J. R. Statist. Soc. A,
(1977), **140,** *Part* 3, pp. 277–297

Rudiments of Numeracy

By A. S. C. Ehrenberg

London Business School

[Read before the Royal Statistical Society on Wednesday, March 2nd, 1977, the President Miss Stella V. Cunliffe, in the Chair]

Summary

Many tables of data are badly presented. It is as if their producers either did not know what the data were saying or were not letting on. Some precepts for improved data presentation are discussed.

Keywords: NUMERACY; TABLES; GRAPHS; DATA PRESENTATION; ROUNDING

People who say they are not numerate usually do not mean that they cannot do arithmetic. Nor should they mean that they cannot do mathematics. Instead, they are really saying that they cannot cope with numerical data—tables, graphs, percentages, and so on. But since such data are often badly presented—requiring much effort even for sophisticated users to understand—the fault is that of the *producers* of the data. That is the starting-point of this paper.

Numeracy has two facets—reading and writing, or *extracting* numerical information and *presenting* it. The skills of data presentation may at first seem *ad hoc* and judgemental, a matter of style rather than of technology, but certain aspects can be formalized into explicit rules, the equivalent of elementary syntax. Such precepts have largely been ignored in statistical practice and teaching.

In this paper I therefore put up for discussion some rules or guidelines for improved data presentation. In doing so my immediate concern is not with the general public but with supposedly numerate people like ourselves—producers and more or less regular users of numerical information. I am not so agitated about the less numerate fringe (e.g. backward school-children or apocryphal company chairmen); they also need help but will not do much with numerical information however well it is presented.

The paper is in five sections. Section 1 gives two examples of how the presentation of data can be improved. Specific rules for doing so are then set out in Section 2, followed by a brief assessment of the relevant literature in Section 3. Possible objections and problems of implementing the rules are discussed in Sections 4 and 5.

1. Seeing the Data

The criterion for a good table is that the patterns and exceptions should be obvious at a glance, at least once one knows what they are. But most tables do not meet that standard.

To illustrate, Table 1 reproduces a small table of data on UK merchant vessels from *Facts in Focus*, a typical publication of statistical information for general use (CSO, 1974, Table 63). The table may at first appear reasonably well laid out. But in forming this view one's attention probably has centred not on the numbers but on the captions, i.e. Dry cargo, Tankers, Gross and Deadweight tonnages, and so on.

The numbers themselves are not as easy to take in. What are their main features? How can they be summarized? How can one tell someone over the phone? What is one likely to absorb or remember? Looked at with these questions in mind the table now appears like a fairly undigested jumble of numbers. But it need not have been like that.

TABLE 1

United Kingdom Merchant Vessels in Service
(500 gross tons and over)

	1962	1967	1973
Number			
All vessels	2,689	2,181	1,776
Passenger*	242	173	122
Dry cargo	1,847	1,527	1,165
Tankers	600	481	489
Thousand gross tons			
All vessels	20,554	20,375	29,105
Passenger*	2,504	1,709	920
Dry cargo	10,562	10,757	13,520
Tankers	7,488	7,908	14,665
Thousand deadweight tons			
All vessels	26,577	27,448	46,763
Passenger*	1,467	919	349
Dry cargo	13,990	14,362	20,115
Tankers	11,120	12,167	26,299

* All vessels with passenger certificates. Source: Facts in Focus

Table 2 gives an improved presentation of the same data. It is easier to see major patterns and exceptions:

The *numbers* of vessels declined over the years by 30 to 50%, but less for tankers.

The *tonnages* jumped dramatically by up to 100% between 1967 to 1973, except for passenger vessels.

Dry cargo vessels accounted for the largest numbers of vessels and also the biggest tonnages, with tanker tonnages overtaking the dry cargo ones in 1973.

Passenger vessels differed from the others in having larger gross than deadweight tonnages.

Few of these patterns seem as clear in Table 1, even now that one knows what to look for. The original table therefore fails both the strong and the weak versions of the criterion for a good table, whilst Table 2 certainly passes the weak version if not entirely the strong one:

The Strong Criterion for a Good Table: The patterns and exceptions should be obvious at a glance.

The Weak Criterion: The patterns and exceptions in a table should be obvious at a glance once one has been told what they are.

The weak criterion is much the more important one. It applies automatically to all situations which are repetitive, i.e. ones where the probable pattern of the new data is known beforehand. It can therefore cover more complex tables and apply to the experienced user.

The strong criterion sounds fine. But it says nothing more than that the naive newcomer should gain instant insight, unaided. This will seldom work. With data that are altogether new, or at least new to the expected reader, the producer of the table cannot merely announce that "the results are shown in the table" and expect every reader to work out the story-line himself. Instead, he should guide the reader by a brief verbal commentary and tell him what he knows.

This is the *weak* criterion in operation again. It is illustrated by Paul Samuelson's *Economics* (Samuelson, 1976) where every table and graph is accompanied by a short paragraph commenting on what it says, as exemplified in Table 3. Whilst Samuelson's

TABLE 2

An "Improved" Version of Table 1

Vessels of 500 gross tons and over	1962	'67	'73
Number			
Dry Cargo	1,800	1,500	1,200
Tankers	600	480	490
Passenger*	240	170	120
ALL VESSELS	2,700	2,200	1,800
Gross Tons ('000)			
Dry Cargo	11,000	11,000	14,000
Tankers	7,500	7,900	15,000
Passenger*	2,500	1,700	900
ALL VESSELS	21,000	20,000	29,000
Deadweight Tons ('000)			
Dry Cargo	14,000	14,000	20,000
Tankers	11,000	12,000	26,000
Passenger*	1,500	900	300
ALL VESSELS	27,000	27,000	47,000

* All vessels with passenger certificates

TABLE 3

A Table with Commentary

(Table 2·4 from Samuelson's *Economics*)

Nation sizes will look different in the future

Any differences in growth rate accumulate into significant changes. Note how the United States and the Soviet Union grow relative to Western Europe. (Source: United Nations.)

ESTIMATED FUTURE POPULATION OF DIFFERENT COUNTRIES IN 1985 (in millions)

	ANNUAL GROWTH (% per year)	1970	1980	1985
United States	1.3	205	226	240
United Kingdom	0.6	55.7	59.5	61.8
France	0.8	50.8	55.3	57.6
Soviet Union	1.0	243	271	287
Sweden	0.7	8.0	8.6	8.8
Italy	0.8	53.7	57.9	60.0
Japan	1.2	103	116	121

format need not be copied slavishly (usually one would simply comment in the main text) and some of his tables are none too good, his care to communicate is no doubt correlated with the book's phenomenal success over the years.

A common doubt about trying to improve the layout of a table is whether the presentation should not depend on the particular use to be made of the data. But an "improved" version like Table 2 is easier for virtually *any* purpose than the original Table 1. The data could perhaps be displayed in a way even more suited to some specific purpose, but that would merely mean taking the procedures of this paper yet further.

The main steps in going from Table 1 to Table 2 (such as rounding and re-ordering the rows, and possible objections to them) will be discussed in later sections. At this stage I only want to illustrate how some marked improvements in data presentation are possible even with a small and fairly simple table: The golden rule is that the next step or two in looking at the figures in a table must be visually easy.

The 10×10 correlation matrix in Table 4 is another small but more analytical table. The variables here are whether people in a sample of 7,000 UK adults said they "really liked to watch" a range of ten TV programmes like World of Sport (WoS), Match of the Day (MoD), Panorama (Pan), and so on (from data in Goodhardt *et al.*, 1975, Chapter 9).

TABLE 4

Adults who "Really Like to Watch": Correlations to 4 *Decimal Places*

(Programmes Ordered Alphabetically within Channel)

		PrB	ThW	Tod	WoS	GrS	LnU	MoD	Pan	RgS	24H
ITV	PrB	1.0000	0.1064	0.0653	0.5054	0.4741	0.0915	0.4732	0.1681	0.3091	0.1242
"	ThW	0.1064	1.0000	0.2701	0.1424	0.1321	0.1885	0.0815	0.3520	0.0637	0.3946
"	Tod	0.0653	0.2701	1.0000	0.0926	0.0704	0.1546	0.0392	0.2004	0.0512	0.2437
"	WoS	0.5054	0.1474	0.0926	1.0000	0.6217	0.0785	0.5806	0.1867	0.2963	0.1403
BBC	GrS	0.4741	0.1321	0.0704	0.6217	1.0000	0.0849	0.5932	0.1813	0.3412	0.1420
"	LnU	0.0915	0.1885	0.1546	0.0785	0.0849	1.0000	0.0487	0.1973	0.0969	0.2661
"	MoD	0.4732	0.0815	0.0392	0.5806	0.5932	0.0487	1.0000	0.1314	0.3267	0.1221
"	Pan	0.1681	0.3520	0.2004	0.1867	0.1813	0.1973	0.1314	1.0000	0.1469	0.5237
"	RgS	0.3091	0.0637	0.0512	0.2963	0.3412	0.0969	0.3261	0.1469	1.0000	0.1212
"	24H	0.1242	0.3946	0.2432	0.1403	0.1420	0.2661	0.1211	0.5237	0.1212	1.0000

TABLE 5

The Correlations for the 10 *TV Programmes Rounded and Re-ordered*

Programmes		WoS	MoD	GrS	PrB	RgS	24H	Pan	ThW	Tod	LnU
World of Sport	ITV		.6	.6	.5	.3	.1	.2	.1	.1	.1
Match of the Day	BBC	.6		.6	.5	.3	.1	.1	.1	0	0
Grandstand	BBC	.6	.6		.5	.3	.1	.2	.1	.1	.1
Prof. Boxing	ITV	.5	.5	.5		.3	.1	.2	.1	.1	.1
Rugby Special	BBC	.3	.3	.3	.3		.1	.1	.1	.1	.1
24 Hours	BBC	.1	.1	.1	.1	.1		.5	.4	.2	.2
Panorama	BBC	.2	.1	.2	.2	.1	.5		.4	.2	.2
This Week	ITV	.1	.1	.1	.1	.1	.4	.4		.3	.2
Today	ITV	.1	0	.1	.1	.1	.2	.2	.3		.2
Line-Up	BBC	.1	0	.1	.1	.1	.2	.2	.2	.2	

Again the patterns and exceptions are not clear. But appropriate re-ordering of the variables, rounding, better labelling and better spacing lead to a marked improvement, as shown in Table 5. Now we can see that there is a cluster for the five Sports programmes, another cluster for the five Current Affairs programmes, and three locally high correlations of ·2 between Panorama and the Sports programmes.

I am not concerned in this paper with how the appropriate ordering of the variables was initially discovered (although this can be greatly helped by good data presentation). What concerns me here is our ability to see, understand and communicate such a pattern once it has been established. Even now that we know the pattern, it is still not apparent in Table 4. In contrast, anyone can see it in Table 5 (especially anyone already familiar with the notion of a correlation matrix). In fact Table 5 is largely redundant, as with all tables which satisfy the *strong* criterion of a good table. Its main pattern could be described in words alone, as consisting of two clusters: correlations of ·3 to ·6 between the five Sports programmes and of ·2 to ·5 between the five Current Affairs, with correlations of ·1 or so between these two clusters. Yet some deliberate redundancy in communication usually helps.

Experience indicates that most people would agree that the improved tables illustrated here are somehow better than the original versions. But more formal assessment of this also seems desirable—not merely to "prove" the difference, but to see in what *ways* the improvements work, for what kinds of people, and under what circumstances. Some exploratory studies in this direction are summarized elsewhere (Chakrapani and Ehrenberg, 1976).

2. Six Basic Rules

The table improvements illustrated so far involved a combination of factors (subsumed by the golden rule that the *next* steps in looking at a table should be visually easy). These factors can be considered separately and in this section I outline six specific rules or guidelines which deal in turn with drastic rounding, marginal averages, choosing between rows and columns in a table, ordering the rows or columns, the use of space, and the differing roles of graphs and tables.

These rules will be illustrated with another example from *Facts in Focus* (CSO, 1974, Table 97) concerning the level of unemployment in Great Britain over four selected years, as reproduced in Table 6. Although this is again a small and simple table (chosen for conciseness of exposition here), the numerical details are once more not obvious at a glance.

Table 6

Unemployment in Great Britain—Original Version

	1966	1968	1970	1973
Total unemployed (thousands)	330.9	549.4	582.2	597.9
Males	259.6	460.7	495.3	499.4
Females	71.3	88.8	86.9	98.5

Suppose we now look away from this table. What do we remember having seen, *without looking back*? What can we say about the numbers of unemployed?

Rule 1: Rounding to Two Significant or Effective Digits

Understanding any set of numbers involves relating the different numbers to each other. But in Table 6 this is not easy. For example, mentally subtracting the 1966 total from the 1973 total and remembering the answer is relatively difficult (330·9 from 597·9 = 267·0). Taking ratios mentally (330·9 *into* 597·9) is virtually impossible. Most of us can do such mental arithmetic only by first rounding the figures to one or two digits in our heads.

In Table 7 this rounding has been done for the reader. The general rule is to round to two significant or effective digits, where "significant" or "effective" here means digits which vary in that kind of data. (Final 0's do not matter as the eye can readily filter them out.)

Now we can see that the difference between 330 and 600 for total unemployed is 270, and that 330 *into* 600 is almost 2, i.e. an increase of almost 100%. We can also see that the increase for males from 260 to 500 is again nearly 100%, and that the corresponding increase for females is about 40%, from 71 to 99. Total unemployed up by almost 100%, males up by

TABLE 7

Unemployed in GB—Rounded

000's	1966	'68	'70	'73
Total unemployed	330	550	580	600
Male	260	460	500	500
Females	71	89	87	99

almost 100% and females up by less than 50%: that is something one *can* remember. It is also easier to recall that the range for total unemployed is from about 330 to 600 than that it is from 330·9 to 597·9.

Returning to Table 6 we see how any comparable assessment of the figures would necessarily involve mental rounding. Pocket calculators are not the answer since knowing that 597·9/330·9 = 1·8069 does not greatly help us to see and absorb the patterns in the table. For better or for worse, drastic rounding is necessary if we are to see and assimilate the data. Whether rounding to two significant digits is "going too far" is a possible objection considered in Section 4.

A lesser problem is that the male and the female numbers are shown to the nearest ten thousand and the nearest thousand respectively, by being rounded to two significant digits *in their own context*. This avoids over-rounding when different groups of figures vary greatly in size (as also occurred in Table 2). The consequence is that the figures do not add up exactly. This is an undoubted nuisance, but a lesser one than the perceptual difficulties of the unrounded data in Table 6—anyone who cannot learn to cope with rounding errors will probably not get much out of such statistical data anyway.

Rule 2: Row and Column Averages

The next rule concerns the use of row or column averages to provide a visual focus and a possible summary of the data. Table 8 illustrates this by giving the row averages across the four years. (The column totals in this table serve almost the same purpose as column averages.)

TABLE 8

With Averages

000's	1966	'68	'70	'73	Ave.
Total unemployed	330	550	580	600	520
Male	260	460	500	500	430
Female	71	89	87	99	86

Even with a small table such averages prove useful. Noting that the average male/female ratio is 5 to 1 (i.e. 430/86), we can see more readily how this ratio varies over the years, from less than 4 to 1 in 1966 to just over 5 to 1 in the three succeeding years. Put in statistical jargon, by making the "main effects" explicit (here the row averages and column totals) we can see more easily any "interactions" between rows and columns (here sex by the years). The general rule is to work out row and column averages before scrutinizing the detailed figures.

Rule 3: Figures are easier to Compare in Columns

Figures are easier to follow reading down a column than across a row, especially for a larger number of items. Even for our small example here, Table 9 makes it easier to see that each category of unemployed was substantially lower in 1966 than in the three later years.

TABLE 9

Rows and Columns Interchanged

GB	Unemployed (000's)		
	Total	Male	.Female
1966	330	260	71
'68	550	460	89
'70	580	500	87
'73	600	500	99
Average	520	430	86

We also notice minor variations and sub-patterns more, for example that contrary to the total trend, the female figures levelled off only for 1968 and 1970 (in fact dropping slightly in the latter year), and that the 1973 figure of 99 is markedly high. Compared with Table 8 we are beginning to see more of the data.

The improvement is a perceptual one. To see in Table 8 that the main variation for total unemployed is from roughly 300 to 600, the eye first had to take in and then partially ignore the symbols and gaps in the sequence

330 550 580 600

and it had to travel relatively far to do so. But in Table 9 the hundreds are close together.

The eye can run down the first digit in each column and totally ignore the rest, i.e.

3..
5..
5..
6..

It could also marginally take in the second digits whilst still concentrating on the first

33.
55.
58.
60.

This we tend to do anyway when we read long strings of longer numbers, whether across the page

330·9 549·4 582·2 597·9 261·3 734·6 790·2

or, preferably, *downwards*

330·9
549·4
582·2
597·9
261·3
734·6
790·2

where the blip (a copying or typing error ?) typically stands out more clearly.

Rule 4: Ordering Rows and Columns by Size

Ordering the rows and/or columns of a table by some measure of the size of the figures (e.g. their averages) often helps to bring order out of chaos. It means using the dimensions of the table to enable us to see the structure of the *data* rather than merely reflecting the structure of the row or column labels (which is usually already well known). The tables in Section 1 gave striking examples.

The present unemployment data already have the rows and columns in an effective order of size because the trends happen to coincide with the order of the years. But to illustrate the rule further Table 10 gives the data with the rows in another order, *A* to *D*. Even with such

TABLE 10

Rows in Some Other Order

GB	Unemployed (000's)		
	Total	Male	Female
A	550	460	89
B	580	500	87
C	330	260	71
D	600	500	99
Average	520	430	86

a small table it is less easy to see that the Row *C* (or 1966) figures are generally the smallest. Interactions are even harder to spot, e.g. that the male figures in Rows *B* and *D* are identical at 500 whilst the female ones differ markedly at 87 and 99.

When ordering rows or columns by size a subsidiary question is in which *direction* the figures ought to be ordered. People differ in their predilections here. Some like to have figures running from large on the left (as in Table 9), or from large at the top, whilst others prefer the opposite. With time-series, some like to have time progress from the left or the top of the tabulations, whilst others prefer to have the latest figures there. But these views are usually not held very strongly, nor do they appear to have any marked perceptual consequences when ordering *columns*. But for the *rows* of a table, showing the larger numbers above the smaller numbers (as in Table 11) helps because we are used to doing mental subtraction that way.

TABLE 11

Rows in Decreasing Order of Totals

GB	Unemployed (000's)		
	Total	Male	Female
1973 (D)	600	500	99
'70 (B)	580	500	87
'68 (A)	550	460	89
'66 (C)	330	260	71
Average	520	430	86

The combinations of consecutive numbers in Table 11 happen to be simple and hence fairly easy to subtract either way round. But in other cases the effect is more marked. If the first two numbers had been 640 and 580 instead of 600 and 580, the arrangement would matter more:

$$\begin{matrix}580\\640\end{matrix}\ \text{compared to}\ \begin{matrix}640\\580\end{matrix}$$

With less rounded numbers the effect is even stronger. For example, subtracting 583 from 637 is easier in the form

$$\begin{matrix}637\\583\end{matrix}\ \text{than}\ \begin{matrix}583\\637\end{matrix}$$

Facilitating such mental arithmetic is important when one is scanning large sets of data.

Rule 5: Spacing and Layout

Table 12 illustrates a form of table layout widely used in typed reports and prestigious printed documents. The rows are given in double or triple spacing and the columns are spread right across the page.

TABLE 12

Widely Spaced Figures

		Unemployed	
	Total	Male	Female
1973	600	500	99
'70	580	500	87
'68	550	460	89
'66	330	260	71

Such tables look nice but are counter-productive. The data are not easy to read because the eye has to travel too far. The rule is that figures which are meant to be compared should be placed close together. Single spacing is particularly effective in making the eye read down columns. But there are also need to be deliberate gaps to guide the eye *across* the table (e.g. between groups of 5 or so rows) as was illustrated in Tables 2 and 5.

Rule 6: Graphs versus Tables

Graphs are widely thought to be easier on the reader than tables of numbers, but this is only partially true. Graphs are of little use in communicating the *quantitative* aspects of the data, but they can highlight *qualitative* results (like that something has gone up, is a curve rather than a straight line, or is small rather than large). For example, the bar-chart of the unemployment data in Fig. 1 shows dramatically that

(i) Unemployment increased most from 1966 to 1968,
(ii) Female unemployed were far less than male,

but these are *qualitative* features of the data which can also be conveyed quite well verbally, as in (i) and (ii). But a graph can make the points more "graphic", and hence graphs can be very useful at the beginning or end of an analysis.

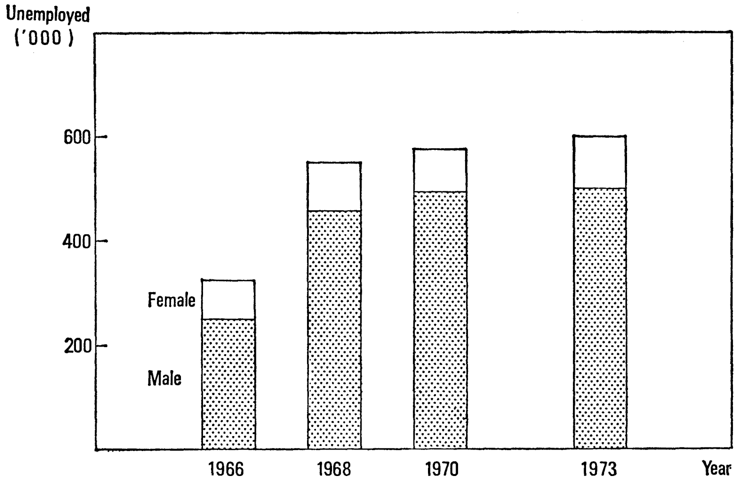

FIG. 1. A Bar-chart of the Unemployment Data.

However, graphs are of little if any use for quantitative detail. In Fig. 1 the size of the increase from 1966 to 1968 is not obvious at a glance (one has to project the blocks onto the vertical scale and interpolate). Nor is it clear just how small a proportion the female unemployed were, nor whether this proportion went up or down over the years, let alone by how much. This quantitative failure of graphs had led to some of the numbers often being shown as well, as illustrated in Fig. 2 (though this is not done in *Facts in Focus*, say).

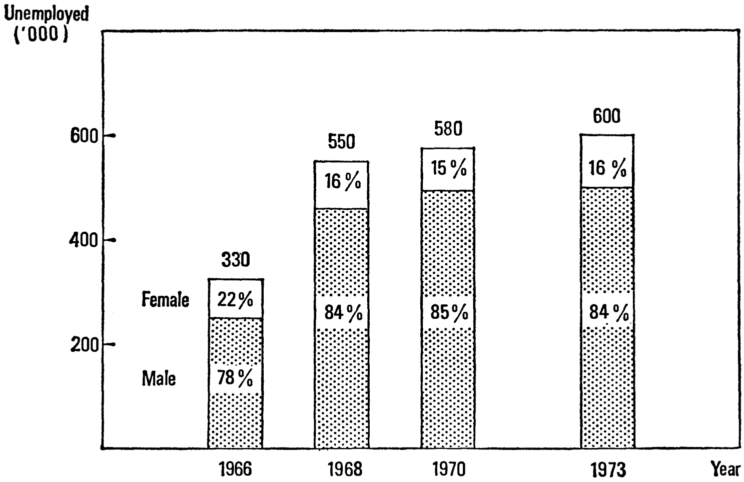

FIG. 2. With Numbers added.

One then mostly looks at the numbers (e.g. to see that the proportion of female unemployed actually went *down*) rather than at the graph, so that graphs with numbers inserted are often little more than badly laid-out tables. Arithmetical manipulation of the readings is made difficult rather than easy (e.g. taking averages, differences, ratios, or deviations from an average or a trend-line). Hence well-designed or "graphic" tables are better than graphs for any detailed numerical analysis, especially with extensive ranges of data.

3. THE LITERATURE

The literature relating to the successful presentation of statistical data seems to be sparse. The possible sources are psychological, typographical and statistical.

A good deal of work on information processing has been reported in psychology (e.g. Schroeder *et al.*, 1967; Lindsay and Norman, 1972). But little of it seems directly relevant to our narrow area of highly structured numerical tables and graphs—not even most of the work on pattern recognition and attention span. However, in writing about formal mathematical rather than *empirical* tables, Wright (1973) has noted that it is helpful to space related columns of figures closer than unrelated ones, and to arrange items so they can be scanned vertically rather than horizontally.

More fundamentally, Herbert Simon (1969) in discussing short-term memory in his Compton lectures noted not only that we can generally recall numbers of up to 7 or even 10 digits correctly if we are not interrupted in any way (i.e. not even by our own thoughts), but that there is also now experimental evidence that if we *are* interrupted by any task (however simple) the number of digits we retain in our short-term memory generally drops to *two*. (I am greatly indebted to Professor David Chambers for drawing my attention to this reference recently.) This would explain our need in Rule 1 to round figures to two significant (or variable) digits if we are to be able to perform mental arithmetic with them, i.e. to keep the figures accessible for immediate recall whilst being "interrupted" though having to relate one figure to another. (More than two significant digits being retained across an interruption can usually be explained parsimoniously either by our having recoded the information into two larger "chunks", or by having taken enough time—about 5 seconds per chunk—to fixate the information in our *long-term* memory.)

The study of *typography* (e.g. Spencer, 1969) has centred on the legibility of type-faces and sizes, on page design and problems of reduction and degradation, rather than the interrelation of different aspects of a numerical table. Yet some precepts apply, like the well-established typographical rule that strings of capitals are relatively difficult to read (BEING ALL OF THE SAME HEIGHT). This is often ignored for headings and captions in statistical tables, especially by many manufacturers of peripheral computer equipment.

As for statistics, our *Society's* original objects were to procure, arrange and publish facts, centering as far as possible on those which could be stated numerically and arranged in tables (cf. RSS, 1974). But there has been little discussion of what makes tables of numbers easier for the reader or even the analyst himself to understand and use.

Some statistical writers have stressed the importance of limiting the number of digits. But they almost invariably continued to use unnecessarily large numbers themselves. Giffen (1913), for example, greatly stressed rounding but used up to nine digits himself—"to the nearest acre" (I am indebted to Professor Bill Kruskal for the reference). Golde (1966) referred to a loss in accuracy of "only 3·41%" when dropping the third digit in a certain number.

Professor Ray Bauer has a splendid diatribe on digits (in Buzzell *et al.*, 1969, Chapter 5):

> "The data should not pretend to be more than they are. One of the most misleading practices indulged in by pretentious researchers is to present complex tables with the percentages carried out to the third, or even fourth, significant figures. Sometimes this is done because a researcher is lazy and he does not want to round out the figures which come out of the computer. Sometimes he is afraid to round his percentages because they won't add to precisely 100 per cent. Other times he actually believes that the third figure is important: this is virtually never true."

Yet nobody has taken much notice of this because Bauer has still only stressed the pointlessness and lack of precision of the later digits, rather than the *positive* advantages of eliminating them—that we can see, manipulate, and communicate two-digit numbers much better.

Graphics are currently attracting a good deal of attention (see Beniger and Robyn, 1976, for a bibliography), but there seems to be nothing new to help in communicating quantitative information effectively. The *Council for Social Graphics* in Washington, DC, has recently been assessing people's "graphicacy" (how well one can read graphs), but their target audiences are mainly the innumerate or inexperienced (e.g. school-children) rather than professional or regular users of numerical information.

On the whole, the presentation of numerical data to facilitate their use has been a relatively neglected area. Perhaps people have not realized how unnecessarily incomprehensible their supposedly competent tables usually are.

4. SPECIFIC OBJECTIONS

Whilst using and teaching the present approach to data presentation during the last few years (e.g. Ehrenberg, 1975, Part I) a number of objections and problems have been raised. I now discuss these, taking the six rules in turn.

Rounding (*Rule 1*)

Rounding is the rule which tends to raise the most (or the most heated) objections. It is the only rule where information is actually discarded and many people seem to feel that observed data should be treated as sacrosanct, e.g. that if some clerk or computer happens to have recorded the data to five digits, that is how the data should perhaps always remain.

Yet rounding is readily accepted in graphical presentations and also in fitting mathematical models to the data. Nor would most people object to reducing statistical data to *three or four* significant digits. But they often feel that rounding to only two significant or effective digits is overdoing it. Unfortunately such rounding is necessary to facilitate mental arithmetic. For example, few of us can divide 17·9% into 35·2% in our heads (most percentages are reported in effect as "per mille" rather than as "per centum"). Of several thousand people asked to do this over the years only two US mathematicians at Purdue have claimed success. But they got different answers, so at least one of them was wrong. In contrast, dividing 18% into 35% is obviously about 2. Thus two digits are better.

The finding noted earlier that our short-term, quick-access memory is limited to manipulating two-digit numbers applies even to a simple arithmetical task like scanning a column of more or less equal figures against their average, and even when all the figures remain in front of us as in the following examples:

549·2	550	549
582·2	580	582
601·9	600	602
621·3	620	621
734·6	730	735
617·9	620	618

With *four* digits in the first column we can hardly recall the average of 617·9 as we run down the column from one entry to the next (unless we go in for mental rounding). With *two* significant digits there is no problem—we can check the trend whilst readily holding the average of 620 in our head. With *three* digits there is still quite a problem of recalling the average as we go down the column, although we can cope better by transforming the 618 into two "chunks" like 6 18 (six eighteen) or into a "sing-song" *six*-one-eight. (With three digits we could also take the time to transfer the average of 618 to our longer-term memory, but this would hardly work when scanning a table with many such columns.)

Instead of asking for any particular data "Can we possibly round them to two digits?" we need to check only whether there is some specific reason why we should *not* do so. One can think of exceptional situations, e.g. where large multipliers might be involved, where rounding errors can build up as in compound interest calculations, or where we are analysing deviations from a model. (One would then often keep a third digit for working purposes in calculating averages, slope-coefficients or other parameters to avoid minor problems with rounding errors, but round again when actually reporting the results.) Again, one can round the sequence 186, 97, 93 to 190, 97, 93 without effective loss, but with our decimal system one might not round 106, 97, 93 to 110, 97, 93 because the error of rounding the 106 is large compared with the range of 13. But such cases are exceptions.

One safeguard is that no information need be *completely* lost by rounding. The two-digit rule is a guideline for statistical working tables and the final presentation of results, not necessarily for basic data records. One can put the more precise data in an appendix or, better still, in a filing cabinet or other data bank just in case somebody should want them sometime.

The more precise data are however unlikely ever to be used. When would the earlier unemployment figures really be needed to the nearest 100 people as in Table 13, rather than

TABLE 13

Unemployment in Great Britain

(Table 6 repeated)

	1966	1968	1970	1973
Total unemployed (thousands)	330.9	549.4	582.2	597.9
Males	259.6	460.7	495.3	499.4
Females	71.3	88.8	86.9	98.5

Source: Facts in Focus

rounded to the nearest 1,000 ot 10,000? The degree of precision that might be required can be judged against the range of the observed variation that has to be explained, the size of the residuals in any formal model-building, and the likely requirements of any deeper analysis (not to mention any inherent inaccuracy of the data).

For example, in Table 13 the average error in rounding the female unemployed to two digits would be 300. This is trivial when assessed against the overall increase of almost 30,000 in the female figures from 1966 to 1973, and the contrary drop of 2,000 from 1968 to 1970. The rounding errors are also trivial when compared with the residuals from a mathematical model like $F = 0{\cdot}1M + 43$. This represents the relationship between female and male unemployed quite well, the correlation being ·85; but the residuals average at 3,000.

Finally, the rounding errors are trivial in the context of any fuller analysis of unemployment. This would never mean digging deeper into the eight selected readings in Table 13. Instead, it would necessitate taking account of vastly *more* data: for other years, different regions of the country, different industries, different age-groups (treating school-leavers and students separately), plus figures for employment, reported vacancies, inflation, investment, stock-piling, dumping, Gross National Product, the money supply, birth rates, immigration, mechanization, business cycles, world trade, unemployment in other countries, and so on, as well as intensive comparisons of figures based on different definitions and measuring procedures (i.e. the whole question of the "quality" of the data).

Each monthly issue of the *Department of Employment Gazette* gives about 8,000 two-to-four-digit numbers on UK unemployment. They may be mostly the same as in the previous month, but the need to see the wood for the trees becomes even more urgent than with the eight figures in Table 13. Hoping to explain variation to the third digit (less than 1%) becomes even more absurd. People who object to rounding to two effective digits because they feel that "there may be something there" can have had no experience of *successfully* analysing and understanding extensive empirical data.

Averages (Rule 2)

Averages are not always directly descriptive or "typical" of the readings in question. The inclusion of such averages in a table is often criticized as useless or even misleading. But they can still provide a visual focus for inspecting the data, and also useful parameters when comparing different distributions of the same type.

To illustrate I use the slightly more complex example shown in Table 14 (a re-presentation of the earlier TV data in Table 5). It says how many adults who "really liked to watch" one

TABLE 14

Duplication Analysis: Percentage of Adults who like one TV Programme who also like Another

	% of adults who really like to watch										
	Wrld of Sprt	Mtch of the Day	Grnd stnd	Prof Box ing	Rgby Spec ial	24 Hrs	Pan ora ma	Ths Wk	To day	Line Up	Av. excl. 100's
	39	38	35	32	16	31	31	27	24	9	28
% who also really like to watch	%	%	%	%	%	%	%	%	%	%	%
World of Sport	(100)	75	80	75	74	49	52	51	47	50	61
Match of the Day	73	(100)	77	72	75	47	48	45	41	46	58
Grandstand	72	71	(100)	68	75	47	48	45	41	47	57
Prof. Boxing	61	60	62	(100)	65	41	49	40	38	44	51
Rugby Special	28	30	32	31	(100)	21	23	19	18	25	25
24 Hours	39	38	40	39	44	(100)	68	61	50	66	49
Panorama	41	38	42	42	42	67	(100)	58	47	59	48
This Week	34	31	34	34	33	53	50	(100)	48	53	41
Today	29	26	28	28	29	39	37	43	(100)	44	34
Line-Up	12	11	12	13	16	20	18	19	17	(100)	15
Average (excl. 100's)	43	42	45	45	50	42	44	42	39	48	44

TV programme also really liked to watch another. Thus 39% of adults really liked to watch World of Sport (WoS), and of these 73% really liked to watch Match of the Day, 72% Grandstand, and so on. (The earlier correlations of Table 5 can be calculated from the data in Table 14: the new table involves no change in empirical content but only in language or model.) Writing in the row and column averages by hand, as one would do with computer output or other working tables, we quickly see that none of the averages represent the data at all well. They are not typical or modal of the individual figures in the corresponding row or column.

Nonetheless, the averages provide a useful focus. By fixing first on the average for each row and ignoring the 100's in the diagonal, we see four above-average figures followed by five below-average figures in each of the first five rows, and the opposite pattern in the last five rows (five *below*-average figures followed by four *above*-average ones). Blocking out the row averages on the right by hand dramatizes their effectiveness in providing this visual focus.

Correspondingly the column averages help us to see that there is no such simple pattern in the columns. But we also note that these column averages are all virtually equal. (Focusing first on the overall average of 44 helps to see this.)

This suggests checking whether the figures in each column follow the same form of distribution (having the same means). Starting with the overall average of 44 and the column of row averages on the right, we see a "High–Low–About average–Low" pattern. Inspecting the individual columns in the body of the table against their averages (with the 100's in parentheses so that the eye can ignore them more easily) shows the same pattern in the first five columns, but a somewhat different pattern for the last five: "About average–Low–High–Low". We now see that the virtual equality of the column averages across all 10 columns was a coincidence.

Thus used critically as a working-tool, averages are of value in getting to know the data even when they do not summarize the data as simple "typical" figures. Calculating marginal averages for a table generally remains a helpful first step to see which way the data goes. But they need not always be retained in the final presentation.

Using Columns for Figures to be Compared (*Rule* 3)

An objection often raised about interchanging rows and columns is the difficulty of fitting long row captions into column headings. This can usually be done by abbreviation, by spreading the headings over two or three lines (as in Table 14), and by relegating detail to footnotes. (Some people are said to object to having to look at the footnotes to a table, but they are probably not the sort of people who would get much out of a complex-looking table anyway.)

It is important to get one's priorities right. The design of a table must be determined by the data that are being displayed not by the logic of the row and column labels. Yet tables often are designed without taking any account of the data itself. For example, some recent proposals by the *Business Statistics Office* (Fessey, 1976) for interchanging rows and columns in its regular *Business Monitor* series were judged unconvincing, but the dummy tables that were prepared contained no numbers. In practice, if a table layout is designed without reference to the data, that is what the final table will probably look like.

To illustrate, Table 15 greatly clarifies the patterns in Table 14 by interchanging the rows and columns. (Omitting the 100's, inserting column averages, and appropriate spacing also seem to help.)

TABLE 15

Duplication Analysis: Rows and Columns Interchanged and Sub-group Averages

Adults who Really like to Watch		% who also Really like to Watch											
		WoS	MoD	GrS	PrB	RgS	AV.	24H	Pan	ThW	Tod	LnU	AV.
World of Sport	%		73	72	61	28		39	41	34	29	12	
Match of the Day	%	75		71	60	30		38	38	31	26	11	
Grandstand	%	80	77		62	32		40	42	34	28	12	
Prof. Boxing	%	75	72	68		31		39	42	34	28	13	
Rugby Special	%	74	75	75	65			44	42	33	29	16	
AVERAGE		76	74	71	62	30	63	40	42	33	28	13	31
24 Hours	%	49	47	47	41	21			67	53	39	20	
Panorama	%	52	48	48	49	23		68		50	37	18	
This Week	%	51	45	45	40	19		61	58		43	19	
Today	%	47	41	41	38	18		50	47	48		17	
Line-Up	%	50	46	47	44	25		66	59	53	44		
AVERAGE		50	45	46	42	21	41	61	58	51	41	18	46
ALL ADULTS	%	39	38	35	32	16	32	31	31	27	24	9	24
							63/32 = 2.0 41/32 = 1.3						31/24 = 1.3 46/24 = 1.9

We now see that the figures in each column tend to be similar within each of the two programme categories, and hence close to the averages shown at the bottom of each block. For example, World of Sport (WoS) in the first column is liked by about 76% of those who liked one of the other Sports programmes (the individual figures varying between 74 and 80%), and by about 50% of those who liked one of the Current Affairs programmes.

Since World of Sport (WoS) is liked by about 39% of *all* adults (as shown in the last row of the table), we can see now that it was about twice as popular amongst those who liked another Sports programme, and about 1·3 times as popular amongst those who liked a Current Affairs programme, than amongst the population as a whole (76/39 and 50/39).

The same pattern holds for the other Sports programmes shown in the next four columns (MoD to RgS). From the averages we estimate "duplication-ratios" of $63/32 = 2{\cdot}0$ within the Sports cluster and $41/32 = 1{\cdot}3$ between the Sports and Current Affairs programmes. The pattern also applies to the Current Affairs programmes in the last five columns of the table. The duplication-ratios here are again 1·3 for Current Affairs versus Sports and 1·9 within the Current Affairs cluster itself.

Table 15 may appear more complex than the earlier correlation matrix in Table 4, but it provides much more insight into the data. It is an instance of the so-called "duplication law", which says here that the percentage of people who like programme P amongst those who like programme Q is directly proportional to the percentage of the whole population who like programme P, the proportionality-factor or "duplication-ratio" being a constant for a particular grouping of programmes. This form of relationship has already been found to occur in a wide range of choice situations (e.g. Ehrenberg, 1972; Goodhardt *et al.*, 1975) and also has strong theoretical backing (Goodhardt, 1966; Goodhardt *et al.*, 1977).

Table 15 may not seem obvious at a glance if one is seeing it for the first time. But it brings out the duplication pattern clearly enough for anyone already knowledgeable in the area, and in particular for anyone involved in using the model in question. This typically involves examining and communicating literally hundreds of thousands of such figures over the years. A form of layout meeting the weak criterion for a good table which allows one to scan and grasp extensive data then becomes essential.

A common query about changing rows into columns is whether all users of the table will want to compare the figures in the columns rather than those in the rows. In practice they must always do *both*. But the main pattern in the data should be looked at first and hence in columns because that is easier. Then, having seen the main pattern, one can look at the rows and at any row-and-column interactions. Again, with a table of time-series one usually looks first at each series on its own (which is easier in columns) and only then correlates the different series.

Ordering by Size (*Rule 4*)

Ordering the rows or columns of a table by some measure of size raises two problems. One is that different measures of size can be used, resulting in different possible orders. The criticism is that readers (especially *other* readers) might be misled by the particular order chosen.

For example, in Table 2 ordering the rows of the shipping data by the numbers of vessels led to the sequence Cargo/Tanker/Passenger. Ordering by the *1973 tonnages* would have led to the sequence Tanker/Cargo/Passenger. But users of a table do not have to accept the chosen ordering as sacrosanct. One order will show up the conflict with another, and *some* visible ordering is always better than none (as in the original Table 1). In Table 2 anyone can see that the 1973 tonnages *were* out of step. (Were you, the reader, in any way misled, or would you only be worrying about the possible effect on *others*?)

The second problem arises when there are many different tables with the same basic format and straight application of the rule would lead to different orders for different tables. In such cases the same order should be used in every table.

A good example occurs with tables giving various social and economic statistics for different countries, or for different regions or towns within the same country. A useful common order for all such tables might be population size. This provides an instant visual rank correlation between the absolute and *per capita* rates for each variable.

Such an ordering is often criticized as departing from the alphabetical listing of the countries, making it difficult to look up the result for a particular country. But a statistical table is not a telephone directory. To use an isolated figure one must understand the context of the surrounding ones and see the general pattern of the data. If there are many such tables and they are large, an alphabetical key will be worth giving. In any case it is probably easier to find an isolated name in a non-alphabetical listing than to interpret an isolated number from an unstructured table.

Spacing and Layout (*Rule* 5)

The basic guidelines for table layout may seem simple: single spacing and occasional deliberate gaps to guide the eye; columns spaced evenly and close together, and occasional horizontal and vertical rules to mark major divisions. But many typists, printers and computers are programmed differently. Double spacing in tables is common, as are columns spaced unevenly according to the width of the headings, and occasional irregular gaps between single-spaced rows because some row captions ran to two lines. One needs not only good typists or printers, but also thoughtful *control* of these facilities.

The traditional printers' embargo on vertical rules is widely accepted but can be side-stepped increasingly through modern off-set and duplicating methods (as illustrated by the tables in this paper). However, ruling off every column routinely, as in Tables 1 and 6, is counter-productive. In contrast, a few well-chosen rules can have a startling effect, as the reader may see by drawing a vertical and a horizontal rule by hand to separate the two sets of five programmes in Table 14.

The niceties of good spacing and layout are many and complex. Some specific ones are illustrated by Table 15. A good general example (not illustrated here) occurs with tables where the data come in pairs of related figures, e.g. "observed" and "theoretical" values, or "last year" and "this year": The use of closely spaced *pairs* of columns is then very effective. More work is needed to make the possible variety of such procedures more explicit.

Graphs versus Tables (*Rule 6*)

A claim often made in supposed contradiction to Rule 6 is that people (especially "other people") find graphs easier to look at than tables, e.g. Fig. 1 rather than Table 13. They probably do, but this is misleading because it does not reflect that people usually extract and retain little information from a graph.

It is no use merely saying "Here are some complex data—let's put them on a graph". What does one learn even from simple graphs (like Figs 1 or 2)? They are supposed to be easy, but suppose one looks away, what does one remember? Some shapes or qualitative features perhaps, but seldom any quantities. In any case, Fig. 3 reminds us that most graphs do not show simple patterns or dominant numbers which can readily be grasped. Success in graphics seems to be judged in producer rather than consumer terms: by how much information one can get on to a graph (or how easily), rather than by how much any reader can get off again (or how easily).

5. Some Problems of Implementation

Perhaps the most frequent comment on the rules of data presentation discussed here is that they are "mere common sense", with the insidious implication that real statisticians do not have to bother with them. But at this stage the rules reflect neither common knowledge nor common practice. In that sense they are decidedly uncommon.

The rules seem obvious only once they have been stated. No particular skills or knowledge are then required to assess them—only one's common senses, i.e. whether they feel, look and sound right. The procedures therefore lack the technical mystique of a Durbin–Watson test or non-Euclidian space that tends to guarantee a certain instant popularity.

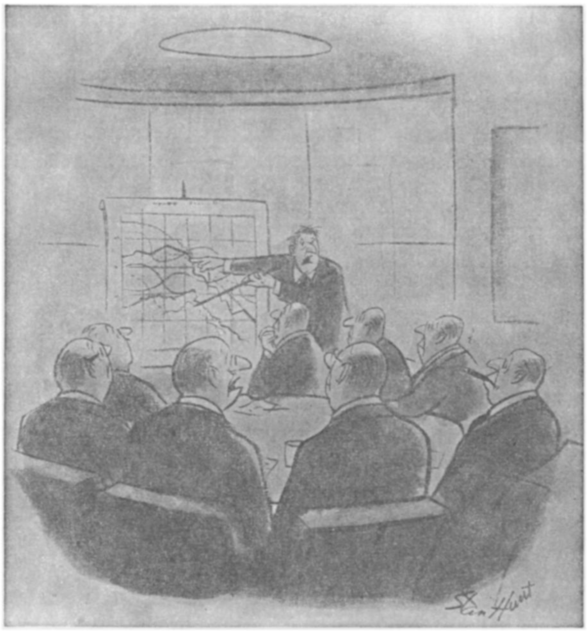

FIG. 3. "McBelding certainly has a gift for making cold statistics come to life."
(Drawing by Stan Hunt; © 1975 The New Yorker Magazine, Inc.)

Implementation of such presentation rules is often thought to depend on two factors:

(a) the data,

(b) one's purpose.

But in this final section I show that the rules generally transcend these factors.

The Nature of the Data

The data to be presented can be classified along several dimensions, namely whether they are

Simple or complex

New or repetitive

Reliable or uncertain

To tell a particular story, or presented only "for the record".

I now consider these four aspects.

Complexity. The tables illustrated in this paper have been mainly small and simple, for reasons of space. Nonetheless, Table 1 was in effect a *three-way* table, and the Table 2 version made it much easier to correlate the three different variables (i.e. the numbers of vessels and the two types of tonnages). Table 15 was another fairly complex example. More generally, my experience is that the six rules also apply to still more complex or extensive data, where they are in fact needed even more. Some examples have been considered elsewhere (e.g. Chakrapani and Ehrenberg, 1976; Ehrenberg, 1975, 1976b, c; Ehrenberg and Goodhardt, 1977).

New or Old? The analytic situations discussed in the statistical literature mostly follow the "exploratory" approach, where a new data-set is to be analysed as if that kind of data were being looked at for the first time ever. Good data presentation then takes a fair amount of work, since one seldom gets a new table completely right the first time round. But the real problem with new data is not that of presenting it well, but of having first to understand it. Luckily this is (or should be) relatively rare. Most situations faced by professional or frequent users of data are *repetitive*, in that they have already seen a good deal of similar data before and therefore know their probable structure (Ehrenberg, 1976a).

To take Table 15 as an example, one's usual task is not to discover the basic duplication pattern for the first time (that can strictly happen at most once, and did so about 10 years ago in this instance—Goodhardt, 1966; Ehrenberg and Twyman, 1967). Instead, one needs to assess these particular data against one's prior knowledge of the duplication law, to establish and understand any apparent anomalies, to communicate the results to others, and to *use* the results (e.g. for theoretical model-building, practical decision-making, prediction or control). In such well-understood repetitive situations the rules of data presentation can be applied routinely. Their use becomes highly efficient.

The Quality of the Data. It is often said that the "quality" of the data should affect how they are presented. This presumably refers to outliers, sampling errors and basic measurement problems. But given that certain numbers are to be reported at all, it is better to present them clearly rather than obscurely, so the rules still apply. Good data presentation makes outliers and misprints stand out: Twyman's Law—that any reading which looks interesting or different is probably wrong—can only be applied if we first see that a reading *is* out of step.

Sampling errors occur if sample sizes are small. Most modern statisticians are of course highly trained to deal with this (if with nothing else) and in a paper before this Society the existence of such issues can, I hope, largely be taken as read. I only add that some analysts' habit of attaching a standard error to every reading in the body of a table is both visually obnoxious and statistically naive. If standard errors or other devices of statistical inference need to be explicitly quoted, this should be done either in a separate display, or in footnotes, or in the text.

The basic problem with data is what the variables in question actually measure. In our unemployment example the figures are for registered unemployed (with a good deal of small print in the definitions), and do not properly represent "unemployment", whatever that may be. Female unemployed tend, for example, to be markedly under-represented, especially at times of high general unemployment. Learning to understand what one's variables mean usually depends on comparing different *types* of measurement (e.g. the official figures of registered unemployed with sample survey data of supposedly "actual" unemployed). This is usually a complex task and the need for effective data presentation remains. Even if our measurements are known to be biased that is no reason for leaving the numerical results obscure.

For the Record or . . . ? Three main types of empirically-based data tables can be distinguished:

working tables, for the use of the analyst and his immediate colleagues, with no wider communication in mind;

the final presentation to a more or less specific audience, to support or illustrate some specific conclusion or findings;

tables set out "for the record" (as in official statistics) in case someone wants to use the data.

In the first two cases the structure of the data needs to be apparent both to the analyst himself and to others. Hence the rules of this paper apply. With data presented "for the record" however it is sometimes argued that the data will contain so many different stories, for different kinds of uses and users, that its presentation must vary accordingly. But few real instances have been quoted and this conclusion seems to be the exception rather than the rule. In any case, it does not follow that the data must be presented to tell *no* story, as is so often the case.

One frequently-cited illustration of the use of such data is the politician who wants to quote a single number in some speech (e.g. the number of doctors in his home town) and who appears profoundly disinterested in gerneral patterns and laws, e.g. of the incidence of doctors in different towns and places. But this is wrong. No meaningful use can be made of an isolated number. What good is it knowing that there are 57 doctors in the town without some idea of whether this 57 is high, low or normal—as *many* as 57, or *only* 57, or what? Without aiming to turn politicians or other users of the odd statistic into fully-fledged statisticians, we need not pander to the mindless misuse of statistical data. In any case, most occasional users would be happy to see, or be told, that 57 *is* high (or low) on a *per capita* basis and after allowing for the local age-distribution, or whatever. Bringing out general patterns in official statistics can do little harm, and may do much good.

Applying the present rules to official statistics will take time and effort, but this will be more than balanced by savings in paper and printing costs, not to mention the fuller and better use that will be made of the data. Yet the practical problems of implementing these rules of data presentation must not be under-estimated. There can be very substantial set-up and upset costs in changing from traditional practices. The methodology is still under-developed. People are not only unfamiliar with the techniques, but also with the fundamental notion that most tables can be improved to communicate better.

The Purpose of the Analysis

It is commonly suggested that one should formulate one's purpose explicitly before tackling the analysis or presentation of some data. Sir Maurice Kendall (1969), for example, has said that if he had some data and wanted in some sense to describe their structure, he would do nothing except store the original observations until someone could specify the object of the exercise. This reads like a denial of the purpose of ordinary science, i.e. to understand a system, and to do so before attempting to make practical applications.

In a recent book review in this *Journal*, Pridmore (1976) extended this view to the novice, complaining that apparently he had not been told to ask himself such questions as "What am I going to do with the results? Why am I doing this? Why are the data wanted?" before starting his analysis. But I doubt if a novice could answer these questions, or should be expected to do so. More generally, I feel that the emphasis on establishing a purpose prior to first studying one's data is unrealistic.

I am not saying that one should not have a purpose, but only that one cannot formulate a realistic purpose if one knows nothing yet about one's data (i.e. no prior knowledge and also no peeking). But, as mentioned earlier, most analyses are of a repetitive kind, so that one usually has prior experience of other, similar data to influence how one approaches the new data.

Formulating a purpose without knowledge of the data would in any case mean that the analyst's uninformed perception of his purpose would determine how he analyses and presents his data. This would be very subjective. The contrary view, which I support, is that the detailed analysis and presentation of the data should be dominated by the facts. One's personal objectives or purpose should mainly determine how one then *uses* the results.

The kind of presentation rules discussed in this paper are themselves often regarded as subjective because the presentation is to be deliberately influenced by one's knowledge of the data. This is anathema to some statisticians, due to a misunderstanding of certain technical problems in statistical inference for small samples.† But if the presentation rules are made explicit, any reasonably experienced person can follow them and obtain more or less the same results, which is the touchstone for achieving objectivity.

† It is, for example, often argued that faced with more than two sample means, one should not pick out the smallest and largest and use a t-test to assess whether they differ significantly. But there is nothing wrong with picking out the largest and smallest means in one's data, and then doing an appropriate test of significance. All that is wrong is using the ordinary *two-sample* t-test in a situation for which it was not designed.

The presentation of data of course involves judgement. But that is true of any form of analysis (e.g. in choosing one's variables, measurement techniques, conditions of observation, sample sizes, cleaning-up procedures, analytic techniques and models, significance levels, etc.). Judgement is largely what distinguishes a good analyst from a lesser one. But this judgement must have knowledge, experience and techniques to bite on, and subsequently be replicable by others. Thus the main aim in this paper has been to discuss rules or guidelines for data presentation which can be applied more or less routinely, with judgement.

ACKNOWLEDGEMENT

I am greatly indebted to the large number of helpful comments I have received on drafts of this paper, in particular from Mr T. P. Barwise, Dr S. Hodges, Mrs H. Lewis, Professor W. H. Kruskal, Dr S. Rosenbaum and Mrs J. Sheppard.

REFERENCES

BENIGER, J. R. and ROBYN, D. L. (1976). The history and future of graphics in statistics. *Social Statistics Proceedings*, 192–197. Washington, DC: The American Statistical Association.

BUZZELL, R. D., COX, D. F. and BROWN, R. V. (1969). *Marketing Research and Information Systems: Text and Cases*. New York: McGraw-Hill.

CENTRAL STATISTICAL OFFICE (1974). *Facts in Focus* (2nd edn). London: Penguin.

CHAKRAPANI, T. K. and EHRENBERG, A. S. C. (1976). Numerical information processing. (Working Paper.)

EHRENBERG, A. S. C. (1972). *Repeat-buying: Theory and Applications*. Amsterdam: North Holland; New York: American Elsevier.

—— (1975). *Data Reduction*. London and New York: Wiley.

—— (1976a). We must preach what is practised. *The Statistician*, **25**, 195–208.

—— (1976b). Communicating market data. *J. of Adv. Research*, **16**, 27–32.

—— (1976c). Annual Reports don't have to be obscure. *J. of Accountancy* (August), 88–92.

EHRENBERG, A. S. C. and CHAKRAPANI, T. K. (1977). Factor analysis or BGA? *Appl. Statist.* (submitted for publication).

EHRENBERG, A. S. C. and GOODHARDT, G. J. (1977). Developing and using a model. In *Proceedings of the 20th Annual Conference of the Market Research Society*. London: The Market Research Society.

EHRENBERG, A. S. C. and TWYMAN, W. A. (1967). On measuring television audiences. *J. R. Statist. Soc.* A, **130**, 1–59.

FESSEY, M. C. (1976). Private communication.

GIFFEN, SIR ROBERT (1913). *Statistics*. (H. Higgs and G. U. Yule, eds.) London: Macmillan.

GOLDE, R. A. (1966). *Thinking with Figures in Business*. Reading, Mass.: Addison-Wesley.

GOODHARDT, G. J. (1966). The constant in duplicated television viewing. *Nature*, **212**, 1616.

GOODHARDT, G. J., CHATFIELD, C. and EHRENBERG, A. S. C. (1977). *The Dirichlet choice model*. (In preparation.)

GOODHARDT, G. J., EHRENBERG, A. S. C. and COLLINS, M. A. (1975). *The Television Audience: Patterns of Viewing*. Farnborough: Saxon House; Lexington, Mass.: Lexington Books (D. C. Heath).

KENDALL, SIR MAURICE (1969). Private communication.

LINDSAY, P. H. and NORMAN, D. A. (1972). *Human Information Processing*. New York: Academic Press.

PRIDMORE, W. A. (1976). Review of "Data Reduction". *J. R. Statist. Soc.* A, **139**, 268–269.

RSS (1974). *Regulations and List of Fellows*. London: The Royal Statistical Society.

SAMUELSON, P. (1976). *Economics* (10th edn). New York and London: McGraw-Hill.

SHROEDER, H. M., DRIVER, M. J. and STREUFORT, S. (1967). *Human Information Processing*. New York: Holt, Rinehart & Winston.

SIMON, H. A. (1969). *The Sciences of the Artificial*. Cambridge, Mass.: M.I.T. Press.

SPENCER, H. (1969). *The Visible Word*. London: Lund Humphries.

WRIGHT, P. (1973). Understanding tabular displays. *Visible Language*, **7**, 351–359.

[The Discussion of Professor Ehrenberg's paper appears on pp. 307–323.]