Comprehension Failures in Educational Assessment

Alastair Pollitt and Ayesha Ahmed

Research and Evaluation Division
University of Cambridge Local Examinations Syndicate

Paper presented at the European Conference on Educational Research, Edinburgh

September 2000

Contact details:

Alastair Pollitt, RED, UCLES, 1 Hills Road, Cambridge, CB1 2EU

pollitt.a@ucles.org.uk tel: +44 1223 553847, fax: +44 1223 552700

Comprehension Failures in Educational Assessment

Alastair Pollitt and Ayesha Ahmed

University of Cambridge Local Examinations Syndicate

Abstract

In this paper a psychological model for the process of answering an examination question is described. The model consists of six 'stages' - Learning, Reading, Searching, Matching, Generating and Writing; it is argued that valid examining requires the question writers to understand how these processes work, and how they are sensitive to interference effects. The value of this model for improving school examinations is illustrated by analysing some questions from science examinations. The questions are all "contextualised" and the analysis, validated empirically, shows how school students can be misled or biased by the improper use of context. The paper offers some advice for the proper use of context in examining.

Question Difficulty and Validity

School examinations are measuring instruments, and they generally consist of a series of tasks, or questions. This means, therefore, that the *ability* measured by examinations is the student's ability to gain marks on these tasks. What is this *ability*?

Traditionally, there have been several approaches to judging validity, some more empirical, others more impressionistic. Our interest in the ability being assessed, rather than the syllabus or the outcomes of assessment, focuses our attention on what is traditionally called *construct validity*, but in a non-traditional way. Rather than studying the students to see what other traits that vary amongst them also correlate with examination scores, we focus on the tasks themselves to see what the students' minds are doing during the examination.

Examinations measure students' success at doing the things the tasks require them to do; these 'things' are the demands that each question makes. To understand, therefore, what the examination is measuring we need to understand two things about each task: what is the process that students' minds follow as they answer it and, as they pursue that process, what are the features of the task that determine how easy or difficult it is. In discussing his paper on construct validation (Stenner et al, 1983) Stenner commented: “If you can't tell me why this question is more difficult than that one then you don't know what you're measuring”.

In our research we have identified many *Sources of Difficulty* and *Sources of Easiness*, known as SODs and SOEs (Pollitt et al, 1985; Hughes et al, 1998 a&b, 1999; Bramley et.al., 1998; Fisher-Hoch et.al., 1998). A question can contain valid SODs and SOEs, intended by the question writer, but it may also contain invalid SODs and SOEs that were not intended. Different students will tackle even a very simple question in different ways, and will bring to it different complexes of abilities and different memories of experiences. One student's SOD may be another student's SOE, and examiners should look out for a range of SODs and SOEs in each question and ask how different students may react to each of them.

Good question writers will deliberately control the sources of both easiness and difficulty in order to ensure that the principal steps in the answering process involve the students in carrying out the kind of cognitive activity that constitutes good practice in their subject. For example in a science examination students should be doing the kind of thinking scientists do. This captures the essence of a cognitive psychology approach to construct validity: examiners manipulate questions to ensure that, as far as possible, the students' minds will be doing the things the examiners want them to show they can do. Then the SODs in the question will determine how well they are able to do these things.

What should students' minds be doing during assessment? That is a question for the curriculum specialists, but assessment specialists have a responsibility to help them achieve construct validity. One issue that has caused us some concern is the very popular use of contextualised questions in science examinations, for we have seen considerable evidence that the contexts can interfere with the intended cognitive processes, and so compromise validity.

The role of context in assessment has recently been discussed widely in mathematics (e.g. Boaler, 1994, Wiliam, 1997; Nickson, 1998;). In this paper we outline a cognitive psychological model of the processes involved in answering an examination question, and consider how context affects these processes in science examining.

The Model of the Question Answering Process

The development of a single general model that attempts to describe the process by which students tackle and answer examination questions is described in more detail in Pollitt & Ahmed (1999). The model grew out of work reported in Pollitt, Entwistle, Hutchinson & De Luca (1985), who identified a number of features that seemed to cause difficulty in Scottish O Grade examination papers in geography, mathematics, chemistry, English and French (as a foreign language).

Answering language comprehension test questions

One outcome of that project was a linear model for the process of answering comprehension questions after reading a previously unseen text.

Largely for logical reasons, six stages were proposed for the comprehension question model:

1 Reading the text
2 Reading the question
3 Searching the text
4 Interpreting the text
5 Formulating an answer
6 Writing the answer

Some statistical evidence supporting the general form of this model is described and evaluated in Pollitt & Hutchinson (1986).

Answering examination questions

The language comprehension model can be generalised to all examinations and all kinds of questions, and can be expressed in terms of general cognitive processes: the key to this generalisation is to see that Stage 1 (Reading the text) can correspond to the larger scale process of reading - that is studying - a whole subject. There are therefore six phases in the general question answering model :

0 Learning the subject
1 Reading the question
2 Searching the memory
3 Matching question to memory
4 Generating an answer
5 Writing the answer

In the earlier model students were said to search the "text" for information relevant to the question. Modern cognitive theory would insist that it should be the mental representation of the text that is searched instead of the text itself (Taylor, 1996), and this enables us to generalise from the assessment of reading ability to the assessment of learning in all examinations. The students will need to search their mental representations of the whole subject, which are stored in their memory, to find learning relevant to the task.

In traditional 'problem solving' tasks, each of the phases in the model may be quite distinct. Consider verbal analogy problems. Sternberg (1982) proposed a model of the psychological processes involved in solving analogy problems, of the form *A is to B as C is to ...?*, which involves six "components":

encoding, inference, mapping, application, justification, and response.

Although described as " components" it is clear that these are intended to form a sequence of distinguishable steps which, in simple cases at least, will represent how the problem is actually solved. Again, it would seem logically necessary, for example, to infer the relationship before you can map it. The similarities between this model and our model are obvious, if "encoding" and "inference" are seen both as parts of Reading (see below for further discussion of the reading process).

However, in many examination questions some of the phases are so automatic that it is difficult to distinguish them as discrete stages in a process. Consider a cloze task. The student reads a text, building up a more or less coherent mental representation of it, in the normal manner of processing written language. As this building process continues the mental representation is made complete using the remaining redundancy in the text, and the gaps in the text are then filled in a usually very automatic way. It appears that the process of completing a cloze test bears more resemblance to the normal speaking process than to reading.

Most examination tasks fall between the two extremes illustrated by ‘problem solving’ and ‘cloze’, and are not so simple; they will involve repeated cycles of Search, Match and Generate, and the outcome of Generating may be another search in memory. Nevertheless, the simple linear model has proved to be a useful heuristic for investigating the various factors identified in the research programme as Sources of Difficulty and Sources of Easiness. If we can understand how they achieve their effects we should achieve better control and be able to ensure that we base assessment on more valid questions.

Understanding the question

The largest source of influences on question difficulty is the first of the phases in the model - Reading the question. We find that students very often misunderstand what they are asked to do, and many of the effects we see as errors in Searching, Matching and Generating have their origin in problems with Reading. The next section therefore begins with a more detailed consideration of the process of reading the question, and of some of the principal ways in which this process may go wrong.

The Effect of Context on Question Processing

Reading comprehension

Most current theories of text comprehension stress the active and constructive nature of the process. Meaning is generated by the cognitive processes of the reader, building a personal representation which is provoked by the text and mostly built from pre-existing knowledge, and which is modified by personal attitudinal characteristics and intentions. Thus, for example, van Dijk (1981) talks of “knowledge structures”, or Anderson (1977) of the “ideational scaffolding” of comprehension; to Johnson-Laird (1981) words are “cues to build a familiar mental model” as we read. In all of these sources the reader is envisaged as constructing a mental model which represents in some way their particular understanding of the text. When we look for evidence that a text has been comprehended, then, it is to this mental representation that we are addressing our questions.

In the case of reading an examination question the same applies, for the question text only provides meaning to the extent that it can be incorporated into the emerging mental model of the task. There are some properties of mental model building, though, that are particularly important when the 'text' is an examination question.

First, each student constructs their own unique representation of the question text, out of their own memories and experiences, emphasising whichever elements seem most salient to them, whether they are relevant or not.

The activation of these elements is automatic; the student will not even be aware of them since they occur *before* the reading of the text reaches consciousness. Thus these elements, although selected unconsciously, will affect the subsequent interpretation of the text. Evans (1989) calls this 'pre-attentive bias':

> *many biases are caused by pre-attentive or preconscious heuristic processes which determine selective encoding of psychologically "relevant" features of the problem. (p 91)*

Further, when a student reads a question they build a representation of a text which is embedded in a very specific and familiar setting, that of examinations. Therefore the student will construct a representation not only of the text itself but also of the purpose for which it is being read. In this case the purpose is the satisfactory completion of a task, where satisfactory means one that will gain as many marks as possible from the examiner. This means the representation will include the student's construing of the task and in some cases, like cloze, the outcome itself. In less automatic cases the representation will include prescriptions for action, which will at least initially direct the student's thinking.

The Reading phase of our model, then, consists of building an appropriate mental representation of the task expressed in the question. If this happens effectively, then phases 2-4 will usually be quite automatic; if it does not, then the subsequent phases will at best require explicit attention and at worst be distorted and lead to a wrong answer or no answer at all. Examinations are stressful experiences, in both affective and cognitive ways. Some of a student's attentional resources will inevitably be directed to coping with the stresses rather than to answering the questions, and anything that increases the attentional demands of the tasks, such as dealing with an “interesting” context, will further jeopardise effective task comprehension.

Context and schemas

How does the use of context for a question influence processing? Two features are significant here, the effect of schemas, and the kinds of cognitive process that students will expect to engage in.

The most powerful effect of context is best considered in terms of schema theory (Bartlett, 1932). Schemas are automatically activated when the student reads the question - indeed some schemas relating to the whole examination process will be activated even before this. Context evokes very powerful schemas, powerful because they are familiar; indeed this is central to the justifications for using context at all. But powerful schemas bring with them strong expectations, of what sort of thinking is likely to be needed in the next few minutes to cope with the demands of the task. These expectations may come from the general examination schema, from more specific schemas for the subject, or from schemas activated as the student reads the question. Students' schemas will resemble each other in important respects. They will construct a representation of the task incorporating common expectations of examination requirements, and of the reality behind the context.

Yet everyone's schemas will also be different to some degree, especially when they are founded in reality and experience. Real world contexts will activate different concepts in different students' minds, creating different conceptualisations of the task. In effect, each student answers a different question and the examiners lose control of the measurement process.

One way to summarise these issues is in terms of the power of *expectations* to influence how students answer an examination question. Long experience of the culture of examinations, through textbooks as well as tests and past papers, causes students to expect to be asked to perform in certain ways, both in general and specifically in certain topics, while their teachers also prepare them to meet specific topics, questions and demands in examinations in their subject. As well as these prior expectations there are also aspects of questions that will lead to certain expectations during the question comprehension process. In particular, some contexts will cause students to expect to be asked about a particular aspect of science in a particular way, whether or not the question really demands it. So strong are the expectations, and so anxious are the students to 'recognise' the tasks, that even the most careful wording and typographical design may fail to make students see what the question is really about.

Examples

The examples that follow are taken from science examination papers designed for 16 year olds. They have been chosen from the hundreds that we have studied to illustrate the dangers of contextualisation and the power of expectations. They are thus a deliberately biased selection, and are not typical of UK science examining.

Ex 1 Period of sulphur

(b) *Put a ring around the name of an element which is in the same period as sulphur.*

fluorine

magnesium

oxygen

potassium

Even though the students were reminded to "Use the Periodic Table on the back page to help you answer these questions" very many chose **oxygen** rather than **magnesium**. Oxygen is in the same *group* as sulphur, not the same *period*; the same column not the same row. But most of a chemistry student's time with the Periodic Table in fact deals with the groups or columns since these share similar chemical properties, a message which was re-emphasised here because part (a) was about the Alkaline Earths and part (c) about the halogens. Periods are much less significant in the students' minds than groups, and they do not expect to be asked about them except, perhaps, in a question explicitly about atomic structure. Here they clearly gave the correct answer to the question they *expected* to be asked.

Ex 2 Ice lolly holder

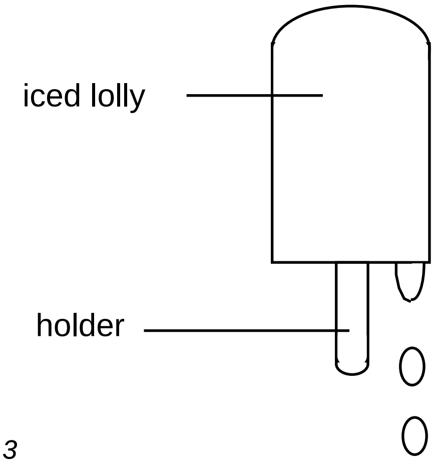

Information on three different materials is shown in Table 3

Table 3

material	**effect of heat**	**effect of water**	**manufacturing problems**
wood	easily burns	absorbs water and rots	may crack or split
plastic	may give off poisonous fumes	very resistant	may discolour
metal	good conductor of heat	may corrode	may have sharp edges

Which material should be used as the iced lolly holder? Give a reason for your answer.

Material ___________________________

Reason ____

____ *[2]*

After some questions about milk, and states of matter, students met the question about Table 3. They were expected to use the information about the properties of materials in the table; the answer expected by the examiner was 'plastic' as it is 'very resistant' to effects of water. However, children's ice lolly schema contains the wrong answer, since "everybody knows" that ice lollies come on wooden "sticks", and many of the students gave this as their answer. Experience of the real world creates a strong expectation; here the schema of a very familiar context led many students to give the wrong answer to the question, whether or not they understood the science. They sometimes went to great lengths to justify choosing wood, such as appealing to rotting as, environmentally, a good thing.

Ex 3 Desert Tower

In desert countries there is a shortage of fresh water.

A scientist has suggested a new way of making fresh water from sea water.

Electricity might be generated at the same time.

The diagram shows the scientist's idea.

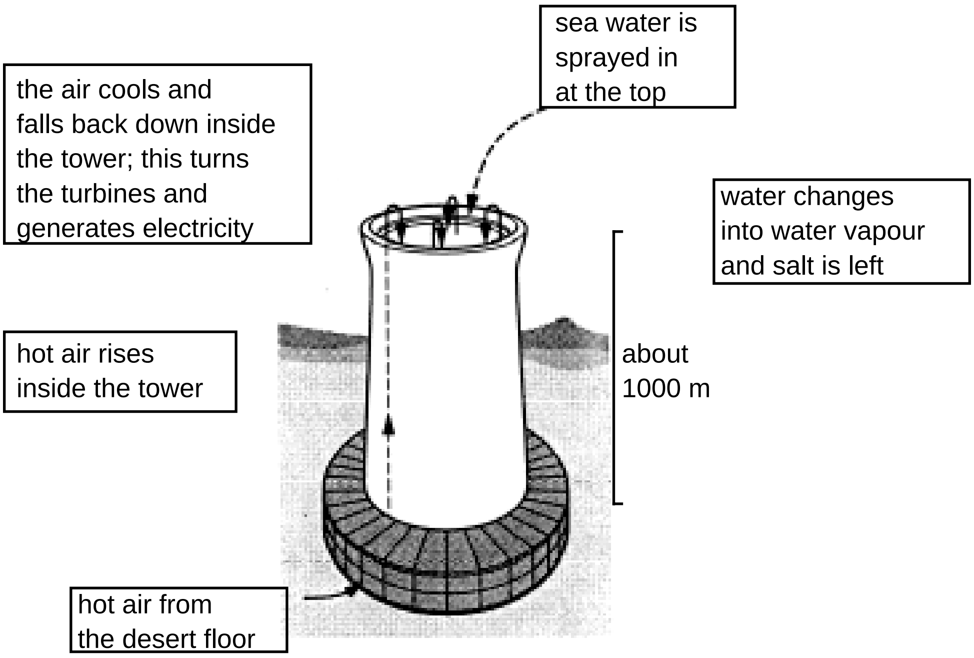

(a) Air rises inside the tower in a convection current.

Explain why convection happens.

____

____

____ *[2]*

*(b) Name **one other** way in which heat is transferred from place to place.*

____ *[1]*

Here again the context created certain expectations. The context of this question involved a tower for making fresh water in the desert and students expected to be asked about the tower. Part (a) asked about convection, but not how it happens in the tower. Questions of this type often begin with a phrase like "Using your knowledge of particles, …" to make sure that students ignore the context when it is irrelevant.

Part (b) asked for other ways heat is transferred, and many students again tried to answer the question using the context, assuming that it was relevant to the answer. In fact this question, like part (a), did not relate to the context at all and students simply had to give a textbook answer. There were 58 correct answers and 61 wrong ones that were based on elements of the context, as the following table shows:

Types of answer (random sample of 199)	**Number**
Correct answer: conduction/radiation	58
Total correct	**58**
Wrong answers containing ‘sun’ ‘solar energy’	3
Wrong answers containing ‘wind’ ‘air’	14
Wrong answers containing ‘desert’ ‘tower’	4
Wrong answers containing ‘radiators/pipes/domestic heating ’	16
Wrong answers containing ‘evaporation’ ‘condensation’	16
Wrong answers containing ‘electricity’	8
Total wrong, in context	**61**
Total other	**80**

Ex 4 Fish Cartoon

Thousands of golf balls end up in the sea.
They are lost by golfers on seaside golf courses and on cruise ships.

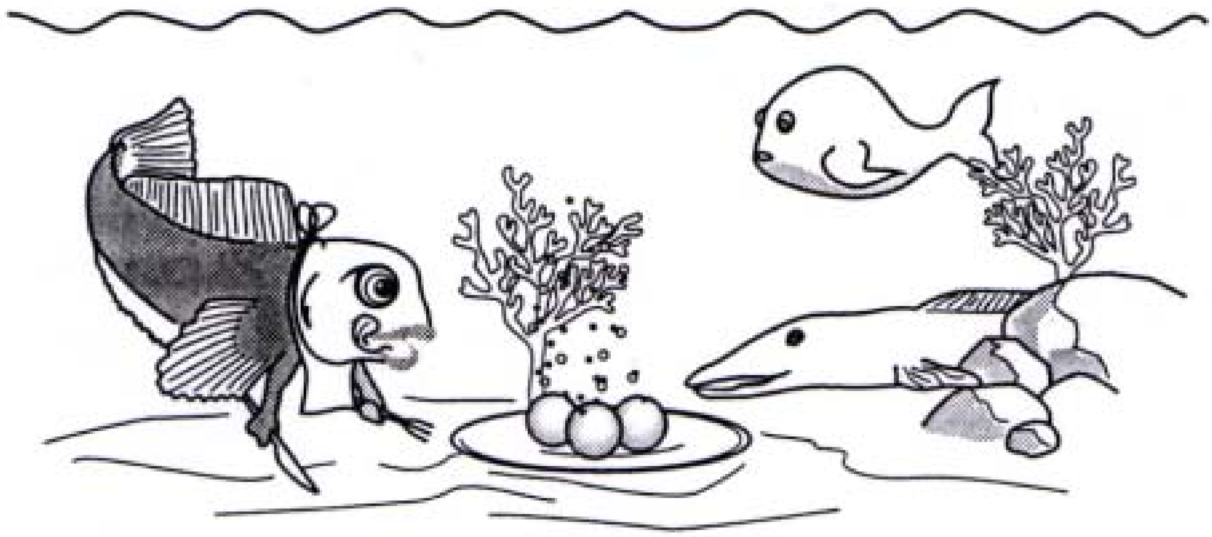

(a) *Fish sometimes swallow these golf balls.*
Explain why this could harm the fish.

____

____

____ *[1]*

The key issue here is that the question is illustrated with a cartoon drawing, which caused the activation of a cartoon schema - in which fish are like humans and eat with a knife and fork in their 'hands', wearing a bib round their 'necks'. This schema would activate the concepts of breathing and choking, and many of the students did write that the fish could choke on the golf balls if they got stuck in their 'throat'.

Contexts evoke schemas, which activate concepts and associations that can mislead. This example misled several of our research staff, and we can safely assume that school students, operating under the stress of exam conditions, would be more rather than less susceptible.

Ex 5 Cleaning the boat

Old paint can be removed from machinery using solvents made from oil.
Suggest ***two disadvantages*** *of using oil-based solvents.*
1 ______________________________
2 ______________________________ *[2]*

The mark scheme read:

any two of	*reject*
flammable, oil is a finite resource, toxic/health hazard	disposal problems, dangerous

Real world contexts evoke rich schemas which contain a complex web of common sense knowledge - most of which will be irrelevant to answering the question. Students have to decide which of this knowledge is relevant and which is irrelevant.

In interview one student commented on her thinking: "It sticks and it's messy. I didn't know if that was scientific or not; whether to use common sense or just scientific principles." She was uncomfortable with the idea of using *oil* to make something *clean*, when she associated oil with *dirt*. The expected answers concerned environmental issues, unrelated to the specific context of machinery.

Ex 6 Why is drinking water in Gibraltar stored underground?

(c) *Suggest why the water is stored underground.*
____
______________________________ *[1]*

Parts (a) and (b) of this question concerned the particle theory of matter, and students were asked to explain evaporation and condensation in warm and windy conditions. For part (c) the examiners therefore expected an answer about how underground reservoirs would minimise the evaporation of stored drinking water. They found, though, that many students suggested that underground reservoirs would keep the water cleaner, or would keep pollution out, and others suggested that Gibraltar had no space or flat surfaces for overground reservoirs. These responses are undeniably true, and were eventually accepted as correct, but it was clear that the examiners had lost control of the students' thinking processes.

It is also likely that many students were uncertain about the acceptability of their answers. This was supposed to be a science examination, and were answers based on hygiene, or geography, appropriate? Were they supposed to use their general knowledge about Gibraltar?

Ex 7 Copper and Silver

> *Suggest a reason why the inventors of this system decided to use a mixture of copper and silver for the electrodes.*

Is real world knowledge allowed or expected? In this case **only** an economic answer was accepted – "silver too expensive" - even though the earlier parts of the question were about electrolysis. Just 10% of the students got the mark, and most of the others tried to give a Chemistry answer, concerning the different charges on the ions or their reactivities, or gave no response at all. When we modified the question by adding … *rather than silver on its own* the success rate went up to 30%. This illustrates the power of students' expectations.

Conclusions

Why is context used in assessment?

Despite the problems caused by contextualisation, three plausible reasons for using context in assessment can be put forward. Note that there may be many other good reasons for using context in teaching, and some might argue that contextualised assessment is necessary to reinforce these pedagogical arguments; the issue here is solely about the direct effects of context on assessment. What is good for teaching may not be good for assessment.

(i) Contexts bring *relevance*, making examination tasks more familiar to students who have not yet developed interest in the more formal aspects of the discipline. At age 16, for instance, students are studying many subjects and cannot be considered as scientists, historians or geographers.

(ii) Context can make abstract concepts *concrete*, giving specific objects to consider in place of textbook generalisations. This reduces the complexity of the task, since the concepts of concrete objects are strongly linked in the brain to specific referents, rather than weakly linked to many referents; this may allow students more easily to show their understanding of the principles in real cases.

(iii) Context allows the examiners to assess the student's ability to *apply* learning, rather than just their knowledge of the content of textbooks. The application of knowledge is generally considered to be a more valuable and indeed more difficult skill, less dependent just on the ability to memorise.

All of these reasons, however, carry with them the risk of bias - relevance is personal, concrete examples will be more familiar to some students than others, and what is intended to be application may even be just recall to those fortunate enough to have considered that example, or a similar one, in their learning. All of these considerations mean, in effect, that different students will construct different mental representations of the question task; when realistic contexts are employed these differences will be very much greater.

Context is a powerful, and therefore a dangerous, device. Its principal effect is to activate many concepts in the student’s mind, and these will be different concepts for students with different experiences. Different classes with different teachers or textbooks will also activate different concepts under the influence of a realistic context.

When questions deliberately cause such variation in the mental processes of the students we can no longer assert that they are all 'doing the same test'; what is inference for some will be recall for others, and a novel context for some will be a familiar one to others. In effect, we may have lost control of the assessment process.

“Relevance” brings with it, necessarily, aspects of natural thinking - both cognitive and affective – that will make it more difficult for the students to stick to scientific thinking and discourse. The context and language problem is particularly severe for school students. They are taught many subjects simultaneously, and their single mind has to learn to distinguish the various discourse domains appropriately. Word meanings alter, sometimes subtly, as the domain changes - think of the varying meanings of "force" or "reaction" in mathematics, physics, chemistry, politics or everyday language. Good educators try to help students make these discriminations, so that they will learn to function appropriately in different domains, without encouraging the view that the subjects are discrete, a difficult balance to maintain. Evoking real world contexts under conditions of stress may confuse the boundaries, leaving students in doubt as to what is acceptable as an answer to an examination question.

Our examples were selected to show some of the worst effects that context can cause. Are there ways to ensure that contextualisation is beneficial?

Using context appropriately

Is there a proper way to use context in educational assessment? Our model suggests that the dangers of contextualising examination questions are considerable. One response is suggested by Brown (1999) in a discussion of mathematics: examination questions should concentrate on the mathematics or science that is the essence of the school subject, not on the outside world. The place for contextualised activities is in projects and coursework, where the relatively relaxed conditions, and the presence of teachers, allow students a reasonable opportunity to explore the complexities of the real world.

If this is too extreme, there are some rules that might ensure that contexts in examination questions do not bias, distract or otherwise distort the assessment processes. We are developing the notion of *focus*, by which we mean the extent to which the most salient aspects of a context correspond to the main issues addressed by the examination question. In such a case the context will help activate relevant concepts, rather than interfering with comprehension and scientific thinking. We see good examples where the bias of the context actually helps the students to concentrate on the required activities, where the natural focus of the context is also the point of the examination question. For instance, in a ‘shopping’ context the natural focus is the final cost of the purchases; the parts of an examination question in such a context should lead clearly towards this final outcome. We would then expect to see fewer comprehension failures of the kinds we have illustrated here.

We think that examiners should identify the focus of any context they wish to use and should then make sure that all the parts of their question lead towards it, rigorously avoiding any side issues, however tempting it might be to divert the student into them. A context in this view is a purpose for analysing reality and applying science, not merely a setting within which interesting scientific phenomena may be explored.

Experimental work to explore this concept of focused context is in process within our research programme.

References

Anderson, R.C. (1977) The notion of schemata and the educational enterprise. In Anderson, R.C., Spiro, R.J. & Montague,W.E. (Eds.) *Schooling and the Acquisition of Knowledge*. Hillsdale, NJ: Lawrence Erlbaum.

Bartlett, F.C. (1932) *Remembering*. Cambridge: Cambridge University Press.

Boaler, J. (1994) When Do Girls Prefer Football to Fashion? An analysis of female underachievement in relation to ‘realistic’ mathematics contexts. *British Educational Research Journal, 20 (5) 551 – 564.*

Bramley, T., Hughes, S., Fisher-Hoch, H. & Pollitt, A. (1998) Sources of Difficulty in Examination Questions: Science. *UCLES Internal Report.*

Brown, M. (1999) One Mathematics for All? In Hoyles. C., Morgan, M. & Woodhouse, G. (Eds.) *Rethinking the Mathematics Curriculum.* Studies in Mathematics Education Series: 10. Falmer Press, London, pp 78-89.

Evans, J. St. B.T. (1989) *Bias in Human Reasoning: Causes and Consequences*. Hove: Lawrence Erlbaum.

Fisher-Hoch, H., Hughes, S. & Pollitt, A. (1998) Sources of Difficulty in Examination Questions: Geography. *UCLES Internal Report.*

Hughes, S., Fisher-Hoch, H. & Pollitt, A. (1998a) Sources of Difficulty in Examination Questions: Mathematics. *UCLES Internal Report.*

Hughes, S., Fisher-Hoch, H. & Pollitt, A. (1998b) Sources of Difficulty in Examination Questions: English. *UCLES Internal Report.*

Hughes, S., Fisher-Hoch, H., Pollitt, A. & Ahmed, A. (1999) Sources of Difficulty in Examination Questions: French. *UCLES Internal Report.*

Johnson-Laird, P.N. (1981) Mental Models of Meaning. In Joshi, A.K., Webber, B.L. & Sag, I.A. (Eds.) *Elements of Discourse Understanding*. Cambridge: Cambridge University Press.

Nickson, M. (1998) What is the difference between a pizza and a relay race? The role of context in assessing KS2 Mathematics. *British Journal of Curriculum & Assessment, 7 (3), 19-23.*

Pollitt, A., Entwistle, N., Hutchinson,.C & de Luca, C. (1985) *What Makes Exam Questions Difficult?* Edinburgh: Scottish Academic Press.

Pollitt, A. & Hutchinson, C. (1986) The Validity of Reading Comprehension Tests: What Makes Questions Difficult? In Vincent, D., Pugh, A.K. & Brooks, G. *Assessing Reading.* Proceedings of the UKRA Colloquium on the Testing & Assessment of Reading. Macmillan, London.

Pollitt, A. & Ahmed, A. (1999) A New Model of the Question Answering Process. *Paper presented at IAEA conference, Slovenia, May 1999.*

Stenner, A.J., Smith, M. & Burdick, D.S. (1983) Toward a Theory of Construct Definition. *Journal of Educational Measurement, 20 (4), 305-315.*

Sternberg, R. (1977) Component processes in analogical reasoning. *Psychological Review*, 84, *353-378.*

Taylor, L. B. (1996) An Investigation of Text-Removed Summary Completion as a Means of Assessing Reading Comprehension Ability. *Unpublished PhD thesis, University of Cambridge.*

van Dijk, T.A. (1981) Review of RO Freedle (Ed) (1979) *Journal of Linguistics, 17, 140-148.*

Wiliam, D. (1997) Relevance as a MacGuffin in Mathematics Education. *Paper presented at British Educational Research Association Conference, York, September 1997.*