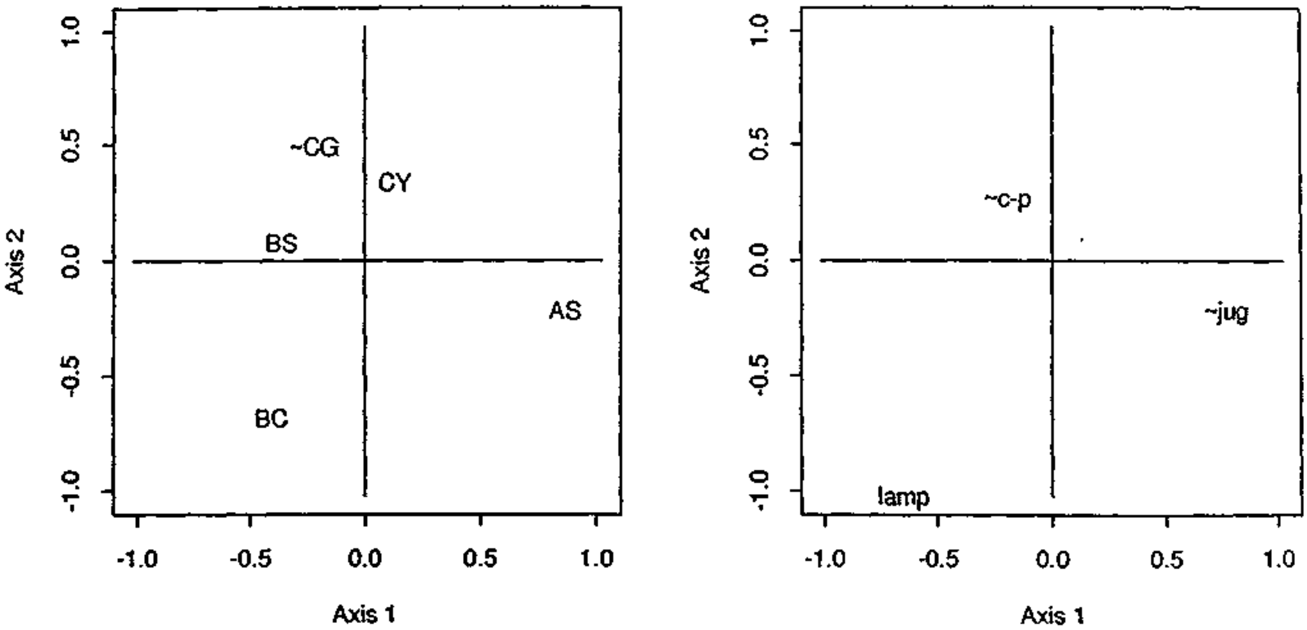

Fig. 2. Correspondence analysis of eve proportions on which the pies used in Figure 1 are based. The left plot is for contexts and the right for forms; labels identify contexts and forms with a ~ indicating that amalgamation of contexts or forms has occurred.

example, context and form are related but in a way that does not depend on fabric. The analysis presented here is of the reduced '5 by 3' table of context by form (Orton andTyers 1993, 20).

The statistical analysis undertaken has already established that context and form are related, and the relationship may be displayed graphically using correspondence analysis. The outcome of this is shown in Figure 1. For a given context, the pies for each form may be re-expressed as proportions of the total pie count for the context. The set of pies re-expressed in this way define the 'profile' of the context. The left plot shows how similar contexts are, in terms of their profile for the forms found within them. For example, the contexts *BS*, *--CG* and *CY* are more similar to each other than to the other contexts, while context *AS* is rather different from other contexts. The right plot shows how similar forms are in their distribution across contexts. If one imagines the two plots being superimposed it is possible to infer, for example, that context *BC* will have a high proportion of the form 'lamp' relative to other contexts, while *AS* will have a higher proportion of the form '~jug' than other contexts. This aspect of interpretation is discussed in more detail in the Appendix. It is possible to assess the reliability of the plots in a variety of ways (see e.g. Orton andTyers 1993).

Figure 2 is based on the eve proportions. There are some differences in the detail of the two analyses but the interpretation would be the same in both cases. This also proved to be true for comparison of eves and pies using glass (Cool 1994). Further empirical comparisons would be useful.

7. SUMMARY

This paper was motivated by several factors, the first of which was our attempt to come to grips with Orton's work on pottery quantification in order to apply the ideas to the quantification of Roman glass assemblages. Not all the ideas and assumptions are straightforward to grasp and we hope our interpretations and explanations may help others who wish to understand the literature. It is evident from conversations that we have had with pottery specialists that at least some perceive a need for discussion at a level intermediate to that which is available in the literature.

A second concern of some we have spoken to relates to the need to use pies rather than eves for quantification. This concern is acute where data *have* previously been recorded as eves, but by type rather than one sherd family per record. We believe that a lot can still be done descriptively with such data, including the use of correspondence analysis. The main limitation compared with the pie philosophy is that inferential statistics, and in particular log-linear models, cannot be used. Practically, this means that if fabric and form are related in a way that does depend on context (see Section 6) then correspondence analysis over-simplifies the data and could be misleading. Where pie-slice is applicable, the extent of this problem can be formally investigated; where it is not applicable and a problem in the data is suspected, more *ad hoc* approaches are possible (e.g. extracting separate marginal tables for contexts where there is enough data, and running separate correspondence analyses for each table to check that conclusions are similar).

Pie-slice, like all such packages, is designed to undertake a specific set of tasks and present the results in a particular way. Potential users should be aware that there are things they might wish to do with their data, including the descriptive inspection of eve values, that may be more readily accomplished using other software. The calculations from eves to pies can be achieved, with relative ease, in other packages which have more analytical flexibility.

Quantification is important. One way of gaining insights into past activity is to compare the detritus left behind on sites, to explore the different types of activities that might have taken place on them. This can only be ddne if the proportions of the different types of artefacts in these assemblages can be reliably measured. The work involved in the development of pie-slice has been important in leading to an awareness of this, and we hope that this paper and those of Cool (1994) and Cool and Baxter (forthcoming) will contribute to further development.

APPENDIX 1 —TECHNICAL DETAILS

Following the notation established by Orton and Tyers (1989, 1990, 1991) let there be m records in an assemblage with total eve value of IF and with w. the eve value for the i'th record. Let the sum of the w^2 be S^2. If pooling the different estimates for a type is legitimate (see Section 4.2) the estimated pie value for an assemblage, $M_{/5}$ is

$$n, = aWIS^2 \qquad \text{(i)}$$

where $a — (m - l)lm$.

Suppose next that a single chunky type with m. records is identified, with the remaining m . records non-chunky. We use the notation $-j$ to identify calculations omitting the m. records. Suppose that the average eve value for chunky types is e_t and for non-chunky types e_2. The chunky eve total must then be redistributed among

$$m_{ej} . = m_{-j} W_J.IW_{-J}. = me_1.le^\wedge_{1\,2}$$

imaginary records to give the same eve average as non-chunky types. For example, if the average eve value per record is twice that of non-chunky types, and there were originally 13 chunky records, then we proceed as if the eve values for chunky types are distributed among 26 records.

Making suitable adjustments the new formula for estimating the pie total then becomes (Orton and Tyers 1990; equation (8))

$$n_2=bWW_{-J}.IS^2_{-J} \qquad (2)$$

where $b = (m_e - l)lm_e$ and m_e is now the estimated number of records. Assuming that the number of records is large enough for $(m -l)lm$ to be approximated by 1 etc. gives

$$njn_{2,1} = m_{-J}\ S^2lmS_{e\ -j}^{\ 2}$$

using $W\ JW - mlm_c$.

Assuming, for simplicity that all chunky records have eve value e_t and non-chunky records have eve value e_2 this becomes

$$njn_t = (k^2 + q)/(k + q)$$

where $k = e_tle_2$ and $q = mlm..$ Since k is bigger than 1 by definition this shows that the adjustment for chunkiness will tend to increase the pie values.

For the example in the text, where pie totals are calculated using each type separately we need the following formula (e.g. Orton and Tyers 1992; equation (3))

$$n_j = aWWWKW'S^2_{i-i} + W^2_iS^2_{-j}) \qquad (3)$$

It remains to discuss aspects of the interpretation of the correspondence analyses of Section 6. A detailed account of archaeological applications of correspondence analysis, including the assessment of how reliable it is, is given in Baxter (1994). Essentially a correspondence analysis of a two-way table results in two 'maps', one for the rows (contexts) and one for the columns (forms) of the table. In a successful correspondence analysis these can be read in a similar way as conventional plots and show, for example, how similar contexts are with respect to their assemblage of forms. Close contexts have similar assemblages; distant contexts have dissimilar assemblages.

It is common, though we have not done so, to superimpose the two plots in order to be able to say something about the relationship between rows and columns. There are, in fact, different conventions that may be used to effect the superimposition that differ in the apparent 'distance' between a row and column point. Usually plots similar to those in Figures 1 and 2 are overlaid. In the examples in the text the general principle is that contexts away from the origin — the (0, 0) point — will have a higher proportion of those forms that occupy a similar region of the graph than other contexts. Thus *AS* is an 'extreme' point in Figure 1 and is most closely associated with the form ~*jug.* What we *cannot* do is interpret distances between points from the two different plots directly. Thus, on overlaying the two plots in Figure 1, context *CY* appears closer to form ~ *jug* than context *AS,* but this cannot be interpreted as meaning that *CY* has a higher proportion of the form. This is because the apparent distance between points from two plots are a function of the overlaying convention adopted. It is the relative extremity of *AS,* coupled with the location *of ~jug* in the same section of the plot, that permits the interpretation given in the text.

Acknowledgements

The research leading to this paper was supported by a grant from English Heritage.

We are particularly grateful to Clive Orton and Paul Tyers for their detailed comments on earlier drafts that have clarified many points, and for the careful reading and suggestions of the editors and referees of Medieval Ceramics. Needless to say any remaining errors, misunderstandings and interpretations are ours alone.

BIBLIOGRAPHY

Baxter, M.J. 1994, *Exploratory Multivariate Analysis in Archaeology,* Edinburgh University Press, Edinburgh.

Cool, H.E.M. 1994, *The Quantification of Roman Vessel Glass Assemblages,* unpublished report for English Heritage.

Cool, H.E.M. and Baxter M.J. (forthcoming) 'An approach to the quantification of glass assemblages', *Proceedings of the 13th Congress of the International Association for the History of Glass.*

Moreno-Garcia, M., Orton, C.R. and Rackham J. 1996, 'A new statistical tool for comparing animal bone assemblages', *Journal of Archaeological Science* 23, 437-53.

Orton, C.R. 1975, 'Quantitative pottery studies: some progress, problems and prospects', *Science and Archaeology* **16,** 30-5.

Orton, C.R. 1980, *Mathematics in Archaeology,* Collins, London.

Orton, C.R. 1982, 'Computer simulation to assess the performance of measures of quantities of pottery', *Wbrld Archaeology* **14,** 1-20.

Orton, C.R. 1992, 'Quantitative methods in the 1990s', in Lock, G. and Moffett, J. (eds.), *Computer Applications and Quantitative Methods in Archaeology 1991,* BAR International Series S577, Tempus Reparatum, Oxford, 137-40.

Orton, C.R. 1993, 'How many pots make five? — An historical review of pottery quantification', *Archaeometry* 35, 169-84.

Orton, C.R. and Tyers, P.A. 1989, 'Error structures of ceramic assemblages', in Rahtz, S. and Richards, J. (eds.), *Computer Applications and Quantitative Methods in Archaeology 1989,* BAR International Series 548, Oxford, 275-85.

Orton, C.R. and Tyers, P.A. 1990, 'Statistical analysis of ceramic assemblages', *Archaeologia e Calcolatori* 1, 81-110.

Orton, C.R. and Tyers, P.A. 1991, 'A technique for reducing the size of sparse contingency tables', in Lockyear, K. and Rahtz, S. (eds.), *Computer Applications and Quantitative Methods in Archaeology 1991,* BAR International Series 565, Tempus Reparatum, Oxford, 121-6.

Orton, C.R. and Tyers, P.A. 1992, 'Counting broken objects: the statistics of ceramic assemblages', *Proceedings of the British Academy* 77, 163-84.

Orton, C.R. and Tyers, P.A. 1993, *A user's guide to pie-slice.* University College London Institute of Archaeology.

Orton, C.R., Tyers, P.A. and Vince, A.G. 1993, *Pottery in Archaeology,* University Press, Cambridge.

Pack, P. and Jolliffe, I.T. 1992, 'Influence in correspondence analysis', *Applied Statistics* **41,** 365-80.

Tyers, P.A. and Orton, C.R. 1991, 'Statistical analysis of ceramic assemblages', in Lockyear, K. and Rahtz, S. (eds.), *Computer Applications and Quantitative Methods in Archaeology 1991,* BAR International Series 565, Tempus Reparatum, Oxford, 117-20.

* Department of Mathematics, Statistics and Operational Research, The Nottingham Trent University, Clifton Campus, Nottingham NG11 8NS.

** 16 Lady Bay Road, West Bridgford, Nottingham NG2 5BJ.

Resume

La quantification d'assemblages mobiliers venant de fouilles, en tant qu'aide ä l'etude comparative soit d'assemblages ou de sites, constitue un aspect de plus en plus important dans l'analyse des resultats de fouilles. En ce qui concerne la ceramologie, une infuence importante provient de l'oeuvre de Clive Orton et de PaulTyers, culminant avec l'apparition recente du Programme d'analyse informatique 'pie-slice' (en francais:' tranche de tarte').

Une grande partie de la litterature publiee est soit hautement technique, soit de nature expositoire necessitant une acceptance quasi-absolue au sujet du materiel technique ou des suppositions sous-jacentes. L'etude presente entend se placer ä un un point intermediaire.

Nous addressons certains aspects les plus compliques et les moins evidents lies ä la mise en pratique de cette 'Philosophie pie-slice'. Un exemple pratique est presente pour mettre en valeur des aspects qui appartiennent aux suppositions ou calculs concernes. Une certaine attention est portee ä ce qui peut etre effectue en dehors du Programme 'pie-slice'. Un probleme concerne l'analyse d'assemblages quantifies ä l'usage d'equivalence estimee de vases ('eves'), mais pas de facon ä permettre la conversion en equivanents d'information ceramique ('pies') dans le Programme 'pie-slice'.

Zusammenfassung

Die Quantifizierung von Fundgruppen bei Ausgrabungen als Hilfsmittel vergleichender Untersuchungen von Ansammlungen und Ausgrabungsstätten ist ein sich entwickelnder und zunehmend bedeutsamer Aspekt für die nach der Ausgrabung stattfindende Analyse. Von großer Bedeutung für Töpfereistudien sind die Arbeiten von Clive Orton und Paul Tyers, die ihren vorläufigen Höhepunkt in dem kürzlich herausgekommenen "pie-slice" Paket für Computeranalyse gefunden haben.

Viele Veröffentlichungen sind entweder hochtechnisch oder darlegend, wobei man einen großen Teil des technischen Stoffes und die zugrunde liegenden Annahmen einfach glauben muß. Die vorliegende Arbeit zielt auf einen Mittelweg.

Wir beschäftigen uns mit einigen der vielfältigeren und weniger offensichtlichen Probleme, die mit der Verwendung der "pie-slice" Philosophie zusammenhängen. An einem ausgearbeiteten Beispiel zeigen wir schlaglichtartig einige der verwendeten Annahmen und Berechnungen. Auch wird der Frage Aufmerksamkeit gewidmet, was man außerhalb des "pie-slice" Pakets tun kann. Eines der Anliegen ist die Analyse von Fundgruppen, die mit Hilfe geschätzter Gefäßäquivalente (eves) quantifiziert wurden, jedoch nicht in einer Art, die eine Umwandlung in Töpferei Informationsäquivalente (pies) im "pie-slice" Paket erlauben würde.