8
Teaching-as-Design and the Ecology of University Learning

Introduction

Teaching in modern universities is very difficult to sustain as a solo activity, conducted by a lone academic. Co-teaching is becoming more common. But even where teachers manage their own individual courses, the need to work in a team arises in a number of situations: for example, in activities such as clarifying assessment standards, planning programs so that they are coherent to students or mapping graduate attributes across a curriculum.

The suspicion that teaching is hard to sustain as a solo activity is reinforced when learning technologies come into the picture. It is rare for an individual academic to have all the knowledge and experience needed to make the best choices among learning tasks, technologies and ways of organising students, and to make sure that these choices are aligned to best effect. We have a strong sense that the increasing complexities of designing for productive student learning – given the changing array of student needs, teaching methods and curriculum demands – will encourage more teachers to do more of their work in collaborative programme teams. In principle, such teams also offer better opportunities for specialist educational development and technology staff to contribute to the improvement of learning opportunities: their contributions are amplified by working with a group, and can crystallise more readily into the improved practices of a course team than into the work of the solo academic. The combined or extended team of academics and educational development and technology specialists can also embrace more complex approaches to educational design, and make use of appropriate design tools and methods – tools and methods that have not been taken up widely by solo academics, in part because of the time involved in mastering their use.

The purpose of this chapter is twofold: to explain the notion of 'teaching-as-design' and to show how the design work of teaching teams fits into the broader ecology of university learning. Among other things, teaching-as-design implies that teaching teams have a lot of work to do before the first

student arrives on campus, and that regular evaluation, reflection and review are needed to close the loop between students' experiences of learning and the (re)design and ongoing enhancement of all aspects of educational provision. Many of the books that offer advice about educational design assume teachers are working on a 'greenfield' site – setting up a new course or programme. In actuality, most educational design work is aimed at improving existing provision. This presents a very different set of problems from those encountered when setting up something new. For example, enhancement of existing courses and patterns of learning and teaching can rarely tackle everything at once: at any one time, some aspects of curriculum, resources, etc., have to be treated as fixed constraints and the enhancement focus has to illuminate other areas on which to work. The other feature of educational design we need to mention at this point is that design work applies at several scale levels, but not everyone can deal with issues at every level. For example, some staff – usually working for a 'central' unit with university-wide responsibilities, will have a remit that includes such things as changing the university-wide learning management system, upgrading the ICT facilities in lecture rooms, or installing wireless access points. Programme-level teaching teams usually have to take these as givens – they are inherited constraints over which they have very little power (Conole & Jones, forthcoming). Conversely, staff in 'central' units, even at very senior management levels, have little or no power to influence the details of how a course is taught. Failure to achieve a reasonable degree of integration up and down the levels can create major problems in an institution, and the uncertainties and miscommunications that arise can be a serious threat to sustainable innovation.

The Idea of Teaching-as-Design

Teaching in higher education is a multi-faceted job. Individual beliefs and preferences, departmental and disciplinary traditions, student numbers, the affordances of the teaching spaces available and access to other resources all combine to shape the ways in which teachers teach. The classic image of university teaching evokes lecturing to a large group, or tutoring a small one. But teaching also includes course and curriculum planning; consulting employers, professional bodies and academic colleagues about course goals, standards and requirements; setting assessment tasks; identifying useful learning resources; preparing lists of recommended reading and links to useful websites; creating or updating course websites; allocating students to working groups; scheduling activities; briefing teaching assistants (where available) and updating lecture notes. Most or all of this goes on before the first student comes onto campus. Once semester begins, in addition to lecturing, running seminars and tutorials, supervising laboratory or other practical classes, providing ad hoc advice about course requirements and counselling students who have academic or personal problems, university teachers also spend a great deal of time marking

assignments and providing feedback (though rarely in enough detail to keep students happy). At semester's end there may be an exam to invigilate, answers to mark, exam boards to attend and course evaluation data to scrutinise. A wise teacher will make notes about what went well and less well, what needs changing next time round, what should be kept and what scrapped.

Within this complex mix, the classic activities of lecturing and taking tutorials can be seen to make a significant, but far from exclusive, call on the teacher's time. And if we look at things from the students' point of view, these supposedly core teaching activities become even less salient. Of course, they are the main opportunities for students and teachers to meet, and they provide a good chance for students to get a better sense of what the teacher values (or says they value). But if we think about how the *students* spend their time, then the work that teachers put in to lectures and tutorials diminishes in significance. Especially in the arts and social sciences, students spend a great deal of time doing things other than attending lectures and talking in seminar or tutorial groups. They devote much more time to tackling the required reading and in preparing essays. Even in the sciences, where there may be much more timetabled face-to-face activity, the quality of what students do depends less on how well their teacher presents a lecture than on how well he or she has designed the activities in which the students are asked to engage. We are *not* saying that lectures are unimportant. Good lectures are very good ways of achieving some of the things that need to be achieved in higher education (Bligh, 2000; Schwartz & Bransford, 1998). The simple point is that students learn a very great deal outside the lecture room. As Tom Shuell pointed out, and as John Biggs has reminded us all, the thing that matters is *what the learner does* (Shuell, 1986, 1992; Biggs & Tang, 2007). The learner's mental activity is the thing that changes what they know: any changes in competence or understanding are dependent on what the learner does.

In thinking about educational outcomes and their improvement, it is therefore essential to take a view that focuses on learner activity. Everything else is at best secondary.

If the quality of students' learning activity is crucial, then much of what teachers do that makes a difference can be thought of as a kind of design work. For sure, significant amounts of teaching involve rapid interaction with students, improvisation, fast thinking and making decisions intuitively and on-the-fly. But much more of the work that teachers have to do actually provides them with opportunities to be more reflective and plan more carefully in their teaching. It is concerned with designing things such that the quality of students' learning activity is as good as it can be.

Focus on Learning: What Needs Designing?

This immediately raises the question of what needs designing and we are quick to say that it is not *learning* that needs to be designed. Indeed, one of our

objections to the term 'learning design' (e.g. Koper & Tattersall, 2005) is that it suggests the impossible – that learning itself can be designed. (It *might* be reasonable to talk about learners designing their own learning. We would nevertheless want to say that one cannot design somebody else's learning, just as one cannot design somebody else's feelings or experience.)

Learning activity is key: what the learner does is what makes a difference to their learning outcomes. As Figure 8.1 points out, such outcomes can be of various kinds – often classified as cognitive, psychomotor and affective. (Cognitive outcomes are concerned with thinking and understanding; psychomotor outcomes with skilful movement and perception; affective outcomes with emotions, attitudes and predispositions.) Giving due recognition to other accounts of learning, we would also want to give space, among candidate learning outcomes, to such things as shifts in one's sense of identity, and the degree to which one has become part of a community or culture, adopting its values, perspectives and practices. These too are reflected in Figure 8.1.)

Figure 8.1 also emphasises the activity to outcome connection. While different kinds of outcome are achieved in different ways, all are dependent on activity. That activity may be mental (e.g., quiet reflection, or practising vocabulary). It may be physical (e.g., using laboratory equipment, or carrying out a dissection). It will often be both. It may involve tools or other artefacts (e.g., in using a computer or reading a book).

Once the centrality of students' activity is recognised and accepted, the teacher's design attention can focus on that which *can* be designed. What is the strongest influence on what students actually do? Usually, it is the tasks which teachers set that have the strongest influence on students' activity. Of course, students rarely do exactly what they are told. Indeed, teachers often set tasks in such a way that some creative interpretation is necessary and so that students can adapt the task to create a better fit with their own needs and interests. In addition, there are likely to be numerous other influences on what students actually do and some of these will – from time to time – prove more powerful

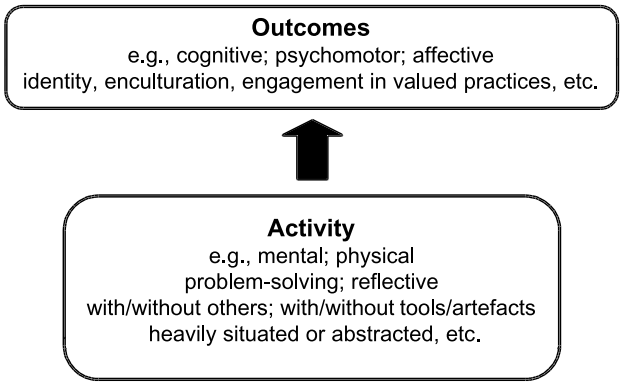

Figure 8.1 Learning outcomes depend on what the learner does.

than what the teacher has set. Nevertheless, a crucial resource for students' learning activity is the task set by the teacher.

Figure 8.2 represents the fact that learning tasks (as set by the teacher) are transformed by the student through their own interpretive and other work, such that it is the student's activity that mediates between the task as set and the educational outcomes achieved. Tasks as set have an *indirect* effect upon learning outcomes (Winne & Marx, 1982; Goodyear, 2000).

There are numerous potential influences on the processes of transformation that turn a task specification into an activity. For activities that last a significant time – hours or days rather than seconds or minutes – the processes that shape and reshape the actual activity can be quite complex, and change over time. As an example, think of a student setting out to tackle an assessment task that involves writing a 3,000 word essay. Imagine there is a deadline six weeks away, that the essay is worth 30% of the total marks for the course, and that the teacher has set a specific essay question. There will normally be an initial stage in which the student tries to make sense of the task specification – what the teacher means by the essay title, what qualities they will be looking for in the final essay, how much scope there is for personal interpretation, what kinds of reading will be needed, what kinds of evidence will need to be deployed? After that, as time goes by, the student will work on the essay: by reading, thinking, making notes, talking to peers, drafting and revising. The quality of this real-world activity will be influenced by a number of factors, including the other competing demands on the student's time, energy and attention. There will be some serendipity in the work, and some slippage and satisficing of goals as the deadline approaches. In short, much may happen to intervene between the task as the teacher set it and the learning outcomes that result from the student's activity. Many of these outcomes may best be thought of as *by-products* of the activity of producing the artefact – in this case an essay – that the task requires.

As we have seen in Chapters 4 and 5, students' conceptions of learning, and approaches to study, are heavily bound up in this task – activity nexus. Indeed, one can argue that an approach to study is a way of describing what is happening when students translate tasks they are set into actual learning activity. When we interviewed them to find out about their approaches to study, what they described was their *intentions* – what they hoped to achieve by acting in particular ways – and their *strategies* – what they actually did. Both of these are retrospective, and somewhat abstracted and tidied up, *accounts* of their translations of tasks into activities, but they nevertheless give us some compelling insights into the task – activity nexus.

Because of this, it must be acknowledged that the connections between tasks as set and learning outcomes are sometimes tenuous. Moreover, it is likely that only a subset of what the student has learnt will be visible to the teacher in the artefact submitted by the student. Does this make task design

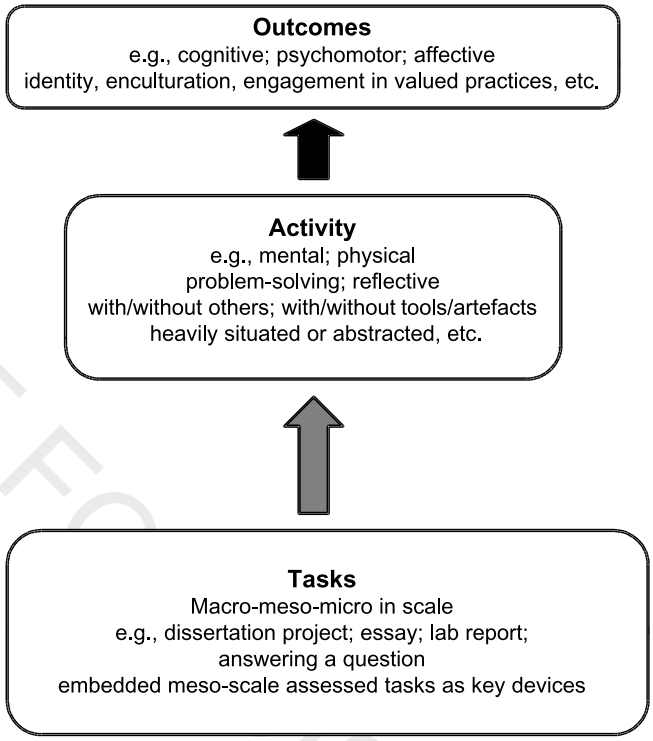

Figure 8.2 Tasks as set have an indirect effect on learning outcomes.

irrelevant? No, it does not. In some circumstances it may be appropriate for the teacher to design a task such that there is a much *tighter* link between task and activity (and thereby, though with less certainty, because the quality of activity will still be variable, to outcomes). Areas of safety-critical training may well require this kind of tight coupling. At the opposite extreme, teachers can sometimes be seen to hand over task design, largely or completely, to the student. Examples would be in programmes that use student-designed or negotiated assessment tasks as key learning activities (see e.g. Brook, Hunt, & Hughes, 1994; Boud, Cohen, & Sampson 1999). But, we would argue, much of higher education practice, for good as well as bad reasons, sits between these extremes and necessarily involves situations in which there is some interpretative work, satisficing and slippage in the task–activity nexus (Goodyear & Ellis, 2008). This being so, task design remains a core responsibility for most teachers.

Figure 8.3 emphasises the point that tasks are not the only influence on student learning activity. As we explained in Chapter 2, learning activity is *situated*. The nature of the activity is influenced by the social and physical context in which it occurs. By pointing out that learning activity is socially situated, we mean that the web of relationships with other people influences the activity: what is done, how it is done and, to an extent, why it is done. We label the social context with the term 'people' in Figure 8.3. In some situations,

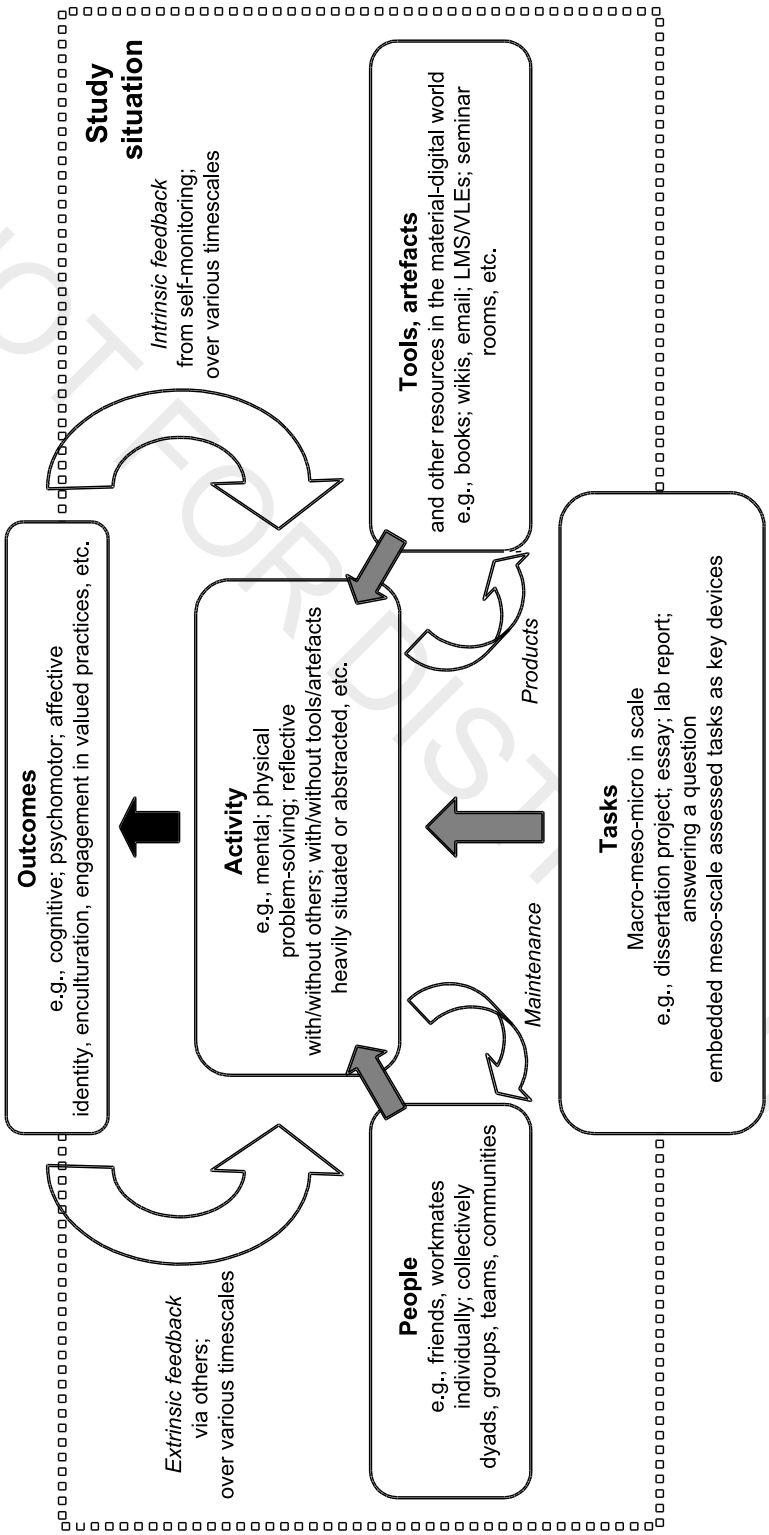

Figure 8.3 Key influences on learning activity and learning outcomes.

it may be that the strongest influence on the learner's activity comes from their relationship with just one other person: a close friend or colleague. In other situations, the influence may come from a small group – perhaps a workgroup to which the student has been assigned by the teacher. Bigger and more diffuse social groupings may also have influence. Examples would be where the student feels themself to be part of a learning community, or an apprentice member of a community of practice. In such cases, community norms, values, practices, language, history, etc., can influence the student, and how they approach their work, in ways that are sometimes powerful and sometimes quite subtle. Among the ways in which relations with others can influence the student's activity, we can think of things such as the manner in which work, effort and academic achievement are talked about, and the images of being a 'good' or 'savvy' student that emerge from such talk. Also, we can point to the ways in which divisions of labour allow some things to be done, while other experiences get closed off. The division of effort within a workgroup has strong implications for the nature of an individual student's activity, and what they gain from it.

In a similar fashion, we need also to think about the ways in which learning activity is *physically* situated. (For simplicity, we include the digital in the physical here.) What someone can do is often influenced quite powerfully by the tools and resources that come to hand (Säljö, 1995; Bowker, 2006), though the effects can sometimes be subtle and difficult to unravel (Salomon, 1993; Hutchins, 1995).

These social-cultural and physical-digital influences on the learner's activity need to be taken into account by the teacher as designer. A key point is that the teacher cannot design the social relationships within which the learner's activity is embedded. Neither can they design all the details of the learning place – the immediate, concrete, physical–digital environment in which a learner's activity is set. Students play a significant role in configuring their learning places. They make choices about the tools and resources they will use. They take responsibility for creating a 'defensible space for learning' – within which they can protect themselves from unwanted interruptions and can bring to hand the things they need. But they do this in an environment that they have not created themselves – whether we are thinking of studying on campus, at home, at work or while travelling. Students, teachers and others *co-configure* the learning place (cf. Nardi & O'Day, 1999; Crook & Barrowcliff, 2001). In a similar vein, teachers cannot prescribe friendships and workmates. Students exercise considerable freedom in choosing who to spend their time with, who they talk to about their work, and whose views they most value. Even in groupwork tasks, where the teacher has determined group membership, students still exercise some freedom in choosing how much or how little effort they will invest in group tasks, how far they will rely on their colleagues to get the job done, and the extent to which they will turn a group task into a

number of parallel individual tasks. But they do this in a social milieu that has, in part, been created by the actions of teachers. Teachers and administrators create such things as classes and year groups. Many teachers are interested in creating the conditions within which a 'learning community' can prosper. They cannot create or design a community, but there is social engineering work they can do to make it more likely that community will evolve and thrive. In this sense, relations between activity and context exemplify Giddens's notion of *structuration*, which we introduced in Chapter 1. Students exhibit agency – a significant level of control over their activity – within a context that influences, but is also created by, aspects of their activity. In Figure 8.3, we have tried to represent this with the feedback loops labelled 'products' and 'maintenance'. The 'products' loop represents the fact that some student activity, either directly, or as a by-product, creates and assembles resources, tools and other artefacts that enrich the students' learning places (Wild, Modritscher, & Sigdarson, 2008). The 'maintenance' loop represents those aspects of student activity that, directly or indirectly, have the effect of maintaining and improving their social/working relationships (with peers, teachers, members of their extended communities, etc.).

To sum up, Figure 8.3 alerts us to the following issues. First, learning outcomes depend on student activity. Tasks are resources for activity, not prescriptions of it. Activity is also strongly influenced by social and physical context. Teachers cannot design these in detail, but they do have a responsibility for setting in place some of the social and physical resources on which students will draw. Design is therefore an indirect process, through which teachers aim to influence but not control the students' activity.

Layers in Design

Figure 8.3 is a very abstract way of thinking about the problem-space of educational design. Lots of very different things are grouped together under the headings 'people' and 'tools, artefacts, etc.' Building a shared understanding of teaching-as-design is easier if we now shift to looking at some of the complexities hidden away by Figure 8.3.

One recognised way of simplifying a complex problem is to break it down into smaller parts. There are lots of models of educational design that take a 'divide and conquer' approach. They take a big problem, such as the design of a whole course, and break it down into more manageable chunks (e.g. planning a lecture at a time). While this can be helpful in tackling problems that are difficult because they are *large*, it is not necessarily helpful in dealing with problems that are difficult because they are *complex* (e.g. ones in which there are many internal dependencies between the elements of the solution). For complex problems that involve lots of interdependent elements, a good strategy is to simplify by separating out the thing that has to be designed into several relatively self-sufficient functional elements or 'layers'. A useful analogy

would be to think about the design of a building and to visualise separate layers that provide different functions. For example, Stuart Brand (1997, pp. 13) suggests that the design of any building can be decomposed into six layers: site, structure, skin, services, space plan and stuff.

- site: the geographical setting and the legally defined building site, having boundaries and context;
- structure: the foundation and load-bearing elements of the building;
- skin: the exterior surfaces;
- services: the communications wiring, electrical wiring, plumbing, sprinkler system, heating, ventilating, air conditioning, and moving parts like lifts and escalators;
- space plan: the interior layout – where walls, ceilings, floors, and doors go;
- stuff : Chairs, desks, phones, pictures, kitchen appliances, lamps, etc.: things that move around inside spaces.

An instance of this approach of simplifying through layering, in the educational design area, is Goodyear's four-layer 'pedagogical framework' (see e.g., Goodyear, 1999, 2005). This offers a way of thinking about tasks and intended learning outcomes by identifying four loosely-coupled design layers.

Goodyear's notion of a 'pedagogical framework' emerged as part of an attempt to compare and contrast the main educational features of a large number of technology-based learning innovations funded by the European Union's Socrates and Minerva programmes. It is a simple attempt to identify and relate four main layers in the design and execution of these innovations (Table 8.1).

The pedagogical framework offers a way of handling some of the complexity of task design, by suggesting that most of the issues involved in enacting a pedagogical approach can be categorised as being concerned with one of the four levels shown in Table 8.1. It helps structure the problem space of deciding *how* teaching will be conducted. This can be particularly helpful when one also acknowledges that the four levels need not be tightly-coupled. While there are usually advantages to having coherence across levels, the relations are not tight, deterministic ones. There is scope for some freedom of action, improvisation, the introduction of variety, etc.

The pedagogical framework is principally concerned with the 'how' of teaching and learning, rather than with *what* is to be learnt. Decisions about what is to be learnt map onto the specification of learning outcomes (Figure 8.3). Considering scale and layering introduces the notion that intended learning outcomes vary in scope, and can be seen as nested. There can be very broad brush outcomes (e.g., 'acquire the ability to communicate effectively with a range of personnel'), nested within which one can find more specific or

Table 8.1 Pedagogical Framework

Philosophy	Beliefs about the nature of knowledge and competence; about how learning occurs; about how people should and should not be treated, etc. Examples might include positivism, realism & interpretivism; constructivism & behaviourism.
High level pedagogy (HLP)	A broad educational approach, such as cognitive apprenticeship, collaborative knowledge-building, problem-based learning or programmed learning. An HLP does not contain direct prescriptions for action, but it foregrounds some kinds of possible action, and backgrounds others. An HLP is a way of turning a philosophical position into a space of commitments and possibilities.
Pedagogical strategy	A broad brush depiction of plans, e.g. about a sequence of tasks that will be proposed. Best seen as a way of communicating and negotiating within a programme team and with students. Best stripped of its commercial and militaristic resonances of outwitting a rival or enemy. A device for joint sense-making, which avoids the fine detail of pedagogical tactics. Examples would include collaboration scripts.
Pedagogical tactics	The detailed moves through which a strategy is enacted, or out of which a strategy emerges.

Source: Goodyear (1999)

finer-grained outcomes (e.g., 'recognise the dangers inherent in writing satirical memos').

Gibbons and Rogers (2008) have also drawn on the layering approach as a way of separating out design issues that can be treated in bundles. Their approach uses seven layers: content, strategy, message, control, representation, media-logic and data management. These are defined in Table 8.2.

In our view, the layering used in Table 8.2 is particularly appropriate if one thinks about educational design as being focussed on the crafting of communication and messages.

Our own preference is for a focus on *situated activity*. Receiving and interpreting instructional messages is one kind of activity and although it is important, we do not see it as necessarily the best focus when thinking broadly about learning and educational design. From a broader perspective, we can use the concept of design layering to help decouple things which are best managed separately, and also to help identify who has responsibility for what. The layering idea can be applied to each of the three main design components we discussed with the help of Figure 8.3: tasks, 'people' and 'tools, artefacts, places'.

Table 8.2 Gibbons's Instructional Design Layers

Content layer.	A design must specify the structures of the abstract subject-matter to be taught, must identify the units into which the subject-matter will be divided, and must describe how elements of subject-matter will be made available to instructional functions performed by other layers.
Strategy layer.	A design must specify the physical organisation of the learning space, social organisations of participants, their roles and responsibilities, instructional goals, allocation of goals to timed event structures, and strategic patterns of interaction between the learner and the instructional experience.
Message layer.	A design must specify the tactical language of message structures through which the instructional experience can communicate content-derived information to the learner.
Control layer.	A design must specify the language of control structures through which the learner expresses messages and actions to the source of the learning experience.
Representation layer.	A design must specify the representations that make message elements visible, hearable, and otherwise sense-able: the media representation channels to be used, the rule for assigning message elements to media channels, the form and composition of the representation, the synchronisation of messages delivered through the multiple channels, and the representations of content.
Media-logic layer.	A design must specify the mechanism by which representations are caused to occur in their designed or computed sequence.
Data management layer.	A design must specify data to be captured, archived, analysed, interpreted, and reported.

Source: Adapted from Gibbons & Rodgers (2008)

One way of doing this is to see what comes into focus when we look at the main areas of educational design at different scale levels. For example, the area of task design changes some of its characteristics when viewed at different temporal scale levels (e.g. tasks that last a few minutes or a few hours; tasks that extend over a whole degree programme). The 'people' area also has different characteristics when viewed at different social scale levels. There are important differences between (a) putting students into pairs to tackle a task, and (b) arranging a programme so that it can support both full-time and part-time students. The design area that we have labelled 'tools, artefacts, etc.' can be seen as a nested hierarchy ranging from a whole university campus down to a mouse, pen, PDA or paper. (The interpenetration of the digital and material

worlds upsets this neat hierarchy, but for now this simplification is not unhelpful.)

Different groupings of university staff have responsibilities for different scale levels or layers in these design areas. The solo teacher may well have a lot of autonomy to design short tasks (e.g. holding a buzz group in a lecture) but will usually just inherit, as a given, some of the constraints set by arrangements at larger scale levels (e.g. the length of a semester). Programme teams can work at micro and meso levels. For example, they will often have some say in the kinds of rooms to be used for face-to-face work with students. They typically have a high level of control over the intended outcomes of a program, though within constraints set by the university's statements about generic graduate attributes, and/or requirements set by professional accrediting bodies. The macro level – the broadest layers of design activity – tend to be the preserve of university managers and staff who have a university-wide remit. We are thinking here of such things as decisions about which learning management system to use, about whether to increase the number of lecture rooms or flexible learning spaces, about whether to create a new programme or close down programs that are under-recruiting, etc. In a well-functioning university, there will, of course, be a great deal of information flow and negotiation across layers and levels. Also, as we have noted above, none of the people working in a design role has complete control over any of the design areas. That said, it is still worth observing that different groups tend to orient their design activity to particular layers in the design space.

Also relevant to this point is that design activity at different levels tends to draw on different knowledge sources. For example, it is much easier to make good use of psychological research arising from detailed studies of short learning episodes if one is thinking about the design of short learning tasks than if one is considering broader issues of course design. Research on knowledge-sharing activities in communities of practice is easier to use if one is thinking about programme design rather than whether to use a buzz group in tomorrow's lecture.

Self-Awareness, Feedback and Self-correction: Iterative Design and Sustainable Improvement

University teaching is a complex activity which a conscious design approach can help to simplify and improve. Thinking about how teaching-as-design fits into the broader ecology of learning is a good way of reminding ourselves that students' learning activity is at the heart of design consciousness – awareness of that central purpose of 'good learning' is key. As we saw in Chapter 6, it is not always the case that teachers adopt a carefully planned approach, or one that is centred on good ideas about student learning.

Another key aspect of this awareness (design consciousness) is that once a task is designed, it develops a life of its own. It becomes part of the students'

experience as a learning activity and they may respond to it in more or less helpful ways for the development of their understanding. Effective designers will approach their work knowing that feedback from student learning can help with the incremental improvement of designed tasks.

Incremental revisions of designed tasks will vary depending on level in the ecology: the micro, for example, might focus on the internal structure of the key aspects of the task; the meso, might be concerned with how teams of teachers, tutors and/or students are likely to interact with the task in light of faculty policies and culture; the macro might consider the interplay of university enterprise technologies and physical spaces to afford learning. Another key aspect of the development of self-awareness about task design is that designing is rarely undertaken on a *tabula rasa*, even when tasks are being designed for the first time (Laurillard, 2002). The ecology in which the design occurs will provide constraints of which the teacher must become aware.

Teachers rarely think that they get the design of a task right the first time. It is reasonable to assume that most designs will go through an iterative process of use, feedback, redesign and reuse. In an ecology of learning, this cycle tightens the links between design and teaching, reinforcing the idea of teaching-as-design.

We find it helpful to think about the traces of student activity as an artefact that can inform the incremental improvement of a design. The use of student feedback, obtained via course experience questionnaires, focus groups, etc., can be very valuable in thinking about the improvement of a course. But it is just one source of data. When students are making use of virtual environments, for example to hold discussions, share work or co-ordinate activity, they leave persisting traces. These may be simply the artefacts they have created – such as discussion logs – or they may also include various representations of group processes – such as graphs showing contributions over time. In short, the whole set of online traces left by students at the end of a course is a potentially valuable, complex artefact. It may sometimes be too complex to use in the time-pressed business of setting up a new virtual space for the course the next time it runs. But routinely to delete such artefacts is wasteful.

There are a number of alternatives. One is to design the virtual space for a course so that parts of it persist from year to year – allowing students to read, re-use and improve upon the work of previous years, while having their own fresh-minted space to operate within. Another is for the teacher to select what they see as most useful and repurpose it for next year, editing and annotating as they see fit.

Our main point is that the 'learning loop' of feedback and incremental redesign means that the quality of a design improves steadily over time. If everything has to be recreated from scratch, the teacher's time is dissipated over an excessive range of jobs: some of which do not need academic expertise and attention. If the teacher's time can be focussed on selective improvements,

then quality ratchets upwards. Having a sharp sense of 'good learning' makes it much easier to prioritise areas for improvement.

From an ecological perspective, feedback on design also needs to flow from other sources, not just from the students. For example, some aspects of a design may have unintended consequences for library staff, IT services or other teachers. In a well-functioning ecology, such information flows to where it is needed, on a timely basis.

Aggregation of feedback on task design across many courses can be particularly illuminating for staff who have faculty-wide or university-wide remits. Not least, being able to spot trends in technology choices, and in other decisions that have consequences for infrastructure, means that planning can stay ahead of the game, rather than merely reacting to crises. We examine some of these challenges more deeply in Chapters 9 and 10.