

Making the Most of the Evidence in Education: A Guide for Working Out What Works Here and Now

Nick Cowen, King's College London and
Nancy Cartwright, Durham University

CHESS Working Paper No. 2014-03
Durham University
28/10/2014



CENTRE FOR HUMANITIES
CHESS
ENGAGING SCIENCE AND SOCIETY

Making the Most of the Evidence in Education: A Guide for Working Out What Works Here and Now

Nick Cowen and Nancy Cartwright

Nick Cowen
Department of Political Economy
King's College London
Second Floor
Strand Building
Strand Campus
London WC2R 2LS
Email: nicholas.cowen <at> kcl.ac.uk

Nancy Cartwright
Philosophy Department
Durham University
50 Old Elvet
Durham DH1 3HN
Email: nancy.cartwright <at> durham.ac.uk

**Making the Most of the Evidence in Education:
A Guide for Working Out What Works
Here and Now**

Nick Cowen & Nancy Cartwright

Consulting research results can help you make better decisions about what will work for your students, here and now. But how do you make the most of this evidence for more reliable decisions?

This is a guide to using research evidence when deliberating about educational policies. It is intended for teachers, for school heads, for boards of governors – for anyone who has to settle on policies, programmes or approaches, whether for a single student, a whole class, a school or an entire school district. It supposes that research evidence can help make for better decisions about what will work for *your* student, *your* class, *your* school or in *your* district but that there's no recipe for how to use research evidence, there's no simple read across from research evidence, no matter how good the quality of it, to what will be likely to work for you here and now. You have to reason that out as best you can. This pamphlet provides some information and some strategies that can make that reasoning easier and more reliable.

We offer some strategies for expressing what you already know about a problem in more explicit ways, and thinking about how to combine that knowledge with evidence from academic research on ‘What Works’. This will help turn the evidence of what works that you find at places like the [Education Endowment Foundation](#) into predictions of what is likely to work here and now.

Making Your Decision

Working out what interventions will help your school typically requires some understanding of the *causes* of success or failure of an intervention. You need to:

- **know** where your school’s weaknesses lie,
- **work** out as best you can what the *causes* of these weaknesses are,
- **predict** what is likely to happen if you carry on without a change of practice,
- **propose** some possible courses of action,
- **think** through
 - what would probably happen under these alternatives
 - what the costs and benefits would be
 - for whom, and then,
- **decide** the best course of action.

Steps and Strategies for Thinking about the Evidence

THREE STEPS

I Identify the challenges for your school or your students and their *causes*. The challenge might be poor attainment in maths, and the cause could be difficulty reading written maths problems.

II Consider *how* the intervention is supposed to work and think about whether it can do that in your case. For example, one-to-one tuition is supposed to work by giving intensive support to help students catch up with their peers; this might not work if the student has a SEN that means that catching up in that way is unlikely.

III Identify what *support factors* are needed for the intervention to work, and evaluate whether those support factors are present. One support factor might be having a sufficiently qualified member of staff to deliver the one-to-one tuition.

SEVEN STRATEGIES

1. *Spider-grams*
2. *Chain maps*
3. *Measuring intermediate outcomes*
4. *Cake diagrams*
5. *Crystal ball method*
6. *Quick exit decision trees*
7. *Intervention score-cards*

STEP ONE – Identify the challenges and causes

How do you identify the challenges your students are facing?
How do you identify their causes?

Often, you will have some data to start with, perhaps internal marks or scores from external assessments. A group of students may be scoring below average in a particular subject. But why that's happened and what's missing can be a challenge to figure out. The school environment is complex.

In this section, we are going to show you sources of bias and sources of complexity that make deliberation about outcomes in schools such a challenge, and some strategies that will help you to overcome them.

Complexity and bias

Sources of Bias

Confirmation bias – holding on to a belief despite new information that tells against it.

Anchoring bias – the first impression of a student or class or a proposed intervention can shape the future interpretation of information about them.

Availability bias – being selective about what information to consider, with most attention being given to information that is vivid, concrete and recent.

The fundamental attribution error – the tendency to explain other people's behaviour as due to internal personality traits with insufficient attention to the context in which they acted.

Sources of complexity

No single cause, no single cure. Students face a variety of challenges. It is rare to find a single problem and it is rare to find a single intervention that can help with a student's complex of problems.

Different values. What is to be achieved, what counts as success, and what are the costs and benefits of different interventions may not be fully agreed. For example, students, parents and teachers may differ on acceptable levels of homework, classroom techniques or overall objectives. It is not always easy to identify what the goal should be.

Multiple participants. Students' success is affected by their family, peers and community. Each may matter to the outcomes. Each may have interests that need to be taken into account.

Overlapping causes. Addressing a cause of several problems at once can be a good strategy if that cause can be tackled. But that cause may not be easily accessible nor the most open to change.


Choose your battles. Sometimes it is better to tackle a problem indirectly. For example, rather than insisting that homework be done at home, you may want to set up an after school club. Sometimes you may have to settle for addressing a different problem. If you can't improve a student's maths scores, perhaps you can improve their history marks (perhaps in order to boost a student's confidence).

Strategy 1: Spider-gram

Spider-grams help you visualise the challenge you face and its causes. Depict your main problem (e.g. 'slow progress in maths') in the centre with possible causes around it.

Mark the ones that tend to increase the effect with +, those that diminish it, with – (or using colours).

A key cause in most urgent need of tackling can emerge. It might be a very different factor from the one you initially identified.



It would be nice if we could in every case draw a spider-gram which is plainly the best. But the story is nearly always *underdetermined*. That means the evidence available is insufficient to identify for sure which conclusion we should reach. In trying to understand why a student is struggling with maths, it may be possible to identify some salient factors. But there may be other factors that are causally important but not immediately clear, e.g. a problem with another sibling at home that stops them from concentrating in class.

STEP TWO – Find and evaluate possible solutions



Education
Endowment
Foundation

Once you identify the problems your students face and their causes, you need to find interventions to tackle them.

Various resource centres can help. The [Education Endowment Foundation](#) (especially the EEF's toolkit) is likely to be most relevant. The Early Intervention Foundation may also be helpful.

Their information is a good starting point. You can learn from the resource centres about what interventions have worked and not worked elsewhere and how much they tend to cost. They may also offer insights about how and why an intervention is supposed to work, or about the kinds of schools or students for whom the intervention has worked in the past. But you will have to evaluate which of these interventions can work in your school.

Suppose that teaching phonics early in primary education has improved outcomes in reading ages in a variety of schools studied. It may be that this intervention does not work for students with dyslexia. Or, suppose that evidence suggests that more time for team sports is a great way to motivate students and it helps them concentrate better during lessons, but in your setting, team sports go against the cultural traditions of a significant group of students, keeping them from participating. Then you might not get the effect hoped for. It could even have a negative effect, for instance if it seems to exclude or stigmatise students. Noting these issues early means you can choose a different intervention that *does* have the power to improve outcomes in your setting.

Thinking about averages

Results reported by What Works sites tend to be *averages*: the average effect of the intervention across individuals in the study. They are averages obtained on trial populations that may be like yours but may be different.

Suppose you are thinking about an individual student. You know that where there is an average there is often a lot of variation around that average. For some individuals who make up the average the outcomes will be very good; for some, much less good; and for others still, the intervention may even make matters worse. You can't just assume that the outcomes for your student will be close to the average.

Nor can you assume that the reported average from a trial is the same average you would see across your students or classes. What a study can establish is an average over the individuals in the study. Even if it is a very large study and very well conducted, the average is still only for the study population, in the circumstances in which the study was conducted.

There are some statistical strategies that researchers follow to deal with this problem. The EEF makes available the results of a meta-analysis across all available studies. That is essentially an average of all the averages. Does that make the effect more robust? That depends on what you know about what it takes to get the intervention to work. It also depends on what you know about your students. Perhaps if you know nothing at all, 'take the average of the averages' is the best you can do. But then you should be cautious in your predictions about what your average will look like.

Strategy 2: Chain Maps



Chain maps are a great way of thinking through step-by-step how the intervention is supposed to produce the desired outcome. Causal chains are only as strong as their weakest link. If any of the intermediate stages look like a leap rather than a step, for instance they are either difficult to achieve or implausible, that is a sign that the intervention won't work. When you have a close connection between each step, you have a strong indicator that your intervention is on the right track.

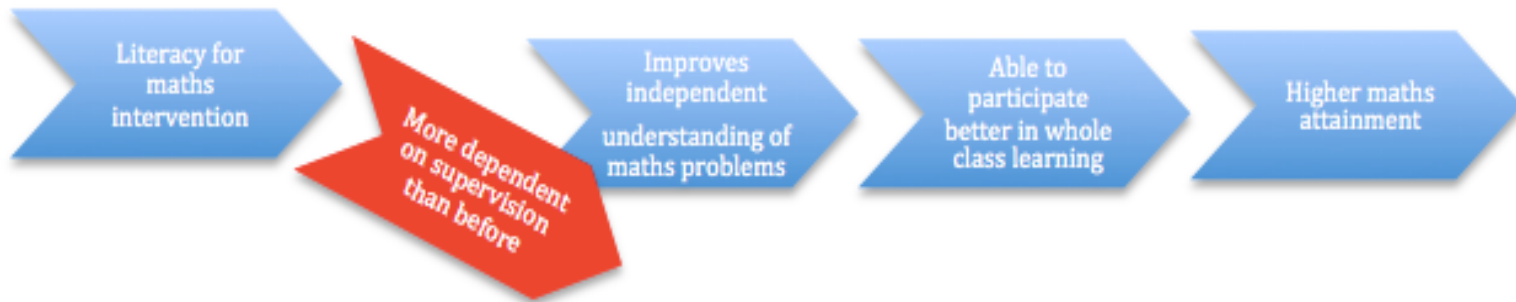
Chain maps also help you work out what you would hope to see happening along the way so you can monitor as you go and sometimes catch problems before it is too late.



Unintended consequences and side effects

Beware. Many interventions have side effects you need to consider. For example, small-group or one-on-one interventions can work, especially if designed to tackle a student's weaknesses. However, sometimes too much use of this can allow students to become too reliant on close adult supervision. They can become less able to cope with learning in class when unsupervised. This can reduce success even if the intervention can work in principle.

But have no fear! Anticipating potential problems and responding to feedback can prevent these problems.



Strategy 3 – Measuring intermediate outcomes

Sometimes the final outcome of an intervention cannot be seen for a long time, perhaps so long that it will be hard to change tack if the approach is going wrong. But you can get a heads up on this by measuring intermediate outcomes that should happen if the intervention is working.

For example, giving students explicit learning objectives is one intervention for retaining knowledge and skills. Retention of knowledge is a long-term process, difficult to observe in the short term. But there is an intermediate stage that is supposed to contribute to the outcome: remembering the learning outcomes. Typically, teachers show the learning objectives on a whiteboard and ask students to write them down in their exercise books. You can then give students a quick quiz on this at the end of the lesson. If they can't remember what was on the board or in their exercise books then the strategy of giving them learning objectives is probably not working as it should. If they can at least remember the learning objectives, that is a positive sign that the rest of the material will be more easily recalled and used in the future.

STEP THREE - What support factors are needed?

Suppose you have decided that an intervention has potential to work in your setting. Still in order for it to work, some other factors might have to be present. Interventions alone are rarely enough on their own to produce the right outcome. They need help – what we call *support factors*.

Striking a match is a good way to get a flame. But not if it is sopping wet or there is no oxygen in the room. It is just the same with causes in education. The cause you concentrate on is often not the only one needed to produce the targeted outcome. There's always a set of supporting factors needed as well.

If you don't have these in your setting – or can't arrange for them to be there – your intervention will not produce the intended outcomes even if it has the capacity to do so in other cases.

Support factors affect which interventions do best: Homework vs. One-to-one tuition

The EEF toolkit summarises one-to-one tuition as offering ‘Moderate impact for high cost, based on extensive evidence’ while homework in secondary schools is ‘Moderate impact for very low or no cost, based on moderate evidence.’ They are both associated with an average of 5 months increased progress. That could be taken to suggest that homework is simply better than one-to-one tuition, achieving the same outcome but for less cost. But as the toolkit explains, the 5-month average hides a wide variation in outcomes.



Suppose two-thirds of your students have English as an Additional Language and also lack the space at home to complete their homework. Giving out more homework may be unlikely to help in this context. On the other hand, the relative disadvantage of the class means that more resources are likely to be available via the pupil premium. This makes extensive one-to-one tuition more affordable, so in this case it may be more likely to be a success than alternatives.


Strategy 4: Cake Diagrams

A cake diagram maps all the factors that must work together if the targeted outcome is to occur. So it helps you focus on the support factors necessary for your intervention to work.

As with baking a real cake, you need all the ingredients in order to get the desired outcome.

This reminds you to see whether you have all the factors you need. Maybe you will be able to get the missing ingredients or substitute an alternative.

For example, a famous California class-size reduction programme failed to increase student attainment because it was rolled out too quickly so that there were too few qualified teachers and no extra classroom space. Both these are necessary support factors if small classes are to help.



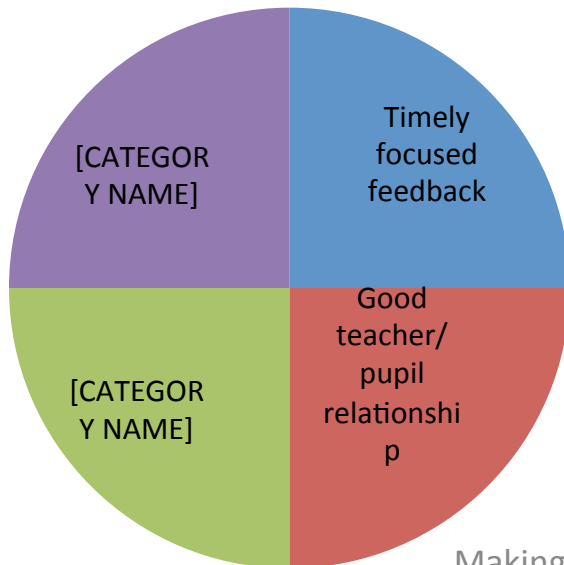
Cake Diagrams - Missing Ingredients

Suppose that providing better feedback helps most when parents know what the feedback means and know how to help their children respond to it when doing homework – it is a support factor for feedback to produce really good outcomes.


Now suppose you face a context where parents are less inclined to be involved with how well their child is doing and to help with things like homework so you cannot expect feedback to have a strongly positive effect. For some students, who struggle to respond positively to feedback without extra help, the effect could even be negative.

One response might be to combine the feedback intervention with an after-school homework club (maybe one that is already in place or one that can be established for the purpose). Explaining how written feedback works to members of staff and volunteers who supervise the homework club can allow them to help students respond when trying to improve. This could mitigate the problem of disengaged parents.

Using Feedback with parent support is present



Using Feedback with parent support absent



Combining Interventions

Usually we think that combining interventions makes the outcome even more likely – their influence is in a sense ‘additive’. On the other hand, two different causes can interact in a way that heightens or lessens the effects of each, or they may together have no effect or even the opposite to the effect that each would have on its own (known as the ‘reversal of effect direction’).

This is something to keep in mind when predicting the effects of interventions. Positive effects can even be reversed, for example, when several interventions are tried simultaneously but in a way that lacks co-ordination. Typically the overall effect of two interventions deployed together is less than the sum of what each can be expected to produce on its own.

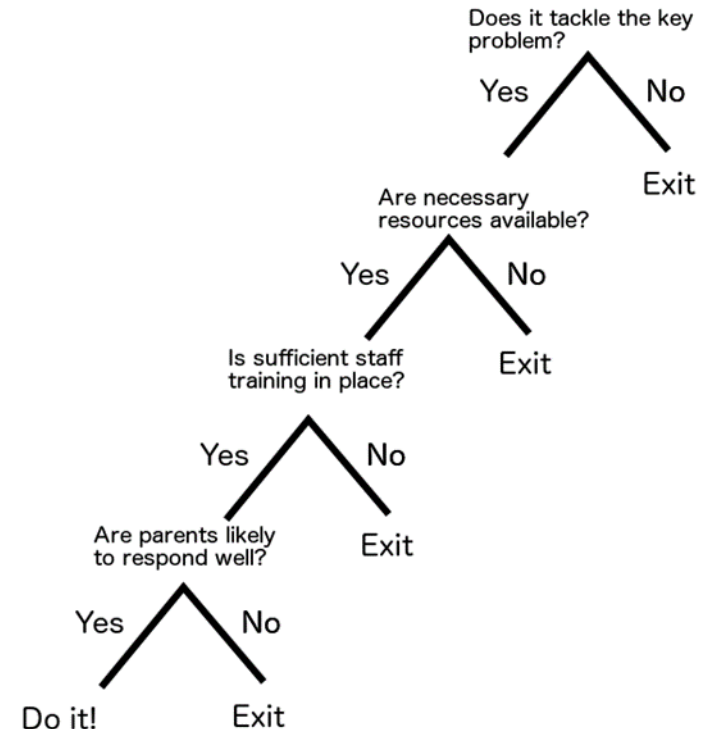
For example, you may identify a lack of effective feedback and a lack of familiarity with word-letter sounds as two factors limiting a student’s reading. Suppose you introduce two interventions, one involving a literacy computer game for the feedback problem, the other an intensive phonics programme for the word-letter sounds. They are each estimated to improve reading ages by 3 months. But combining these might produce only 4 months of improvement, rather than the 6 you would hope for if the interventions were simply additive.

Strategy 5: Quick exit decision trees

Quick exit decision trees are an alternative to cake diagrams. Rather than trying to fill in an entire cake diagram of supporting factors at once, you think of them one-by-one. Ask: can this be put into place? If the answer is YES, go on to the next. As soon as the answer is NO, you know the intervention won't work.

This is a good way to rule interventions out.

It can also be used to rule them in, so long as you are sure enough you have listed all the support factors necessary for the intervention to work.



Thinking together


You can think about evidence and your teaching practice on your own. But it can be even more useful to deliberate about these issues together. Group discussion, especially in team meetings, is likely to be essential for introducing any major intervention successfully. But it can also be useful to discuss approaches in less formal settings with colleagues and fellow professionals. The two following strategies are designed with these sort of settings in mind.



Members of your team will all have different pieces of the picture to help make a decision. But sometimes, especially in meetings, these pieces of information may not come out. One or two individuals, for instance those who are confident in speaking to the whole group, may dominate the discussion. One way of getting round this is to get people to write down their views individually before discussion takes place.

Strategy 6: The 'Crystal ball method'

Imagine a crystal ball reveals that next term the intervention you are now considering has failed. Why would it have failed? What factors would you suspect in hindsight caused it to fail? Those factors might well be the support factors you need here and now to make the intervention a success.



One of the strengths of using this method in a group is that it gives a clear message that it is all right to mention risks, to discuss the possibility of the plan going wrong. It encourages people to use their creativity, to look for difficulties and to be rewarded for finding them.

For instance, you could ask people individually to write down where they find possible sources of failure, then go around the room reading them out. Starting with the person who is in charge of the proposal can set up a clear dynamic of critical appraisal. This contrasts with the different scenario where an enthusiastic senior leader presents the proposed intervention and asks if anyone has criticisms or concerns, which can make people more shy about raising possible snags.


Strategy 7: Score-card for interventions

- Get your team to assess separately the following categories for each intervention under consideration. Give each category a score of 1-10:
 - Size of the predicted effect (10 for very big)
 - Relevance for targeted students (10 for very well targeted at the students the intervention is for)
 - Cost (10 for virtually no cost)
 - Bad side-effects (10 for no foreseen bad side-effects)
 - Supporting factors (10 for all foreseen supporting factors in place).
- Total up each individual's score for each intervention (out of 50 in this case).
- Compare the scores of team members.
- See where you agreed. If lots of people independently give a high score to an intervention, that's a positive sign that the intervention is likely to work in your school.
- You can then look at where you have different scores and discuss why.

Intervention Score-card	
Effectiveness	7
Relevance	6
Cost	10
Side-effects	3
Support factors	9
TOTAL	35

Making the Most of the Evidence: Key Points

- Research evidence is evidence of what has worked in various schools in the past. Whether something that has worked elsewhere will work *here* and *now* requires careful thought.
- Research evidence has its limits. Its strongest findings tend to be the average effects of measurable outcomes. The average may be indicative for your students but cannot be definitive. Individual students will produce very different outcomes. Some important outcomes and some important causes of them cannot be easily measured.
- It is important to understand the causal role of the intervention, or the mechanism through which an intervention is meant to succeed. This helps you know when an approach is likely to work *here* and *now* and to spot when potential problems with the intervention have emerged.
- No evidence-based intervention can be introduced alone. All will have support factors that must be in place for them to succeed.
- Research evidence can be important and very useful but it does not speak for itself. Even the best research requires critical engagement.
- Deciding whether an intervention will work, or is working, requires professional judgement, preferably in the context of free and open deliberation amongst colleagues. Research evidence is always showing what works for a particular context on particular population of students.



www.shutterstock.com · 182996780

Cartwright & Cowen would like to thank the AHRC, the British Academy, the Cabinet Office and the Spencer Foundation for support for research leading to this pamphlet.