

## Copyright Notice

Staff and students of the University of Edinburgh are reminded that copyright subsists in this extract and the work from which it was taken. This Digital Copy has been made under the terms of a CLA licence which allows you to:

- access and download a copy
- print out a copy

This Digital Copy and any digital or printed copy supplied to or made by you under the terms of this Licence are for use in connection with this Course of Study and should not be downloaded or printed by anyone other than a student enrolled on the named course or the course tutor(s). You may retain such copies after the end of the course, but strictly for your own personal use.

All copies (including electronic copies) shall include this Copyright notice and shall be destroyed and/or deleted if and when required by the University of Edinburgh.

Except as provided for by copyright law, no further copying, storage or distribution (including by e-mail) is permitted without the consent of the copyright holder.

The author (which term includes artists and other visual creators) has moral rights in the work and neither staff nor students may cause, or permit, the distortion, mutilation or other modification of the work, or any other derogatory treatment of it, which would be prejudicial to the honour or reputation of the author.

**Course of Study: Research Methods**

**Name of Designated Person authorising scanning:**

**Title: “Quantitative Data Analysis: a user-friendly guide for social scientists”. Chapter 8, pp. 94-112 “Creating Categories”**

**Name of Author: Ian Dey**

**Name of Publisher: Routledge, London**

**ISBN/ISSN: 041505852x**

**Name of Visual Creator (as appropriate): n/a**

## Creating categories

'We think in generalities, we live in detail'

(Whitehead quoted in Peter 1982: 493)

We have to interpret our data in order to analyse it. But analysis can go beyond interpretation. We can try to create conceptual tools to classify and compare the important or essential features of the phenomena we are studying. This involves a process of abstracting from the immense detail and complexity of our data those features which are most salient for our purpose. For example, the chemist focuses not on 'water' but on H<sub>2</sub>O, stripping away the many connotations of the term to isolate those characteristics essential for analysis (Klein and Warren 1967). Abstraction is a means to greater clarity and precision in making comparisons. We can focus on the essential features of objects and the relations between them. However, it is important to remember what we are abstracting from. The Taoist Chuang Tzu wrote:

Fishing baskets are employed to catch fish; but when the fish are caught, the men forget the baskets; snares are employed to catch hares, but when the hares are got, men forget the snares. Words are employed to convey ideas, but when the ideas are grasped, men forget the words.

(Quoted in Capra 1983: 36)

Abstractions are powerful means for making comparisons, but we must also remember their origins and limitations.

In making comparisons, it is helpful to distinguish between two forms of relation between objects or events (cf. Sayer (1992: 88). On the one hand, we can identify 'substantial' relations of connection and interaction. When we laugh at a joke, for example, there is a substantial connection between the joke and our laughter. On the other hand, we can identify purely 'formal' relations of similarity and difference between things. Thus we can distinguish between jokes and laughter, as different types of phenomena. This type of comparison involves categorizing phenomena according to their similarities or differences. In this and the following two chapters, I

focus on how we can categorize qualitative data, before considering issues raised by substantive connections.

In this chapter, I shall look at the problems of generating categories for the analysis. In the following chapter I turn to the issues raised in assigning these categories to the data. Then in Chapter 10, the last dealing with categorizing data, I consider how we can 'split and splice' categories – in other words, subdivide or integrate categories as ways of refining or focusing our analysis. Naturally, we have to create categories before we can assign them, and we have to assign them before we can split or splice them. But despite this logical precedence, in practice we may find ourselves shifting backwards and forwards between these different aspects of categorizing data.

The very quality of qualitative data – its richness and specificity – makes for problems when we try to make comparisons between observations. For what are we comparing? There are no standard categories in terms of which to compare observations. Indeed, there are no clear boundaries as to what constitutes an observation.

To compare observations, we must be able to identify bits of data which can be related for the purposes of comparison. How can this be done? The answer is deceptively simple. In principle, we could organize the data by grouping like with like, so that any observation which seems similar to or related to others can be grouped with those observations. We can put all the bits of data which seem similar or related into separate piles, and then compare the bits within each pile. We may even want to divide up the items in a pile into separate 'sub-piles' if the data merits further differentiation. We can then compare observations within each pile or sub-pile, looking for interesting similarities or differences within the data. We can also make comparisons between the different piles or sub-piles, again looking for patterns or variations in the data.

However, this procedure begs two important questions. First, what is an observation? We referred above to 'bits' of data, but how are these bits to be identified or distinguished from the rest of the data? There must be some criterion or criteria which allow us to distinguish one bit or observation from another. Second, how can an observation be judged similar to or related to some other observations? Why put a bit of data into one pile, but not into another? Because they are alike, or related? But things are not just alike or related – they are alike or related in some respect or another. Although we may say that observations are alike or related without explaining why this is so, nevertheless there must be some respect or other in terms of which this judgement is made. If we distinguish between an employer and an employee, for example, we implicitly refer to a variety of social and economic features which characterize the difference between them. Distinctions are always conceptual as well as empirical – they reflect some criterion or criteria in terms of which observations are distinguished

and compared. In data analysis, we can try to make explicit (or at any rate, as explicit as possible) the conceptual criteria in terms of which distinctions are made between observations.

Grouping data in this way therefore involves developing a set of criteria in terms of which to distinguish observations as similar or related. Typically, this is done through the development of a set of categories, with each category expressing a criterion (or a set of criteria) for distinguishing some observations from others similar or related in some particular respect(s). The development of a set of categories allows the data to be organized through a variety of different distinctions. Data within each category can then be compared. If necessary, further distinctions can then be drawn within each category to allow for a more detailed comparison of data organised within a set of sub-categories. Conversely, data assigned to different categories can be compared and interrelated to produce a more encompassing analysis of the data. This process can continue until the analyst is satisfied that all relevant distinctions between observations have been drawn, and observations can be compared effectively in terms of an established category system.

In categorizing data, we are not simply bringing together observations which are similar or related. A comparison is already implied in the adoption of a particular category. The data is being classified as 'belonging' to a particular group and this already implies a comparison between this data and other observations which do not 'belong' to the category in question. Categorizing involves differentiating between the included and excluded observations. The process of categorization may seem akin to naming observations, and in the literature is sometimes referred to as 'labelling' data and categories are sometimes referred to as 'labels'. However, this may be confusing and misleading if we think of 'naming' in terms of proper names rather than classes of objects. We may name a baby 'Rebecca' and this name distinguishes her from her siblings 'Katie' and 'Paul'. However, naming here simply aims to provide a label sufficiently unique (for practical purposes) to designate an individual person. It is not a label which stands for a class of objects – it is not a concept. There is a role for labelling in this sense of designating names for unique bits of data, as we shall see later. But where we are 'labelling' in order to group data, it may be less confusing to use the term 'categorization' for this process of making and assigning distinctions within the data.

Creating categories is both a conceptual and empirical challenge; categories must be 'grounded' conceptually and empirically. That means they must relate to an appropriate analytic context, and be rooted in relevant empirical material. Categories which seem fine 'in theory' are no good if they do not fit the data. Categories which do fit the data are no good if they cannot relate to a wider conceptual context. We could say that categories must have two aspects, an internal aspect – they must be meaningful in

relation to the data – and an external aspect – they must be meaningful in relation to the other categories.

It is not by accident, therefore, that we refer to creating categories in the plural. A category cannot be created in isolation from the other categories we want to use in the analysis. When we devise a category, we are making decisions about how to organize the data in ways which are useful for the analysis – and we have to take some account of how this category will 'fit' into this wider analytic context. It is usual, therefore, to think in terms of generating a set or list of categories through which to organize comparisons between observations. The formal relations between the categories are important in defining the relation between any particular category and the data. In generating categories, therefore, we have to think systematically and logically as well as creatively.

How does one begin to generate a category set? This is a question which researchers have had some difficulty in answering, in part because there is no single or simple answer. Obviously the methods of generating a category set will reflect the type of data being analysed, and also the aims, inclinations, knowledge and theoretical sophistication of the researcher. The theoretically-inclined participant observer with voluminous field-notes and the policy evaluator with a set of open-ended interview responses to analyse may have quite different starting points and quite different resources upon which to call. The theorist may be able to draw upon existing theoretical perspectives. By contrast, the policy evaluator is more than likely to generate a category system around an established set of policy issues, and specific categories may already have been anticipated in the methods used to collect data. That said, there are some considerations which may apply to the generation of categories, whatever the specific aims and circumstances of the analyst. This becomes clear if, rather than trying to characterize and contrast different approaches, we consider instead the common resources which can be utilized in any approach to generating a category set.

One source of ideas for generating categories is the data itself. Qualitative methods often involve the classification of data which cannot be accommodated within pre-existing categories. This is, indeed, usually part of the rationale and justification for using a qualitative approach. It is assumed that the researcher cannot establish all (or perhaps any) of the important categorical distinctions at the outset of the research. In some forms of research, such as participant observation or ethnography, the analyst may be reluctant to adopt any prior conceptions before entering the field, and may therefore depend almost entirely on inferring distinctions from the data. In less unstructured research, though some categories may be established in advance, these may still require confirmation in the data, while other categories or subcategories may be derived from distinctions suggested in the data. Even with a relatively structured technique, such as a

structured interview schedule with open-ended questions, all the responses produced cannot be assigned to categories in advance of analysing the data. At the very least the adoption of a pre-existing set of categories requires confirmation that these are indeed the important distinctions within the data. In any case, these distinctions are more than likely to be preliminary rather than exhaustive. Further differentiation within each category will almost certainly draw upon distinctions made within the data.

Thus distinctions within the data can generate new categories, or contribute significantly to refining or modifying the original categories. Amongst those using qualitative techniques, there is usually a strong emphasis on creating categories based on distinctions in the data, most especially where these are recognized or used by the research subjects themselves. Qualitative research is often concerned to elucidate the ways in which subjects experience and perceive situations and events. It would certainly be difficult if not impossible to convey these experiences and perceptions without taking account of how the subjects themselves distinguish situations and events.

At the same time, qualitative researchers often employ observational methods which may produce data inconsistent with how subjects experience, perceive or explain events. A subject's 'explanations' of events may involve assumptions or preconceptions which s/he only dimly recognizes, if at all; and 'explanation' can serve a variety of purposes, including self-justification and exculpation, which have little to do with providing an accurate or truthful account. Thus some distinctions may be drawn by the subjects of the research, while others may be suggested by the data, though not recognized explicitly or implicitly by the subjects themselves.

Categories should not be imposed upon the data arbitrarily; the categories adopted should reflect the data. The distinctions established through categorization should be meaningful in terms of the data being analysed. However, reflecting the data does not mean that categories merely reproduce distinctions which are made or are apparent in the data, although these distinctions can sometimes provide some useful ideas for categorization. A 'reflection' (i.e. mirror image) involves a new view of the data – that of the researcher; and this view can only emerge through 'reflection' (i.e. thought) on the part of the analyst, since the distinctions drawn must be those of the analyst, and related to the overall direction and purpose of the research. Distinctions must serve some analytic purpose which the analyst brings to the data. The actor acts; the analyst analyses – this is integral to their respective roles as subject and researcher. This dictum is not the whole truth, for its inversion may also be applicable – the actor may also analyse and the analyst may also act. But the analyst cannot escape responsibility for the analysis, which must be based on his or her own ideas about the data.

Categorization of the data requires a dialectic to develop between

categories and data. Generating and developing categories is a process in which one moves backwards and forwards between the two. It is this interaction of category and data which is crucial to the generation of a category set. To try to generate categories in the absence of both these resources would be premature. Although we can consider these as separate resources, in practice the generation of categories is an interaction between the two. For example, even if we have not read a single line of data, any ideas we have prior to the analysis must still anticipate the kind of data we will want to analyse.

A rich source of ideas for categories can be found in the questions in terms of which the research originated and developed. These questions, perhaps vaguely formed and poorly articulated at the outset, may already have been considerably redefined and reformulated by the time the final stage of data analysis has been reached. The process of research, whether through interviewing, documentary analysis, or observation, inevitably involves selecting data. This selection is made by the researcher in terms of what seems significant, puzzling or problematic, and the criteria used in selecting data can provide a rich source of ideas for generating a category system.

In documentary analysis, for example, the criteria for selecting documents, or for focusing on particular extracts, should reflect the issues on which the researcher is seeking evidence. There must be some criteria for inclusion and exclusion of documentary data, even if these are broadly defined and refer mainly to the boundaries rather than the substance of the subject being researched. With interviewing, the researcher will have some idea in advance of what questions to ask and which topics to pursue – no matter how non-directive the interviewer may be, the interview has to be conducted with some research purpose in mind. With observation, the observer must make decisions about which sites, situations and events to observe. Often these decisions will be affected by the data emerging from observation and new priorities will develop through the course of the research. In each of these approaches, initial or emergent issues, more or less explicitly defined by the researcher, will provide some guidance to the categories worth developing in the analysis of data. While, in the nature of qualitative research, such questions are not likely to be either comprehensive or exhaustive, they may nevertheless provide a vital starting point for generating categories for analysis.

The process of finding a focus for the analysis, and reading and annotating the data, leads on naturally to the creation of categories. In practice, a sharp distinction cannot be drawn between these processes. The techniques we discussed earlier for generating categories also provide fertile ground for the generation of categories. However, compared with browsing and annotating data, creating a category set requires a more disciplined and systematic approach. Ideas must be sifted, their import assessed, their

relevance evaluated. Some may be discarded. Others may suggest key concepts through which to understand the data.

**Resources for generating categories**

- Inferences from the data
- Initial or emergent research questions
- Substantive, policy and theoretical issues
- Imagination, intuition and previous knowledge

Any or all of these resources can be brought to bear in the task of generating categories for analysis. It is not possible to predict in advance which will prove most useful in developing a category set. That will depend both on the richness and complexity of the data being analysed, and the range and relevance of the experience and ideas which the researcher can bring to the analysis.

Suppose we want to derive a category set from our ideas on humour. Through our initial efforts to find a focus for the analysis, we identified some key ideas to use in the analysis:

- Incongruity
- Catharsis
- Values
- Victims
- Stereotypes

These ideas were inspired by a prior reading of the relevant literature, and their relevance to our analysis has been established through our review and annotation of the data. Do they provide a useful basis for distinguishing differences and similarities in the data? We have found sufficient elements of the different styles of humour to be confident in the general utility of these ideas. We can treat them tentatively as the first approximation of a category set through which to analyse the data.

At the same time, we may want to modify the categories in the light of our knowledge of the data. We might note an element of exaggeration in the sketch 'In the Office', in the references to tarmacking the drive eight times, or receiving seventeen phone calls. There is also an element of transposition, where the women discuss dress instead of the news. We may want to modify our categories in the light of these observations, by adding to our list or perhaps by incorporating them as elements of existing categories. For instance, we could see both exaggeration and transposition as interactive elements in incongruous humour, the latter confirming our expectations only for the former to confound them.

- Incongruity (exaggerating, transposing)

- Catharsis
- Values
- Victims
- Stereotypes

While annotating the data, we noted the role of stereotypes in creating humorous effects. The humour invoked stereotypical views of women's problems with sleeplessness, diet, drugs, their concerns with appearance and body odour, and an inclination to frivolity and gossip. We may therefore want to develop some categories to capture the kind of stereotypes being used. There is no one way to categorize this data: we must choose between different alternatives. **we create a category based on the data, we address the general question: 'What is this data about?'** We also address more specific questions inspired by our analytic concerns: e.g. 'How is this funny?' 'What kind of stereotypes are used here?'. **Categories express ideas about the data. The categories we use will reflect the kind of questions we ask.** For example, we could try either of the two category sets listed in Table 8.1

Table 8.1 Alternative category lists

Detailed category list	Broad category list
Sleeplessness	Health
Odour	Appearance
Diet	Character
Drugs	
Dress	
Gossip	
Frivolity	

The first list of categories is longer and more refined; the second is shorter and more general. One list stays close to the data; the other is at one remove, already implying an implicit categorization (Figure 8.1)

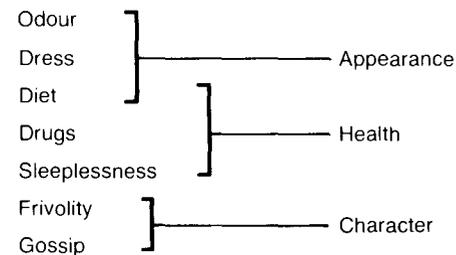


Figure 8.1 Alternative category lists for analysing female stereotypes

We might amend our category set to incorporate these ideas, for example by including the shorter, more general categories in our category list:

- Incongruity (exaggerating, transposing)
- Catharsis
- Creativity
- Values
- Victims
- Stereotypes (appearance, character, health)

There are at least three questions we might ask of this list of categories. What do these categories mean? Is this list sufficiently refined? And what about relationships between the categories? Let us take each of these points in turn.

First, let us take the question of what the categories mean. We noted above that categories must denote bits of data, and relate this data conceptually to the wider analysis. The meaning of a category is therefore bound up on the one hand with the bits of data to which it is assigned, and on the other hand with the ideas it expresses. These ideas may be rather vague at the outset of the analysis. The meaning of a category is something that evolves during the analysis, as we make more and more decisions about which bits of data can or cannot be assigned to the category. To make these decisions, we need not just a category but also a developing set of criteria in terms of which we can decide when and where to assign the category to the data. In other words, we should always try to refine and redefine our categories by specifying and modifying the criteria for assigning them to the data. Without such criteria, our analysis may seem arbitrary and impressionistic. At the same time, we must recognize that any definitions we develop at the outset are liable to be rather general and contingent in character. In defining categories, therefore, we have to be both attentive and tentative – attentive to the data, and tentative in our conceptualizations of them.

Take the category ‘catharsis’ as an example. We introduced this category to refer to humour which involves some sort of emotional release, and related it to topics which are anxiety-laden or even taboo. Now we need to ‘operationalize’ this category in terms of observations in the data. We might start with sex and violence as subjects which are often sensitive issues. ‘Often’ is not ‘always’ and not all references to sex and violence will be ‘cathartic’. We might also start to think about what ‘sensitivity’ involves. Suppose we try to make preliminary ‘definitions’ of the main categories we have created so far (Illustration 8.1).

This is only a starting point. Our ‘definitions’ are notably vague, reflecting the rather abstract nature of the categories we have chosen. We

Illustration 8.1 Preliminary definitions of categories

Incongruity	Include any data where expectations are disrupted.  Disrupted expectations are not always humorous. Identify elements involved in humorous disruption: farce, absurdity, exaggeration.
Catharsis	Include any data which can be seen as sensitive in some way.  Possible topics: sex, violence.  Sensitivity: associated with discomfort, embarrassment, humiliation, guilt.
Values	Include any data which relates to the affirmation or subversion of social values.
Victims	Include any data where humour has an identifiable human target, but only where this target may be representative of a wider group or institution.  Include any such targets, regardless of whether or not the humour ridicules or criticizes with empathy.
Stereotypes	Include any data which seems to invoke stereotypical images.  A stereotype is an ‘unduly fixed’ image which may or may not be accurate but is applied regardless.  Note that not all fixed images are stereotypes.

must be ready to extend or modify our criteria as the data demands. We shall therefore return to the problem of defining categories when we consider the process of assigning them to the data.

Meantime let us turn to the question of whether our initial category list is sufficiently refined. Our category set is short, and the categories are very general. Is this satisfactory? Should it be nearer eight categories, than eighty, or eight hundred? There is no single answer to this question. Data analysts who emphasize the importance of ‘grounding’ categories in the data sometimes advocate a ‘line-by-line’ approach to generating categories (see Strauss 1987). Perhaps more accurately, this is a ‘bit-by-bit’ approach where each bit of data can be as small as a single word. Each bit is considered in detail, to identify aspects which may be relevant to the analysis. The significance of a bit of data can be considered by contrasting it with other bits, by imagining this bit in alternative contexts, or by drawing on relevant theoretical or policy issues. In this way a variety of distinctions may emerge, some of which may eventually prove fruitful in analysing the data. The aim is to generate theory which is fully grounded in the data. Once categories have been developed in this detailed way, the analyst can identify the most relevant categories for further elaboration, and finally

proceed to a more integrated analysis around the core categories which emerge from this process.

By contrast, it is possible to begin with categories which are based on a general comprehension of the data and proceed to a fuller and more detailed categorization (Jones 1985). The emphasis here is on a 'holistic' approach, attempting to grasp basic themes or issues in the data by absorbing them as a whole rather than by analysing them line by line. Broad categories and their interconnections are then distilled from a general overview of the data, before a more detailed analysis fills in and refines these through a process of subcategorization. This approach is more feasible where the analyst already has a fair idea of what s/he is looking for in the data.

Most data analysis probably falls some way between these two extremes. Perhaps the most flexible approach is to develop 'middle-order' categories, which draw some broad preliminary distinctions within the data (cf. Becker and Geer 1982). Often these distinctions may be based on fairly common sense categories, around which the data can be organized quite effectively, without implying commitment to any particular theoretical approach. This approach also fits well with policy-oriented research where a policy agenda already provides a source of categories for analysis. Once the data has been organized into broad categories, the analysis can move in either direction, towards more refined distinctions through subcategorization or towards a more integrated approach by linking and integrating the 'middle-order' categories.

A middle-order approach offers a flexible compromise which allows the analysis to develop in a more detailed or holistic way as time and inclination permits. A middle-order approach is also attractive if the data, although qualitative, is not entirely lacking in structure. Policy issues and programme conditions in evaluative research, for example, can provide a framework for generating a middle-order category set which can already be anticipated in the identification of 'key issues' used in collecting data. Our choice may be as much influenced by our confidence in the organizing power of our initial categories as by any considerations of time or purpose, significant though these may be. For example, if our ideas from other sources at the outset of data analysis are very limited, we may find a bit-by-bit approach applied to at least part of the data is most useful in generating categories for the analysis.

None of these approaches has a monopoly of virtue, and whether one takes a holistic view, begins with middle-order categories or starts with a bit-by-bit analysis  be a question of pragmatism rather than principle. There is no sense  undertaking a bit-by-bit analysis, for example, if time does not permit such a detailed and laborious approach, with the result that some parts of the data are never properly analysed. Where time is tight, a middle-level or holistic approach may provide a better method of selecting

data for categorization. On the other hand, a bit-by-bit approach may be better suited to data analysis where the overriding aim is to use the data to generate theory rather than bringing it to bear on existing policy or theoretical concerns.

In any case, the contrast between these different approaches can be overdrawn, since we are only discussing where to begin. In categorizing the data, the analyst has to work at each of these levels. A holistic approach has still to be rooted in the data, through middle-level categories and bit-by-bit analysis. A middle-level approach has to be geared from the start to the development of a more detailed and integrated analysis. With a bit-by-bit approach, the analyst must become more selective and integrative in subsequent phases of the analysis.

One good reason for adding categories to a set is to ensure that our category list is sufficiently comprehensive. Obviously the length of a category list will depend in part on the range of issues and the breadth of the data being analysed. However, we can also add categories by developing a more refined category set. Would this not save time going over the same material again later to make further distinctions within it? Or would a more extensive category set depend on making further subdivisions than seems desirable at this point in the analysis?

Why not make distinctions now, if it saves us work later on? If categories are too broad, too much data would be included within each category for useful comparisons to be possible. The data would cover too many distinctive topics, and still require further differentiation along the broad lines which we can already identify as relevant at this point in the analysis. Why not assign the data immediately to the more refined category, reducing two operations to one? Why produce an enormous and unwieldy mass of data under a single category, if the data can already be more differentiated amongst several?

Against a concern with efficiency, however, we must balance the issue of confidence. The introduction of more refined distinctions at this stage should only be contemplated if we are confident that these distinctions are sufficiently grounded conceptually and empirically to form a useful framework for analysis. In general, we may prefer to use broad categories at the outset to avoid prejudicing subsequent analysis and perhaps even precluding particular lines of development. It is important not to close off options at this stage by making distinctions which are not based on a thorough review of all the relevant data. This will only create more work later on. 'Errors' in initial categorization also exact a price in terms of subsequent efficiency. 

There are advantages  in giving the task of refining categories until later. One is that all the data for a category can be brought together and examined as a whole before deciding upon any refinements. This may be useful in identifying the weight attached to particular issues and establishing

the connections between them. A detailed inspection of the data may suggest a rather different way of refining the analysis. Rather than make precipitate judgements at this stage, it may be better to wait until all the data has been categorized under the broad category. The main purpose of a middle-order category set is to make possible a more detailed inspection of the data by extracting and ordering observations through some broad preliminary distinctions.

Another difficulty with using too many detailed categories is simply remembering all the relevant categories in terms of which the data can be distinguished. So long as categories are broad, this need not be a major problem, as it will generally be fairly obvious whether the data can be distinguished under one heading or another. But if a category list becomes lengthy through being excessively detailed, it may be difficult to recall all the relevant distinctions when working through the data. An unduly long and cumbersome category list would be difficult to apply to the data.

The degree of refinement at this stage in the analysis may therefore reflect the volume of data to be analysed and the degree to which categories can be readily identified. In developing an initial category set, we may as well take account of distinctions which are already clearly relevant, providing always that this does not preclude the possibility of developing the analysis later on in a variety of directions. Obviously the degree of refinement required is difficult to determine, since we have to balance some competing considerations in devising an initial category set. How this balance is struck may affect the reliability, efficiency and flexibility of the analysis. Too few categories, and later flexibility may be ensured, but at a high price in terms of efficiency, since distinctions still have to be made which could have been applied in the initial categorization. Too many categories, and efficiency may be enhanced, but at the expense of reliability and later flexibility. Striking a reasonable balance is a matter of judgement, reflecting the range, complexity and volume of data to be differentiated (Figure 8.2).

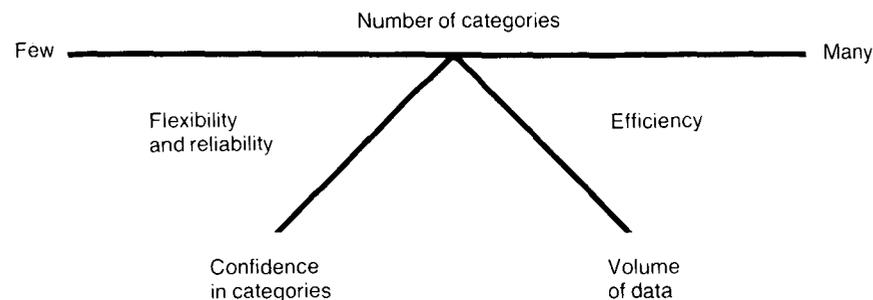


Figure 8.2 Weighing up the degree of refinement in initial category set

Returning to our categories for stereotypes, we must weigh up the virtues of using a more or less refined list. The more refined, the more categories we shall use. A single sketch has generated several possible categories – drugs, diet, etc. – and we have to beware of being overwhelmed by the sheer number of similarly refined categories generated by the data, which includes many other sketches. We may benefit by using the more general categories – health, appearance and character – which also express more interesting conceptual distinctions. Even these distinctions could be dispensed with, if we opted instead for a very general category ‘stereotypes’. But at this level of generality, we may be overwhelmed by the sheer volume of data which would be categorized under such a general category. At a minimum we may want to differentiate between different kinds of stereotypes, and we may also start to refine other categories to produce a much more extensive category list (Illustration 8.2).

At this point, we can turn to our third question and consider relationships between the categories we use. This raises two basic issues. First, are categories inclusive or exclusive? And second, how many levels of classification do we want to use?

Illustration 8.2 Developing a more extensive category list

- Catharsis–sex
- Catharsis–suffering
- Catharsis–other
- Incongruity–exaggerating
- Incongruity–transposing
- Incongruity–other
- Values–confirming
- Values–subverting
- Victims–ridiculing
- Victims–empathizing
- Stereotype–appearance
- Stereotype–health
- Stereotype–character

Categories can be either inclusive or exclusive. If two categories are inclusive, then we can assign them both to the same bit of data without being inconsistent. If two categories are exclusive, then we can only assign one or other to the bit of data. Categories which are exclusive are always related in some way to an underlying concept or overarching category. In Figure 8.3 I have differentiated between inclusive and exclusive categories by using different ways of depicting the way categories interrelate. I have used a long bar with short arms to bracket those categories which are inclusive, and a short bar with longer arms to indicate those which are exclusive. Thus the category ‘victims’ includes two exclusive categories ‘ridiculing’ and ‘empathizing’, while the category ‘stereotypes’ embraces three categories which are inclusive. Because the categories ‘victims’ and

'stereotypes' are inclusive rather than exclusive, we could assign 'ridiculing' and 'appearance' to the same bit of data. It doesn't matter how you choose to depict these different relations between categories, so long as they are consistently noted in some way.

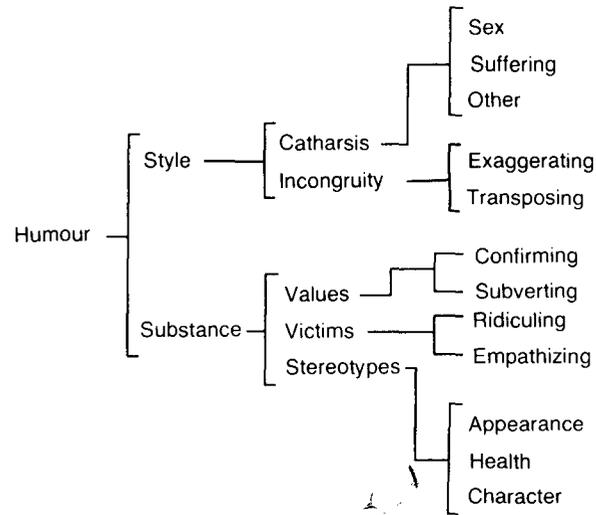


Figure 8.3 Developing a more refined category list

Second, we have to consider levels of classification. Figure 8.3 already involves four levels of classification, with some categories 'nested' within others. But in terms of our category set, some of these levels can be deemed redundant when it comes to categorizing the data. There is simply no point in using the category 'humour' – nor of distinguishing between 'style' and 'substance'. These are distinctions which may be useful conceptually, but have little analytic power when it comes to organizing the data. They do not discriminate sufficiently between different aspects of the data. At most we might opt for a couple of levels, using the subcategories of style and substance, and their own subcategories. It is important to keep track of these different levels of classification, and the easiest way to do this is graphic representation. Some computer packages provide facilities to support this aspect of analysis.

Although conceptually our classification includes different levels, in practice we can operate as though we have an undifferentiated category set. This is more efficient since it saves categorizing the data at different levels. We do not want to have to assign overarching categories, where we have already used the relevant subcategory. For example, if we have used the category 'appearance' then we don't need to assign its overarching category, 'stereotype'. We know anyway that the data assigned to 'appear-

ance' also belongs to the category 'stereotype'. Suppose later on we want to compare all the data implicitly assigned to the category 'stereotype'? How can we if we have not assigned the category? The computer provides a simple answer, since it allows us to locate or combine all the bits of data assigned initially to the separate subcategories.

In developing our category list, therefore, we can include only the most refined categories. To proceed in this way using a single category list is simpler and more efficient, but there are two requirements we must meet. Combining categories later on is only possible if we use a unique specification for each category in our category set. Our initial category set must therefore contain no duplicates. We have to designate each subcategory uniquely, and not assume that the computer will differentiate between two identical subcategories which belong to different overarching categories.

The second requirement is that our categories must be exhaustive. Where the subcategories are exhaustive, there is no need to include the over-arching category in the list, since it is entirely included within the subcategories. Where the categories are not exhaustive, a residual category is required to pick up data which cannot be assigned to the appropriate subcategories. This decision may be taken in establishing an initial list, but the question of 'exhaustion' is one which really can be answered only in relation to the data. At any point, we may have to amend our list if we encounter data which does not fit within the existing subcategories. In practice, therefore, it may be convenient to use an overarching category, such as 'stereotype', as a residual category for data which cannot be readily assigned initially to one of its subcategories. This allows us to retain flexibility in the development of the analysis and avoid premature judgements for the sake of efficiency.

So far, we have looked at the resources upon which we can draw in creating categories, and the issues which arise in creating an initial category list. To conclude this discussion, let us consider some common injunctions applied to the process of creating categories. This process can take a variety of forms, depending in part on the purpose of the research and in part on the time and resources available for the analysis. While qualitative data analysis is not an endeavour subject to 'rules', there are several points worth considering in developing a category system. These provide some general guidelines which can be adopted (or adapted) as appropriate, rather than an established set of procedures to be followed mechanically.

Category development requires the analyst to become thoroughly familiar with the data, and it is worthwhile acquiring this familiarity at an early stage in the analysis. Although this may seem a rather innocuous point, in practice becoming thoroughly familiar with the data can prove quite an onerous undertaking. When observations are voluminous, as is often the case with qualitative data, the temptation to take short-cuts is considerable. Given the constraints imposed on research budgets and

timetables, and the limited aims of some data analysis, some short-cuts may also be to some extent unavoidable. However, time spent becoming thoroughly absorbed in the data early in the analysis may save considerable time in the later stages as problems are less likely to arise later on from unexpected observations or sudden changes in tack.

Perhaps the most common injunction to would-be analysts is that data must always be considered in context. One of the major concerns of qualitative analysis is the observation of opinion or behaviour within a 'natural' setting. From this perspective, meaning depends upon context, and the interpretative action or opinion must take account of the setting in which it is produced. An observation cannot be fully understood outwith the context in which it is made. To consider this fully, it is often essential to regard the researcher as part of the context being studied. This is obviously relevant in interviews, where the respondent is responding to some sort of stimulus on the part of the interviewer. It is also relevant in observational research where the researcher interacts socially with the subjects of the study. How subjects perceive and respond to the observer can then have a significant effect on what they say or do. The researcher's own actions and perceptions therefore become part of the social interaction, and need to be observed and analysed as such.

The injunction to consider context may seem somewhat paradoxical, however, since for the purposes of comparison, it is necessary to abstract data from its immediate context, and consider it from a point of view which transcends that context and allows the data to be compared with observations made in a different context. For example, a stereotypical comment may be made in the context of specific sketch, and yet I may want to compare these stereotypical comments across a range of different sketches. These observations must therefore be abstracted from their immediate context. However, confirmation is required that these comments can be meaningfully compared. In practice, this confirmation can be established partly through comparing how stereotypical comments have been used in different contexts. We saw earlier how much observation depends upon implicit classification, or, in other words, implicit comparisons. Thus comparison is itself a useful method of identifying and understanding the specific context within which observations occur. This does not remove the tension between analysing meaning in context and analysing it through comparison, but it does imply that both processes are necessary for an adequate elucidation and interpretation of the data. This is why categories have to be meaningful both internally, in relation to the data understood in context, and externally, in relation to the data understood through comparison.

A third point concerns alternative ways of categorizing observations. Our category set cannot be entirely arbitrary, for it must make sense of the data. But there is no single set of categories waiting to be discovered. There

are as many ways of 'seeing' the data as one can invent. Any distinction has to be considered in relation to the purpose for which it is drawn. With respect to that purpose, it may be more or less useful, but one distinction cannot be considered more or less valid than another independently of the reasons why it is made. It is better to be profligate in producing alternative categories than to foreclose analysis by adopting one particular set too early on in the analysis.

A related point is that flexibility in extending, modifying and discarding categories is important in developing an effective category system. The fit between data and categories is subject to continual adjustment. Flexibility is required to accommodate fresh observations and new directions in the analysis. Categories may be extended or modified to cope with data which does not quite fit, but at some point categories which cannot cope with the range and complexity of the data must simply be discarded in favour of more promising alternatives. It is also likely that the analysis will shift in emphasis or direction as initial assumptions are modified in the light of the data. New categories may be needed which more accurately reflect the changing aims of the analyst.

Categories can be considered, not just in terms of the data but also in terms of their connections with other categories. While the same observations may be categorized in several different ways, reflecting different aspects of the analysis, too much overlap between categories leads to an inefficient and cumbersome analysis. As far as possible, categories reflecting the same dimension of the analysis should not overlap unduly. If categories do overlap, then this should reflect significant differences in the distinctions being made about the data.

Developing categories usually involves looking forwards towards the overall results of the analysis as well as looking backwards towards the data. It is worth working towards a holistic view, even if this may not be feasible at the outset. It is not the case that a more holistic view will somehow simply emerge from an accumulation of detailed categorization. Perhaps if we plant a sufficient number of individual trees, we do create a wood, but to see the trees as a wood still requires a shift of vision. To pursue the analogy, it is also necessary to plant trees whose type and location are related, for otherwise there will be no wood but simply a jumble of scattered trees. Therefore it is as important to consider the relation between categories as it is to consider the relation between a particular category and the data.

The process of developing categories is one of continuous refinement. Inevitably, the criteria for including and excluding observations, which may be rather vague at the outset, become more precise during the course of the analysis. These criteria need to be set out as clearly as possible if observations are to be categorized in a reliable way. It is worth trying to spell this out as far as possible at the outset. If nothing else, this will indicate where

ambiguities exist which one can try to clarify or remove during the process of categorizing observations. Decisions about what to include and exclude are themselves part of the process of clarifying the criteria involved. In developing categories, therefore, it is useful to keep track of the criteria which are being adopted in making such decisions. Recording these criteria provides a running commentary on how a category is being used and provides a basis for developing a more precise statement of what distinctions are being drawn. Where the research is a cooperative activity and more than one person is involved in analysing the data, this is essential if categories are to be applied on a consistent and uniform basis. But much the same point applies even where only one analyst is involved, since even here there is a need to secure a consistent approach and avoid arbitrary and ad hoc decisions.

Since categorizing observations is a crucial phase in data analysis, there is some virtue in regarding this as a public rather than private activity. In other words, we may consider how decisions about categorizing data can be explained and justified to others. This can help to sharpen our approach and ensure sufficient rigour in decisions about categorization. It may also help to clarify the relationship between observations, interpretation and the overall results of the analysis. We may be better placed to explain why a particular tack has been taken in analysing the data, why some distinctions have come to be regarded as crucial while others have been discarded, and just how comparisons, resulting from categorization are rooted in the available evidence.

**Some common injunctions in creating categories**

- Become thoroughly familiar with the data
- Always be sensitive to the context of the data
- Be flexible – extend, modify and discard categories
- Consider connections and avoid needless overlaps
- Record the criteria on which category decisions are to be taken
- Consider alternative ways of categorizing and interpreting data

The categories that we create become the basis for organizing and conceptualizing our data. Categorizing is therefore a crucial element in the process of analysis, and as well as considering how to create categories, we also have to consider the issues involved in assigning them. This will be the focus of the following chapter.