Introduction

This paper sets out to address some apparently simple questions:

- *What is good pedagogy?*
- *What kinds of frameworks or tools could help us to capture it?*
- *How could this promote better learning?*

In focusing on these questions, we recognise that it may seem more obvious to start thinking about teachers' professional learning and development by focusing on the necessary conditions for such learning to occur. For example, we might argue that teachers need to feel trusted and valued, that their experiences and perspectives are acknowledged, that the culture of the schools in which they work should promote critical questioning and innovative approaches, with space and encouragement for discussion and sharing of ideas. We will return to these issues, but first we focus on what that learning should be. Again, it might seem obvious that this is already well known: we surely know what great teaching looks like; we just need to create the culture in which teachers feel empowered and free to do it.

In fact, there is some evidence that an understanding of what constitutes effective pedagogy – the method and practice of teaching – may not be so widely shared, and even where it is widely shared it may not actually be right (Strong et al, 2011; Hamre et al, 2009). Hence it is necessary to clarify what is known about effective pedagogy before we can think about how to promote it. Unless we do that there is a real danger that we end up promoting teaching practices that are no more – and perhaps less – effective than those currently used.

We also review research that has shed some light on what works in terms of the practices of professional learning – whether it is the frameworks used to define teaching effectiveness or observing peers, entering into dialogue and feedback and helping to improve practice.

This study presents a brief review of the existing research evidence that is relevant to these questions. The original research questions we set out to address are given in full in Appendix A.

What is good pedagogy? Elements of teaching effectiveness

Defining 'good pedagogy'

Defining effective teaching is of course problematic. Ideally, we might define effective teaching as that which leads to high achievement by students in valued outcomes, other things being equal. We acknowledge that available assessments – and particularly those that have been used for high-stakes accountability or in existing research studies – may not fully capture the range of the outcomes that we might specify as desirable aims for education (Popham and Ryan, 2012; Muijs et al, 2014; Polikoff, 2014).

We also acknowledge that 'other things being equal' may be open to different interpretations about what factors should or can be taken into account. A number of factors will influence students' achievements, for example, pre-existing student characteristics (both of individual students and collectively), characteristics of the school and of the teacher (some of which may be alterable, others not), and of the context. In practice, the attribution of an 'effect' to an individual teacher or school is generally determined by what cannot be explained by factors that are judged to be outside the control of that individual (Raudenbush, 2004). This kind of 'residual attribution' – interpreting value-added simplistically as the effect of the teacher – is, of course, problematic (Newton et al, 2010; Hill et al, 2011; Dumay et al, 2013).

Despite these limitations, wherever possible, it makes sense to judge the effectiveness of teaching from its impact on assessed learning. If the assessments and value-added models available to us are not good enough, we need to improve them. In the meantime we must exercise some caution in interpreting any claims about teaching effectiveness.

A further concern is that in practice, any kinds of observational measures provide at best poor approximations to how much students actually learn. Whether they are based on classroom observation, student surveys, book scrutiny or other sources, their predictive power is usually not high. For example, even in a high-quality research study such as the Measures of Effective Teaching Project (Mihaly et al, 2013, Table 3, p24), the median correlation between a range of value-added and observation ratings was only 0.3. Although a correlation of 0.3 will often be presented as 'highly significant' by researchers, in practice it means that if we were to use classroom observation ratings to identify teachers as 'above' or 'below' average in their impact on student learning we would get it right about 60% of the time, compared with the 50% we would get by just tossing a coin. It is better than chance, but not by much; there is information in classroom observation, but not enough to base important decisions on it. And of course, this is a best-case: with regular teachers or principals using un-validated observation protocols and no quality assurance process to check judgements are aligned, the correlation will be much less, perhaps even negative (Strong et al, 2011).

Developing indicators of good pedagogy that can be used reliably

There are at least two kinds of problems we could encounter in trying to 'operationalise' good pedagogy - that is developing a set of measures of good (and great) pedagogy that can be reliably used to assess teacher effectiveness. One is to be too specific: to define it in terms of a checklist of observable, effective practices or skills. A potential problem with trying to reduce great teaching to constituent elements is that the whole may be greater than the sum of its parts. The choices a teacher makes in orchestrating their skills may be an essential part of what makes them effective. Focusing on the behaviours themselves will always be too limited. Instead we need to think in terms of a professional pedagogy in which judgement is an essential component. Nevertheless, evaluating the quality of such choices is unlikely to be straightforward.

The other problem is not to be specific enough. Although it is important to be clear about the principles that underpin pedagogy (James and Pollard, 2011), we must also relate them to something that is observable. Theory must be specific enough to be empirically testable and a guide to well-defined actions.

Shulman (1988, p38) has written of the need for "a union of insufficiencies, a marriage of complements, in which the flaws of individual approaches to assessment are offset by the virtues of their fellows". His argument was that although each individual measure of some aspect of teaching effectiveness may be flawed and inadequate, when our view is informed by a varied collection of such measures their failings can be overcome. However, this view seems not to take into account how we might assess the teacher's role in selecting and orchestrating these 'effective' approaches, nor does it address the practical difficulties of turning an array of insufficient indicators into a meaningful whole. Indeed, Shulman himself seems later to have retracted this view (Shulman, 2009). Before we can think about the validity of any measures of teaching effectiveness we need to be clear what those measures are intended to be used for. On some wish-lists will be requirements: for use in selection for initial professional entry; for awarding certification as a qualified teacher; for recognising professional progression, perhaps linked to probation, tenure, promotion, retention, or performance-related pay; for identifying under-performing teachers, with associated support or firing. Unfortunately, the evidence seems clear that our best currently available measures of teaching effectiveness are not adequate for most of these kinds of purposes (Gitomer, 2009).

Our purpose here is a little different. We take the view that low-stakes, formative use of teaching effectiveness indicators, with an emphasis on feedback, support and challenge, and professional learning, may lead to improvements in student learning, even if those indicators are in many ways 'insufficient'. In this we echo Shulman's (2009) distinction between assessment *of* teaching and assessment *for* teaching. However, where Shulman emphasises creating measures for which 'the very act of preparing for and engaging in assessment would be a powerful form of professional development' (p241), we also stress the role of feedback from and discussion about the results of an assessment in professional learning, and the

role of a clearly specified framework of performance indicators to focus teachers' attention and effort on things that are important.

With this approach, our criterion for validating a measure of teaching effectiveness is not 'Does it produce a complete, unbiased and accurate measure of a teacher's impact on student learning?', but 'Can using it as part of a system of self-evaluation, feedback, dialogue and re-assessment lead to improvements in student learning?'. In technical terms, we value consequential validity over criterion-related validity. This perspective also allows us to acknowledge that quality teaching is multidimensional: a profile of multiple, independent strengths and weaknesses may be more useful – and a better fit to reality – than a single, unidimensional measure.

Types of evidence relevant to 'effectiveness'

There are a number of sources of evidence about the skills, knowledge, behaviours, qualities and competences required to be an excellent teacher. A key feature of the current review is that we try to limit our attention to well-defined, operationalisable behaviours, skills or knowledge that have been found to be related, with at least some justification for a causal relationship, to measureable, enhanced student outcomes. Following Rosenshine (2010, 2012) and Muijs et al (2014), these sources of evidence include:

- Evidence from educational effectiveness research about teacher behaviours associated with learning gains

- Evidence from intervention studies about what can be changed, and its effect on outcomes

- Evidence and theory from cognitive science about learning: how our brains acquire, make sense of and use information

There are two key requirements for the inclusion of a teaching approach as 'great teaching' in this review:

- There must be a clear, well-specified and implementable intervention associated with promoting the approach. It has to be something we can change. For example, the knowledge that 'great teachers have high expectations' is of no use to us unless we have a strategy for encouraging teachers to raise their expectations

- There must be some evidence linking the approach with enhanced student outcomes. There is not necessarily any assumption that such outcomes should be limited to academic attainment: whatever is valued in education should count.

One of the features of research on effective practices is that there are a number of reviews available with quite different claims about what characteristics of teacher practice are associated with improved outcomes. For example, a review by Husbands and Pearce (2012) contains 'Nine claims from research', of which the first is that 'Effective pedagogies give serious consideration to pupil voice' (p3). A

good definition of ‘pupil voice’ is given, but as far as we can tell, none of the studies cited contain robust evidence to link it causally to improvements in pupil outcomes. There is some evidence of a link to changes in teachers’ practices and perceptions, and to more positive attitudes for both teachers and students, though many of even these studies would not meet basic quality standards for robust support of such claims. Using pupil voice may indeed be an effective pedagogy, but we believe that the evidence currently available does not support this claim, so have not included it.

However, we acknowledge that the question of what teaching practices are shown by research to be effective remains contested. An example from England is Brown et al’s (2001) analysis of different views of the research basis of the National Numeracy Strategy. From the US an example is Boaler’s (2008) critique of *The Final Report of the National Mathematics Advisory Panel* (National Mathematics Advisory Panel, 2008).

It is also clear that a lot of the research which has set out to discover the elements of effective teaching have simply asked the wrong questions. As Good and Biddle pointed out more than 25 years ago, looking back then over at least 20 years of this kind of research,

> At various times educators in this century have advocated as answers large-group instruction, smallgroup teaching and individualised teaching!...However it seems clear that simple characteristics of instruction have never predicted instructional effectiveness...The issue is not individualised instruction or small-group instruction, but rather the quality of thought and effort that can occur within these structures...(Good & Biddle, 1988 p.116)

A salutary example is from Brown et al (2001), who confidently identified a list of instructional practices that empirically distinguished effective from less effective teachers, as determined by their students’ learning gains. They then tested the predictive power of an observation schedule based on evaluating these practices for a different group, but found the results rather disappointing:

> We are therefore left with the perhaps rather happy conclusion that the behaviour of effective teachers and less effective teachers are not easily characterised; much depends on the particular way that teachers and classes as people relate together. There are signs that certain types of behaviour may often lead to higher gains, but there are always exceptions in both directions.

A final caution is from the US *National Mathematics Advisory Panel* (2008):

> Unfortunately, little is known from existing high-quality research about what effective teachers do to generate greater gains in student learning. Further research is needed to identify and more carefully define the skills and practices underlying these differences in teachers’ effectiveness, and how to develop them in teacher preparation programs.

Examples of effective practices

In this section we present a collection of teacher behaviours, approaches, classroom practices and skills that meet our criteria of being well-defined, implementable and linked to gains in student outcomes. We have sought to include here some practices that are counterintuitive, or that challenge the accepted orthodoxy about what is effective teaching, on the grounds that these examples may have value more as a prompt to critical questioning rather than a checklist of desirable behaviours. Teachers may need to have clear understanding of why, when and how each of these practices can be effective, and exactly what it means to demonstrate them in a way that is optimal to promote students' learning. Good summaries of the wider evidence about effective practices can be found in Muijs et al (2014) and in Ko et al (2013).

Some important caveats are required before presenting these examples of 'effective practice'. All of them are open to interpretation. All of them could be done well or done badly. All of them could be inappropriate in some contexts and appropriate in others. For these reasons it may be unproductive or even harmful to treat them as if their meaning is unproblematic or to require them as a recipe or formula. Nevertheless, they are all supported by robust evidence of positive impact on student learning, so may be seen as offering at least a 'starter kit' for thinking about effective pedagogy.

Danielson's Framework for Teaching

The use of this framework as a classroom observation instrument is discussed in more detail below (p31), but for now we present an outline of the elements that are evaluated.

1. **Planning and preparation**
    a. Demonstrating Knowledge of Content and Pedagogy
    b. Demonstrating Knowledge of Students
    c. Setting Instructional Outcomes
    d. Demonstrating Knowledge of Resources
    e. Designing Coherent Instruction
    f. Designing Student Assessments

2. **Classroom environment**
    a. Creating an Environment of Respect and Rapport
    b. Establishing a Culture for Learning
    c. Managing Classroom Procedures
    d. Managing Student Behaviour
    e. Organizing Physical Space
3. **Instruction**
    a. Communicating with Students
    b. Using Questioning and Discussion Techniques
    c. Engaging Students in Learning
    d. Using Assessment in Instruction
    e. Demonstrating Flexibility and Responsiveness
4. **Professional responsibilities**
    a. Reflecting on Teaching

b. Maintaining Accurate Records
c. Communicating with Families
d. Participating in the Professional Community
e. Growing and Developing Professionally
f. Showing Professionalism

The Classroom Assessment Scoring System (CLASS)

CLASS (Pianta, La Paro, & Hamre, 2008) is an evaluation framework for classroom observation that identifies three main domains and a number of dimensions within each:

Emotional Support

Classroom climate (positive and negative) – warmth, respect, enjoyment, enthusiasm
Teacher sensitivity to student needs
Regard for student perspectives – respect for student autonomy, interests, motivations

Classroom Organization

Behavior management
Productivity – time management, maximizing opportunity to learn
Instructional learning formats – activities that maximize engagement

Instructional Support

Concept development – focus on higher order thinking
Quality of feedback
Language modelling – questioning, expanding, use of vocabulary

Rosenshine's Principles of Instruction

Rosenshine (2010, 2012) has summarised at least 40 years of research on effective instruction with a key set of principles that maximise its impact. The starting point for this evidence base is a set of correlational studies linking particular observed classroom teacher behaviours with higher student outcomes. For each of these principles there is also experimental evidence showing that attempts to train teachers in adopting these behaviours can result in changes in teacher behaviours and improvements in student outcomes.

In outline the ten principles are:

1. Begin a lesson with a short review of previous learning
2. Present new material in small steps, with student practice after each step
3. Ask a large number of questions and check the responses of all students
4. Provide models for problem solving and worked examples
5. Guide student practice

6. Check for student understanding
7. Obtain a high success rate
8. Provide scaffolds for difficult tasks
9. Require and monitor independent practice
10. Engage students in weekly and monthly review

Creemers and Kyriakides' Dynamic Model

A huge body of research in the educational effectiveness tradition has focused on the characteristics of schools and teachers that are associated with high learning gains. Much of the evidence is correlational, cross-sectional and lacking a strong theoretical foundation (Scheerens et al, 2001). However, the Dynamic Model (Creemers & Kyriakides, 2006, 2011) is empirically grounded, well enough specified to be testable and has indeed been subjected to considerable testing and verification.

The model identifies 21 particular teaching practices, grouped under eight headings. Creemers & Kyriakides (2011) have also developed a set of instruments for capturing these practices, consisting of two low-inference classroom observation instruments, a high-inference observational instrument and a student questionnaire, together with a teacher questionnaire for measuring school factors.

Table 1: The dynamic model of educational effectiveness (Creemers & Kyriakides, 2006)

(1) Orientation	**(a) Providing the objectives for which a specific task/lesson/series of lessons take(s) place** **(b) Challenging students to identify the reason why an activity is taking place in the lesson.**
(2) Structuring	(a) Beginning with overviews and/or review of objectives (b) Outlining the content to be covered and signalling transitions between lesson parts (c) Drawing attention to and reviewing main ideas.
(3) Questioning	(a) Raising different types of questions (i.e., process and product) at appropriate difficulty level (b) Giving time for students to respond (c) Dealing with student responses.
(4) Teaching modelling **.**	(a) Encouraging students to use problem-solving strategies presented by the teacher or other classmates (b) Inviting students to develop strategies (c) Promoting the idea of modelling
(5) Application	(a) Using seatwork or small-group tasks in order to provide needed practice and application opportunities (b) Using application tasks as starting points for the next step of teaching and learning.
(6) The classroom as a learning environment	(a) Establishing on-task behaviour through the interactions they promote (i.e., teacher–student and student–student interactions) (b) Dealing with classroom disorder and student competition through establishing rules, persuading students to respect them and using the rules.
(7) Management of time	(a) Organizing the classroom environment (b) Maximizing engagement rates.
(8) Assessment	(a) Using appropriate techniques to collect data on student knowledge and skills (b) Analysing data in order to identify student needs and report the results to students and parents. (c) Teachers evaluating their own practices.

Evidence from cognitive psychology

Because of the fragmentation of academic disciplines, a parallel source of evidence can be found in research in cognitive psychology that has investigated the nature of learning, the conditions under which it occurs and the role of memory in this process. A good summary can be found in Bransford, Brown, & Cocking (2000).

One paradoxical finding is that some approaches that may appear to make learning harder in the short term, and less satisfying for learners, actually result in better long-term retention. Emphasising the difference between short-term *performance* and long-term *learning*, Bjork and Bjork (2011) call these 'desirable difficulties', and give four specific examples:

- **Varying the Conditions of Practice**: Varying the learning context, types of task or practice, rather than keeping them constant and predictable, improves later retention, even though it makes learning harder in the short term.

- **Spacing Study or Practice Sessions**: The same amount of time spent reviewing or practising leads to much greater long-term retention if it is spread out, with gaps in between to allow forgetting. This "is one of the most general and robust effects from across the entire history of experimental research on learning and memory." (Bjork and Bjork, 2011, p59).

- **Interleaving versus Blocking Instruction on Separate To-Be-Learned Tasks**: Learning in a single block can create better immediate performance and higher confidence, but interleaving with other tasks or topics leads to better long-term retention and transfer of skills.

- **Generation Effects and Using Tests (Rather Than Presentations) as Learning Events**: Having to generate an answer or procedure, or having to retrieve information – even if no feedback is given – leads to better long-term recall than simply studying, though not necessarily in the short-term. Testing can also support self-monitoring and focus subsequent study more effectively. "Basically, any time that you, as a learner, look up an answer or have somebody tell or show you something that you could, drawing on current cues and your past knowledge, generate instead, you rob yourself of a powerful learning opportunity" (Bjork and Bjork, 2011, p61).

A recent and comprehensive summary of the impact, strength of evidence and generality of conditions under which a number of learning techniques have been shown to be effective is presented by Dunlosky et al (2013).

Table 2: Effectiveness of ten learning techniques, from Dunlosky et al (2013)

High utility	**Practice testing**
	Self-testing or taking practice tests on material to be learned
	Distributed ('spaced') practice
	Implementing a schedule of practice that spreads out study activities over time
Moderate utility	**Elaborative interrogation**
	Generating an explanation for why an explicitly stated fact or concept is true
	Self-explanation
	Explaining how new information is related to known information, or explaining steps taken during problem solving
	Interleaved practice
	Implementing a schedule of practice that mixes different kinds of problems, or a schedule of study that mixes different kinds of material, within a single study session
	Summarization
	Writing summaries (of various lengths) of to-be-learned texts
Low utility	**Highlighting**
	Marking potentially important portions of to-be-learned materials while reading
	Keyword mnemonic
	Using keywords and mental imagery to associate verbal materials
	Imagery use for text learning
	Attempting to form mental images of text materials while reading or listening
	Rereading
	Restudying text material again after an initial reading

Examples of teacher characteristics

As well as observable behaviours, there are also some teacher characteristics that may not be directly observable in classroom behaviour, but which have been found to be related to students' learning gains.

(Pedagogical) Content knowledge

A number of studies have found a relationship between measures of a teacher's knowledge of the content they are teaching and the gains made by their students. It seems intuitively obvious that 'Teachers cannot help children learn things they themselves do not understand' (Ball, 1991, p5). However, the search for a relationship between characteristics such as academic qualifications or general ability and student performance has been rather disappointing: correlations are typically very small or non-existent (Rockoff et al, 2011). Nevertheless, there seems to be an emerging body of work that can link more specific measures of

content knowledge, and in particular the kinds of content knowledge that are relevant to teaching, to student gains.

For example, Sadler et al (2013) tested a group of volunteer, experienced middle school (seventh and eighth grade) science teachers on their understanding of the content they were teaching and on the kinds of misconceptions they expected students to show. Generally, their understanding of the content was good, though there was enough variation to give some predictive power to teachers' subject knowledge: overall, teachers answered 83% correctly, compared with 38% by their students. However, the teachers' ability to identify common misconceptions was hardly above chance. Overall, there was a positive but modest relationship between teachers' understandings and their students' gains. However, an item-level analysis of the relationship between teachers' and students' understanding of specific concepts had considerably more predictive power. This suggests that targeting support for teachers at particular areas where their understanding or their knowledge of student misconceptions is weak may be a promising strategy, a claim that is supported by reviews of the impact of teacher professional development in these areas (Timperley et al, 2007; Blank and de las Alas, 2009).

Hill et al (2005) investigated the importance of teachers' pedagogical content knowledge in mathematics. They cited a number of studies that have found that teachers' level of understanding of the mathematics they are teaching is related to how effectively students learn it. In their own analysis, they found that the difference between high and low scoring (a 2 SD gap) teachers on their Content Knowledge for Teaching (CKT) was associated with more than a month's additional learning for students in a year. Although this is not a huge effect, it is of similar order to the strength of the relationship between socioeconomic background and attainment, for example. Interestingly, most of the difference was between the lowest scoring teachers and the rest: once their CKT score was into the third decile there was no further relationship with student learning.

Beliefs about learning

Askew et al (1997) found that highly effective teachers of numeracy were characterised by a particular set of beliefs, which in turn led to a corresponding set of teaching approaches. They claim that "The mathematical and pedagogical purposes behind particular classroom practices are as important as the practices themselves in determining effectiveness" (p5). In other words, simply describing or defining observable practices or approaches is not enough to characterise teachers as more or less effective; it matters *why* the teachers adopt them.

In particular, Askew et al (1997) identified beliefs about the nature of mathematics and what it means to understand it, along with teachers' beliefs and theory about how children learn and about the teacher's role in promoting learning, as important distinguishing factors between those who were more and less effective (see table 3). Given the potential significance of the need to focus on teacher beliefs, it seems surprising that these findings do not seem to have been extensively tested by further research; although there is extensive research on teacher beliefs, links with pupil progress are much less common. A study by Higgins and Moseley (2001) of teacher beliefs about Information and

Communication Technology failed to find any convincing relationships between beliefs and pupil progress.

However, some corroboration can be found in the evidence from Timperley et al (2007) that the professional development programmes with demonstrable benefits for learners mostly included some attempt to engage with teachers' existing theories, values and beliefs (p196). Such a claim is also consistent with a view of effective pedagogy as consisting of more than just a set of classroom techniques, but depending on the ability to make complex judgements about which technique to use when.

Table 3: Characteristics of highly effective teachers of numeracy, from Askew et al, 1997

Highly effective teachers were characterised by beliefs about

What it means to be numerate:

- having a rich network of connections between different mathematical ideas
- being able to select and use strategies, which are both efficient and effective.

  They used corresponding teaching approaches that:
  - connected different areas of mathematics and different ideas in the same area of mathematics using a variety of words, symbols and diagrams
  - used pupils' descriptions of their methods and their reasoning to help establish and emphasise connections and address misconceptions
  - emphasised the importance of using mental, written, part-written or electronic methods of calculation that are the most efficient for the problem in hand
  - particularly emphasised the development of mental skills.

How children learn:

- almost all pupils are able to become numerate
- pupils develop strategies and networks of ideas by being challenged to think, through explaining, listening and problem solving.

  They used teaching approaches that:
  - ensured that all pupils were being challenged and stretched, not just those who were more able
  - built upon pupils' own mental strategies for calculating, and helped them to become more efficient.

The role of the teacher:

- discussion of concepts and images is important in exemplifying the teacher's network of knowledge and skills and in revealing pupils' thinking
- it is the teacher's responsibility to intervene to assist the pupil to become more efficient in the use of calculating strategies.

  These teachers used teaching approaches that encouraged discussion, in whole classes, small groups, or with individual pupils.

Less effective teachers believed in the importance of either

- pupils acquiring a collection of facts and standard methods, and that pupils varied in their ability to remember these. They used teaching approaches that:
  - dealt with areas of mathematics discretely
  - emphasised teaching and practising standard methods and applying these to abstract or word problems without considering whether there were alternative more efficient ways of solving a particular problem.

or

- developing numeracy concepts using practical equipment and waiting until pupils were ready to move onto more formal methods. They used teaching approaches that emphasised pupils working things out for themselves, using any method with which they felt comfortable.

Other characteristics

A large number of studies have set out to find links between a variety of other teacher characteristics and student achievement gains. Wayne and Youngs (2003) conducted a review of the available literature and concluded that there were positive (though often inconsistent and probably small) associations between

student learning gains and teacher characteristics such as the status of the college they had attended or their scores on certain kinds of tests, such as licensure or reasoning tests, or specific tests of the material they were teaching.

For mathematics teachers, having a higher degree in maths, or a better class of degree, was associated with more student learning, but the same relationship was not found in other subjects. Similarly, being certified (qualified) in maths or science teaching was associated with greater effectiveness, but there was no relationship between certification and effectiveness in other subjects. Ball and Hill (2009) review some of the later literature on the relationships between teacher certification, qualifications and level of study with student learning, and conclude they are generally inconsistent and hard to interpret.

Interestingly, a number of teacher characteristics (such as teachers' self-reported self-efficacy, extraversion and conscientiousness) were found by Rockoff et al (2011) to be related to supervisor ratings of effectiveness but not to actual student achievement gains.

Examples of ineffective practices

It may seem unduly negative to focus on things that do not work, but there are a number of reasons for wanting to do this.

One is that it provides a challenge to complacency. A potential problem with lists of 'best practice' is that they can be susceptible to confirmation bias. If the list of effective practices is long enough, and contains descriptions of practices that are open to a bit of interpretation, most teachers will be able to identify some they think they are doing. Such lists can also seem, like motherhood and apple pie, to be good, but predictable, obvious and nothing new. Including some examples of 'worst practice' is likely to provoke a stronger reaction, which we hope can be challenging in a constructive way. Clearly, bluntly telling a teacher that some aspect of their practice is wrong may not be a good way to get a discussion going, however.

A second reason is that many of these ineffective practices seem to be quite popular, though most evidence here is anecdotal and selective. It may be that as well as telling us 'what works', an important contribution of research is to tell us what doesn't work. By stopping doing things that are either ineffective or inefficient, we should allow more time to focus on thing that will make more difference.

The following are examples of practices whose use is not supported by research evidence:

Use praise lavishly

Praise for students may be seen as affirming and positive, but a number of studies suggest that the wrong kinds of praise can be very harmful to learning. For example, Dweck (1999), Hattie & Timperley (2007).

Stipek (2010) argues that praise that is meant to be encouraging and protective of low attaining students actually conveys a message of the teacher's low

expectations. Children whose failure was responded to with sympathy were more likely to attribute their failure to lack of ability than those who were presented with anger.

> "Praise for successful performance on an easy task can be interpreted by a student as evidence that the teacher has a low perception of his or her ability. As a consequence, it can actually lower rather than enhance self-confidence. Criticism following poor performance can, under some circumstances, be interpreted as an indication of the teacher's high perception of the student's ability." (ibid)

Allow learners to discover key ideas for themselves

Enthusiasm for 'discovery learning' is not supported by research evidence, which broadly favours direct instruction (Kirschner et al, 2006). Although learners do need to build new understanding on what they already know, if teachers want them to learn new ideas, knowledge or methods they need to teach them directly.

Group learners by ability

Evidence on the effects of grouping by ability, either by allocating students to different classes, or to within-class groups, suggests that it makes very little difference to learning outcomes (Higgins et al, 2014). Although ability grouping can in theory allow teachers to target a narrower range of pace and content of lessons, it can also create an exaggerated sense of within-group homogeneity and between-group heterogeneity in the teacher's mind (Stipek, 2010). This can result in teachers failing to make necessary accommodations for the range of different needs within a supposedly homogeneous 'ability' group, and over-doing their accommodations for different groups, going too fast with the high-ability groups and too slow with the low.

Encourage re-reading and highlighting to memorise key ideas

This finding has already been mentioned in summarising the review by Dunlosky et al (2013). Re-reading and highlighting are among the commonest and apparently most obvious ways to memorise or revise material. They also give a satisfying – but deceptive – feeling of fluency and familiarity with the material (Brown et al, 2014). However, a range of studies have shown that testing yourself, trying to generate answers, and deliberately creating intervals between study to allow forgetting, are all more effective approaches.

Address issues of confidence and low aspirations before you try to teach content

Teachers who are confronted with the poor motivation and confidence of low attaining students may interpret this as the cause of their low attainment and assume that it is both necessary and possible to address their motivation before attempting to teach them new material. In fact, the evidence shows that attempts to enhance motivation in this way are unlikely to achieve that end. Even if they do, the impact on subsequent learning is close to zero (Gorard, See & Davies, 2012). In fact the poor motivation of low attainers is a logical response to repeated failure. Start getting them to succeed and their motivation and confidence should increase.

Present information to learners in their preferred learning style

A belief in the importance of learning styles seems persistent, despite the prominence of critiques of this kind of advice. A recent survey found that over 90% of teachers in several countries (including the UK) agreed with the claim that "Individuals learn better when they receive information in their preferred learning style (for example, visual, auditory or kinaesthetic)" (Howard-Jones, 2014). A number of writers have tried to account for its enduring popularity (see, for example, a clear and accessible debunking of the value of learning styles by Riener and Willingham, 2010), but the psychological evidence is clear that there are no benefits for learning from trying to present information to learners in their preferred learning style (Pashler et al, 2008; Geake, 2008; Riener and Willingham, 2010; Howard-Jones, 2014).

Ensure learners are always active, rather than listening passively, if you want them to remember

This claim is commonly presented in the form of a 'learning pyramid' which shows precise percentages of material that will be retained when different levels of activity are employed. These percentages have no empirical basis and are pure fiction. Memory is the residue of thought (Willingham, 2008), so if you want students to remember something you have to get them to think about it. This might be achieved by being 'active' or 'passive'.