MRS Reports/2018

Deconstructing bias

Lessons from 70 years of research and insight

MRS Delphi Group

MRS Evidence Matters™

Contents

About MRS Delphi Group

The MRS Delphi Group is led by a collection of respected thinkers in the marketing and research sectors. The Group delivers valuable insight across a range of important business, social and political issues.

The Steering Board includes: Dr. Nick Baker (chair) Morar HPI; Phil Sutcliffe, Kantar TNS UK; Colin Strong, Ipsos; Kevin Evans, Pepsico; Zoe Ruffels, Samsung; Tim Britton, Springer Nature; Cat Wiles, VCCP; Jake Steadman, Twitter; Chet Henderson, Unilever; Jane Frost, CEO of MRS. Additional thanks are due to Caroline Hayter and Alison Camps for their contributions.

Foreword
Jane Frost, Chief Executive, MRS

Researchers know a huge amount about unconscious bias – how it affects people's behaviour and the responses they give, and how to recognise the warning signs when bias starts creeping into research design and analysis.

Bias is insidious and impossible to strip out or completely neutralise. It can be a force for good or discriminatory. It can be unconscious, or bias can be consciously applied to redress an unfair imbalance.

The purpose of this report is to share the knowledge that resides within research and insight professionals with the other parts of the business. One of the founding principles of the MRS Delphi Group is to make more of the collective knowledge of the research sector. Behind every insight presented and every project delivered is a treasure trove of information and learning which is, more often than not, discarded. Just as businesses have become acutely aware of the value of their data, similarly this sector can use to its advantage the combined intelligence of 70 years of research practice and experience.

Recognising where bias occurs and having some tools to redress the balance is critical for decision making. This is particularly true in those parts of the business where the impact of bias is critical to the overall health of the business – from HR to data analytics; from customer services all the way up to the senior leadership team.

I have long been concerned that algorithms may be directing our lives to a degree that doesn't allow for serendipitous discovery – whether its new music on Spotify or diverse viewpoints on Facebook. To rectify this we may need a positive bias to aid us in discovery – to put new ideas, products and perspectives in our purview.

As experts in their field, researchers, data analysts and insight professionals are paid for their viewpoint. An interpretation of research findings will be nuanced and informed by personal experience as well as by hard fact. You could say we pay experts for their bias.

Increasingly, more parts of the business are accessing customer data and analytics. Some of us may lack the skills to interpret this evidence and make the appropriate decisions. I think we all need some practical tips to help us adopt a critical role and challenge our analysis; to ensure we are applying controlled and consistent measures, for example sample sizes and survey design.

As a senior decision-maker I hope you find this report as useful as I have.

Jane Frost CBE,
Chief Executive, MRS

Introduction

Researchers are familiar with unconscious bias. In fact, over the 70 years since the sector was established researchers have created methods specifically designed to remove bias where possible. Whether it's the problems of sampling, making decisions on how to phrase the questions we ask and interpret the answers, or questioning the assumptions that we (or our clients) make.

As such the research sector has a valuable contribution to help businesses become more effective and efficient in how they deal with bias.

This paper considers three aspects of unconscious bias in business, and how market research expertise can help solve the problems that it creates:

1. Cultural bias

The way we unintentionally mistreat staff and customers, and miss out on opportunities.

2. Management bias

The flawed decision-making that undermines both strategy and day-to-day decisions.

3. Data bias

The danger of unquestioningly "baking in" the bias in our data and assumptions to the algorithms that organisations use.

"Debiasing" has become big business – one oft-quoted figure is that, in the US, it has become an $8 billion business.[1] In this paper we will also consider the approaches that organisations could make, and what we know about their impact, using the available evidence. The conclusion is that we appear to have a bias to action: that is, we do something, whether or not that course of action is likely to create a positive outcome.

Indeed, there is plenty of evidence that superficial efforts to offset unconscious bias, though well-meaning, may backfire.

There are, however, interventions that have more chance to help. Where those exist, we point them out.

We believe it is time to look again at the way in which we evaluate and respond to the world, at how we make judgements and evaluate them, because it could be that the largest mistakes are the ones we don't yet know we are making.

[1] McKinsey: Focusing on what works for workplace diversity (2017) https://www.mckinsey.com/featured-insights/gender-equality/focusing-on-what-works-for-workplace-diversity

Why bias matters

Unconscious bias has both a personal and an organisational cost – as seen in these three stories about recruitment:

1

In 2003, in Boston and Chicago, Emily and Brendan replied to job advertisements with a CV and covering letter. So did Lakisha and Jamal. So did many other people, all of whom (like these four) had been invented by two economists. They randomised the qualifications and experience of their imaginary applicants, and found that those with white-sounding names were 50% more likely to be called back for interview than Lakisha and Jamal, across all occupations and company sizes.[1] This experiment has since been replicated in different countries, with different groups.

2

When US symphony orchestras changed policy to have their auditions performed behind screens, so that it was not clear if they were male or female, the result was a 25% of the increase in the percentage of orchestra musicians who were female. In the 1970s, 5% of musicians in orchestras were female, and it's now almost 40%.[2]

3

When Amazon attempted to avoid bias in hiring by building an AI that it trained using CVs it had received in the previous 10 years, the AI taught itself to penalise applications from women. The problem was partly that the overwhelming majority of its applicants, used the train the AI, had been male. Development on the tool was abandoned.[3]

The three stories show that unconscious bias is real and that its effects are widespread. In the case of hiring, if you unintentionally do not select well-qualified candidates, then that damages both them and your organisation. But Amazon's experience shows that it's not easy to overcome this problem. (As you will see, technology can be biased too.)

On the other hand, doing nothing is not an option, because the bias we are examining is not a conscious choice. We may have good intentions, but we also know that good intentions are not enough.

We commonly describe bad decision-making as suffering from an absence of evidence. But that's only part of the picture. We can also suffer from using narrow, skewed or incomplete evidence, or from testing too narrow a range of ideas, or ignoring interesting new ideas because we are not looking hard enough for them. And we are at a historical inflection point in how decisions are made; as we trust more to machine learning and AI, we risk systemising this failure in our algorithms.

Research purposefully creates balanced samples to ensure that decision making is based upon data that is less likely to suffer from bias of this type and that the source data is from as wide, and as balanced, a sample of 'inputs' as possible. In short, properly structured research is important as a check against bias.

[1] Bertrand, Marianne and Sendhil Mullainathan (2004), "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination", The American Economic Review 94(4): 991–1013.

[2] Goldin, Claudia and Cecilia Rouse (2000), "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians", American Economic Review 90(4): 715–741.

[3] Dastin, Jeffrey (2018), "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, 10 October.

What is unconscious bias?

There's no formal definition of unconscious bias, but we can say that unconscious biases are states of mind that inform our everyday decisions.

Note that there's nothing here that implies that what we do will be bad, just that they affect our choices. The groundbreaking scientific research on this was done by Amos Tversky and Daniel Kahneman 50 years ago. They first brought knowledge of our flawed decision-making to widespread attention.[4]

"People rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations," their most well-known paper, that draws together many of the behaviours that we now call cognitive biases, tells us, "In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors."

Behind this is what Kahneman has immortalised as "System 1" and "System 2" thinking in his book *Thinking, Fast and Slow*.[5] System 2 decision-making uses time and is considered. But many of our decisions use System 1, which is instinctive, and uses shortcuts. System 1 associates new information with existing patterns or thoughts to reach a conclusion.

In business, we all make hundreds of System 1 decisions every day. Most are unremarkable (or at least pass unremarked). But severe and systematic errors can be devastating, not least because we are likely to be unaware of the problem until it is too late. We usually find out about the destructive nature of these biases by experiencing their impact on us, or on our business.

Research can help us understand the different influences on decision-making, deconstruct the drivers of behaviour and therefore direct organisations on how to drive culture or behavioural change. The challenge to the research sector is to do this in a manner which is applied and commercially savvy, not theoretical.

Unconscious, unintentional, or misguided??

The way in which we describe this problem isn't settled either. Some dislike the use of "unconscious", as it encourages us to think about this is terms of neuroscience, a sort of faulty wiring to be fixed. Kahneman (and many other researchers) dislikes this interpretation, because it's not always useful. It's almost impossible to reliably fix our System 1 errors by tinkering with our heads. But, if we have systems in place to spot when they are misleading or destructive and the processes to correct those missteps, then that is maybe the best we can hope for.

There is also the temptation to categorise all discrimination in this way. Colin Strong, global head of behavioural science at Ipsos points out that we can make similar judgements, with similar outcomes, with varying degrees of consciousness. A recruiter who refuses to interview women because he argues that they would not be able to cope in a crisis is consciously discriminating based on sex, not showing an unintentional bias. A team leader who tends to allocate riskier tasks disproportionately to men is also treating women unequally – but he may be unaware that he is doing it until it is brought to his attention. As a result, the intervention you choose in these two cases may be entirely different.

[4] Tversky, Amos and Daniel Kahneman (1974), "Judgment under Uncertainty: Heuristics and Biases", Science 185(4157): 1124-1131.

[5] Kahneman, Daniel (2011), *Thinking, Fast and Slow*, Farrar, Straus and Giroux.

Market research at its best is at the heart of debiasing decision-making, and opening our minds so we can hear the voices of others.

Personal or organisational?

This, however, is both a personal and organisational challenge to solve. On the personal side, we can decide to reflect more on how we make decisions, or spot the warning signs of the assumptions we do not question, or the biased ways we behave in the work environment, and moderate them as a result. But it is also the duty of the businesses to put in place both guard-rails (clear expectations and policies, explicit values) and to help (offer training when it is useful, incentivise and reward progress) that may provide those cues. The two can be complementary, and the evidence is that one without the other is rarely successful in the long term.

Is the problem ethical or financial?

It is both. When an organisation that acts unfairly to its staff, the fact that it didn't set out to do this is no excuse. When it treats different groups of customers unequally, it cannot complain that its data or its algorithms didn't warn it. There is clearly a legal responsibility not to discriminate, but also an ethical one.

But if this narrative isn't compelling enough, organisations that are slaves to unconscious bias are at a disadvantage. They will fail to seek out new types of customer. They will recruit from a narrow pool of people who "fit in", and miss out on creative or innovative ideas as a result.

How can market research help?

Market and social research at its best is at the heart of debiasing decision-making, and opening our minds so we can hear the voices of others. As a sector we are unique in that we make a conscious effort to try to correct for bias, for example in our sample and questionnaire design, and our statistical rigour.

But we have learned from our mistakes. The data we collect has bias, because some people are harder to reach than others, or are more honest than others, or shout louder than others. The people whose opinions we report have their own unconscious biases. They may be unaware of their true feelings or motivations, or wish to appear more virtuous than they are.

But we are also victims of our biases. Our research sometimes fails to uncover new ideas, because it focuses on the same questions, asked to the same groups. Or we make subjective assumptions about what facts are relevant, or that we think are important, and don't test or communicate those assumptions thoroughly enough.

But, day by day, we improve. The techniques we use, and have been using for many years, to mitigate bias have a wider application. It has never been more important to take a considered look at the sources and outcomes of bias, because as many of our services become automated, so we risk encoding that bias in the algorithms that make so many of the decisions that affect our lives.

The application of derived measures for understanding the choices individuals make using advanced analytics (e.g. regression analysis etc) are incredibly important. The use of big and small data, the combination of transactional or customer data produced by the operational processes of running business alongside attitudinal research data offer us unique and powerful ways to understand the 'what' and the 'why' – in essence the holy grail for organisations in determining strategies to influence future decisions people make whether about purchase, service design or other such requirements.

A practitioner's viewpoint

Bias can be a double-edged sword writes Caroline Hayter, MRS main board member and co-founder and strategist, Acacia Avenue.

We put processes and systems in place to make our lives easier. They become behavioural defaults and their advantage is that they alleviate us from having to expend huge amounts of mental effort for every single decision that needs to be made. So, there's a clear benefit to this.

But of course, there's a downside too. Which is precisely that we don't revisit assumptions or processes put in place, often for years. But the fact is that cultural tides are constantly changing and challenging our established norms yet many business practises shift far less easily.

As a research agency, we see this in the briefs we receive. They often call for creative research approaches to solving a particular commercial issue but when we read the brief in detail, there are language leaks that betray an underlying lack of desire to change at all. There may be phrases along the lines of 'we'd like to view the groups' or 'we assume groups in London, Birmingham, etc'. This is a classic unconscious bias and I suspect that all too often these sorts of deeply baked in biases feature across the board in business.

There are many ways in which we can challenge these biases but they need to be systemic, rather than processes that are bolted on. Here are two that we've embedded at Acacia Avenue:

1. Multi-disciplinary teams – our teams come from the worlds of marketing, communications and research rather than just research. This enables us to interrogate every facet of a project from multiple angles.

2. Multiple strands of enquiry – similarly, looking at a problem from more than one angle also helps, using multiple workstreams against one objective, or in research, more than one strand of enquiry.

While these sorts of practises are conducted in pockets of businesses, they're not nearly common enough.

The other aspect to this double-edged sword is that we don't live in a binary world that's either prone to, or free from bias. Without bias, our worlds would become exceedingly difficult – we wouldn't be able to manage the cognitive load.

Bias is context dependent – it alters based on where we are, what we're doing and who we're with. Most tools that focus on bias, such as the original Harvard IAT, focus on individual response when we know full well that we are social animals and that we're influenced by broader cultural expectations and themes. There is no categorical diagnosis for bias – it's more elastic than that.

We need to recognise that bias isn't always the assumed negative that its name implies, and that sometimes it's acceptable, even helpful. As business consultants, we're paid for our points of view, not for our objectivity. This means being able to identify how to interpret bias in all its facets, recognise where and when it needs to be addressed and make brave recommendations that may fly in the face of popular opinion.

With the demand for agility, rapid-response and laser-sharp thinking in business today, the real challenge is how to ensure that our thinking in getting there isn't so laser-focused that we miss what's really going on.

1

Cultural bias

Aviva's internal pension holders who identified as female used to find that their pension documentation defaulted to talk about the benefits to "your husband". All male pensions referred to the benefits to "your wife". Efforts on behalf of Aviva Pride, the internal LGBT advocacy group, to get it changed would be pushed to the back of the queue, as the IT department had other priorities. "The difficulty we had: it doesn't affect many people," recalls Jan Gooding, who from January 2016 to August 2018 was the global inclusion director for the Aviva Group.

Whoever created the system was unlikely to have been homophobic, just careless. But once the system was in place, it persisted.

All organisations have a unifying culture, some elements of which are explicit (our contracts, regulation), and some are implicit. But Gooding's priority wasn't what the implicit part of that culture creates, it was what it inadvertently prevents. "That's where I was able to harness Aviva's interest. It was less about equal opportunities from an ethical point of view, and more commercial: the productivity and happiness of our workforce because they were reaching their potential, and then the ability to do more for our customers, because we were more aware of them."

Unconscious bias, if unchecked, can become a part of that implicit culture. As Gooding points out, this has an impact both internally and externally.

Recruitment

One internal impact is that people are not recruited. The dominant culture is unchallenged because we choose people who are like us. The potential to create groupthink through recruitment and the desire to "fit in". The external impact is that it is difficult to create products or services for groups outside your traditional focus.

Behaviour at work

The clearest example of this is the way in which we treat co-workers. There is plenty of evidence that women, especially, are treated differently. They tend to be interrupted more often, or their opinions given less weight. This is not just about the habits of the speaker, it's about how those speaker differentiate their audience. For example, one experiment[6] found that when randomly placed in pairs for a discussion, for example, women were interrupted more often. In a three-minute conversation, men interrupted men twice, and women 2.6 times. But women, on average, interrupted men once, and other women 2.8 times.

Cognitive diversity

The work of Andrew Tenzer at Reach (see Focus on: Cognitive Diversity p.14) is just one example of how we may unwittingly foster groupthink. Tenzer and Gooding both point out that, when we say we want cognitive diversity, we often prefer people who think differently, but in ways we find amusing or valuable.

[6] Hancock, Adrienne and Benjamin Rubin (2014), "Influence of Communication Partner's Gender on Language", Journal of Language and Social Psychology 34(1): 46–64.

What can we do to offset these biases? As researchers, we would advocate setting up objective measurement when possible. Recruitment is a good example, for which standardising sets of questions or having someone unconnected with the recruitment process to audit the hiring process helps. Iris Bohnet, a professor at Harvard University, advocates using blind CVs, structured interviews, or data-driven evaluations to measure hiring.

Jake Steadman, senior director of insight and analytics at Twitter, points out the important of encouraging a broad mix of recruits at entry level as a way to create a diverse, inclusive workplace, especially in teams that are doing research or creating the insights that guide the organisation's decision-making. "Some of the basic ways in which we tend to screen applicants can narrow the pool of thought," he warns. An example is that market-leading organisations may decide to recruit their insight teams only from highly-ranked (and perhaps less diverse) universities. At Twitter, Steadman has deliberately looked for talent outside narrow academic parameters.

An example is the approach adopted by Goldman Sachs, which in 2016 stopped doing on-campus interviews, based only in a narrow set of universities, for its first round. Candidates from any universities can now use a pre-recorded interview application called HireVue instead to make a first impression remotely. In the second round, Golman has adopted a structured interview format.

Russell Horwitz, co-chief operating officer of the securities division, said the previous system, in which recruiters went back to the universities from which most of them had graduated, and recruited "based on instinct and feel ... does not produce the best outcome."[7]

Richard Thaler, one of the pioneers of researching our unconscious biases in decision-making who won the 2017 Nobel prize for economics, goes further: he suggests recruiting by giving people tests. Traditional interviews are "useless" he argues, as a process of unbiased evaluation. "You wouldn't hire a race-car driver by giving them an interview," he says, " We'd put them in a car or, better yet ... behind a video game."[8]

But countering day-to-day unconscious bias in the workplace can be more problematic. The standard approach is to engage employees in a session of unconscious bias training. Studies show that isolated training sessions can have little effect, or even a negative effect. If they are compulsory, they may be resented (former Google employee James Damore wrote his famous memo criticising diversity programs that went viral after one of these bias training sessions). Another unintended consequence is that showing employees that everyone is unconsciously biased serves to unconsciously legitimise bias. An example: an experiment in which subjects underwent training to discourage age-related bias increased System 1-based age discrimination afterwards.[9]

[7] Noonan, Laura (2016), "Goldman Sachs to drop on-campus interviews", Financial Times, 23 June.

[8] Javetski, Bill and Tim Koller (2018), "Debiasing the corporation", McKinsey Quarterly, May.

[9] Kulik, Carol, Elissa Perry, Anne Bourhis (2000), "Ironic evaluation processes: Effects of thought suppression on evaluations of older job applicants", Organizational Behavior 21(6): 689–711.

Research shows that effective training isn't a one-off. Unsurprisingly, the evidence is that you can't "fix" unconscious bias with a single training session any more than you can change a shopper's habits with a single advertisement. You may open their minds to the ways in which they can change their behaviour, but as Kahneman pointed out, expecting to rewire the brain is probably too ambitious. Our values change slowly, if at all.

Instead, it has been successful to target our habits and associations, which may begin to moderate our beliefs. An example of this is what Twitter has put in place to reduce bias in its culture. Candi Castleberry Singleton is the Vice President of Intersectionality, Culture, and Diversity, has supported the creation of many internal groups to give platforms to diverse groups and interests, and all employees are encouraged to join in their activities.

Research teaches us that the way in which a proposition is framed can have a large effect on how well people respond to it.

The aim is to go beyond diversity to achieve inclusion, so that the diverse groups are valued, and their points of view contribute to the success of the business. "We do have targets for inclusion, but what is more important is the approach we take to foster this, and that everyone in our staff feels included," says Steadman. He adds that, for an organisation like Twitter to hold the behaviour of its users to account, it also has to hold itself to account, and so part of the drive to create an organisation in which everyone feels included will be to publish statistics on its progress to these goals in meeting these targets.

Research teaches us that the way in which a proposition is framed can have a large effect on how well people respond to it. Tinna Nielsen and Lisa Kepinski, two experienced trainers in unconscious bias, point out that a session called "Unconscious bias awareness training" or "Inclusion and diversity training" tends to trigger unconscious associations for employees that may undermine the training.[10] They may assume that they are going to be "fixed", or accused, or lose status as a result.

"Instead, why not motivate by designing a 'title' focusing on the 'meaningful destination or outcome', such as better performance, innovation, engagement," they say, rather than focusing on the means.

Sensitising the people in your organisation to the effects of unconscious bias and then continuously encouraging employees to change the way they think and act, by examining their own assumptions and seeking out people who belong to groups unlike your own, is more likely to be successful.

[10] Nielsen, Tinna C and Lisa Kepinski (2016), "Unconscious bias training is hot, but the outcome is not. So what to do about it?" *Medium.com*, 23 May.

Recommendations:

How do we identify cultural bias?

- Be positive about diversity and inclusion– create a culture of celebration and encourage everyone to get involved.
- Inclusion is your goal, and that requires action. Listen to all voices to create an environment that doesn't marginalise particular groups, just because they are less numerous or powerful.
- Checking that your research sample is representative ensures that inclusivity is taken into account from a 'data input' point of view.

Should we change the way we recruit?

- If your recruitment focuses on a particular set of universities or degree types, you are recruiting from an arbitrarily narrow group.
- Consider practical tests and anonymised CVs rather than traditional selection.
- Involve people from outside the department of organisation in the interview process.

How should we approach anti-bias training?

- Avoid a box-ticking exercises: compulsory and one-off sessions are much less effective than tailored, long-term programmes.
- Measure your progress. Set goals, survey employees before and after (but beware of social desirability bias, the desire to be seen as a good person by others).
- Frame the training by describing its positive outcome, not the problem.
- Consider training leaders in moderation skills – see MRS training.

Focus on: Cognitive diversity

"People experience the world, and interpret it, differently," says Andrew Tenzer, head of group insight at Reach plc, formerly Trinity Mirror. "Because people who work in advertising are young and affluent and live in a metropolitan city, do they have a different thinking style? This matters, because or thinking style drives behaviour at an unconscious level."

As head of insight for a company that has Tenzer spends a lot of his time working with advertising agencies, especially media buyers. But as a researcher with a particular interest in behavioural economics, he was fascinated by whether the young people at these agencies thought about the world in the same way as the group defined as the modern mainstream – that is, the middle 50%, who represent more than 50% of brand buyers.

Tenzer, and Ian Murray, partner and co-founder of agency House51, did this using a nationally representative sample of 2,415 people, compared to a sample of 150 people drawn from London-based media agencies.[15] The results were startling. They found that people who work in media agencies thought differently about concepts like fairness. Mapping their responses to a framework of culture and values, and "they massively overestimate power and hedonism, and massively underestimate some of the more positive aspects of the way we live. You can see that people working in media agencies and people in the modern mainstream have different value systems," Tenzer says.

But, even more startling, they found that the two groups also perceive the world differently. They reran some experiments first done by Richard E Nisbett, a psychologist whose experiments led to the groundbreaking discovery that people from different parts of the world think – and even see – the world in different ways.[16] Nesbitt focused on the difference between eastern cultures that emphasise culture and relationships, and western cultures that were more analytics, "straight line" thinkers. In this case, the modern mainstream had far more of a community focus, avoided risks and strong emotions, and had less need for belonging than people in advertising.

"Brands and advertising have lost relevance with large swathes of the UK".

This means that, as a group, advertising people are more interested in new things, over-emphasise brand relationship. "Brands and advertising have lost relevance with large swathes of the UK," the report concludes.

Tenzer draws a lesson too for diversity: that while we may be successful in recruiting or assembling teams that look diverse, this still limits diversity if we instinctively recruit people whose worldview we feel comfortable with. "The diversity that doesn't get spoken about nearly enough is cognitive diversity," he says, "All the research suggests that when we recruit, we recruit people who are culturally quite similar to us. I just wonder if the non-conformists in our industry are really as non-conformist as we think."

[15] Tenzer, Andrew and Ian Murray (2018), "Why we shouldn't trust our gut instinct", available from Reach plc (https://www.reachplc.com).

[16] Nisbett, Richard (2005), The Geography of Thought: How Asians and Westerners Think Differently – And Why, Nicholas Brealey Publishing.

2

Management bias

While cultural bias affects how we respond to, and treat, other individuals or groups, for people in positions of leadership our unconscious biases can also show up as flawed decision-making.

For 70 years the members of the MRS have advised businesses on how to make better decisions by using the best available evidence. But many decisions made day-to-day are based on what we call "gut feel". Our gut is where unconscious bias lives.

Using statistical techniques to mitigate bias in the data (for example, by trying to ensure that the sample represents the population from which it was taken) is the bread and butter of quantitative market research. But the challenges of political polling (see Focus on: Polling p.22) show this is a journey, not a destination. This also applies to qualitative research. We all bring our biases with us, but we can do our best to be transparent about the judgements we make when doing research.

Our gut is where unconscious bias lives.

There is now a long list of cognitive biases, some of which are more common (or perhaps more realistic) than others. Many are familiar:

Confirmation bias

As the head of insight at one multinational technology company we spoke to told us: "Decisions are made and then data is presented to reinforce them. If we uncover data that contradicts the prior point of view, its validity will be challenged." Research generated to support a point of view has bias baked into it.

False-consensus bias

In groups, we tend to overestimate the degree of consensus that exists, or how much others share our views. Research is well-placed to challenge this, but only if we look beyond the group, or search for sources of disagreement or conflict, either inside or outside the organisation. Alfred Sloan, the chief executive of General Motors from 1923 to 1946 and one of the pioneers of market research as a way to improve decision-making, can inspire us in this. At the end of a meeting he once told the group "I take it we are all in complete agreement ... I propose we postpone further discussion of this matter until the next meeting to give us time to develop disagreement."

Hindsight bias

"We're all geniuses after the fact", Thaler points out. He encourages the people he works with to write down their decisions when they make them, agreeing on the goals, and the assumptions behind the decision. This disentangles good decision-making from luck, because these assumptions can be researched. The decision may still go well or badly, but it avoids the temptation to rationalise a lucky gamble.

The list is long: hyperbolic discounting, the sunk cost fallacy, insensitivity to sample size, availability bias. While we will never truly eliminate them from decision-making, Kahneman's prescription is the same: avoid System 1 thinking when it is practical.

"We would all like to have a warning bell that rings loudly whenever we are about to make a serious error, but no such bell is available," Kahneman writes in Thinking, Fast and Slow. The best we can hope for, he argues, is to have cues, which prompt us to pause and engage System 2. Research, whether by introducing statistical evidence, or new voices, or creative ideas, can be a way to engage that thinking.

Recommendations:

Can we moderate gut feel?

— Make important indicators continuously available, for example as dashboards, across the a business – ensure these dashboards include a wide range of sources.

— Use quick bursts of research in decision-making.

How do we encourage cognitive diversity?

— Formally encourage challenges to ideas to avoid confirmation bias. For example, Psychologist Gary Klein advocates conducting a pre-mortem before you begin. This assumes your project has failed: encourage the team to discuss why that would have happened.

How do we ensure our research is representative?

— Always allocate time to challenge methodologies, especially when the result is the one you expected or wanted and encourage researchers to explain them.

— Check whether the results you achieved were what research predicted, and use this to learn and improve.

Focus on: The IAT

One of the most widely-used, but most controversial, strategies for organisations to reduce unconscious bias is the implicit association test (IAT), developed at Harvard University 20 years ago. One of the reasons you might be familiar with it is that it is so easy to experiment with: chances are that you have either seen a link, or been sent one, to take the test.[11]

The IAT attempts to access our personal System 1 shortcuts for race, gender, sexual orientation, and other possible sources of unintentional bias by timing the delay with which we associate positive and negative attributes to particular groups when they flash up on a computer screen.

The theory is that the instinctive response overcomes social desirability bias which is the tendency to respond in ways that make us look good to others. Researchers have dealt with it for many years – for example, in research on charity donations, or drinking habits. And of course, few staff would choose to focus on feelings of racism or sexism, especially in front of co-workers.

The creators of the test have in the past made strong claims for the IAT's usefulness. In 2013 they claimed that it:

"predicts discriminatory behaviour even among research participants who earnestly (and, we believe, honestly) espouse egalitarian beliefs."[12]

This claim amplified via thousands of Facebook shares and mandatory training sessions, has stuck. The IAT is clearly is potentially a very powerful source of insight into ourselves, which can be the foundation of bias training within an organisation. But it also comes with some health warnings.

To be valuable, any intervention based around the IAT should have a clear goal that we can reasonably expect it to help to achieve. The second is that the IAT has to give reliable insight. The third is that this insight has to be actionable.

A clear goal. It can be eye-opening to consider that we have implicit biases. Used as part of a programme to help us mitigate their effect, the IAT may be one element that can help us reflect on how we see the world. There is evidence, however, that the IAT's ability to predict how we actually behave is extremely weak[13], and it is behaviour that we aim to affect. This is not a reason by itself to ignore IAT studies, but to treat them at best only one element training that seeks to reduce the impact of unconscious bias.

Reliable insight. This is more problematic. Many studies have shown that the results are not stable. That is, our score on the test is affected by all sorts of contextual factors, even the time of day. Whether those results remain private or are shared with an employer, research into, say, preference for packaging design that was as unstable as the IAT would rarely be considered useful.

Is it actionable? Therefore, an individual's IAT score isn't a reliable indicator either of that person's unconscious bias, or that person's likely actions. There is also the problem that being confronted with our IAT encourages us to believe that bias is purely a personal failing, a problem to solve. Many experienced trainers emphasise that improving company structures and values are often more successful in challenging the cultural biases of staff, especially if those habitual responses have been incentivised or encouraged in the workplace.

[11] If you haven't tried it, and you're curious, Project Implicit is here: https://implicit.harvard.edu

[12] Banaji, Mahzarin and Anthony Greenwald (2013), Blindspot: Hidden Biases of Good People, Delacorte Press.

[13] Oswald, Frederick L, Gregory Mitchell, Hart Blanton, James Jaccard, Philip Tetlock (2103). "Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies", Journal of Personality and Social Psychology 105(2): 171–192.

A practitioner's viewpoint

The research sector should take on a stronger role in tackling bias writes Alison Camps FRSA, Deputy Chairman, Quadrangle.

In my view, the reason bias is still a problem in the marketing services sector is not just because it is unconscious.

It's because it is institutionalised.

I have 30 years' experience under my belt and am the Deputy Chairman of one of the UK's leading independent insight agencies, but I *still* find myself contending with men who think it perfectly fine to interrupt and talk over me every time I open my mouth.

Earlier this year, one of my friends who is a gay British Asian man working in strategic communications was so frustrated by the lack of diversity he saw around him that he set up his own agency, The Unmistakables, to help clients create better workplaces and campaigns for minorities. He is being the change he wants to see happen in the world and I salute him for it.

I think there's a lesson here for us.

One of the things I love most about research is that it can reveal and challenge bias, by giving voice to the views of those who are often overlooked by brand marketers and decision-makers, too many of whom – let's be honest – are still too white, too male and too straight.

But for it to do that, we must design it into the process. Because it seems to me that for all our meticulously specified Nat Rep samples, our conscientious boosts amongst minority groups and so on, there is *still* a tendency in our industry to design research which reflects our own 'lived experience' (or that of our client) more than the realities of those whose views and opinions we seek to understand.

Nowhere is this more evident than in the language we use. Language and culture are intertwined so it stands to reason that if we want to understand a group of people, we need to use language that is meaningful and relevant to them. We should literally speak their language.

And we need *deliberately* to give greater visibility to minority or under-represented voices (which can often include women, by the way, who last time I looked represent over 50% of the population). We need to ask the awkward questions of our clients, draw attention to it in our proposals and make it a focus in our analysis and reporting including the way we use imagery in deliverables.

We're all on a journey here. If we want to rid ourselves of unconscious bias I think we need to achieve a state of conscious competence.

In my experience, big brands are already doing great work to address the issue internally. As a sector, I'd hope that we have a big role to play in helping our clients do the same *externally*.

3

Data bias

In 2015, T J Fitzpatrick, who is black, tried to use the automatic soap dispenser in the public bathroom of the Atlanta Marriott hotel. If didn't work. When his friend Larry tried it, it did. They filmed the phenomenon, and the "racist soap dispenser" became a YouTube hit.

Automatic soap dispensers work by detecting reflected light from our hands. In this case, Fitzpatrick's dark skin didn't reflect enough light to trigger the dispenser.

It's not the only example of bias in product design. Currently the most obvious is in facial recognition software that uses machine learning to train its algorithm. At the beginning of 2018, Joy Buolamwini, a researcher at the MIT Media Lab, built a dataset of 1,270 faces, using the faces of politicians, and found that gender was identified correctly in more that 99% of light-skinned males, and 93% of light-skinned females. But more than a third of dark-skinned females were misidentified. Buolamwini now leads a project called the Algorithmic Justice League to fight bias in machine learning.[17]

"When there is bias in data, we can end up making some very bad decisions," says Tom Evans, director of data science at Kantar.

"When there is bias in data, we can end up making some very bad decisions."

What is behind this? Increasingly, algorithms – automated sets of decision rules – are part of the products and services we create. Broadly they split into two categories.

The first are deterministic, in which the data scientist uses existing data, for example transactional data, customer experience, research into preferences, economic trends, and so on, and creates a model based on how an outcome (for example, propensity to purchase) depends on a set of variables. The data sets can now be as large as the population they are applied to.

The second take a limited set of data points, and use it to train an artificial intelligence, which then applies its learning to a larger data set, and constantly refines its results based on its experience of the world.

But both suffer from bias. The reason is that unintentionally biased people made them.

A product built from data provided by a limited, unintentionally similar, group of people will be optimised for that group. It doesn't matter whether that product is a soap dispenser, a financial service, or a television show. We may assume that large datasets will remove any bias, but this is not so, warns Harry Davies, head of marketing measurement and attribution for the UK at Google:

"Some people have lost the skills set and the knowledge we had when we didn't have so much data, which is to look for a representative sample. But even in our largest datasets there is still a lot of bias. Not everybody is walking into your store, for example, and so you end up viewing people who are doing things you can see, not the other people you want to reach."

[17] Read about the Algorithmic Justice League here: https://www.ajlunited.org/, or watch Buolamwini's Ted talk on the topic: https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms.

This has two unwelcome outcomes. When using biased data, your conclusions will be in error. But also, the data might not even investigate important preferences or attributes, because the people who collected it didn't consider them to be important.

Machine learning offers a similar, but potentially much harder to spot, problem. The algorithm may learn based on inadequate training data, and then magnify that problem. This was the case in the earlier example of Amazon's failed recruitment AI, which "learned" to discriminate against women, or facial recognition systems, for which the training data can often be a very narrow sample – for example, taken from the researchers who created it. When we consider that, currently, eight out of nine AI developers is male for example[18], this can create serious problem.

It is not easy to fix these problems, but it is possible. Evans explains that data science should not be about throwing data at a problem to see what correlates: he recommends that good practice means having a clear idea of the business problem and the regulatory environment in which the algorithms exists. If discriminating on a characteristic is undesirable or illegal, then building a deterministic algorithm without that data will make it less likely that it will be biased (it is far from impossible, because some unregulated data, for example postcode, may be correlated with attributes such as race).

But an AI is, largely, a black box: it cannot explain why it does what it does. Algorithmic bias is currently the focus of regulatory discussions in many countries (not least because they are now making important decisions in regulated industries, or that affect protected categories, and so need to be accountable). We argue that, if you wait for regulation to solve your bias problem, this may be inadequate for the practical reason that customers and citizens value fairness highly.

There are, in their early stages, "fairness tools" which can be used to test the attributes that AIs are using to make judgements. This relies on a carefully thinking through and communicating what "fairness" will be. This is a subjective judgement, but one that organisations cannot avoid. Rumman Chowdhury, Accenture's global responsible AI lead, explains that her company's tool looks at the error rates for each variable that it uses to make a decision, and whether errors are higher for one group rather than another.[19] Adjusting the parameters to equalise errors makes the algorithm less effective, and so businesses that will use it will become less biased, but potentially less profitable. Where this trade-off is made will be an important consideration for decision-makers, and one that should involve the voice of the customer.

"This is nothing new, " Evans adds, "we have been using algorithms for years to make decisions, for example on insurance pricing. But data bias has meant that we are making decisions that, we are beginning to see, are pretty scary."

[18] Mantha, Yoan (2018), "Estimating the Gender Ratio of AI Researchers Around the World," Medium.com, 17 August.

[19] Peters, Adele (2018), "This tool lets you see–and correct–the bias in an algorithm", Fast Company, 12 June.

Recommendations:

Is my data biased?

— Never assume that just because a data set is large, it is representative of the population. Demand sense checks – how are minority groups represented in the data?

— Be cautious when using freely available data sources – they may not be representative enough to ensure a sound data set.

Are my algorithms fair?

— Using data for decision-making needs a hypothesis. Build a model that makes sense to the business and test it, rather than fish around in the data for correlations.

— The law prevents discrimination against some groups, but consider an ethical framework that goes beyond regulation, and communicate this to all stakeholders.

— Test your outcomes: formally, the emerging category of 'fairness tools', for example from Accenture may help reveal discrimination, but allocate time to investigate both whether there is clear discrimination and whether your customers perceive bias.

Focus on: Polling

In the 2015 general election, the final estimates of polling companies had the Conservatives and Labour neck-and-neck on 34%. The Conservatives won 38% of the vote, Labour 31%. In the 2016 Brexit referendum, polls famously predicted remain. In 2017, some polls one day before voting were predicting a double-digit lead for the Conservative party. In the event, the Conservatives remained the biggest party, but lost their overall majority as voters returned a hung parliament.

No group of professionals puts such effort into creating unbiased samples and accurate predictions as pollsters. What can go wrong, and what can it teach us about the impossibility of completely escaping from bias?

We tend not to notice bias until we get things wrong. This, it turns out, is also a cognitive malfunction known as outcome bias: if the result of our is what we predicted, we don't look hard at the methodology. The MRS and British Polling Council produced a joint report into what happened in 2015[14], which pointed out that "the polls have been nearly as inaccurate in other elections but have not attracted as much attention because they correctly indicated the winning party." So looking back so say "why didn't we see it coming can inform us about our blind spots.

We tend not to notice bias until we get things wrong.

Unrepresentative samples

"Even careful data collection contains bias", the report concluded. In 2015, the mistake was to not give enough weight to people who didn't want to participate in research, or couldn't participate because they did not use mobile phones. The problem is that any adjustments to do are informed guesses. Being transparent about how you interpret raw data encourages constructive debate.

Confirmation bias even affects experts

When making adjustments or deciding which results are outliers, says Anthony Wells, director of YouGov's political and social opinion polling, it's natural to give more credibility to those that are closer to your prior belief of the truth. "We don't try to make those decisions but we can't help it. In 2016, everyone expected the remain vote to win. In 2017, everyone expected Labour to be hammered," he recalls. Seek out other opinions. They will often be wrong, but give them attention.

People lie, even if they don't mean to: "People are very poor at explaining their own behaviour," says Joe Twyman, co-founder of Deltapoll, "People may agree with something on a logical level, but the heart says something different." Look for emotional engagement or actions, rather than rationalisations.

[14] Sturgis, P, N Baker, M Callegaro, S Fisher, J Green, W Jennings, J Kuha, B Lauderdale, and P Smith (2016), Report of the Inquiry into the 2015 British general election opinion polls, Market Research Society and British Polling Council.

Download all Delphi reports here: mrs.org.uk/reports

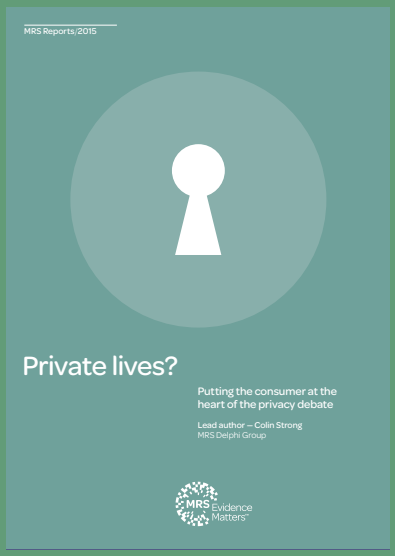

Private Lives?

A look at privacy issues through the lens of the consumer.

The Politics of Persuasion

What influences voters and how to improve democratic engagement.

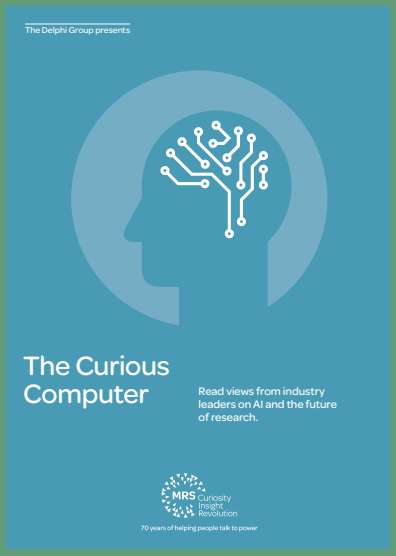

The Curious Computer

The impact of AI on the research and insight sector.

Towards an Insight Driven Organisation

The people, skills and processes that enable insight to drive business growth.

Prediction and planning in an uncertain world

How insight and research can help organisations connect with the future.

Great Expectations

How technology impacts consumer trust.