



21 APRIL 2020

## Critical Thinking

PROFESSOR ALEX EDMANS

### Confirmation Bias: Why We Might Not Think Critically

- a. “Equality is better for everyone”
- b. “CEO remuneration packages actively discourage innovation in UK’s top companies”
- c. “The outperformance of ESG strategies is beyond doubt”
- d. “Professors on corporate boards increase profits”

These four statements have three things in common:

1. Those making them argue that they’re backed up by research.
2. Most people would like these statements to be true. Indeed, I have a particular interest in them being true:
  - a. In last year’s lecture, “Executive Pay: What’s Right, What’s Wrong, and What’s Fixable”, I argued that shares should be given to all employees, not just the CEO, so that the benefits of corporate success are distributed fairly.
  - b. The same lecture argued that executive pay should be reformed to encourage long-term thinking
  - c. My inaugural lecture, “Purposeful Business: The Evidence and the Implementation”, argued that companies that perform well on Environmental, Social, and Governance dimensions may also deliver higher returns to shareholders.
  - d. If this statement were true, I would be more likely to be offered board seats.
3. None of the statements is actually backed up by rigorous research.

Importantly, (2) and (3) are linked. Since most people would like these statements to be true, they accept them uncritically, even if the underlying evidence is flimsy. This is an example of confirmation bias: accepting a statement uncritically if it confirms what we’d like to be true, and rejecting a statement that contradicts it.

This lecture is on how we can think critically and avoid succumbing to confirmation bias. Note that “critical thinking” includes other topics such as logical reasoning (e.g. deduction, induction, and abduction), but I will focus on confirmation bias since it is so pervasive. My TED talk “What to Trust in a Post-Truth World” also addresses confirmation bias; I’ve deliberately chosen to focus on different aspects in this lecture to avoid duplication. Thus, the reader interested in this topic should find the TED talk of additional value.

To illustrate my points, I need to use examples. These examples are *not* chosen in order to “bash” particular studies. Indeed, as mentioned above, I have a personal interest in these studies being true. If my goal was to bash a study, it would be easiest for me to choose a non-impactful (and thus likely low-quality) one. I’ve deliberately chosen studies that are highly influential – in some cases I used to quote them myself because of my own confirmation bias – to illustrate how easy it is to think uncritically.

The purpose of this talk is to be entirely constructive. Companies underperform, economies stagnate, and societies malfunction. If improving performance was as easy as claimed by some viral talks, influential studies, and best-selling books, this would not be the case. Instead, the underperformance may be due to following conclusions that aren’t actually backed up by evidence. Any elements of this talk that may appear sharp are only to highlight the severity of confirmation bias, and thus the importance of critical thinking.

### The Neurological Basis for Confirmation Bias



A 2016 study, published in *Scientific Reports*, featured human subjects with liberal political views.<sup>1</sup> The researchers read out either a political statement that the subjects agreed with (e.g. “The death penalty should be abolished”) or a non-political statement (e.g. “The primary purpose of sleep is to rest the body and mind”). The researchers then read out contradictory evidence and measured the subjects’ brain activity. There was no effect when non-political statements were contradicted. When political statements were contradicted, the amygdala of subjects’ brains was triggered. This is the same part of the brain that is activated when a tiger attacks you and induces a “fight-or-flight” response. Thus, some people respond to their views being contradicted as if attacked by a tiger. It overrides the prefrontal cortex, the rational part of the brain. Nobel Prizewinner Daniel Kahneman, in his book *Thinking, Fast and Slow*, referred to the impulsive thought process as System 1 and the rational process as System 2.

We’ll now look at various forms in which confirmation bias may manifest, and then discuss practical tips to avoid such a bias.

## 1. Statements Without Evidence

Many views are expressed in the form of simple statements. When making a statement, you have the incentive to make it as one-sided as possible. For example, Senator Bernie Sanders claimed that “Wall Street CEOs who helped destroy the economy, they don’t get police records. They get raises in their salaries.” Such an extreme quote will be widely shared by opponents of capitalism, since it illustrates capitalism’s alleged ills in a very stark manner. Moreover, an extreme quote reduces the need to back it up with evidence, by giving the impression that the point is so obvious that no proof is needed. Indeed, extreme quotes are often made without referring to evidence at all. The CEOs of Bear Stearns and Lehman each lost nearly \$1 billion of personal wealth when their firms collapsed.

You may think that you didn’t need to attend a Gresham lecture to hear the advice “check the facts”. It seems obvious that you should check the facts. However, due to confirmation bias – people are willing to believe that capitalism is crooked – System 1 is in overdrive even though System 2 might tell you that you should check the facts. Thus, many people shared or believed the quote without verifying it.

## 2. Statements Citing Evidence That Doesn’t Support the Statement

Sometimes, statements are backed up by references. However, these references may not actually show what the statement claims that it shows. But, if the statement confirms what readers would like to be true, they won’t bother checking the reference.

For example, the UK House of Commons Select Committee Report on Executive Pay stated that “the evidence is at best ambiguous on the impact of individual CEOs on company performance”, with a footnote. Many people may have seen that there was a footnote, said to themselves “there must be evidence behind it”, believed it and moved on without checking the footnote – who actually provided the evidence. Others may have gone further and checked the footnote, and seen that it referred to the evidence submitted by “Professor Alex Edmans” to the Executive Pay inquiry. Perhaps people might have kindly thought that I am a semi-reliable source of evidence, and similarly moved on. However, even that is insufficient – one has to check whether the evidence I submitted actually said what the report claimed it did. In fact, nowhere did my evidence state this. The closest sentence I wrote on the “impact of individual CEOs” was “CEOs with high equity incentives outperform CEOs with low equity incentives by 4-10% per year, and the researchers do further tests to suggest that the results are causation rather than correlation”, which suggests that CEOs do have a significant impact on firms.

Similarly, Malcolm Gladwell’s book *Outliers* famously introduced the 10,000 hour rule – if and only if you spend 10,000 hours on any cognitively complex task, you’ll become an expert. He quotes research Anders Ericsson as evidence. I must admit that I viewed this best-selling book as authoritative, and thus used to quote it in the orientation address I gave to LBS MBAs each year. It confirmed what I wanted to be true – you can do anything

---

<sup>1</sup> Kaplan, Jonas T., Sarah I. Gimbels and Sam Harris (2016): “Neural Correlates of Maintaining One’s Political Beliefs in the Face of Counterevidence.” *Scientific Reports* 6, 39589. For a humorous cartoon summarising this study, and confirmation bias more generally, see <https://theoatmeal.com/comics/believe>.



you put your mind to – so I didn't bother to look up the original study by Ericsson. But I've since done so, and it makes no such claim. Indeed, Ericsson's own book, *Peak*, has a section "The 10,000 hour rule isn't really a rule". I'll explore what Ericsson's research actually shows on how to become an expert in Lecture 6, [The Growth Mindset and the Abundance Mentality](#).

A related issue is that a statement may be quoted out of context. Former GE CEO Jack Welch is frequently cited as saying "Shareholder value is the dumbest idea in the world". Two points that we've already discussed apply here. First, the extremity of this statement should ring alarm bells – even if shareholder value is not the best idea, it's unlikely to be the worst idea in the world. Second, the statement is not backed up by evidence. While Welch was an extremely successful CEO, he did not conduct a large-scale study comparing companies which pursued shareholder value with otherwise identical companies which pursued a different objective. But there's a third, new point – this statement is quoted out of context. The full sentence is "On the face of it, shareholder value is the dumbest idea in the world", which has a very different meaning. Welch then goes on to clarify his point and explain why shareholder value is not as "dumb" an idea as it may initially seem. Many people quote Welch, conveniently missing out "on the face of it", even though it significantly changes the meaning of the sentence – indeed, you can Google to find many articles that do so. This is bad practice.

Even more worrying than others mis-citing evidence that the authors never claimed, authors themselves may misportray their own evidence. For example, a TED talk used to be titled "Want a more innovative company? Hire more women." The study correlated the share of revenues that stems from new products with a composite of 6 measures of diversity, including not only gender but other factors such as country of origin and age. Ignoring the fact that correlation does not imply causation (which we'll shortly discuss), and assuming that "share of revenues from new products" is a good measure of innovation (which it may not be – this measure could be high simply because you've failed to maintain your existing products), this study does not even document a correlation between innovation and gender diversity – the correlation is with a composite measure of diversity which contains five other measures than gender. It could be that one or more of the five other components drove the relationship.<sup>2</sup>

Moreover, some studies make claims without correlating anything at all. A study claimed that "CEO remuneration packages actively discourage innovation in UK's top companies" yet does not have any measures of innovation, so there are not even correlations.<sup>3</sup> It finds that UK companies' remuneration packages have certain features (e.g. earnings-based bonuses), which they believe to be detrimental to innovation. However, without showing that these features actually deter innovation, this conclusion is invalid – it is a case of "sentence first, verdict afterwards." The paper simply *assumes* that earnings-based bonuses deter innovation because innovation expenditure decreases earnings. But the effect of incentives on behaviour is extremely complicated, which is why there is an entire literature on it. Indeed:

- The same logic implies that earnings-per-share incentives will encourage share buybacks, because share buybacks increase earnings-per-share. But there is no evidence of this, either in the US<sup>4</sup>, or in the UK study PwC and I conducted for the UK government.<sup>5</sup>
- Some incentives can have the opposite effect to what intuition would suggest – e.g. incentives can *deter* activity by undermining intrinsic motivation.
- Earnings-based incentives may encourage the company to "get its act together" and improve in many dimensions, including innovation. For example, by reducing injury frequency, Alcoa got its act together and improved in many other dimensions including productivity. Indeed, evidence shows that takeover threat (which might cause a company to focus on earnings to avoid takeover) increases innovation.<sup>6</sup>

---

<sup>2</sup> The title of the talk has since been changed to "How Diversity Makes Teams More Innovative", perhaps due to complaints that the original title was misleading. However, this new title still claims a causal relationship.

<sup>3</sup> "The Invisible Drag on UK R&D: How Corporate Incentives Within the FTSE 350 Inhibit Innovation"

<sup>4</sup> Bennett, Benjamin, J. Carr Bettis, Radhakrishnan Gopalan and Todd Milbourn (2017): "Compensation Goals and Firm Performance." *Journal of Financial Economics* 124, 307–330.

<sup>5</sup> "Share Repurchases, Executive Pay and Investment". BEIS Research Paper Number 2019/011.

<sup>6</sup> Atanassov, Julian (2013): "Do Hostile Takeovers Stifle Innovation? Evidence from Antitakeover Legislation and Corporate Patenting." *Journal of Finance* 68, 1097–1131.



### 3. The Narrative Fallacy

Let's say that neither problem (1) or (2) apply – we have facts that seem to be true. For example, Simon Sinek claims that Apple used to have the “why” statement “Everything we do, we believe in challenging the status quo.” And Apple has indeed been extremely successful. (It turns out that nowhere did Apple say “everything we do, we believe in the status quo”, nor even similar language, but people took this for granted and indeed you can find 23,000 articles on Google quoting this phrase because they accepted Sinek's claim uncritically. However, I've already discussed the importance of fact-checking, so let's assume that this claim is true and Apple did have this statement.)

The *narrative fallacy* is the temptation to see a set of facts and ascribe a cause-effect relationship to these facts, even if there's no true link. We have this temptation because we want to be able to make sense of the world, and thus form links between phenomena even if they're unrelated. Indeed, when something happens that's difficult to explain (e.g. a tragedy), we tell kids “everything happens for a reason.”

Continuing the above example, Sinek's best-selling book *Start With Why* claims that Apple's success is due to it being driven by its “why”. I very much like Sinek's TED talks and my own book is on the importance of purpose. However, it is impossible that Apple's “why” drove its success. There could be very many other contributing factors, e.g. Steve Jobs's ideas or his network of relationships. However, the narrative that success was due to Apple's “why” is particularly appealing to readers as anyone can adopt a “why” – whereas not everyone can suddenly think of a novel idea or has a network of relationships.

Moreover, often interpretations are imposed after-the-fact, reverse-engineered to fit the facts. Indeed, as Jobs himself pointed out in his 2005 graduation speech at Stanford University “You can't connect the dots looking forward; you can only connect them looking backwards”. With hindsight, it may appear that Apple started with “Why”, but it may have not actually pursued this strategy at the time – just as a snooker player may unintentionally pot a ball when aiming for a different pocket. As discussed above, Apple never said “everything we do, we believe in challenging the status quo” - instead, this seems to be a “why” statement that Sinek ascribed to them after the fact, because it fits his narrative.

A solution might be to study many companies. If we look across many successful companies and find that they all have a common characteristic (e.g. humble leadership), then isn't this proof that the success was due to this common characteristic? Indeed, this is the approach taken by many best-selling business books, such as *In Search of Excellence*, *Good to Great*, and *Built to Last*, which studied 43, 11, and 18 companies respectively. However, even ignoring issues of correlation vs. causation (which I'll get to shortly), “all successful companies have humble leadership” *does not mean* that “all companies with humble leadership are successful” (just as “all company CEOs are humans” does not imply that “all humans are company CEOs”). This is an error in logical reasoning. To show that all companies with humble leadership are successful latter, you can't just focus on companies that ended up successful – you'd have to study *all* companies with humble leadership (regardless of whether they turned out successful) and compare them with companies without humble leadership.

Indeed, force-fitting an explanation after the fact – whether to one or to many companies – is an example of *spurious correlation*. If there is no true relationship, then it will disappear in subsequent time periods (economists refer to this as “out-of-sample”). Indeed, many of the companies in the above books have subsequently underperformed, such as Xerox, Philip Morris, GE, Fannie Mae, Circuit City, Wells Fargo, and Ford.

As Daniel Kahneman points out, “regression to the mean” is to be expected if the true cause of the relationship is luck. He writes: “The basic message of *Built to Last* and other similar books is that good managerial practices can be identified and that good practices will be rewarded by good results. Both messages are overstated. The comparison of firms that have been more or less successful is to a significant extent a comparison between firms that have been more or less lucky. Knowing the importance of luck, you should be particularly suspicious when highly consistent patterns emerge from the comparison of successful and less successful firms. In the presence of randomness, regular patterns can only be mirages.”



The standard way to address “luck” is to test for statistical significance. A significant test gives the likelihood that a statistical relationship arose from pure chance. The conventional significance level is 5% - we only conclude that a relationship exists if there is less than a 5% chance that the data could be due to luck. However, many studies have no tests of statistical significance. To show statistical significance, you typically need hundreds of companies, not 43, 11, or 18.

#### 4. Alternative Explanations

So if you find a statistically significant correlation between two variables (A and B), does this prove a relationship? For example, A is whether a company has a professor on the board, and B is performance. Does this mean that professors on the board cause superior performance? No, for two reasons:

- *Reverse causality.* B might have caused A. It’s good performance that leads to professors being on the board. Professors might be risk-averse, hence preferring the ivory tower of academia to the cut-and-thrust of business. So they’ll only join the board of a company that’s already performing well.
- *Omitted variables.* There’s a third variable, C, which causes both. For example, a forward-thinking board might want to increase its diversity of thought and appoint a professor. Thus, C causes A. Also, a forward-thinking board causes the company to perform well due to injecting it with good ideas. Thus, C causes B as well; there’s no relationship in either direction between A and B.

Now the phrase “correlation does not imply causation” is very well-known. Again, you probably didn’t need to listen to a Gresham lecture to hear this phrase – but we often forget it because it’s only System 2 that gives this warning, and it may be overruled by System 1. Accepting correlation uncritically as causation is extremely common in business, particularly since people like to believe that there are certain practices that will improve performance. Indeed, many studies do contain a few sentences acknowledging that correlation doesn’t imply causation, but as a perfunctory disclaimer – the authors typically market the report (through press releases, launch events, and adverts) as showing causation. For example, statements such as “Equality is better for everyone” imply causal relationships even though the study contains brief caveats.

Sometimes authors will claim to have separated causation from correlation. There are indeed valid methods of doing so, but often these claims are unwarranted. For further detail, see the article “A Layman’s Guide to Separating Causation from Correlation ... And Noticing When Claims of Causality are Invalid” ([www.alexedmans.com/correlation](http://www.alexedmans.com/correlation)).

To illustrate how correlation is often interpreted as causality, despite the “correlation does not imply causation” phrase being so well-known, consider the book *The Spirit Level: Why Equality is Better for Everyone*. I have chosen this example as the book has been extremely successful and influential, and I commend it for injecting the importance of inequality into the public debate; this discussion will be purely on its evidence base. The book gives scatter plots of country-level inequality against various country-level outcomes (e.g. obesity, happiness, health and social problems) and argues that there is a relationship to the naked eye. In addition to problem (3) (there are no tests of statistical significance), there are also no control variables. There are very many other country-level factors that could cause the outcomes, the most obvious being poverty. Countries with high inequality also have low average incomes, and it could be poverty, rather than inequality, which leads to these negative outcomes. The authors address this explanation by drawing a scatter plot of health and social problems against poverty and arguing there is no relationship to the naked eye (again, without statistical tests).

However, elementary statistics tells you that it’s invalid to rule out a “poverty” explanation by doing *separate* regressions on poverty and inequality. Instead, you need to do a *single* regression on inequality controlling for poverty (and for the myriad of other factors that may drive obesity, e.g. national diet, amount of PE taught at schools, availability of fitness facilities). Just as seriously, the authors forget about poverty explanation for the rest of the book; almost all their future graphs (with different outcome variables to “health and social problems”) study inequality only. Again, confirmation bias have led to System 1 being in overdrive, which is why respected newspapers and other outlets have referred to the book’s evidence being very rigorous and hard to dispute. However, any A-level statistics student would be able to dispute it.



Reverse causality is also an issue. It could be that inequality is the result of the various “outcomes”, not the cause. For example, ill health may cause people to be unable to study or work, in turn leading to inequality. The authors acknowledge on rare occasions that correlation is not causation but they “believe” that this relationship is causal. Their arguments for this belief are invalid. One is that the strength of the relationship suggests that it’s causal, but a relationship could be strong yet be in the other direction (or there could be a strong omitted variable). Another is the fact that a relationship has been found in many settings. Again, causality could be in the other direction in all of these settings, or an omitted variable could be present in all of these settings (and it need not be the same omitted variable in each one).

## 5. Universality Bias

Up to now, we’ve discussed *internal validity* – whether a study is internally coherent and shows what it claims within the dataset that it uses (e.g. data on firm performance and CEO humility across 100 companies). Internal validity is violated if we ascribe a relationship when there is none – for example, if there is no causal effect from CEO humility to firm performance within those 100 companies.

This discussion concerns *external validity*. Even if we could convincingly nail a relationship with a dataset, it may not apply to companies outside that dataset (i.e. outside those 100 companies). I am coining the phrase *universality bias* to describe our preference for “universal” explanations that apply elsewhere. We like to have a theory of everything. This allows us to make sense of the world more easily. But very rarely is there a theory of everything which applies to every company or ever person. Even if we could prove that Apple’s “why” caused its success, a “why” might not help other companies. Perhaps those in low-margin, hyper-competitive industries should focus more on cost control rather than finding their “why”. There are books, articles, and talks on “Five ways to ...” that claim to apply to everyone. But, “Five ways to nail a job interview” might be misleading, as the best advice might differ depending on whether you’re an introvert or extrovert, or the type of job you’re interviewing for. “Five ways to lose weight really fast” should depend on your level of fitness, age, exercise regimen, and so on. But, a book which claims that its advice applies to every reader is likely to sell more copies than one that acknowledges that the evidence was only collected on a small subset of companies or individuals.

Universality bias illustrates the danger of overextrapolating from single stories. Stories are vivid, they’re powerful, they bring a topic to life and get re-told. They’ve thus been used successfully in business schools (as case studies), books, and TED talks. But stories tell you little unless they’re backed up by large-scale data, because you can always hand-pick a story to support any viewpoint. Supporters of an exclusive focus on profit might use the story of GE under Jack Welch to show it can succeed. Opponents might use the story of Enron to show it can fail. Indeed, both GE and Enron are major business school case studies, but neither story tells us whether running a company for profit works in general.

Even if you show a relationship in a large sample of data, it may not apply outside that sample. Gladwell claimed that “The 10,000 hours rule says that if you look at any kind of cognitively complex field, from playing chess to being a neurosurgeon, we see this incredibly consistent pattern that you cannot be good at that unless you practice for 10,000 hours”. However, the Anders Ericsson study was on violinists. He and his coauthors did not consider chess, neurosurgery, or other cognitively complex activities. Limited external validity applies even to my own work. A paper I published in 2011 showed that US companies with high employee satisfaction outperform their peers.<sup>7</sup> Since then I’ve extended the study to non-US countries and found that the relationship only holds in those with flexible labour markets; it doesn’t hold in countries with regulated labour markets.<sup>8</sup>

The above concerns the application of a result to fields that you haven’t even analysed (e.g. neurosurgery). A related manifestation of the universality bias does involve analysing data in different fields, but imposing your “theory of everything” on it and not considering alternative explanations – as captured by the *Alice in Wonderland* phrase “sentence first, verdict afterwards”. If authors have a hammer, everything looks like a nail. For example, the blurb

---

<sup>7</sup> Edmans, Alex (2011): “Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices.” *Journal of Financial Economics* 103, 621-640.

<sup>8</sup> Edmans, Alex, Lu Li and Chendi Zhang (2020): “Employee Satisfaction, Labour Market Flexibility, and Stock Returns Around the World.”



for *The Spirit Level* reads: “Why do we mistrust people more in the UK than in Japan? Why do Americans have higher rates of teenage pregnancy than the French? What makes the Swedish thinner than the Greeks? The answer: inequality.” But there may be very many other explanations for those phenomena, and the authors don’t consider (let alone test) them.

### Practical Tips: What to Trust

I recognise the above examples paint a very bleak picture. They seem to suggest that we cannot trust anything – even statements in official government documents or in best-selling books. Even if a statement by an influential person may not be backed up by evidence (1). Even if it refers to evidence, the statement may be misquoting what the researchers actually claim (2). Even if the statement correctly conveys what the researchers claim, this claim may be invalid – the researchers may not have even studied the variable in question (6), they may be ascribing a relationship where none exists (3), they may not have considered alternative explanations (4), or the results may be overextrapolated beyond the initial context (5).

Unfortunately, the picture is indeed bleak. Very influential studies and books have a much weaker evidence basis than they claim. Given the neurological basis for confirmation bias, there is a huge incentive to convey results that others would like to be true. Your policy will get widespread support, your book will become a best-seller, and your report, blog, quote, or talk will be widely shared.

So, what are practical tips we can use to ensure we think critically and overcome confirmation bias?

1. For a statement, check if it is backed up by a reference. If it is a quote, check the context.
2. If it is, skim the reference and see whether it actually claims what the statement states that it claims. (Even if the reference is an academic paper behind a paywall, the abstract will always be available free of charge. Moreover, the full version of most social science papers is available on [www.ssrn.com](http://www.ssrn.com)).
3. If the authors’ claim agrees with the statement, check that the authors’ actual analysis supports the claims. Think whether there might be plausible alternative explanations, such as omitted variables, reverse causality, or luck.

It’s true that most published papers are behind a paywall – you need a subscription to the journal to access it. This might seem to preclude point (3). However, for research on business and other social sciences, almost always the pre-publication will be available for free on the Social Science Research Network ([www.ssrn.com](http://www.ssrn.com)). This is exactly the same as the published version except the latter will incorporate some minor copyediting comments to comply with house style. Thus, it’s indeed possible to have a quick check of a paper without having a journal subscription.

Nevertheless, point (3) might be difficult for non-academics to verify, particularly if the paper is technical. So, the reader should check if the paper is published in a top peer-reviewed academic journal. Practitioner studies have significant value. They often have better access to data than academics, and are often a superior source for statistics. But academics have particular expertise in drawing relationships *between* statistics – teasing apart causation from correlation and addressing alternative explanations. Importantly, academic studies have to undergo rigorous peer review to check their scientific accuracy. The very top journals have the highest standards, using the world’s leading specialists to scrutinise a manuscript, and reject up to 95% of manuscripts. The 5% not rejected aren’t immediately accepted either; instead, their status is “revise-and-resubmit”. The reviewers highlight concerns that the authors need to address, and the paper can still be rejected at the next round. It’s not unusual for a paper to take five years to be published after its first draft. A hard slog for the authors, but it helps ensure that the published results are correct.

Here’s an example of the importance of evidence quality, to show it’s not just an “academic” issue but can transform a study’s implications for real-world practice. In the UK House of Commons’ 2016 inquiry into corporate governance, a witness quoted research which “found that firm productivity is negatively correlated with pay disparity between top executive and lower level employees”, referencing a January 2010 work-in-progress draft. The finished



version had actually been published three years prior to the inquiry.<sup>9</sup> Having gone through peer review and tightened up its methodology, it found completely the opposite result:

- *We do not find a negative relation between relative pay and employee productivity.*
- *We find that firm value and operating performance both increase with relative pay.*

I also gave evidence at the inquiry, and I submitted follow-up evidence explaining that the study's results had been completely overturned in peer review. Despite publishing my follow-up evidence, perhaps due to confirmation bias, the final report of the inquiry still wrote that "The [witness] states that "There is clear academic evidence that high wage disparities within companies harm productivity and company performance."<sup>10</sup> " This is an example of why the picture is indeed bleak, since even House of Commons reports include evidence known to be tainted. Thus, it is essential that we exercise critical thinking; we cannot even take the conclusions of official reports at face value.

I stress the importance of a paper being published in a *top* peer-reviewed journal. That a journal calls itself "peer-reviewed" is far from sufficient to guarantee its rigour, since there's a vast range in the quality of reviewing standards. The analytics company Cabell's has a blacklist of 8,700 journals that claim to be peer-reviewed but actually aren't. Journal quality can easily be checked by looking at one of the freely-available lists of the best ones, such as the *Financial Times* Top 50. One doesn't need to be an academic insider to do this. Peer review isn't perfect – mistakes are made. Sometimes sloppy papers get accepted and good papers are rejected. But it's better to go with something checked than something unchecked.

Of course, every paper starts out unpublished, but we can gauge the quality of a new paper by considering the authors. Similarly, not every communication of information is through an academic paper – it may be through an article, interview, or even a Tweet – but again we can gauge the reliability by considering the authors. We should particularly ask ourselves:

1. Do they have expertise *in the relevant field*? This helps avoid the issue of "halo effects", where a person is seen as a guru due to expertise in one field and thus is trusted as an authority in different fields. Liverpool FC manager Jürgen Klopp honourably declined to give his opinion on the coronavirus crisis, fearing that people might believe him because he is famous. While System 2 might tell you that you should clearly not believe a football manager on a scientific issue, actors have persuaded people to not vaccinate their children. A less obvious error is that, because Jack Welch was a successful CEO at one firm, he is viewed as an authority on the effect of a shareholder value on performance in firms in general – which is a question that warrants academic study rather than business leadership. Some doctors are proclaiming miracle cures for the coronavirus, even if their expertise is not in pathology or epidemiology. People ask me whether we should invest in the stock market after it has fallen due to the coronavirus, because I'm a Professor of Finance, although my expertise is in individual companies rather than macroeconomic conditions.
2. Do they have a hammer? Are they known for a "theory of everything" which may have blinded them to alternative explanations?
3. Relatedly, if they had found the opposite relationship, or no relationship, would they have still released the study? Or are their incentives strongly skewed to finding the relationship they are claiming? For example, authors who claim that inequality explains every social problem would be unlikely to publish a paper showing that inequality does not cause a particular problem. In 2019, PwC and I released a study, commissioned by the UK government, showing that share buybacks are generally not misused in the UK.<sup>11</sup> I have other work suggesting share buybacks are sometimes misused in the US.<sup>12</sup> This hopefully gave further credence to the study, as any bias might have led me to wanting to show that share buybacks are similarly misused in the UK.

---

<sup>9</sup> Faley, Olubunmi, Ebru Reis and Anand Venkateswaran (2013): "The Determinants and Effects of CEO-Employee Pay Ratios." *Journal of Banking and Finance* 37, 3258–3272.

<sup>10</sup> The goal of the report is to summarise the main points, rather than repeat every point. Thus, there was no obligation for the report to include a point made that was later shown to be false.

<sup>11</sup> "Share Repurchases, Executive Pay and Investment". BEIS Research Paper Number 2019/011.

<sup>12</sup> Edmans, Alex, Vivian Fang, and Allen Huang (2020): "The Long-Term Consequences of Short-Term Incentives."





## Is it Practical to Check Every Source?

Are these practical tips indeed practical? Even if we don't read the actual analysis in every paper referenced, it takes time to check the quality of the journal it's published in, or the credentials of the authors. So we should only scrutinise the references for the most important points that an article makes. Moreover, we should be particularly discerning if:

1. The article or research draws conclusions that you would like to be true (your own confirmation bias)
2. The article or research draws conclusions that the author would like to be true (the authors' or researchers' confirmation bias).
3. The article of research is one-sided, claiming "clear evidence" or universality. There are nearly two sides to almost every social science issue, and claiming that an issue is unambiguous suggests that the authors may not have seriously considered the other side and blundered into their conclusion – or are aware of the evidence on the other side but buried it. For example, the TED talk on diversity claimed "The answer was a clear yes, no ifs, no buts. The data in our sample showed that more diverse companies are simply more innovative, period." *The Spirit Level's* subtitle is "Equality is better for everyone." An article claimed "The outperformance of ESG strategies is beyond doubt" when meta-analyses (reviews of multiple studies) find that the evidence is mixed at best.<sup>13</sup> This misportrayal can have serious implications. The leading UK broker Hargreaves Lansdown claims that "Study after study has shown that businesses with ESG (environmental, social and governance) characteristics have outperformed their peers"<sup>14</sup>, when this is not true. This could lead to citizens investing their financial future in funds that they believe will outperform but do not.

### An Example: Test Your Skills

To test your critical thinking skills, consider the following opening paragraph of a *Harvard Business Review* article.

"Joe Bower and Lynn Paine "had me at hello" (to quote *Jerry Maguire*) with their new HBR article, "The Error at the Heart of Corporate Leadership." Laying out their data, they find that long-term oriented companies create more financial value and more jobs. In fact, if more American companies were focused on the long term, they estimate, investors would have an additional \$1 trillion, workers would have an additional 5 million jobs, and the country would have more than an additional \$1 trillion in GDP."

Should we take this article at face value?

1. Consider the author. Here, we don't even need to examine his credentials. The fact that the Bower-Paine article "had me at hello" is a strong signal of confirmation bias – he was willing to accept their conclusion uncritically.
2. Check the reference. It turns out that the Bower-Paine article does not lay out any data or make any such claim. There is, however, a separate article in the same *HBR* issue as Bower-Paine (by different authors) that makes the claim. It has various measures of long-term behaviour (e.g. investment) and finds that long-term companies subsequently perform better.
3. Check for alternative explanations (internal validity). Reverse causality is an issue. A basic finance principle tells you that when companies' future outlook is better, they should invest more. Omitted variables are an issue. It could be that a good CEO invests more (as she has better ideas) and also improves long-term performance.
4. Check for external validity. The study states "Extrapolating from the differences above, we estimate that public equity markets could have added more than \$1 trillion in asset value ... companies could have created five million more jobs in the United States – unlocking as much as \$1 trillion in additional GDP." However, such an extrapolation to other companies beyond the sample is invalid. It will not make sense for all companies to invest more, e.g. those in declining industries. If Xerox had invested more, it wouldn't have performed better. Moreover, if all companies within an industry make "arms-race"-type investments, such

<sup>13</sup> Renneboog, Luc, Jenke Ter Horst, and Chendi Zhang (2008): "Socially Responsible Investments: Institutional Aspects, Performance, and Investor Behavior." *Journal of Banking and Finance* 32, 1723–1742.

<sup>14</sup> <https://www.hl.co.uk/features/sustainable-approaches-to-a-profitable-future>



as some types of advertising, this split the pie (taking value or jobs away from competitors) rather than growing the pie and adding to aggregate welfare.

© Professor Alex Edmans, 2020

Some of this summary is adapted from my book, *Grow the Pie: How Great Companies Deliver Both Purpose and Profit* (Cambridge University Press).