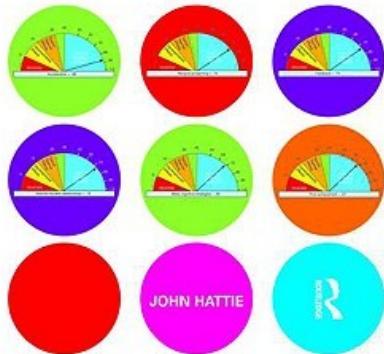


VISIBLE LEARNING
A SYNTHESIS OF OVER 800 META-ANALYSES
RELATING TO ACHIEVEMENT

"Reveals teaching's Holy Grail"
The Times Educational Supplement



Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement

John A.C. Hattie

[Download now](#)

[Read Online ➔](#)

Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement

John A.C. Hattie

Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement John A.C. Hattie
This unique and ground-breaking book is the result of 15 years research and synthesises over 800 meta-analyses on the influences on achievement in school-aged students. It builds a story about the power of teachers, feedback, and a model of learning and understanding. The research involves many millions of students and represents the largest ever evidence based research into what actually works in schools to improve learning. Areas covered include the influence of the student, home, school, curricula, teacher, and teaching strategies. A model of teaching and learning is developed based on the notion of visible teaching and visible learning.

A major message is that what works best for students is similar to what works best for teachers - an attention to setting challenging learning intentions, being clear about what success means, and an attention to learning strategies for developing conceptual understanding about what teachers and students know and understand.

Although the current evidence based fad has turned into a debate about test scores, this book is about using evidence to build and defend a model of teaching and learning. A major contribution is a fascinating benchmark/dashboard for comparing many innovations in teaching and schools.

Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement Details

Date : Published March 1st 2009 by Routledge (first published November 14th 2008)

ISBN : 9780415476188

Author : John A.C. Hattie

Format : Paperback 378 pages

Genre : Education, Teaching, Nonfiction



[Download](#) [Visible Learning: A Synthesis of Over 800 Meta-Analyses ...pdf](#)



[Read Online](#) [Visible Learning: A Synthesis of Over 800 Meta-Analys ...pdf](#)

Download and Read Free Online Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement John A.C. Hattie

From Reader Review Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement for online ebook

Atila Iamarino says

Um livro que qualquer educador deveria ler. Sério. Pode acabar com 80% da licenciatura, 50% do curso de pedagogia e só dar este livro para as pessoas lerem.

Hattie fez uma meta-análise de meta-análises de evidências sobre ensino. Uma meta-análise é um sumário de análises de uma área. Por exemplo, uma comparação de estudos que avaliam a eficácia de se adotar uniformes em uma escola. Hattie comparou, resumiu e nivelou os achados de 800 meta-análises sobre educação. São mais de 50 mil artigos e milhões de estudantes embaixo disso comparando desempenho antes e depois da aplicação de alguma mudança. Uma cacetada de métodos e intervenções em estudantes, professores, escola, métodos de ensino e o que mais pensar.

Para completar, ele apresenta um índice muito intuitivo para comparar os resultados. Um efeito d que dá as chances de um estudante do grupo sob intervenção ter um desempenho escolar melhor do que quem não participou. Ele dá os pontos positivos e negativos de usar esse tipo de índice, o que fica de fora e afins. Mas algum tipo de métrica tem que ser adotada. E o que acho que mais contribuiu para as comparações é que o baseline que assume para ver a eficácia dos métodos não é 0. Segundo ele, um $d=0,2$ é equivalente ao que alunos desenvolveriam naturalmente ao longo de um ano, simplesmente por amadurecerem. Ou seja, a não ser que uma mudança atrasse o desenvolvimento dos alunos, ela vai dar algum resultado positivo. Um $d=0,3$ ou um pouco mais é esperado simplesmente porque alunos estão participando de um estudo e mais bem acompanhados. Então intervenções compensam (em especial as mais caras ou trabalhosas) principalmente quando têm um efeito acima de 0,4.

E este é o maior ponto. Muito do que pensamos, de modas e de ideologias sobre o ensino não têm esse efeito nem de perto. Construtivismo, uniformes, aprendizado baseado em problemas e mais um monte de ideias já foram testadas em milhares de estudantes e não chegam nesse valor. Podem até compensar serem adotadas, mas não são necessariamente o mais prioritário. Nem mesmo diminuir o número de alunos por professor – dentro de números como reduzir de 30 para 15. Enquanto repetir de ano é uma das intervenções que fazem mais mal para os alunos (um dos poucos efeitos negativos). E eu mesmo aceitava como importante ou positivo antes de ler.

Na verdade, quase tudo o que Hattie levantou com um efeito mais alto são atitudes dos professores e dos alunos. Mudanças de comportamento, motivação, planejamento, metas, cobranças claras e outros do tipo. Uma série de medidas relativamente baratas de adotar (em termos de infra-estrutura), mas bastante trabalhosas por demandarem muito de professores e da escola. Mudou muito como vejo o papel do professor. Vai ser um livro para o qual com certeza vou voltar muitas vezes.

Andrew says

Although an extremely dry read (what else should you expect from a synthesis of education meta-analyses?), I feel like it will be an extremely useful reference during my graduate studies and a terrific supplement for newer research that I will read throughout my career. I have high hopes that teaching as a profession will become more evidence based in my lifetime.

Mariah says

Many parts of this book are interesting & certainly useful for discussion. School libraries are not mentioned once in this entire book, and there are studies about the importance of adequately staffed school libraries for academic success. This makes me wonder what other important things the author has omitted or forgotten among his seemingly thorough analyses.

K Love says

It's phenomenal: like all data and meta-analysis, interpretation is still from the perspective of human experience and interaction.

Lauren Gifford says

Every educator should own a copy of this book. I loved the visuals for a quick glance at high effects strategies.

Diz says

This book compares all kinds of variables in education by measuring their effect on achievement on a single scale, which provides an opportunity to compare these variables. It's an interesting idea, but there are some dangers in this. As the author points out in the conclusion of his book, there may be objectives other than achievement. However, an administrator who skims this book may miss that point and make important decisions for the school based on an incomplete understanding of what is covered here. They may see that math classes have a greater effect on achievement than art classes and then decide to cut art classes to add more math classes. That would be a terrible loss. I do wish that the author had been more clear about there being more objectives to education than achievement.

Nathan says

John Hattie's a genius: looking at over 800 meta-analyses of over 50,000 studies, he's isolated the effect of different variables on the achievements of kids in school. There's so much here, I can dig in again and again for parenting and teaching inspiration. Here's what I took away before my brain ranneth over. ($d=0.3$ is a measure of how effective the intervention is; almost every intervention works, but 0.4 is the mean benefit so it's only a really GOOD intervention if it scores more than 0.4).

one of the major results presented in this book relates to increasing the amount of feedback because it is an important correlate of student achievement. However, one should not

immediately start providing more feedback and then await the magical increases in achievement. As will be seen below, increasing the amount of feedback in order to have a positive effect on student achievement requires a change in the conception of what it means to be a teacher; it is the feedback to the teacher about what students can and cannot do that is more powerful than feedback to the student, and it necessitates a different way of interacting and respecting students (but more on this later). It would be an incorrect interpretation of the power of feedback if a teacher were to encourage students to provide more feedback. As Nuthall (2007) has shown, 80% of feedback a student receives about his or her work in elementary (primary) school is from other students. But 80% of this student- provided feedback is incorrect! It is important to be concerned about the climate of the classroom before increasing the amount of feedback (to the student or teacher) because it is critical to ensure that “errors” are welcomed, as they are key levers for enhancing learning. It is critical to have appropriately challenging goals as then the amount and directedness of feedback is maximized. Simply applying a recipe (e.g., “providing more feedback”) will not work in our busy, multifaceted, culturally invested, and changing classrooms.

Marzano (2000) argued that 80 percent of the variance in achievement could be accounted for by student effects, 7 percent by school effects, and 13 percent by teacher effects. He then used these estimates to evaluate the effects on student achievement of an ineffective, an average, and an exceptional teacher in an ineffective, an average, and an exceptional school respectively. Average schools and average teachers, although he said they did little harm, also did little to influence students’ relative position on the distribution of achievement for all students; ineffective teachers, no matter how effective the school, had a negative impact on the standings of all students, whereas students of exceptional teachers, even in ineffective schools, either maintained or increased achievement, many quite substantially. “Exceptional performance on the part of teachers not only compensates for average performance at the school level, but even ineffective performance at the school level” (Marzano, 2000, p. 81).

Indeed, one of the fascinating discoveries throughout my research for this book is discovering that many of the most debated issues are the ones with the least effects. It is a powerful question to ask why such issues as class size, tracking, retention (that is, holding a student back a grade), school choice, summer schools, and school uniforms command such heated discussion and strong claims. Such cosmetic or “coat of paint” reforms are too common. So many structural claims involve the parents (more homework), lead to more rules (and therefore more rule breakers), have hints of cultural imperatives (quietness and conformity is desired), and often include appeals to common sense (reducing class size is obviously a good thing!). However, the most powerful effects of the school relate to features within schools, such as the climate of the classroom, peer influences, and the lack of disruptive students in the classroom—all of which allow students and teachers to make errors and develop reputations as learners, and which provide an invitation to learn.

What a child brings to the classroom each year is very much related to their achievement in previous years—brighter children tend to achieve more and not so bright children achieve less. This should not be surprising given that the correlation between ability and achievement is very high. Hattie and Hansford (1982) reported an average correlation of $r = 0.51$ between measures of intelligence and achievement (an effect size, $d = 1.19$). This high relationship accounts for what many researchers call (usually with a sense of surprise) the “Matthew effect”, which is based on the biblical notion that the rich get richer and the poorer get poorer or do not gain as much. Prior achievement predicts success from preschool to the first years of schooling (Duncan et al., 2007; La Paro & Pianta, 2000; Schuler, Funke, & Baron-Boldt, 1990), between high school and college or university grades (Kuncel, Hezlett, & Ones, 2001), between college and adult success (Bretz, 1989; Samson, Graue, Weinstein, & Walberg, 1984), and between grades in school and later job performance (Roth, BeVier, Switzer, & Schippmann, 1996).

About 70–80 percent of families have two parents in most Western countries, about 10–20 percent of families are single-parent, and about 2–10 percent are other than these structures. Pong, Dronkers, and Hampden-Thompson (2003) found that single parenthood is associated with lower mathematics and science achievement (although the effects are quite small). They also noted that countries with more generous welfare policies, like Austria, showed the smallest gaps. The greatest gaps were in countries such as the United States and New Zealand, who, they claimed, lagged behind other industrialized countries in providing financial assistance, universal child benefits, tax benefits and maternity leave benefits to single and poorer families. They concluded that “to some extent the investment in national family policies explains why Australia ranks at the top but the United States and New Zealand rank last in the academic resilience of children from single-parent homes” (p. 695).

Hong and Ho (2005) concluded that parent aspirations were the most important influence on their children’s achievement, whereas parental supervision in the forms of monitoring students’ homework, time watching television, and time going out with friends appeared to have a negative effect on the educational aspirations of adolescent students. Similarly, Rosenzweig (2000) noted that the relationships between student achievement and parental participation ($d = 0.56$) and supportive parenting ($d = 0.43$) were much higher than with homework supervision ($d = 0.19$), participation in school activities ($d = 0.14$), communication with school and teachers ($d = 0.14$), monitoring school progress ($d = 0.12$), providing structure in the home ($d = 0.00$), and controlling and disciplining parental style ($d = -0.09$). These effects were the highest in high SES families, in elementary compared to high schools, and in Asian and Latino compared to white and African American families. Of as much interest are those family variables that negatively relate to achievement. These factors included external rewards, homework surveillance, negative control, and restrictions for unsatisfactory grades. Overall, “the higher the hopes and expectations of parents with respect to the educational attainment of their child, the higher the student’s own educational expectations and, ultimately, the greater the student’s academic achievement” (Hong & Ho, 2005, p. 40). These high expectations are assisted by greater parent-student communication and the student’s control over their own studies (see also Fan and Chen, 2001).

As an example of its use, Konstantopoulos (2005) found that a substantial proportion of the variation in student achievement lies within schools and not between schools. If the variance is within, this means that factors such as teacher variability have a relatively larger effect on student achievement than do school effects. “It appears that the teachers students are assigned to may be more important than the schools they attend” (p. 36).

New Zealand, as an example, has among the lowest percentage of between-school variance (about four percent and thus the within school variance is much greater.) Using data from the Second International Mathematics Study, Scheerens, Vermeulen, and Pelgrum (1989) found that school effects were undetectable as a source of variance in New Zealand, whereas between-teacher or between-class variance was 42 percent. Harker and Nash (1996; Nash & Harker, 1997) found that the school effect in New Zealand high school performance accounted for between five to ten percent of the variance in mathematics, nine to ten percent of the variance in English, and five to seven percent of the variance in science. The message is that, if you take two students of the same ability, it matters not which school they attend, but it may matter greatly who their teacher is. It is not so much that teachers matter, as that the variance within schools indicates that some teachers matter more than others!

A major reason difficult goals are more effective is that they lead to a clearer notion of success and direct the student’s attention to relevant behaviors or outcomes, whereas “doing your best” can fit with a very wide range of goals. It is not the specificity of the goals but the difficulty that is crucial to success. There is a direct linear relationship between the degree of goal difficulty and performance. There are five meta-analyses relative to this contention (Table 9.2) and the overall effect size is a large $d = 0.67$ (these are not all achievement outcomes and so are not included in the Appendices of this book). The performances of the students who have the most challenging goals are over 250 percent higher than the performances of the subjects with the easiest goals (Wood & Locke, 1997).

Also, difficult goals are much better than “do your best” or no assigned goals. Any school with the motto “do your best” should immediately change it to “face your challenges” or “strive to the highest”. The following five meta-analyses relate to this contention (Table 9.3). This is because “do your best” goals are easily attained—in one sense, anything you do can be defined as your best. Instead, teachers and learners should be setting challenging goals.

Goals have a self-energizing effect if they are appropriately challenging for the student, as they can motivate students to exert effort in line with the difficulty or demands of the goal.

Commitment to the goals helps, but is not necessary for goal attainment—except for special education students, where commitment makes a major difference. Klein, Wesson, Hollenbeck, and Alge (1999) found a high relationship ($d = 0.47$) between goal commitment and subsequent performance, and the effect between commitment and outcome increased as a function of goal difficulty. Donovan and Radosevich (1998) found lower effects of commitment to goals than they expected, but these were still quite high ($d = 0.36$).

When I completed the first synthesis of 134 meta-analyses of all possible influences on achievement (Hattie, 1992) it soon became clear that feedback was among the most powerful influences on achievement. Most programs and methods that worked best were based on heavy dollops of feedback. When I was presenting these early results in Hong Kong, a questioner asked what was meant by feedback, and I have struggled to understand the concept of feedback ever since. I have spent many hours in classrooms (noting its absence, despite the claims of the best of teachers that they are constantly engaged in providing feedback), worked with students to increase self-helping (with little success), and have tried different methods of providing feedback. The mistake I was making was seeing feedback as something teachers provided to students—they typically did not, although they made claims that they did it all the time, and most of the feedback they did provide was social and behavioral. It was only when I discovered that feedback was most powerful when it is from the student to the teacher that I started to understand it better. When teachers seek, or at least are open to, feedback from students as to what students know, what they understand, where they make errors, when they have misconceptions, when they are not engaged—then teaching and learning can be synchronized and powerful. Feedback to teachers helps make learning visible.

Programmed instruction, praise, punishment, and extrinsic rewards were the least effective forms of feedback for enhancing achievement. Indeed, it is doubtful whether rewards should be thought of as feedback at all. Deci, Koestner, and Ryan (1999) have described tangible rewards (stickers, awards, and so on) as contingencies to activities rather than feedback because they contain so little task information. In their meta-analysis of the effects of feedback on motivation, these authors found a negative correlation between extrinsic The contributions from teaching approaches—part I 175 rewards and task performance ($d = -0.34$). Tangible rewards significantly undermined intrinsic motivation, particularly for interesting tasks ($d = -0.68$) compared to uninteresting tasks ($d = 0.18$). In addition, when the feedback was administered in a controlling manner (e.g., saying that the student performed as they “should” have performed), the effects were even worse ($d = -0.78$). Thus, Deci et al. concluded that extrinsic rewards are typically negative because they “undermine people’s taking responsibility for motivating or regulating themselves” (Deci et al., 1999, p. 659). Rather, extrinsic rewards are a controlling strategy that often leads to greater surveillance, evaluation, and competition, all of which have been found to undermine enhanced engagement and regulation (Deci & Ryan, 1985).

Another form of feedback is repeated testing, but this is only effective if there is feedback from the tests to teachers such that they modify their instruction to attend to the strengths and gaps in student performance. Although performance is increased with more frequent testing, the amount of improvement in achievement diminishes as the number of tests increase (Bangert-Drowns, Kulik, Kulik, & Morgan, 1991). Students taking at least one test during a 15-week term scored about half a standard deviation higher in criterion examinations than students taking no tests. When two groups answered identical test items, superior performance was obtained from students who answered the questions on a large number of short tests rather than on a small number of long tests. The caution is that it may not be the frequency of test taking but that frequent test taking made the learning intentions and success criteria more specific and

transparent. Clariana and Koul (2006) found that multiple-try feedback was less effective for surface outcomes ($d = -0.22$) but more effective for higher-order outcomes ($d = 0.10$).

Many states in the United States have high-stakes testing and there is also much testing embedded in the No Child Left Behind imperatives. There have been arguments that such frequent testing is akin to a coaching effect, whereas others consider that any gains are because of narrowing the curriculum, teaching to the test, and because too many students are excluded who may not perform so well. Amrein and Berliner (2002) raised much debate with their analysis of the performance of 18 states with high-stakes testing systems and found little effect of these systems on student achievement. This conclusion was contested (e.g., Braun, 2004; Raymond & Hanushek, 2003; Rosenshine, 2003). Lee (2006) used meta-analysis to compare different state policies on the National Assessment of Educational Progress examination. He found six studies favored high-stakes testing states, five were mixed, and one favored low-stakes testing states. The effects were extremely varied ($d = -0.67$ to $d = 1.24$), although it made no difference as to the focus of the accountability—that is, whether the focus is a combination of schools and students ($d = 0.38$), for schools alone ($d = 0.39$), or for students alone ($d = 0.31$). The effects on mathematics ($d = 0.38$) are slightly higher than on reading ($d = 0.29$), and higher for elementary ($d = 0.44$) and middle schools ($d = 0.35$) than for high schools ($d = 0.03$).

Wendy says

... Because I'm a dork and read things like this for fun. :)

I'm going to go ahead and review this one even though I technically haven't finished it. It's not really the kind of book you read all at once. It'll be something I dip into off and on. I don't know much about statistics, so I can't explain all the terms, but this is obviously John Hattie's life's work. He has synthesized thousands of educational research meta-analyses into one place -- an impressive task. As I understand it, a meta-analysis uses statistics to combine and compare results of many studies on the same topic. It provides a common language to discuss these different studies so we can tell what's really effective and not.

The end product is fascinating -- he ranks 138 different topics within education related to school policies, family, and teaching from most to least effective. Some are common sense -- students suffer when they move around between multiple schools regularly. Some are surprising. For example, Hattie could find no meta-analysis that supports using retention as an intervention method for students. It apparently has all kinds of negative effects on kids but no positives. However, schools use this all the time. Conversely, he could find no negative effects from accelerating learning for gifted students (allowing them the flexibility to move through school more quickly), but schools rarely allow this. (Actually, more frustrating than surprising.) There are also, of course, a lot of strategies and policies that really make no difference whatsoever, positive or negative.

Hattie definitely has his own opinions about what he's found in the research and doesn't hesitate to express

them, which actually makes this book a lot more interesting to read than you would expect. He offers nice explanations and summaries to tie everything together and explain the numbers more clearly. So far I've read through about 30 of the meta-analysis summaries, and I have a lot of things I need to look into more. Once I've done that, I'll come back to this resource to find my next ideas for research. It's not exactly a high-budget production -- some of the pages are poorly printed and kind of a mess, but it's a tremendously important addition to the field of education, and I expect to come back to it frequently over the coming years.

Trevor says

In a world full of presumed shortages, but one of the things that is never in short supply is opinions on how to 'fix' the education system. Just about everyone has been to school and therefore immediately thinks that this alone amply qualifies them as experts able to pontificate on how best to educate the young and to hold forth on everything that was wrong with their education and by extension, everyone else's. Some of these people even become teachers only to develop ideas concerning what 'works' and what doesn't when educating students. Hattie's point is that a lot of what we 'know' about what works in the classroom is based on very flimsy 'evidence'. In fact, he would be prepared to say it is based on virtually no evidence at all. This is because humans are quite strange creatures, strange in that they learn. They learn often despite how they are taught. So, if your measure for success is – the kids know more now than they did at the start of the year – then that might be really setting the bar far too low. As he says at the start of this book, if the teacher is breathing the students will probably come out the other end knowing more than they knew when they went into the class. It would be hard (though, not impossible) for students to come out the other end knowing less than they did when they went in.

Now we need to talk about statistics, sorry. We all know what a standard (normal) distribution curve looks like. It is that bell shaped graph you see all of the time. It says that most people cluster around the average and only a few people are a long way above or below average. But let's say you do something to improve the learning of students in your class and you want to see if it has been effective or successful – what do you do? Well, you could measure how well your students do before and after the intervention and see what change it makes, or you could find another class that isn't going to have the same intervention and measure the two classes before and after the intervention and compare how much change there has been in the class you are interested in. But, even if there has been a change beyond your control class, how do you know if the change is enough to get excited about? This is where this book starts to talk about effect sizes and this is really quite important. Effect sizes are basically measures of how much you have been able to improve averages. It is a measure of a change in the mean (the distance between the means) of two normal distributions. And it measures this change as a ratio of a standard deviation. This means that an effect size of 1.0 would say that students having the measure intervention would have really shot ahead by an entire standard deviation. Most effect sizes are smaller than 1.0. Hattie says that we shouldn't be comparing our teaching interventions against an effect size of 0.0 – that is just silly, because as we said before, if the teacher is breathing we will expect an effect size of greater than zero just because the students are growing older. No, to be a valid improvement Hattie suggests an effect size of 0.4. That is, a little less than half a standard deviation.

He then number crunches over 800 meta-analyses of educational research projects from around the world into what really make a difference to students' learning. These include things the students bring to the classroom with themselves, how schools are organised, students' home environment and finally differences in their teachers and the type of instruction they provide.

And the findings? Well, basically that of the things we can actually affect, teachers are the main contributor

to student success and finding out what good teachers do in their classrooms should be our priority.

Hattie argues that what teachers really need to do is to learn how to become reflective in their teaching – that is, how they make their teaching visible to their students and how they can see how well their students are learning. Effectively, he is saying that most standard models don't actually work in the classroom, because what happens in the classroom is very complex – so to be successful a teacher needs to take this complexity into consideration. A good teacher will start by figuring out where their students are up to – what they already know. They will then work out where they want their students to get to and then they will work out what they need to do next to start their students on the path to getting to where they need to go. They will do this by coming up with some sort of intervention. All learning is about change, and so the teacher is trying to work out how to change their students in a positive direction. However, not all interventions work, at least, they don't all work equally well. But they all 'work' a little bit, perhaps less than an effect size of 0.4, but enough to convince a teacher that change is happening. So, what a good teacher needs to do is to constantly challenge their view of what is working by seeking to prove it isn't working. This is Popper's 'proof through seeking disconfirmation' writ large.

And what did the 800+ meta-analyses tell us? Well, lots of things. That teacher education doesn't make any difference to teacher effectiveness – if you have received six weeks or four years of teacher education training it will not predict how good a teacher you will be. That whole language is a terrible way to learn to read and phonics is great. That direct instruction is the key to learning and constructivist models are trendy, but ineffective. But mostly, that teachers need to make learning visible, make success criteria obvious to both themselves and their students and to constantly provide feedback on how students are going and what they need to do next to improve.

Some of what Hattie suggests is pretty well obviously true. Other bits I feel very uncomfortable with. Still other bits I think are just plain wrong. The problem here is that he has numbers, but as he says himself, numbers are far from the whole story.

One of the things that drove me nuts while reading this book was that at no time did he really define what 'achievement' was. He could measure it to two decimal places, as he self-deprecatingly says of himself at one point, but just what it is that is being measured when we say we are measuring achievement is at no point made clear. As he also says at one point, there is surface knowledge, there is deep knowledge and there is integrated knowledge – and most of the assessment practices used in these meta-analyses are okay at measuring surface learning, a bit worse at measuring deep knowledge, but hopeless at measuring integrated knowledge. My problem is that you have something like direct instruction – which we are constantly being told is the most effective means of improving student learning – but how are you measuring that exactly? If this involves drilling kids for weeks that two plus two equals four and then giving them a test that asks, 'what does two plus two equal?' then I can see the kids that have been drilled on the test questions will almost certainly do better on the test. Now, Hattie does claim this is not what direct instruction is really about and that the knowledge gained can be transferrable, but other than his statements to the fact there isn't much else to go on in this book.

The strangest part of the book for me was the bit saying whole language was a failure in teaching kids to read. I just need to say that my understanding of whole language is that it is teaching kids that texts are meaningful and engaging and doing this by focusing on their comprehension of texts much more than on decoding and phonics. My view would be that no one really reads out of the joy of knowing that sh- or th- or ch- make interesting sounds or that those letter combinations often work together to produce those sounds. You know, C-A-T doesn't really sound out to cat... Rather, people read to gain meaning from texts – comprehension might not be the first thing you need to learn when you are learning to read, but it had better

be the last thing you learn. What was utterly bizarre was that he got higher effect sizes for programs that taught through comprehension than he did with whole language. I tend to think there are ideological biases against whole language, so this might also cloud the effects.

And this is my main problem with all of this book. Which is something he mentions in his concluding chapter. He says that he has given a list of interventions and their effect sizes that come about by implementing them. But he points out that these may well not be additive. So, if a teacher wearing a red shirt improves student learning by 0.4 and the teacher wearing green shorts improves student learning by 0.2 that might not mean that wearing both will improve student learning by 0.6 – why not? Well, because these things might not be independent.

And there is my problem. I assume that many of the studies on which these meta-analyses were based were comparing students with other and different schools – one school that tried an intervention and another that did not. But the schools that did try an intervention may be the kinds of schools that already have advantages and these might be masked by what we think we are measuring.

For example, a Catholic school in Australia claims in its advertising material that Catholic education is really good, a fact it can prove, because in the USA women are 100 times more likely to make it to the Senate if they went to a Catholic school than if they didn't. See the problem? There is a clear social class thing happening here over and above the religion of those attending the school. And that social class thing might just have much more to say about who gets into the US Senate than whether they were taught by particularly scary nuns. Hattie even says at the start of this book that social class might even be more important than teaching in deciding achievement at school – which begs the question, why wasn't social class measured or considered?

I don't want to dismiss this book outright – the problem is that I've no idea what it is that I'm reading. It all looks terribly scientific, but if I'm not really told what achievement is, then telling me that it has improved by 0.63 tells me next to nothing.

As he also says, 'achievement' isn't the be all and end all of education – we also send children to school to socialise them. Schooling is a kind of indoctrination of which achievement is only one part of the process. As he also says, achievement is easy to measure when it is least meaningful. But what if achievement was measured in our ability to think independently. How would you go about measuring that kind of achievement? What sort of five-minute test would you create for that skill? I'm not providing an answer to that – but if you only measure what is simple to measure maybe you are thereby missing what is important to measure.

I'm really not suggesting this book can be ignored, but I don't know I can be as confident in its conclusions as Mr Hattie seems to be. He does a very good job in telling the story behind the data – I just don't understand the data enough (and this after 300 pages of his explanations) to know if the story he is telling is supported by the data or if the data is really capable of supporting any story.

Lindsay Bowley says

I recommend this to anyone in education and anyone making educational decisions or policy. What powerful data!

Geert Hofman says

Well researched but somewhat boring.

Janis Taylor says

The statistics and research in this book are fascinating. The major shortfall in this book is the lack of elaboration on the #1 rated correlation with achievement...."self-reported grades." I did some more research to figure out what self reported grades are because Hattie didn't give much of an explanation. I now have my own ideas about what this is and why it is important. I just think it would have been great to elaborate on the top ten a little more than he did.

Caroline says

I was really excited by the premise of this book. Someone read over 800 meta-analyses related to education and wrote a book summarizing their conclusions? Sign me up! (Also, does that make this review a meta-meta-meta-analysis?) 40 pages in, though, and I am giving up, as the book is too full of egregious errors. Even for the statements that don't have obvious problems, I don't think I can trust this author to represent the info accurately.

My first concern was when Hattie defined the common language effect (CLE) as the probability that a score sampled from one distribution will be greater than a score sampled from some other distribution. The problem? He immediately translated a positive effect into a CLE less than 50%! Students in classes where homework is assigned out-perform those without homework (average $d=0.29$). Hattie translates $d=0.29$ to a CLE of 21%, saying "in 21 times out of 100, introducing homework into schools will make a positive difference." What? To see if I was missing something, I did some googling and found this page. Apparently all the CLEs in the book (which were reported alongside every result), were bogus, and some were even negative or greater than 100%!

I probably could have stopped reading then, but I kept going. Later in the book, the CLE struck again! Discussing the impact of a student's prior achievement on educational results, Hattie says, "The overall effect size of 0.67 is among the highest effect sizes in this synthesis of meta-analyses, although the common language estimate (CLE, see chapter 2) should remind us that, on average, prior achievement will lead to gains in achievement on 48% of the occasions, although there is much that is unexplained beyond prior achievement ($100-48 = 52$ percent that is unexplained)." Not only has he again translated a positive effect to a CLE less than 50%, but he seems to be interpreting it as something like percent of variance explained. As far as I can tell, this is nonsensical. I'm not even sure how / why Cohen's d was calculated for this question, where something like a correlation coefficient seems more sensible. Did they split students into two groups, one with "high" prior achievement and one with "low"? I didn't look into the studies to find out. (After all, reading all the studies myself would defeat the purpose of reading this book.)

There were also several presentations of results that seemed dubious to me. Early on, Hattie said that the "average effect size" across all the studies covered was $d=0.4$ He thus suggests we evaluate interventions by

comparing the effect size to $d=0.4$ rather than 0. Depending on the study in question, Cohen's d may have been calculated as $[\text{mean_treatment} - \text{mean_control}]/\text{SD}$ or $[\text{mean_endoftreatment} - \text{mean_beginningoftreatment}]/\text{SD}$, and I don't think those are directly comparable. The second would include effects from the students getting older during the treatment, so the effect size would depend on the time period of the treatment. For the first, the effect size should presumably be smaller overall, and would also depend on the quality of education received by the control group. Maybe comparing to $d=0.4$ still has value as a rough heuristic, but I'm not convinced yet.

Hattie also points out that almost all studies reported a positive impact, and says "virtually everything works", and "the null hypothesis is not the question of interest" since "nearly any innovation is better than its absence". This sounds more like publication bias to me!

Throughout the book, Hattie includes several tables showing several major types of contributions to learning, along with average or overall effect sizes. It's not clearly stated how he calculated these overall effect sizes. Take the average Cohen's d value from each meta-analysis? From each individual study? Weren't there any that didn't report a Cohen's d ? Maybe those were dropped? Not knowing how these values were calculated, I'm not sure how much weight to put on them. Should I really conclude that the types of interventions with higher "average" effect sizes tend to have a bigger impact? I'm not willing to assume this is true, so I think it's time to put the book down.

Faith Hurst says

Very useful for the practice of teaching. It was interesting but ultimately over described for my purpose. However, I did enjoy the background of the research and the statistical method of placing this all on a continuum. Now to just learn what all of these strategies are!

Beth says

So boring -2*

If you're interested in statistics then this book is for you. If you're sane (like myself) then reading it is probably like watching paint dry.

SO BORING.

A few references but nope don't do it to yourself.
