

JAMES  
BALL

POST-  
TRUTH

HOW BULLSHIT  
CONQUERED  
THE WORLD

\B<sup>b</sup>\

Biteback Publishing

# INTRODUCTION

The US government stockpiled 30,000 guillotines, stored in internment camps – including one in Alaska large enough for two million people – ready to wipe out Second Amendment supporters at a rate of three million an hour.<sup>1</sup> Trump supporters at a New York victory rally chanted, ‘We hate Muslims, we hate blacks, we want our great country back’.<sup>2</sup> Denzel Washington endorsed Donald Trump<sup>3</sup> – and Trump actually won the popular vote in the US election, despite the mainstream media telling you otherwise.<sup>4</sup>

Everything above is grabby, easy to understand, easy to share – and false. All serve as examples of the long-existing but newly discussed phenomenon of outright fake news: easily shareable and discussable stories, posted to social media for jokes, for ideology, for political reasons by groups connected to foreign nations, such as Russia, or – most commonly – to make a bit of money.

These examples are classics of the genre: possible to invent in minutes, but taking hours to debunk. Even the most obvious nonsense claim takes time and effort to prove false. Take the internment camp supposedly prepared to jail two million opponents of Hillary Clinton. If the world’s largest manufacturing building – a Boeing aircraft factory in Washington<sup>5</sup> – had the entirety of its internal space converted to confinement cells, it could only house one million people, and that’s without any corridor space, kitchens, room for security or anything else. An actual site for two million would need to be three or four times larger, constructed entirely in secret, and somehow hidden from any kind of passer-by, whether by land or air. But none of this matters to someone already convinced. What actual proof do I have the site doesn’t exist? Of course the answer is none.

For the determined debunker, just battling outright and obvious falsehoods, from anonymous blogs and hoax sites, would be a losing battle. But there’s a far wider problem than these actual hoaxes – the

whole range of stories that are essentially untrue, but arguable to people who believe them or can convincingly pretend to.

The UK's debate over whether to leave the European Union – the Brexit debate – was littered with such claims. The UK pays £350 million a week to the EU, and voting to leave would mean this money could be given to the NHS. The then Chancellor, George Osborne, would raise income taxes by 2p should the UK vote to leave. Voting to stay in the EU would open the UK to uncontrolled immigration from Turkey, from where twelve million people plan to migrate.<sup>6</sup>

These claims are, to most who dig into them, just as false as the first group, but with two main differences. The first is that there's enough core of truth to each to make them essentially arguable: the short version of the claim put on a leaflet may be an outright lie, but once they drilled down into the detail, two politicians arguing in the media could run the argument to a draw. The second is that these claims aren't made by anonymous figures – they're made by the politicians and the staffers at the centre of the rival campaigns.

Needless to say, this is a problem the US has had plenty of time to grow familiar with: Donald Trump can generate more political nonsense in an hour than most of his rivals can produce in a year. Trump's versatility in generating half-truth, untruth and outright spectacular mendacity borders on genius.

The subjects range from the trivial to matters of major national policy, and no statement is bound by anything that came before it. Take Trump's evolution on his flagship policy of building a border wall, which Mexico would pay for – the country's statement that it would do nothing of the sort was easily ignored. After his election, Trump acknowledged he'd be going to Congress for money for his wall – yet still insisted Mexico would pay for it.

Trump has accused the media of lying about the crowd size for his inauguration by quoting it at 250,000 rather than the 1,000,000+ he claims; pollsters of fabricating his low approval numbers; the CIA of fabricating evidence that Russia intervened in the election in his favour; and unknown authorities of allowing millions of fraudulent votes to be cast in an election he nonetheless won. Trump can even comfortably and casually lie about incidents captured on video: when caught during the campaign imitating the disability of *New York Times* reporter Serge

Kovaleski,<sup>7</sup> Trump routinely states the incident – which happened at the front of a televised rally – didn't happen, or that he'd never met the reporter concerned (he had, repeatedly).

In markedly different ways, the world was reshaped in 2016 by two contests typified by anger against elites, a breakdown in trust in the media, widespread (and wrong) belief among pundits that the contests were foregone conclusions – and the routine use of what for the rest of this book we're going to call bullshit.

Britain's vote to leave the EU ends a relationship of more than forty years between the UK and the world's largest trading bloc. It will involve reforging the country's security partnerships and trading relationships with new and existing partners, and will leave the EU reassessing its own future.

The US election outcome is, if anything, even more significant for the world. At the start of his term, Donald Trump said that he'll try to reshape the country's healthcare system, redefine its relationship with Russia and with NATO, consider ripping up its trade deals, change the USA's long-standing China policy, end Obama's climate change measures and deport far more 'illegal immigrants' than his predecessor.

The consequences of each vote could hardly be more serious, and yet the campaigns that decided them – and masses of the media coverage – were based on trivia, half-truths and lies. It would be a gross oversimplification to claim that either electorate was tricked into their vote, but nor can we rule out that bullshit swung votes, especially as both were relatively close: had just 55,000 voters (out of more than 130 million nationwide) in three states voted differently, Hillary Clinton would be President.<sup>8</sup> Fake news stories alone – leaving out poor-quality information, biased coverage or mainstream media repeating dubious Trump claims – reached orders of magnitude more people than that.

The Brexit vote is less clear-cut, as it was less close: Leave won by a margin of 52 per cent to 48 per cent, or about 1.3 million votes. But analysis of who actually voted shows that the crucial margin of victory came from left-behind, low-income people who don't usually turn up to vote.<sup>9</sup> The question of what urged this group to vote in the referendum when they stayed at home in the previous year's general election remains an open one, but in a contest where one side offered complex economic forecasts for 2030 and the other gave clear-cut messages on handing

money to EU bureaucrats versus the NHS, messaging is an obvious possibility.

Had bullshit been confined to the fringes, to fake news sites which didn't reach significant audiences, such questions could be ignored – but what happened in reality was that mainstream coverage became dominated by repeating and regurgitating claims which were often entirely untrue. The long-standing media habit of leaving campaigns to duke it out over who was telling the truth worked in favour of the liar: make a claim, have it echoed in print, on TV and online, and then get further coverage as the rival campaign challenges its truth.

What effort the major outlets make towards challenging the truth of political claims tends to be confined to specialist fact-check columns, or dedicated political shows – rather than leading the main news broadcasts, or shorter mainstream radio bulletins. The result is – superficially at least – bullshit works: if challenged, it provokes a story about the row that repeats the claim for days at a time; if unchallenged, the claim seems unanswerable.

Before we go further, it's worth explaining why this book talks about 'bullshit' rather than lies or untruth or some other term. One reason is simply that we need a catch-all word to cover misrepresentation, half-truths and outrageous lies alike. The other stems from the Princeton University philosopher Harry Frankfurt, who wrote a full book defining the term in 2005: *On Bullshit*.

Frankfurt's argument, roughly speaking, runs as such: to tell a lie, you need to care about some form of absolute truth or falsehood, and increasingly public life is run by people who don't care much either way – they care about their narrative.

'One who is concerned to report or conceal the facts assumes that there are indeed facts that are in some way both determinate and knowable,' he argues. 'His interest in telling the truth or in lying presupposes that there is a difference between getting things wrong and getting things right, and that it is at least occasionally possible to tell the difference.'

If someone rejects that idea then there are two options: to never again claim anything as fact, or to bullshit – say things are so, but with no recourse to reality.

Frankfurt concludes:

Someone who lies and someone who tells the truth are playing on opposite sides, so to speak, in the same game. Each responds to the facts as he understands them, although the response of the one is guided by the authority of the truth, while the response of the other defies that authority, and refuses to meet its demands.

The bullshitter ignores these demands altogether. He does not reject the authority of the truth, as the liar does, and oppose himself to it. He pays no attention to it at all. By virtue of this, bullshit is a greater enemy of the truth than lies are.

In other words, a bullshitter will say what works to get the outcome they want, and care little whether it's true or not. To many (this author included), this serves as a relatively fair description of many modern political campaigns, and its effect seems to be as damaging as Frankfurt's philosophical text would surmise.

The resultant mass-produced bullshit is too much even for the earnest media outlets who try to report fairly and accurately to attempt to deal with: their culture and norms simply cannot keep up with the onslaught, especially given their bone-deep habit of trying to give a hearing to both sides of a political argument. When it comes to dealing with bullshitters, the mainstream media may be bringing a knife to a gun fight.

But there are plenty of large outlets making no such effort. Many could easily be accused of being part of the bullshit machine themselves, some clearly intentionally. Just as outlets rage at fake news, stories on a wide range of issues are routinely angled to suit the prejudices of the audience – in the UK, right-wing tabloids have been made by regulators to apologise time and again for distorted reporting about Muslims,<sup>10</sup> refugees and immigrants.

Some outlets run front pages which it's almost impossible to believe they could ever think were true. In the run-up to the EU referendum, the *Daily Mail* ran a front page showing a lorry full of people smuggling themselves into the UK. 'WE'RE FROM EUROPE – LET US IN', the headline said. As the *Daily Mail* would surely be aware, EU citizens have a right to live and work in the UK and would have no need to travel illicitly into the country – the person quoted had in fact said 'We're from Iraq'.<sup>11</sup>

Some newspapers will go still further on their websites, to pick up any kind of traffic. The *Daily Express* – which sells 400,000 copies and has 1.5 million unique browsers a day – routinely runs online headlines like 'Chemtrails "will wipe out humans" causing biblical-style floods, says

expert’,<sup>12</sup> referring to a widely discredited conspiracy theory that planes leave behind chemicals designed to keep populations docile.

Such a culture is hardly confined to the UK. US supermarket tabloids have run front pages including ‘HILLARY: 6 MONTHS TO LIVE!’, ‘HILLARY FAILED SECRET FBI LIE DETECTOR!’ and ‘HILLARY HITMAN TELLS ALL!’<sup>13</sup>

Concerns about media accuracy are hardly new: relatives of the ninety-six football fans killed in the crush at Hillsborough stadium in 1989 were faced in the immediate aftermath by a front page in *The Sun* falsely stating – based on untrue accounts from police – that fans had attacked police, robbed the wounded and dying, and had urinated on police officers.

So it goes too with fake news online – a phenomenon the journalist John Diamond had spotted twenty-one years before it hit the mainstream. ‘The real problem with the internet is that everything written on it is true,’ he wrote in 1995.<sup>14</sup> ‘Or rather, there is no real way of discerning truth from lies. The net is a repository of facts, statistics, data: unless anything is palpably wrong, we tend to give all facts on our computer screens equal weight.’

What broke in 1995 has not been fixed in 2017. If a site has a plausible name and a design which looks roughly like a mainstream news site, we tend to believe it – one now-defunct site called the Boston Tribune, a plausible newspaper name for a non-existent paper, ran articles claiming Obama had bought a retirement home in the Middle East, had given his mother-in-law a lifetime pension for babysitting, and that an elderly man had been arrested for shooting a man who was attempting to abduct a seven-year-old child.

What Diamond, who died in 2001, could not have predicted was how the effect he already saw (that everything looks equally credible online) would be compounded by what’s routinely called the ‘filter bubble’. In short, we tend to click on things that suit what we already think, and we’ll rarely try to fact-check a story that suits our preconceptions. In other words, a liberal will likely Google for a fact-check of a claim that Obama was born in Kenya (he wasn’t), but is much less likely to do the same for a claim that Trump once called Republicans ‘the dumbest group of voters in the country’ (he didn’t).<sup>15</sup>

Given that most of us are friends with people with a broadly similar worldview to ours, we see more and more unchecked news we're predisposed to agree with. The result? Where once right- and left-wing partisans disagreed over their interpretations of a roughly shared narrative, now a portion of each side see different, polarised and largely untrue narratives about the other – and each thinks the other is uniquely afflicted by 'fake news': those on the left point to Breitbart or pro-Trump hoax sites, while those on the right flag The Canary or the hoax sites designed to catch anti-Trumpers.

Politicians – in general – have not suddenly become more mendacious. The media have not suddenly become more inclined to lie. And despite suggestions otherwise, the public have not become more stupid or distracted (on the contrary, on average, we're more educated than we've ever been). So why is bullshit now in the ascendency?

In the US, Trump is exceptional in his repeated and tenacious disconnection from reality – but he hardly acts against a political trend on both sides of the Atlantic for focusing on messages that cut through, rather than getting too bogged down in boring-but-important details.

The big and systematic reasons for bullshit's triumph lie in large part on the media side of things, both with traditional outlets and with the new economics of the internet. Most of the time we discuss such things, we focus on the new technologies and platforms and their effect on us. That misses another seismic shift: the economics.

The business model of 'serious' outlets is under sustained pressure, especially so in the case of print media. Circulation numbers are falling, which simultaneously starves outlets of both circulation revenue (cash from the cover price) and advertising revenue, as people pay less as they reach an ever-smaller audience. This advertising drop is compounded by companies switching their ad spend to digital outlets – meaning numerous papers are seeing print revenue fall by 15 per cent or more each year. Less money means fewer reporters, each doing more work than ever with lower budgets, making regurgitating what politicians say a much more cost-effective proposition than digging into what they're saying.

There are knock-on effects too. Some outlets shift their coverage to suit their dwindling and ageing reader base, giving up entirely on ever reaching a younger audience in print.



Most have turned to their online coverage to compensate for their revenue, often compounding the bullshit problem. The biggest and most specialised outlets – places like the *New York Times* or *Financial Times* – have the prospect of getting enough people to pay online subscriptions for their news, allowing something akin to their traditional business model to continue.

Many others are instead in the game of reach: each visit to a news story generates fractions of a penny through display advertising, not nearly enough on its own to fund a news story. The way to make such small amounts pay is to generate huge audiences – millions a day – and to try to make each story as low-cost as possible.

This is not a business model designed to combat bullshit, but rather to propel it as far as possible across the world. Assigning a reporter to spend hours looking into a claim, then writing a cautiously worded article on its truth leads to more cost for fewer clicks. The easier and more lucrative alternative is to write up the original claim, unchecked, within minutes, followed by any angry reaction to it, alongside rebuttals. Any debunk can simply be cribbed from another outlet. The result can be six or more stories – some of them directly contradictory – with zero original reporting.

This model lends itself too to new media and fake news. If the goal is to maximise an audience (and therefore ad revenue) already facing a glut of down-the-middle serious news, then the trick is to hype and promote any kind of row to get a huge influx of partisan readers. While the hyper-partisan right-wing site Breitbart may be the poster child of this movement, the left-wing UK blog The Canary works on just the same model – which also benefits from sowing doubt as to the veracity of the mainstream media, thus promoting sharing and future clicks: ‘Here’s what you won’t read’ is a strong sell.

Fake news sites are the logical conclusion of this particular business model: if a story is going to be unchecked, or exaggerated, why not make it up entirely and reduce the costs even further? One ultra-successful fake story may generate money through advertising revenue, but may also be used to promote affiliate schemes such as casino sign-ups, get-rich-quick schemes, dubious health products or similar programmes. These generate a much more generous income than ad views to fake sites for each new

person who signs up to the product or service, providing another lucrative revenue stream to hoaxers.

There's a twist to this: almost every major news site profits from these fake sites, too, even while wondering how to tackle them and warning of the risk they present. The 'sponsored links' present at the foot or side of posts on almost every major site give the outlet a small amount of revenue for each click, but almost universally link to fake or hyped news. Traditional media boosts and profits from fake news, even as it tries to fight it.

The above is the short version of the mess this book is trying to untangle. Part I will set out how bullshit – in different forms – shaped two of 2016's central political campaigns, tracking down the detail of how the Brexit battle was fought and how Trump rose to take the White House. These accounts will dissect the outright falsehoods, but also show how just the right dose of bullshit – coupled with a credulous media response – helped some narratives run for weeks, while others withered on the vine.

Part II then looks in turn at each of the key players involved in the process: politicians, old media, new media, fake media – and us, the consumers of news. What contribution has each made to what's happening? What are the limits of what each group will do? And, crucially, what reasons does each group have for acting as they do?

Then in Part III, we turn to looking at why bullshit works as a tactic – the toxic mixture of bullshit often playing well into the psychology of the audience, feeding our existing beliefs and reinforcing our social groups, furthering the goals of political actors, and serving the business models and long-standing culture of the media groups standing between the two.

Finally, we look at what's already being done to challenge fake news – if not bullshit – and why it's not working, and what else we might do to tackle the underlying issues that can perhaps turn back the tide.

Fact-checking won't be nearly enough. The media theorist Clay Shirky said in July 2016 that 'we've brought fact-checkers to a culture war'. This isn't to doubt the good that fact-checking can do, but all too often the people reading the debunks are not only far fewer in number, but also nothing like the same people who read the initial false claim. Not only do debunks of this sort do little to heal divides, they can

inadvertently enhance them. We will need to go outside of our comfort zone to tackle bullshit.

Is bullshit an issue we even need to tackle, though? ‘Fake news’ is nothing new, and while ‘post-truth’ may have been the word of the year in 2016, there’s plenty of seemingly bigger things going on: fears about the rise of populism, nationalism, a growing partisan gulf and accompanying erosion of the political centre.

I don’t think the rise of bullshit coming at the same time as a rise in populist sentiment is a coincidence – each feeds the other. A corrosive effect of our casual attitude to truth is that there’s no agreed way to test our conflicting narratives against one another: all we can say is that those who disagree with us are malicious, corrupt or liars. Donald Trump slams the media; left- and right-wing outlets call rival politicians (and each other) liars or ‘fake news’; and supporters of each group turn on one another as dupes or traitors.

Fake news is more a symptom of this vacuum of trust than a cause: bullshit is indeed the enemy of the truth, and without a sense of truth we have no way to debate across the political fence – we can only shout our conflicting narratives. The end result of such an environment gives no more weight to the BBC or the *New York Times* than to a Facebook status or [AmericanPatriotDaily.com](http://AmericanPatriotDaily.com). Such an environment cannot help but be corrosive to the long-term health and stability of a democracy.

One theme of this book is that we all have our biases and we read, share and respond to news in accordance with them – whether we acknowledge them or not – which means it’s only reasonable to share my own, especially as I’ve worked across a range of outlets which are mentioned throughout.

While writing this book, I’m employed as a special correspondent by BuzzFeed News. Prior to that, I’ve worked at *The Guardian* in the US and the UK, at the *Washington Post*, at the Bureau of Investigative Journalism, and for a time, during the 2010 Chelsea Manning leaks, at WikiLeaks. Additionally, I’ve freelanced for a range of UK newspapers and collaborated on projects with the *New York Times*, International Consortium of Investigative Journalists and ProPublica. If you’re looking to call me an MSM shill, the evidence is all there. Beyond that, I’ve tried to source and evidence all claims in this book – if there’s anything you’d like to pick up on from it, do get in touch via Twitter: I’m @jamesrbuk.

This book can't make any claim to have all the answers to tackling the rise of bullshit. What it hopefully does do is set out the scale of the problem, why it's happening, what motivates those engaged in it and why what's being done so far is inadequate, and to suggest some first steps in tackling the problem.

## **PART III**

# **WHY BULLSHIT WORKS**

## CHAPTER NINE

# WHY WE FALL FOR IT

**W**e've been focusing up until now on how bullshit spreads, and some of its consequences. The remaining piece of the puzzle is to look at *why* it spreads: why we tend to believe it, why people and organisations have incentives to spread it, and why stopping or slowing it down proves difficult for others. We're tackling each in turn, and this chapter looks at the various quirks of our psychology that make us amenable to believing and spreading different types of bullshit.

These tie into quirks in how we think and act. These include, in short, that we find information more believable if it aligns with our current worldviews, and that most of us find anecdotes more convincing than statistics. When we're in a group environment – as we are on social media – how we act around each other emphasises these effects: we like to show we're part of the group, we attack people who aren't part of it, and groups which start out moderate can easily pull themselves into more extreme positions, even without anyone intending for this to happen. This chapter will focus, largely in the abstract, on the combination of psychological factors that deal with how we process and produce information – but we should keep in mind what the modern media landscape looks like as we read on. Online, we have our natural groups we create by accident (our Facebook friends and similar) as well as divisions we create deliberately, such as the political parties or causes we sign up to. We can regularly see what other members of those groups see and share, and they can see what we do too – the architecture and infrastructure of the modern internet could almost have been designed to trigger the instincts that make us most likely to believe things which aren't true.

One of the simplest things we're all inclined to do is look for and accept information which supports our current beliefs – for example, if

we're against the death penalty, we're more likely to look for and share articles which seem to show that the death penalty is ineffective or cruel, and even possibly more likely to remember material in these kinds of stories (though this is contentious). This is known as 'confirmation bias': we look for and retain information that confirms our beliefs, and we struggle to accept information that goes against them.

This might seem to be just stating the obvious: no one likes admitting they're wrong about a dearly held belief, and no one wants to admit they're acting irrationally – but we act like this when there's nothing much at stake. The book *Irrationality* by Stuart Sutherland lays out how this can work<sup>1</sup> – here's one such case which I'd invite you to try yourself before reading on. Here's a sequence of three numbers: 2, 4, 6. The task is to work out what the underlying rule is governing that sequence. The participants in the actual experiment were invited to offer alternative three-number sequences. The researchers would then say whether or not that sequence followed the rule. Once they felt relatively sure of what the rule was, they were invited to guess it. So – stop for a moment and consider which sequence you might guess next.

Some people offered sequences like '8, 10, 12', or '16, 18, 20' next, and were told that these followed the rule. This might lead you to conclude the sequence is 'three successive even numbers'. But in reality, you've gained no new information from those sequences versus the first one – and this isn't the rule. Others might try '8, 12, 16', and assume it's 'three ascending even numbers'. Another option would be '1, 3, 5', assuming it's 'any three numbers increasing by two each time'. Both sequences follow the rule, but neither rule is correct. Here's another sequence which follows the rule: '-1, 8, 29', and one that doesn't '5, -2, 14'. What's the rule?

In reality, the rule the researchers used was 'any three numbers in ascending order' – but rather than challenging the apparently stronger pattern we see in the example we're given, we tend to offer numbers which confirm it. But we only get to the right answer by testing contradictory sequences. Our instinct even when we're faced with abstract mathematical sequences is to look for confirmation of an obvious belief or hypothesis – and this effect seems only to get stronger for beliefs that we have an emotional connection to. One study looked at strong supporters of John Kerry and George W. Bush, showing them

apparently contradictory statements made by each candidate, and some statements explaining the apparent contradiction. Predictably, partisans on each side were more likely to acquit their own candidate while feeling their rival had contradicted himself – but the study also found that in these cases an MRI scanner showed little activity in areas of the brain connected to ‘cold reasoning’.<sup>2</sup> We’re already not good at assessing information we disagree with, and when it’s a strongly held opinion, we may even bypass logical assessment anyway. *Irrationality* sets out the most extreme consequences of these cognitive biases – using the example of Admiral Kimmel, the Commander-in-Chief of the US Pacific Fleet in the run-up to the attack on Pearl Harbor, who didn’t cancel shore leave or increase alert efforts even after a submarine was sunk near the base.

‘We have now established two ways unconsciously used to maintain beliefs – refusal to act on contradictory evidence and refusal to believe it or act on it if it is brought to one’s attention,’ Sutherland concluded.<sup>3</sup> ‘Kimmel was guilty of both faults. He failed to seek evidence from Washington to clarify an ambiguous message, and he refused to believe the submarine sunk outside Pearl Harbor was Japanese.’

The US lost four battleships and 188 aircraft in the attack, and 2,403 people were killed. Kimmel was relieved of his command within a fortnight, and retired within months.

Confirmation bias alone is difficult enough to tackle, but some studies suggest there’s an even stronger effect at play: when presented with evidence that contradicts one of our most closely held beliefs, it may actually serve to *reinforce* that belief rather than challenge it. This apparent effect, discovered in experiments by Brendan Nyhan and Jason Reifler at Dartmouth College, has come to be known as the ‘backfire effect’. In a series of experiments, Nyhan and Reifler showed students of different political persuasions an untrue statement, and then information from a credible source correcting the error.

The first experiment, conducted in the autumn of 2005, showed participants a statement from (then) President Bush stating that Iraq had WMD, immediately followed by a correction – cited to AP – quoting a report which said there was no evidence that Iraq had WMD. The correction reduced the likelihood of liberals and neutrals to say they believed Iraq had a WMD programme – but conservatives were actually *more* likely to say Iraq had WMD after the correction. ‘If subjects simply



distrusted the media, they should simply ignore the corrective information,' the researchers noted.<sup>4</sup> 'Instead, however, conservatives were found to have moved in the "wrong" direction – a reaction that is hard to attribute to simple distrust.'

The researchers repeated the experiment on different issues, again triggering backfire effects. Students were shown an untrue claim that President Bush's tax cuts had stimulated the economy so much that they had increased government revenues, then were immediately shown evidence that this wasn't true – but were more likely to believe the untrue claim after the correction than before it. Another study, this time on whether or not President Bush had banned stem cell research, did not itself provoke a backlash – perhaps showing that people's views on this issue are less strong than on the others – but groups opposed to Bush did still exhibit a resistance to believing the information in the correction. One important caveat, though – when another group of researchers tried to test for the backfire effect on a range of thirty-six issues, they only detected a backfire effect on one.<sup>5</sup> How strong and how common the backfire effect is remains a subject of debate.

The consequences of the backfire effect when it does kick in, though, could be huge. 'Just as confirmation bias shields you when you actively seek information, the backfire effect defends you when the information seeks you, when it blindsides you. Coming or going, you stick to your beliefs instead of questioning them,' an article by David McRaney concluded.<sup>6</sup>

What should be evident from the studies on the backfire effect is you can never win an argument online. When you start to pull out facts and figures, hyperlinks and quotes, you are actually making the opponent feel as though they are even more sure of their position than before you started the debate. As they match your fervor, the same thing happens in your skull. The backfire effect pushes both of you deeper into your original beliefs.

When we reach across the political fence online and offer up evidence supporting our arguments, it rarely goes well – we feel affronted that our political opponents are rarely convinced by our reasonable approach, and perhaps even get angry that people won't see sense. We can fire up our own base, but arguing from evidence – as the evidence, ironically enough, suggests – likely won't work.

Compounding these effects is our poor general understanding of statistics: for all sorts of reasons, we both struggle to understand statistics

in news and also tend to disbelieve them if they contradict our anecdotal experience. This is compounded by journalists and others – whether due to their own poor grasp of statistics or in order to push an agenda – often distorting how statistics are presented to the public, with serious and detrimental effects. One fallacy with severe real-world consequences was to confuse correlation and causation: assuming that because something happens shortly after something else, the one caused the other. This was one of the main drivers of a huge outbreak of public concern that the vaccination for measles, mumps and rubella (MMR) was causing autism in children. This claim was not only fuelled by a fraudulent doctor, Andrew Wakefield, but also spread across the media – but it was never supported by a single piece of high-quality evidence.

However, it seemed plausible simply because the two events were correlated: children receive their first jab around the age of one, and autism generally first manifests in those children with the condition when they're toddlers<sup>7</sup> – shortly afterwards. The whole controversy led to thousands of parents deciding not to give their children the MMR vaccine, reducing the level below what's necessary for 'herd immunity', leading in turn to a number of outbreaks across the country. Nine years after being discredited in the UK, Wakefield is near the centre of US politics, pictured at inauguration balls,<sup>8</sup> with Trump apparently endorsing his baseless anti-vaccination views.<sup>9</sup>

When we get into stories that involve actual figures, things get even trickier. Week in, week out, we're inundated with stories that tell us various habits either raise or reduce our risk of cancer – to pick one at random, in 2015, outlets across the world warned that eating two rashers of bacon per day raises the risk of bowel cancer by 18 per cent.<sup>10</sup> This tells us less than it seems, though: without knowing what our risk of getting bowel cancer was if we *don't* eat bacon, we don't know how bad this is. As it turns out, over the course of their lifetime, about five in 400 people will contract bowel cancer. If every single one of those 400 people began eating two rashers of bacon per day, this would rise to six in 400.<sup>11</sup> This second way of looking at the numbers may be less alarming, but also shows how useless the relative figures are: let's say some other foodstuff increases the chances of a more common disease by 10 per cent, but would mean ten extra cases in 400. This 10 per cent risk sounds

less bad than the 18 per cent, but would actually affect far more people. The way the media talks about statistics, and the way we generally understand them, doesn't help us understand the real world.

This problem is heightened when we're dealing with risks that scare us: statistics just don't really reassure us. A sizeable portion of us are afraid to fly,<sup>12</sup> whereas fear of driving is much lower. This is despite the fact that flying is by far the safest mode of transport: only one in 3,000,000 flights have a crash that results in fatalities (and almost every plane crash with fatalities has some survivors). Mile-for-mile, travelling by train is twice as deadly as flying; travelling by car is 100 times more deadly than flying; and taking a motorbike is 3,000 times more deadly than flying.<sup>13</sup> And what's vastly more likely than any of the above to kill you is either heart disease or cancer – but even knowing this, planes remain scarier than, say, obesity.

If we're talking about things that scare us, terrorism comes incredibly high on the list, and is used by the media for headlines and by politicians to push through political agendas and win elections – Trump's promise of a hardline position on terror has served as justification for his travel ban from seven countries, and more. But our fear of terror is beyond all proportion with its risk. Michael Rothschild, emeritus professor at the University of Wisconsin, calculated an extreme scenario to illustrate this risk. He imagined a world in which terrorists successfully hijacked and destroyed one plane every month in the US – dramatically worse than any real-world scenario in any country in the world. Someone who took four flights each month in this world would have a one in 540,000 chance of being killed in any given year. Someone who only took one flight a year would have a risk of dying of less than one in 6,000,000. By contrast, the chances of dying in a car accident on any given year is one in 7,000, of cancer is one in 600, and heart disease is one in 400.<sup>14</sup> Even if terror attacks got unimaginably worse than they are, it would be a much smaller risk than just going about our everyday lives. Our fears, and our inbuilt resistance to being reassured by evidence, makes us susceptible to bullshit.

So far, we've only looked at factors which affect us as individuals, able to assess and react to information on our own – but of course none of us live in total isolation, and in the online era we see our news alongside our social groups. We see the news our friends share, and we

know what our friends think of particular speeches or candidates. For those of us who strongly identify with a political cause or campaign – something which gives us a sense of group identity – this becomes an even bigger factor, and all of these influences affect how we judge and share information, and so shape what we believe.

Generally, we like to fit in with our friends and colleagues. This kind of conformity doesn't easily fit into how we think about ourselves, but a majority of us will act to fit in with a group even when it's a group full of strangers and the question at hand is simple and entirely apolitical. In her book *Wilful Blindness*, Margaret Heffernan sets out a study showing how dramatic this effect can be. Students were brought into a room with a group of fellow students apparently participating in the same experiment and were asked to say which out of three vertical black lines was the same length as a fourth line on the whiteboard. The task was not difficult: the lines were clearly very different in length, and the correct answer was obvious. However, in each room only one student was an actual subject: every other student in the room had been told in advance to choose the same wrong answer, with the actual test subject answering last. Even in a case like this, with the group comprising total strangers and the question at hand being both obvious and uncontroversial, 40 per cent of subjects chose the obviously wrong answer given by everyone else rather than trusting the simple evidence of their own eyes. The experiment was repeated with different variations, and found that under one condition or another, 58 per cent of people would confirm the popular answer in at least one of the tests given.<sup>15</sup>

This kind of conformity isn't about obeying a leader or following instructions – this is just something many of us do to ourselves, choosing to go with the flow rather than make our own assessments. When a majority of us will choose conformity over truth in such simple conditions, it's not hard to see how this habit can affect us in our lives: becoming embroiled in bad corporate cultures, not speaking out against practices we disagree with, reflexively defending politicians or simply not stating that we don't share a particular opinion of a group we generally agree with.

Conformity is little more than a desire to be polite or to fit in – but arguably it can even in its own right be dangerous. Heffernan warns that

conformity can serve to explain why many of us believe climate change is real but take very little real action.

We live with people like ourselves, and sharing consumption habits blinds us to their cost ... we're obedient consumers and we might change if we were told to, but we're not. We conform to the consumption patterns we see all around us as we all become bystanders, hoping someone else somewhere will intervene. Our governments and corporations grow too complex to communicate or to change and we are just left where we do not want to be.<sup>16</sup>

There's evidence that we'll act to show our belonging inside a group, so it's worth thinking about how that works in practice on social media. If one of our concerns is how a group we like and want to feel part of will react, what kind of story will we share? If we're a Trump supporter and have read two articles, one questioning whether his team could have executed his order banning travellers from majority-Muslim countries better, and another attacking the media for misreporting Trump's presidency, we might share the second because we know for sure it'll get a good reaction from our group.

This kind of sharing behaviour, particularly when carried out by those on the left, was given the slightly derisory name 'virtue signalling' – the right-leaning *Spectator* magazine gives itself credit for coining the phrase, which has rapidly become a term of abuse used by those on the right against the left. 'People say or write things to indicate that they are virtuous,' James Bartholomew wrote in the magazine.<sup>17</sup> 'Sometimes it is quite subtle. By saying that they hate the *Daily Mail* or Ukip, they are really telling you that they are admirably non-racist, left-wing or open-minded.'

The way 'virtue signalling' is used as a term almost exclusively relates to the right, but signalling behaviour – people trying to send cues that they're still part of a group – occurs on both sides of the political divide, with tropes such as disliking political correctness, supporting free speech or attacking metropolitan elites working as signalling devices of the political right. One consequence of this kind of signalling can be to convince partisans that public support for their position is much broader than it really is, especially if it's something few people would want to publicly oppose.

Signalling, argued Helen Lewis in the *New Statesman*, 'had many social media users convinced that Ed Miliband could squeak the election; after all, their friends seemed to be lapping up the mansion tax and the

action against non-doms. No one seemed enthused about taking £12bn off the benefit bill, or reducing the help given to disabled people'.<sup>18</sup> Ed Miliband's Labour Party lost the 2015 election, which polls had suggested they'd narrowly win. The combination of the filter bubble and signalling within it can give us false expectations of people's views – who on the centre or left would want to express reservations on the deficit or support for austerity? Many stayed quiet, then voted Tory. Our instincts on conformity fuel a desire toward 'virtue signalling', and distort what all of us see.

When we're talking with people we largely agree with, though, our views actually change. Imagine taking a group of people who by-and-large agree on a topic, but with differing degrees of intensity, and having them discuss that topic for a few minutes. You might imagine that the consensus opinion of the group would settle somewhere around the middle point of the individual pre-discussion opinions. That's not what happens.

The academics Cass Sunstein and Reid Hastie gathered groups from two cities, the relatively liberal city of Boulder and the relatively conservative city of Colorado Springs. Participants were asked to privately give their views (on a scale) on three issues beforehand: climate change, affirmative action and same-sex marriage. People were then matched in groups with other people from their own city (so that the participants in each group were somewhat in agreement with each other before the discussion) and asked to debate each issue, then asked for their opinion – both privately and as a group – on the issues once more.

The results were stark. The Boulder groups became more liberal on all three issues, and the Colorado Springs groups more conservative – not only in terms of the group verdicts, but also in the views the participants expressed privately: the discussion made each individual stronger in their opinions than before. Equally, the groups became much more similar: before the short discussions, there had been a fairly wide range of opinions on each issue – after even a short chat, opinions in the groups had become much more homogenous. Finally, the Boulder and Colorado groups had polarised: where before the discussions there was overlap between the groups, afterwards the liberals and conservatives were much more divided.<sup>19</sup>

These effects were noticed in groups of relatively like-minded people over a period of just a few minutes. With people who interact day after day and month after month, without a lot of outside interaction it's not difficult to see how substantial this shift can become. The effect persists across countries and cultures, and affects even professionals whom we might expect to be alert for such influences. People will polarise in their group assessments on questions of judgement as well as on questions of opinion: studies of business people asked to predict the likelihood that a sales team would sell a certain number of units showed that a group discussion resulted in a much more confident opinion than the average guesses of participants before the chat began – people polarised their factual assessments. Perhaps most dramatically, the effect even extends to federal appellate judges, who sit on panels of three. One study found that while you can try to predict how an appeals judge will rule by finding out whether they were appointed by a Democratic or a Republican President, there's something that will predict their vote even more reliably: which party appointed the panel's other judges. Inevitably, the effect also manifests itself with juries. The amount in damages that juries award against plaintiffs is typically much higher than the average of what jurors individually say they would award before the discussion.

If group polarisation can affect even small groups of professionals trained to watch for biases – like judges – it's not hard to imagine the effect of certain internet subcultures, or even what our Facebook friends or Twitter followers may do to our own opinions. The filter bubble doesn't just shield us from dissenting opinions – it also carries us further away from the centre.

We conform with our groups, we signal our belonging in our groups and we are polarised by our groups. These factors begin to explain why bullshit information that supports our group identity may be more welcome than accurate and verifiable information which challenges it. More than this, though, we build an identity around the opposition of our groups to others – a habit referred to as 'in-group, out-group' behaviour, or 'realistic conflict theory'. Part of what makes us feel like a real member of one group can be rivalry, or even hatred, of another.

The most famous study of this kind of conflict – the details of which have been challenged by later studies – dates back to the 1950s, when researchers arranged an experimental summer camp for twenty-two boys

in Robbers Cave State Park. The boys were split into two separate groups, not meeting each other for several days, instead having time to bond and establish themselves as groups, catching only the occasional glimpse of the other. After a few days, the groups named themselves the ‘Rattlesnakes’ and the ‘Eagles’. Even before any attempt to introduce rivalry between the groups had begun, there was tension: the Rattlesnakes came within earshot of the Eagles, who were playing ball. Their immediate reaction was to ‘run them off’ or ‘challenge them’, going on to make territorial claims to different parts of the park,<sup>20</sup> marking the baseball diamond with a team flag – and both teams said they wanted to challenge the other.

The pieces were in place for rivalry, and it took almost no effort from the researchers to bring it up to fever pitch: they simply offered a prize of some Swiss Army knives as a reward for the winners of a baseball match. The teams taunted one another before the game – fairly standard for sports – but after the end of the match, which the Rattlesnakes won, things escalated rapidly: a few hours after the game, the Eagles tore down and burned the Rattlesnakes’ flag. In retaliation, the Rattlesnakes raided the Eagles’ camp, shredding mosquito netting and stealing personal belongings. The Eagles planned (but didn’t carry out) a retaliatory raid on their rivals, plotting to attack them with ‘socks filled with rocks’. The Eagles referred to the Rattlers as ‘bums’, ‘poor losers’ and ‘cussers’, the Rattlers dubbing the Eagles ‘sissies’, ‘cowards’ and ‘little babies’.<sup>21</sup>

This was a conflict between two very similar groups which had been arbitrarily created and which lasted only for a few days. Both groups were of boys of the same age from the same area – and yet it took several days of very deliberate effort on behalf of the researchers to reverse the divisions they’d created and at least partially reunite the groups. The adults changed the nature of the tasks given to the boys to require both teams to work together, rather than in competition, in order to get any kind of reward.

De-escalating conflict between short-lived, homogenous groups with no real differences required sustained effort from researchers. As such, it’s not hard to see how difficult avoiding polarisation online could be. Trump supporters have Democrats and the media as out-groups. Corbyn supporters’ out-groups often include centrists, Conservatives and (again) the media. Once an entity becomes an out-group, it’s not only less likely



to be taken seriously, but opposition to what it says also becomes part of group identity – and in the real world there are no researchers sitting waiting to try to reunify the people concerned.

Given all of these factors, it shouldn't come as a surprise that online radicalisation is a concern for governments across the world: even without anyone's deliberate effort, our membership of groups can move our views away from the political centre towards the extremes and increase our antipathy towards those with different views. If these natural trends are then being deliberately exploited by people to promote a cause, whether that's Islamic extremism or the far right, then social media has the potential to become a fertile source of recruits.

The UK think tank Demos connects the risk of online recruitment to ISIS, and even of lone-wolf terrorists becoming radicalised through a mishmash of online half-truths and conspiracies, with people's poor ability to distinguish between truth and lies on the internet. A prescient 2011 report concluded:

Our research shows, however, that many young people are not careful, discerning users of the internet. They are unable to find the information they are looking for or trust the first thing they do. They do not apply fact checks to the information they find. They are unable to recognise bias and propaganda and will not go to a varied number of sources...

The potential consequences of this on society as a whole are unknown. One danger is that young people are more likely to be seduced by extremist and violent ideas ... many terrorist groups are fed by bogus online material circulating unchallenged on online echo chambers. Anders Breivik, the recent Oslo terrorist, is a devastating but only the most recent example of the power of internet material to radicalise.<sup>22</sup>

A UK parliamentary inquiry said social media – specifically Facebook, Twitter and YouTube – had become the 'vehicle of choice in spreading propaganda' and 'the recruiting platforms for terrorism'. The MPs called for the companies to hire more people to monitor and delete extremist content, and for greater efforts to proactively report people posting such material to law enforcement.<sup>23</sup> 'The modern frontline is the internet,' said Keith Vaz, then the chair of the parliamentary committee which ran the inquiry.<sup>24</sup> 'Its forums, message boards and social media platforms are the lifeblood of Daesh [ISIS] and other terrorist groups for their recruitment and financing and the spread of ideology.'

Blocking the most extreme content may be part of a solution, but given our ability to radicalise ourselves – and given that even content that may be within the realms of what's permitted to be published by the

social networks could radicalise people – Demos, and other experts, warn this alone isn't nearly enough.

Demos researcher Louis Reynolds argues:

The government is putting more and more pressure on social media companies to censor content, on schools to restrict internet access, on the justice system to become more involved in policing social media, in order to counter jihadi content online. Yet ultimately trying to censor the internet can't be the principal solution to this problem.

There are next to no effective extremism-related efforts to improve young people's critical thinking ongoing in the UK ... We have to teach young people, from an early age, how to spot manipulation, to counter grooming efforts, to challenge extremist views and spot falsehoods, online and offline. The recruitment of foreign fighters to the Islamic State movement represents a failure of reason and a victory for manipulation. If there is a counter-narrative, this should be the message.<sup>25</sup>

Online radicalisation is not something that only happens through sustained and deliberate effort. Instead, it's built into the architecture of how we interact online – meaning that ongoing efforts to take the most extreme content offline can only take us so far.

Much of how we deal with and process information online comes down to how we choose to think about it – and it's a decision that we often don't even realise we're making. The Nobel Prize-winning psychologist Daniel Kahneman set out a thorough explanation of how we can often deal with things reflexively rather than through deliberate consideration in his bestselling book *Thinking Fast and Slow*.

The idea is best illustrated with an example. Consider this maths problem:  $2 \times 2$ . For almost everyone reading this paragraph, the answer to the question will have immediately popped into your head – with no deliberate effort or action required. Now think about another maths problem:  $19 \times 23$ . Unless you're a polymath (or didn't even bother trying to solve the problem), tackling this one will have required some deliberate thought. You may have decided to try to work out  $20 \times 23$  and then subtract 23 from the answer. You may have considered looking for a calculator. But unlike  $2 \times 2$ , the problem will have required some thought and some effort on your part.<sup>26</sup>

These are the two types of thinking Kahneman lays out. The first, effortless form of thinking is System One – our gut feelings, instincts and reflexive actions. System One thinking is easy and beyond our control: we don't decide to automatically solve  $2 \times 2$ , it just happens. System Two is our deliberate thought, requiring conscious mental effort. Thinking in

this way feels much more intentional, like we're making a choice or working something out.

The two systems often fight with each other, as this famous puzzle illustrates: 'A bat and ball together cost £1.10. The bat costs £1.00 more than the ball. What does the ball cost?'

Many people reading this paragraph will have an instant answer of 10p jump into their head – it just looks right when you read the puzzle. But if you stop to consider the puzzle, that answer is wrong: if the ball costs 10p, the bat must cost £1 more, meaning the total cost of the two would be £1.20 rather than £1.10. The correct answer is that the ball costs 5p and the bat £1.05. Our System One thinking guides us towards a quick and obvious answer which is wrong. Our System Two thinking – if we give it time to kick in – lets us ignore that and work out the correct answer.

This gives us a good way to think about how we read and share online. When we see a quote that suits our political beliefs, it's easy for us to believe it and almost reflexively share it to our social circle – to stop and consider the source of the quote and maybe even check whether it's genuine requires a degree of effort. The former is, for social media addicts at least, a System One action. The latter requires us to make a considered decision to engage System Two to assess the information more carefully before sharing. This isn't something we can do without effort: we have to make a decision to think in this way, and that's not as relaxing or straightforward as the way we often browse. Engaging System Two is an act of self-control, and this is something we find tiring – a process called 'ego depletion'. Without deliberate forethought and self-control we will leave ourselves open to bullshit – and in a sense, we only have a certain stock of self-control in the tank each day,<sup>27</sup> and have to decide whether this is where we want to spend it.

There are many psychological reasons that bullshit works on us, even if we consider ourselves well-educated and able to discern good information from bad. Material may fall into our current worldview, suit our social norms, suit something we wish to signal, or reinforce a group identity. None of these mean we're doomed to believe bullshit, but they do mean we're predisposed to do so without deliberate effort. It shouldn't come as a surprise that bullshit works – if it didn't work, it wouldn't be so prolific – but knowing the mechanisms by which it catches us should

help. That takes us some way towards understanding the demand side – why bullshit is an effective strategy. The next step is to tackle the supply: why it's financially worth spreading.

## CHAPTER TEN

### WHY IT'S PROFITABLE

**B**ullshit thrives when there are people and organisations around with an interest in spreading it, and an audience willing to accept it. The reasons for some politicians to spread bullshit are fairly self-explanatory, and have been tackled in previous chapters – they can help bring election victories, can mire opponents in nonsense scandals, and can crowd out unhelpful media narratives. As we've said elsewhere, bullshit can pave the path to power (which doesn't bode well for politicians' footwear).

The motivations for others tend to be less straightforward, but usually come down to money. For some outlets, letting through the occasional bullshit story is just a side-effect of shrinking newsrooms and cost-cutting. Other outlets, pushing for mass traffic whatever the consequences, are happy to run questionable but cheap stories alongside better-sourced information if that will boost their traffic. But behind the news outlets we visit as consumers lies a supply chain with an interest in passing on unverified stories, hoaxes and stunts and feeding them into the news ecosystem – and it's this array of business models, and the incentives of the players across the industry, that we're going to explore in more depth in this chapter.

Working out what gets spread by who and why is important: it's easy to just point at bad viral stories, or outlets credulously running sensational pieces, and purse our lips and say they're doing something wrong. That's also pointless – unless we work out why organisations have financial reasons for acting as they do, we'll miss many of the real solutions that could actually help us tackle the problem.

The place where most of us as consumers encounter questionable – or bullshit – stories is through media outlets. These are facing a huge number of challenges to their business model, virtually all of which encourage them to favour low-cost content that will appeal to huge

audiences. The challenges of the shift to digital are most obvious for news outlets which are used to print revenues.

Print advertising is worth vastly more – in total and on a per reader basis – than digital advertising to most legacy newsrooms. Analysis of five of the largest US print groups by the Pew Research Center in 2016<sup>1</sup> found that digital made up around a quarter of total advertising revenues, much more than it had a decade before. But this apparent growth is exaggerated by decline elsewhere – digital’s share of the total is growing because print is shrinking so fast. Not only are print revenues falling, at rates often at 10–20 per cent a year, but newspapers have suffered from losing job advertising and classified adverts to online rivals. We tend to focus on the challenge of newspapers getting less money for display adverts (banners, pop-ups etc.), and ignore all the other types of advertising they lost out on even more sharply.

Digital advertising makes only a tiny amount of money per click: a banner advert in a prime spot on a high-quality and big-name news website will make only a fraction of a cent, something around 0.5¢ to 0.7¢.<sup>2</sup> This is a premium price – other sites receive far less for each advertising spot. This gets even more difficult still when we start thinking about how we actually browse: a computer screen can show several large adverts around a news story, but a mobile screen cannot – but now around half of all news is viewed on mobile, and this share is growing.<sup>3</sup> This means digital revenues, which were already barely a fraction of print revenues, are themselves dwindling as desktop browsing is replaced by less lucrative mobile traffic. Advertisers are also sceptical of visits on many web pages, due in part to unscrupulous middlemen gaming the system by creating ‘bots’ to generate fake traffic (and therefore ad impressions) – and due in part to simpler concerns that many readers never actually saw an advert halfway down a page.<sup>4</sup>

The result is increasing desperation on behalf of sites to prove that visitors have definitely seen the adverts on the site. This includes buttons which require you to click to continue reading a story (which guarantees you’ve seen the advert that appears as you click), adverts that pop-up and appear before you can read a story, or annoying adverts which divert you to app stores on mobile. All of these are, if we’re honest, pretty irritating to readers – and the result is that around one in five users now use software to block adverts, up from fewer than one in ten a few years

ago.<sup>5</sup> People using ad blockers generate no revenue whatsoever for the sites they visit. The advertising problem just keeps getting worse.

Some sites – like the *New York Times*, *Financial Times*, and *Wall Street Journal* – have shown they can build subscriber bases to rebuild their revenues and become less reliant on advertising, but for the majority of sites that rely on adverts, the situation is difficult. There is a need for constant content to generate large traffic numbers, with not much money around for fact-checking or scrutiny – and the need to get large numbers of visitors may also incentivise running sensationalised versions of a story. It's a great habitat for bullshit to thrive.

Mark Thompson, the CEO of the *New York Times* and former director-general of the BBC, set this out in stark terms in a 2016 essay for the Reuters Institute at the University of Oxford.<sup>6</sup> 'Winter really is coming for many of the world's news publishers,' he wrote, with a nod towards *Game of Thrones*.

Many publishers have responded to this challenge by putting their faith in a model based on audience scale and digital display advertising ... most of the new digital news providers were launched with business models which were parasitic versions of the same idea. They aimed to rewrite and repackage other people's journalism for much less money than it cost to originate it, and then to use superior technology to out-compete the legacy companies in distribution and advertising monetisation ... These models now look suspect.

The media outlets that we read and share are an obvious player in the bullshit game, and their difficult economics make it easy to see why they're amenable to running stories with minimal checking or with exaggeration. But there are layers of agencies and companies acting within this system that aren't visible to us as consumers who play a significant role in how unverified information spreads.

An apparently innocuous news story from February 2017 helps bring some of these players to light. The story centred on footage of a female cyclist in London, pulling up next to a white van at a traffic light. A man inside the van calls out to the cyclist asking for her number, gets rebuffed and then hurls abuse at her, asking if she's on her period. Furious, she pursues the van – which has now taken off as the light changes to green – pauses, and rips off one of its wing mirrors.

The story and footage were picked up by numerous major outlets: the *Daily Mail*, *Metro*, *The Independent*, *The Sun*, the *Sydney Morning Herald*, and more – attracting millions of video views on some. The issue

is that none of them took any action to verify whether or not the video was real before running it, and when *The Sun* did later track down an eyewitness, he said that he thought the video was staged. Elena Cresci, a digital journalist at *The Guardian*, set out why there were obvious warning signs the outlets could have spotted without needing to speak to anyone. ‘Even before this witness came forward a few things made the video seem suspect,’ she wrote.

It’s extremely convenient that the motorcyclist managed to capture the whole thing – though, as many motorcyclists wear helmet cameras, it’s not impossible. Some people also pointed out that it’s a bit difficult to pull off a wing mirror so easily. The video itself seems to follow a model often seen on prankster sites that create outrageous situations and film them with the intention of going viral.<sup>7</sup>

Where had the video come from, to appear across numerous news sites within hours? Like a lot of content designed to spread virally on the internet, it had been promoted and sold by an agency specialising in finding and monetising video going viral on the internet. These agencies work by finding anything taking off on YouTube or similar sites, and acting as a middleman to help the creators monetise their (often accidental) hit – but such agencies can easily be passed content from hoaxers or others.

In the case of the cycling video we’ve been discussing, the agency concerned, Jungle Creations, was selling the video for £400 to be embedded on newspaper websites, or £150 for use on social media – netting thousands of pounds as dozens of outlets picked up the story. However, Jungle Creations said it had done no work to verify the footage. ‘We do not have the resources available to do in-depth checks on every piece of content we get sent,’ it said in a statement. The company later issued a second statement saying the video ‘may be factually incorrect’, and removed it from its own social feeds.<sup>8</sup>

This didn’t provoke a wave of retractions or corrections in the outlets that garnered millions of views from the original dubious video. Some added ‘– but was it real?’ to their original headlines, some did nothing and, audaciously, the *Daily Mail* left its original account untouched, but added a second story reporting on people’s doubt about the video.

Women are routinely harassed during their commutes in London – whether the video is genuine or not, the issue it is highlighting is real. The fakery is not especially serious in this instance, but the compound



effect of stories like it is damaging the news outlets which often demand their readers trust them more than politicians, and certainly more than ‘niche’ or ‘fake’ news outlets. If major sites routinely run unverified stories and videos fed by middlemen with no kinds of checks, why should readers trust them on other issues? How is the casual reader supposed to know which stories are carefully checked and which aren’t?

The reason such videos are spread and run is that they’re a great kind of content for news sites. Video advertising is one of the rosier areas for digital advertising, creating far more revenue per user than standard banner adverts. Thanks to agencies searching and promoting videos which are already going viral, sites can buy near-guaranteed winners at a relatively low cost, boosting their traffic and hopefully their revenues. But no one in the food chain has any incentive to do too much checking – providing the claims in the video aren’t likely to get anyone sued. The agencies are looking to turn a profit on the grabby videos they’ve sourced, and the news sites are looking for easy traffic. The financial incentives for both act against doing too much checking.

Other players in the industry have even stronger incentives to plant stories in the media – for the PR profession, that’s long been part of the job, and it’s one that has adapted to the new news environment. The joy of creating dubious – essentially fake – viral videos for such agencies is that the hoax itself often serves to promote the client, and then any admission that the video was staged promotes the client once again, and promotes the agency itself to new clients. They win both on the initial viral story and again on the debunk.

Such was the case with a story in 2016 where Twitter and Facebook users noticed a high-end white Range Rover parked outside Harrods and daubed with ‘Cheater, it’s over’ and ‘Hope she was worth it!’. The car, left in such a high-profile location, quickly and naturally began to be noticed and go viral on social media and was in turn quickly written up by multiple news outlets, with various degrees of scepticism as to whether or not there was more to the tale than met the eye.

A day later, as the initial PR hit faded, the car workshop behind the stunt admitted everything, securing a second set of write-ups which served their interests even better than the first. An *Evening Standard* article on the hoax referred to ‘the brand new £90,000 white Range Rover’, ‘an advertising stunt for the new Revere Range Rover Vogue’

and to the workshop behind the stunt as ‘a luxury auto design and enhancement workshop’.<sup>9</sup> The article was also accompanied by a second video showing behind the scenes of the stunt and filming the reactions of passers-by as they examined the apparently trashed car. The news sites concerned got two stories for the price of one, the client of the PR company received generous mentions of their product (including yet another here), even in the story revealing they’d tricked the media, and the PR company got to plug itself as somewhere that can plant stories and net traffic – a great way to secure future work.

Specialist agencies and divisions now exist producing just this form of publicity. One Australian agency specialises in shock videos, with examples including ‘Man Fights Off Great White Shark In Sydney Harbour’, ‘Crazy Guy Runs Into Outbreak Tornado To Take Selfie’, ‘Stormtrooper Falls Down Stairs On The Way To Star Wars Premiere’ and ‘Lightning almost strikes girl in Sydney!!!’

The agency began creating the videos to experiment in what kind of material would take off online, but quickly secured clients willing to promote their brand through the stunts, as the company’s managing director Dave Christison told the Huffington Post: ‘Our second video was client commissioned and it was on the same account. Roadshow Films was the client and we created it for an upcoming film called Into The Storm, on the release date, we revealed it was a hoax.’<sup>10</sup>

Despite posting the second film using exactly the same account as their first faked video, no one sussed out the PR agency was behind the films until they revealed themselves after six more viral videos and a total view count well in excess of 200 million – and the reveal itself was, of course, an effort to secure more clients.

These viral stunts are the latest manifestation of a years-long habit of PR agencies trying to find news hooks to secure coverage for their clients, which can result in good and high-quality stories but which can also lead to stories that are baseless or downright damaging. The science writer Ben Goldacre wrote for several years about ‘Blue Monday’ – the idea that there’s a particular Monday in January which is ‘the most depressing day of the year’ – which has now been covered as an annual event in the UK for more than a decade, despite a lack of evidence for seasonal variation in mood, and an obvious total lack of evidence for an arbitrary date marking the low point.<sup>11</sup> The day began as a PR hook for

Sky Travel, before being seized upon by other companies and then, eventually, mental health charities who presumably reasoned if you can't beat 'em, join 'em.

Goldacre, though, noted in 2011 that not all of this bullshit is harmless. A start-up offering counselling via webcam for £50 a pop claimed it had run a survey which found that 34 per cent teenagers faked serious mental health problems to gain attention – a serious allegation which could have negative consequences for vulnerable young people. Goldacre asked the PR agency behind the story for details on how they'd carried out this research, how many teens they'd contacted, how they'd checked demographics, and other details – and received no response.

'Bullshit is a slippery slope,' he concluded on the matter.

At the top is an okay charity endorsing fatuous PR nonsense as a one-off, and normalising it. At the bottom, the health correspondent of a major national newspaper is standing in a river of sewage, shouting at a million readers that teenagers are malingerers who invent serious mental health problems. All I suggest is that you should think a bit before you step onto the crest.<sup>12</sup>

There's every reason to expect more and more bullshit from this source: one way of measuring its impact through the years has been a metric referred to as 'advertising equivalent exposure', a way agencies try to show clients how much they would have to advertise to reach an audience the same size as that reached by the PR coverage – and these days directly through YouTube, Facebook and similar. As brands become sceptical of advertising, with audiences using ad blockers and seeing news directly on social sites, the PR route becomes more appealing – especially as it works. PRs are definitely a significant player in the economics of bullshit.

In the mainstream media, the business models of bullshit focus on advertising and PR, and the same is true for some on the fringes – particularly those who make conspiratorial or hyperpartisan videos, many of whom receive generous funding from YouTube (owned by Google) to do so. As Joseph Bernstein noted for BuzzFeed News,<sup>13</sup> while Facebook receives most of the public attention for the spread of bullshit, YouTube actually funds it. Videos produced by people pushing the pizzagate conspiracy theory (see [Chapter Eight](#)) attract millions of views, and draw advertising from brands including Uber and Quaker Oats – placed by YouTube's ad network. The internet's largest ad network – Google, a

company with a stated mission to ‘organise the world’s information’<sup>14</sup> – and big brands are helping to make conspiracy theories and misinformation pay.

Elsewhere on the fringes, though, the financial models work quite differently – most notably with InfoWars, once at the very furthest of the political extremes, but now with a clear line into the White House. InfoWars is a site and show helmed by Alex Jones, a right-wing conspiracy theorist closely followed by people in the survivalist movement. Jones was a prominent backer of the theory that 9/11 was an inside job, that the Sandy Hook school massacre was a hoax, that the government was behind the Boston bombings, and more. All of these, Jones claims, were part of a sinister agenda of ‘globalist’ forces to take over the world.<sup>15</sup> He’s claimed that Pepsi contains baby parts, that Obama is a member of al-Qaeda, that same-sex marriage is an attack on God and that the government plans to kill off half of the US population.

You may be imagining a man producing a show from a garden shed, but InfoWars has professional studios, syndication of its shows, and revenue in the millions, and has interviewed Donald Trump after his election win – Jones even claims Trump occasionally calls on him for advice.<sup>16</sup>

InfoWars pushes a hyper-masculine, paranoid agenda in which the audience are urged to ‘be part of history’, to man ‘the new battleship of the fight’ – its slogan is ‘You are the resistance!’<sup>17</sup> Its income source is not standard banner adverts or subscriptions or any kind of outside advertiser: instead, the site and the shows relentlessly advertise an InfoWars range of survivalist products – goods designed to help its audience survive an apocalyptic event, or to help them bulk up and prepare for such a conflict.

Products include ‘Super Male Vitality’, produced by ‘InfoWars Life’, which it is claimed ‘has been used by Alex Jones in order to maximize vitality when working up to twelve hours a day or more in the fight for freedom’, for a mere \$44.95 for a month’s supply.<sup>18</sup> Other products tap into the survivalist mindset – people who stock up and prepare to survive for months or years off-grid following a war or disaster. InfoWars sells iodine – a chemical which offers some protection against radioactive fallout – under the brand ‘SURVIVAL SHIELD X-2’, noting the

supplement is ‘tested for radiation’.<sup>19</sup> InfoWars even partners with the ‘patriot pantry’ to offer ‘emergency storable foods’ with a 25-year shelf life, in stockpiles designed to provide enough food for three, six or twelve months, for \$499, \$895 and \$1,797 respectively.<sup>20</sup>

A cynic could summarise InfoWars’ business model as constantly feeding its audience a diet of stories warning conflict and disaster is coming, while relentlessly advertising bespoke products useful only in the case of that kind of apocalyptic event – reminding them all the while that such purchases are necessary to further the struggle. The business model, as InfoWars points out to its audience, is a robust one, resistant to pressure from advertisers, search engines and others: brands can’t pull their advertising from its site because they’re not there in the first place.

If an outlet can gather an audience devoted enough, it can create its own economic system – which has even survived the transition to a friendly President. The struggle has shifted from the need to resist the government to the need to defend the President: audiences are told Obama has ‘a new residence, some are calling it a compound’, two miles from the White House, with ‘a secret army in a secret bunker’ – 30,000 people in service of a ‘globalist deep state’.<sup>21</sup> ‘They’ll probably just kill him,’ Jones says of their plans for Trump. Whoever is in the White House, the paranoia must be fuelled – without it, the InfoWars business model falls apart.

This kind of direct commerce is a big factor in the places which produce straightforwardly fake news. Leaving aside the hoaxers who set up sites for advertising clicks (covered in Chapter Six), there’s a blurred mass of fake news versus other forms of fake sites designed simply to sell products or get-rich-quick schemes. These kinds of things – ‘make \$\$\$ in your spare time’ – long pre-date the internet, when people would try to recruit unwitting targets into the schemes with leaflets and posters scattered around in streets. With the easy reach of the internet, these are everywhere: anyone who’s spent more than ten minutes online will have seen pop-ups, adverts and more promising ‘Millionaire Exposes Her Secret to Earning £127/Hr From Home’ or ‘How to Get Paid £387 Every Day Without A Job in the UK’.<sup>22</sup>

Following the link takes you to an apparent news article at a fictitious news outlet, ‘Online Wealth News’. The site is sophisticated, using data from your web browser to modify the headline so it appears the subject

of the story is from your local area – for me, it reads ‘Millionaire Mum from Islington Tells All’ – followed by an ‘As seen on: Telegraph.co.uk, FT.com, BBC’. The story, a dummy interview, then explains how ‘Leah Williams’ got rich.<sup>23</sup> The trail eventually leads to a site facilitating the trade of ultra-high-risk ‘binary options’ – a way to gamble on market movements – which have led to ‘thousands of British savers losing millions of pounds’, according to reporting by the Bureau of Investigative Journalism.<sup>24</sup>

That the binary options trading site promotes itself with a fake news article isn’t itself especially significant – a few years ago the same site may have promoted itself on a website not posing as news, or simply on an email thread. However, it does serve to warn that while some fake news is relatively harmless to its readers, misinforming them but not hitting their wallets, others could lead to serious financial loss.

But what’s most significant of all when considering the business model of bullshit is where the first grabby headline that led to the fake news story appeared – and that was at the foot of an article on a mainstream news site.<sup>25</sup>

This is the bit of the fake news ecosystem that mainstream outlets tend not to mention when they rail against fake news, demanding that Google, Facebook or national governments take action to tackle it: the fact is that the overwhelming majority of major news sites make money from advertising networks which include either outright fake news or articles of very dubious quality.

These stories come from links you often see at the foot or at the side of an article, often headlined in small font ‘sponsored links’, ‘promoted links from around the web’, ‘promoted stories’ or ‘more from the web’. The links differ from regular adverts as they match the design of the site they appear on – appearing to the casual viewer like a standard, unsponsored, related story link.

On the *Daily Mirror*, for example, sponsored links include ‘Nick Berry Was Gorgeous In “Eastenders” ... But What He Looks Like Today Is Incredible’ and ‘After Losing 170lbs Susan Boyle is Unbelievably Gorgeous’. Clicking through on the Susan Boyle link leads to a site crammed with adverts offering a generic list of ‘Incredible Celeb Weight Transformations’ – showing just one celebrity on the page, requiring a click to see who’s next, handily reloading the adverts and increasing the

site's revenues. Anyone wanting to actually see what Susan Boyle looks like would have to click seventy-one times, seeing only 'after dropping a few pounds ... Boyle looks better than ever', accompanied by a dated photo of Susan Boyle looking much as she usually does – she has not, of course, lost 170lbs.<sup>26</sup>

The highbrow Foreign Policy site offers sponsored links including 'Here's How Spoiled Barron Trump Actually Is and He's Only Ten' and 'Republican Clint Eastwood Revealed Who He Voted For, And Fans Are In Complete Disbelief' (in reality, Eastwood said he supported Trump, to no one's surprise).<sup>27</sup> *The Guardian's* sponsored headlines include 'You Will Not Believe The Cars That are Coming Out This Year', 'The Furniture Site Other Retailers Don't Want You To Know About!' and 'What "Mary Smith" From Eastenders Looks Like Today Is Jaw Dropping'.

These sponsored content boxes are operated by a small group of specialist ad agencies, with Outbrain, Taboola and Revcontent dominating the market. Why are major news outlets who are supposedly committed to fighting fake news willing to let misleading and sometimes downright false headlines appear on their own pages? The short and predictable answer is revenue: even though these links often go to anonymously registered and little-known advertisers, they pay better than regular display adverts.

'The widgets, which hawk dubious dietary supplements and a wide variety of clickbait, have become an unfortunate staple of websites – both news-focused and otherwise – thanks largely to the ad rates they pay publishers, which tend to be higher than standard banner ads,' noted the Nieman Lab.<sup>28</sup> 'They are, to many, a necessary evil for publishers desperate for more revenue from wherever they can get it.'

The same piece quotes an anonymous executive at a major online company saying: 'Publishers hate these companies but make too much money from them to stop working with them.' This is backed up with statistics: a 2016 report found forty-one of the top fifty news sites on the internet used at least one of the major 'promoted stories' advertising networks – exceptions include the *New York Times*, Slate, BuzzFeed and the *Boston Globe*. Others may find it difficult to wean themselves off this revenue, though, if Outbrain is to be believed. 'We have been told from

major, major publishers that we have become their No. 1 revenue provider,' the company's VP of product marketing told Nieman Lab.<sup>29</sup>

Publishers sit on both sides of the 'promoted stories' ecosystem, though – UK news sites including the *Express* and the *Daily Star* appear in Taboola-sponsored content boxes on other sites, driving traffic towards stories about politicians trying to block Brexit in the *Express* and an article on woolly mammoth remains in the *Daily Star*. These kinds of content promotion networks are also used on occasion by some outlets to drive traffic towards sponsored content their sales teams have made for clients – to generate the minimum viewing figures promised to clients.

The dilemma for the publishers running the adverts is set out most clearly by Matthew Ingram in *Fortune*:

If you just published a long investigative piece of journalism on an important topic like immigration or the U.S. political landscape, what message does it send when the reader gets to the bottom of that story and sees links to cheesy sites using photos of scantily-clad women and other gimmicks?<sup>30</sup>

At the minimum, any outlets campaigning or lobbying for action to be taken against fake news while running these adverts has a bullet to bite – like it or not, it's on both sides of the fight. Can they demand that others tackle fake news while not taking steps entirely within their own control do so?

The final source of financial tension comes from the way the priorities of technology companies clash with those of media companies. Facebook, like news sites, wants to keep its 1.8 billion users logging into the social network every day, and spending as long as possible on the network – the more engaged the audience, the better the revenue. Where once Facebook would link to videos, it now hosts them on its site and in its app. Facebook encourages publishers to offer their articles as 'instant articles' within its walled garden – and has the power to prioritise posts of these kinds in its algorithms, helping them reach a bigger audience. Facebook then gets a share of the revenues from such pieces – though it gives publishers at least 70 per cent of proceeds.<sup>31</sup> What publishers lose is control: they have to reshape their plans, their design and parts of their business model to suit big technology. Trickier still for traditional publishers is the uniformity of instant articles – if they homogenise the



look of content, then the source (and its credibility) becomes even less obvious to a casual reader.<sup>32</sup>

Facebook's is not an ecosystem that tends to foster a culture of truth. Pages on the site benefit from trying to build huge communities of fans – and many of the largest work by lifting content from elsewhere, often without payment or attribution, and framing it to provoke nostalgia or humour. Some of the most egregious examples include pages stealing photographs of disabled children 'to gain followers and get shares', to the distress of their mothers.<sup>33</sup> Like it or not, publishers on Facebook are routinely competing for attention in an environment where small and unscrupulous players trade off sentiment and stolen content with little or no regard to truth – and where reputation matters less than you might hope.

Other decisions of the major technology companies can have negative consequences for the post-truth era. Google has constantly refined its search results page, putting more and more information directly on its own site where once it just offered a list of search results – resulting in the risk of falling traffic to outlets which once traded off searches asking, for example, who played a particular character in a TV show, or how old a politician is.

These algorithmic answers can fail badly: Google seems to take minimal steps to prioritise established outlets over niche ones. Peter Shulman, an associate professor at Case Western Reserve, set out an example that he'd encountered after a bad class.<sup>34</sup> He recounted how a student had announced in his history class that at least five Presidents had been members of the Ku Klux Klan. Searching 'Presidents in the Klan' showed in Google's info box – offering supposedly reliable information from sites Google in which has a high degree of trust – a list of Presidents including McKinley, Wilson, Harding and Truman, cited to a Nigerian news outlet and not backed up with any evidence. Just as Facebook's trending news feature backfired when it was automated (see [Chapter Seven](#)), Google's information box does the same. The technology companies end up, perhaps accidentally, competing with news outlets on their own turf and adding to the misinformation mess.

The biggest economic issue for news outlets with Facebook and Google is that the technology companies on which they rely to reach an audience are also their competitors. In many instances, it's easier and

more effective for brands to spend their digital advertising budget directly with Facebook and Google than it is to spend it with news outlets: the sites have such scale that they have the ability to make adverts that work well and are easy to target. That's hard for news outlets to compete with.

The dominance of Facebook and Google in digital advertising provoked the UK's News Media Association – the trade body of UK newspapers – to call on the government to take action. In a briefing to ministers, it warned that 'the online news environment is characterised by aggregation of news stories by third party players who repackage, serve, link to and monetise that content', adding that 'Google dominates these activities in search and Facebook dominates in social'. It warned: 'Significant value is being captured by companies who do not invest in original journalism at the expense of those who do.'<sup>35</sup> Google and Facebook make up around 75 per cent of the US digital display advertising market and more than 50 per cent of the UK market – and the share is increasing.<sup>36</sup>

None of this is to try to paint either Google or Facebook as pantomime villains. Among other journalism projects, Google funds the Digital News Initiative, helping fund efforts to 'support high quality journalism' and 'encourage a more sustainable news ecosystem' in Europe, which has funded €24 million of projects from newsrooms big and small.<sup>37</sup> Facebook has launched a journalism project looking to work with newsrooms to develop new products, support local news and help build new business models.<sup>38</sup> Cynics may say the internet giants are trying to act on the problem so as to stave off attempts by governments to add regulations or introduce levies to fund news outlets – but senior figures in both companies are aware that their sites rely on high-quality content, and so it's in their interest that the media survives. But these relatively small efforts don't change the current underlying reality. On the internet, it's not creating original content that pays off – it's having the scale and the ability to exploit it.

These are the various business models of bullshit, and when they work together in tandem they're corrosive to mainstream outlets and to trust. News organisations face a funding crunch from all sides: ad revenue is low and squeezed by consumer backlash, the shift to mobile and the dominance of search and social media giants. The resulting

newsroom economics encourage outlets to run dubious viral stories and staged videos alongside the articles we're meant to value as high-quality journalism – and surrounding both are sponsored links to questionable advertisers and, on occasion, outright fake news.

The mainstream media is facing simultaneous crises in trust and in business models – and the two pressures are pushing many outlets in opposing directions. Elsewhere in the bullshit ecosystem, outright fakery can pay off, either through advertising – it's so cheap to make that digital adverts pay enough to make it profitable – or by promoting dubious products and services. Fringe sites can pay off by feeding the very paranoia and division the mainstream would like to tackle. Looking at the financial models of the internet might not be a heartening experience, but understanding why places act as they do is the first step to working out how to tackle it.

You might notice that we haven't talked much about sites trading off subscriptions or other business models this chapter, and that's because for a few sites that can generate enough revenue in this way, this model can serve as a counterweight to producing mass-market, low-cost news. There are still well-resourced newsrooms which work in this way – but if this alone was enough to tackle bullshit, the crisis would not have arisen in the first place. That takes us to the question of our next chapter: why isn't it?

## CHAPTER ELEVEN

# A BULLSHIT CULTURE

The previous chapters of this section have looked at why cash-strapped outlets may spread bullshit, and why we might buy into it, but we haven't looked much into why the remaining well-funded and high-quality outlets are failing to stop the tide – until now. First, though, this chapter risks creating a misconception: that there are outlets which only produce bullshit and outlets which only produce quality news. In reality, almost everyone produces some proportion of both.

Even without any bad intent, we've seen how news outlets run headlines that can mislead, struggle to challenge dubious claims from Trump, Brexit, or other campaigns, and rush out viral stories. It would be a nonsense, though, to suggest that places like the *Daily Mail* do no original or investigative journalism – even if it doesn't suit our political agendas. The *Mail* was the driving force in revealing the people behind the killing of a black teenager, Stephen Lawrence, in 1993 and, in a campaign and coverage spanning more than a decade, it was the *Mail* who spurred inquiries into police failings, uncovered police spying operations on Lawrence's family as they campaigned for justice, and reported on the trials of those who were ultimately found guilty of the murder.<sup>1</sup> This chapter is not intended as an example of 'good' outlets versus 'bad' outlets.

However, some sections of the media have fewer financial incentives to publish bullshit than others. The BBC, for example, receives its funding from an annual license fee of just under £150 per UK household, showing no adverts within the UK. Another handful of newspapers and magazines – particularly but not exclusively in finance – have built up substantial revenue from digital subscriptions, filling the gap left by falling advertising sales. These produce no shortage of high-quality journalism – but that's not what we're here to discuss. Our key question

is: given that these outlets don't have the same kind of pressures as many others in the industry, why can't they stem the tide? Partly, of course, it's that asking just a handful of outlets – even huge and powerful ones – to shape the rest of the industry is expecting a lot, but there are also aspects in how these places produce their content that may help bullshit spread unchecked.

The subscription business model is a pretty simple one, and returns to a business model that print newspapers are familiar with: readers pay for their news. The challenge with trying this online is that there are dozens of news outlets immediately available for free, which are able to re-write versions of any original story within minutes. For a successful subscription model, news sites need to be able to attract a dedicated and engaged audience with an attachment to the brand. This favours sites which produce distinctive, high-quality and original content – a good counterweight to other business models which favour quick turnaround and questionable stories.

The *New York Times* has exploited this model more successfully than most, and benefited from a huge bump in interest and sign-ups during the rise of Donald Trump. Digital subscriptions at the *NYT* rose by more than 500,000, to 1.6 million – for all Trump likes to call the paper 'failing', it's doing better than most online.<sup>2</sup> The *Washington Post*, *Wall Street Journal* and *Financial Times* are all pursuing similar strategies, all relatively successfully. This might now seem like the easy and obvious fix to the news ecosystem, but the reality is a bit more complex. The *New York Times* is the second largest newspaper website in the world, and had a huge boost in digital subscriptions thanks in large part to a huge and unique running news story – but even this failed to compensate for falling print revenues. Overall revenue at the *NYT* fell around 2 per cent in 2016. For the very biggest outlets, subscriptions are helping make falling print revenues survivable, but they're still facing decline. For outlets without those huge audiences and large newsrooms to produce enough original content to encourage large numbers of subscriptions, the revenue will provide only a small fraction of the total funding mix.

This is perhaps illustrated by *The Guardian*, another of the world's biggest online newspapers, which offers membership as an alternative to subscription. Readers are urged to support the paper: 'More people are reading the Guardian than ever. But far fewer are paying for it. And

advertising revenues are falling fast. So you can see why we need to ask for your help,' says a call to action at the foot of many stories. Members are asked to contribute at least £5 a month to secure *The Guardian's* future. The initiative has had some success: fuelled by Brexit and Trump, membership in 2016 rose from 15,000 to 200,000, implying up to around £15 million a year in total revenue from the initiative. That's no bad thing, but *The Guardian's* print revenues are around £120 million, digital makes around £80 million, and the paper loses around £60 million a year.<sup>3</sup> The membership money helps, but is dwarfed by advertising – leaving *The Guardian* aiming to increase subscriptions fivefold by 2019, a vastly ambitious target.<sup>4</sup>

For places that opt for a strict paywall which bars nonsubscribers from reading their journalism, the model can also serve against the interests of countering low-quality information across the internet with better material. If the best outlets all hide their content behind paywalls, the mass audience doesn't see it, as Mathew Ingram set out in 2012:<sup>5</sup>

'One thing is for sure: while metered paywalls may help to save specific newspapers, they aren't going to save journalism – if by "journalism" you mean the process of informing as many people as you can about news that is important to their lives,' he argued. 'By definition, paywalls restrict the reach of a newspaper and that has very real implications for the social aspects of journalism as opposed to the business of one specific newspaper. In the end, we have to answer the question: What is journalism for?'

Paywalls and subscriptions work well for some sites, but they aren't a saviour of journalism, and they risk hiding the best-quality material away from the mass public conversation they might hope to influence. But many of the barriers that prevent quality news outlets having the impact they would hope for stem from factors inside the culture of journalism, rather than its economics.

One aspect of this is the culture of objectivity held sacred by UK broadcasters (who are required to behave this way by law) and most US newspapers (who aren't). This rule holds that news outlets and their journalists should act as if they hold no opinions of their own, instead seeking to reflect and question all sides of a debate while never taking a side of their own. The rationale is that this creates outlets that are trusted

by everyone, giving the public the facts they need to debate issues and make their own decisions.

In a world where not everyone follows these values, though, this can look sterile and anodyne: partisan outlets and politicians speak passionately, launching attacks in clear and human terms – and outlets respond in a stilted on-the-one-hand, on-the-other-hand kind of way, which may sound authoritative but is far less relatable. The habit of this form of journalism to question all sides carries other risks, too. Outlets often confuse objectivity with balance: if one political campaign makes an untrue or niche claim, there's a temptation not to judge it on its merits, but instead to leave it for opponents to challenge – not reporting the facts, but instead reporting the argument. The other risk reflects the old Margaret Thatcher quote: 'Standing in the middle of the road is very dangerous; you get knocked down by the traffic from both sides.' Outlets start to look like an enemy to partisans on both sides, rather than a trusted voice for either.

A final and substantial risk of this 'objective' style of journalism is that it becomes, largely accidentally, a reflection of the view of establishment sources – perhaps most egregiously in the run-up to the 2003 war on Iraq. In an admirable confession of its mistakes, the *New York Times* acknowledged that it had been over-reliant on administration sources and Iraqi defectors, both of whom were determined to sell a particular narrative on Iraq having and being ready to use weapons of mass destruction.<sup>6</sup> And while the BBC faces relentless (and often unfair) accusations of bias from all sides, there's reasonable evidence to suggest that if it does have a bias, it leans towards the status quo and the current government of the day<sup>7</sup> – simply by virtue of the fact that they create policy and statements, and so by their actions shape much of the resulting coverage.

The culture of objectivity has been challenged by New York journalism professor Jay Rosen for more than a decade. Rosen, who refers to it as the 'view from nowhere', says it came to dominate quality journalism culture in the twentieth century as the craft professionalised and displaced the campaigners and pamphleteers who became the first newspapermen. Thus partisan or campaigning journalism is referred to by Rosen as 'Old Testament' journalism, and 'objective' reporting is 'New Testament'. But, as Rosen set out in a 2013 speech, New Testament

journalism may not be the right tool for a post-truth era marked by collapsing trust in public institutions and experts, and a collapse of the old media business models.

‘In new testament journalism, the media’s financial security is the norm, made possible by high barriers to entry and large capital costs required to deliver news,’ he said.<sup>8</sup>

The new testament style is risk-averse because the news delivery franchise is so valuable. The mission is not to move public opinion but to maintain trust or, to put it another way, to protect the brand. Audiences tend to be stable. The authorities learn to regularize their relationship with the journalists. Professionalism in journalism and broadcasting interlocks well with professionalism in politics and other knowledge fields. Thus, the rolodex of reliable experts.

This argument is not a no-brainer: plenty of courageous journalism comes from outlets with the culture of objectivity, and its proponents say it helps make sure reporting is trusted as fair and accurate. But there’s a counterpoint: this model takes passionate and engaged reporters who become experts on their fields of coverage year after year but who spend those decades muzzled from giving their own assessments or opinions, risking appearing to be automatons. A final question: as the internet brings us closer together, do we still think the audience believe that reporters at the BBC or the *New York Times* have absolutely no opinions of their own? Social media and comment feeds across the internet suggest otherwise.

There’s another risk factor among the outlets that consider themselves to be ‘quality’ journalism – that writers become more concerned with the opinion of other journalists than with that of the audience. This is a concern that dates back at least to the 1970s. ‘Journalists write for other journalists, the people they have lunch with rather than the reader,’ an unnamed journalist said at the time, leading one academic to conclude: ‘Their image of the audience is hazy and unimportant ... they care primarily about the reaction of the editor and their fellow-reporters.’<sup>9</sup>

This tendency is exacerbated in US awards culture, where the most prestigious prizes favour journalism written at great length – often 10,000 words or more and presented in a dense, discursive fashion. These pieces are often, for a journalist like me, a joy to read and are produced over the course of months. They are often the very best articles their outlets produce – but it’s not hard to argue that they’re not as accessible or impactful as they could be. UK journalism culture favours relatively



short, punchy write-ups of investigations with a sharp top news line. The US format of long stories, often with their top line buried a dozen or more paragraphs in, seem designed for the devoted reader but leave the mass audience behind. An open question: if the Pulitzers, Polks, Scripps Howard and other awards were scrapped tomorrow, would newsrooms present their best stories differently?

Sometimes some of the ways mainstream outlets present stories are restricted by laws or regulations, in a way that journalists understand but which audiences may not. The UK's environment gives a good outline of how this can happen. Reforms to the UK's libel laws made it easier for outlets to run public service investigative journalism. Previously, outlets would have to be able to prove every claim in a story as true – a hugely difficult bar when journalists were trying to find out complex stories with limited information and contradictory sources.

The shift in the law lets news outlets defend a story on the basis that it was fair, based on information which they had good reason to believe was accurate, and where everyone concerned had a chance to respond to the information and allegations therein. This change lets far more stories make their way to the public, but comes with a sting: the journalists concerned can't adopt the allegation – the piece has to be approached in a neutral, 'view from nowhere', kind of way – and any correspondence and memos sent during the reporting of it must be similarly free of bias: if the journalists concerned looked enthusiastic at the prospect of discrediting the politician or company they were writing about, they can't appeal to this libel defence. These restrictions account for the reporting style of quite a lot of modern UK investigative journalism: phrasing is often careful; the newspaper or TV show carefully avoids (in its news section at least) expressing an opinion on its revelations; and any comment on them is left to politicians or other outlets. The article also typically contains quite extensive denials or explanations from the targets of a story – which may be good for fairness, but risks leaving the article looking confusing to a casual reader.

For broadcasters, these obligations are even more onerous, as TV news and channels are bound to comply with rules set by the regulator Ofcom, which requires coverage to show 'due impartiality' and – especially in the run-up to elections – give fair, measured amounts of airtime to all major political parties.<sup>10</sup> The rules aren't intended to be too

rigid – they don't require equal time to be given to all opinions, for example – but they do risk leaving broadcasters leaning towards caution. This presents TV news with a dilemma in the run-up to elections or referenda, when there is an obligation to give airtime to, and reflect the views of, all major parties and campaigns. Navigating the lines of those legal obligations, especially if a candidate or campaign says things that are untrue, is a substantial challenge – and often a source of frustration to those watching from the other side of the political divide. To take Vote Leave's £350 million-a-week NHS claim (covered at length in Chapter Two) as an example: broadcasters are unable not to talk about this, as it was a central and daily plank of an official campaign's platform during the campaign. But calling it out as inaccurate on every mention would not feel impartial to many listeners. The result tends to be that longer broadcasts, like half-hour current affairs shows, note that the claim is disputed, but shorter broadcasts and clips, which hit a mass audience, often let dubious claims run unchallenged. Keeping broadcast impartial can help keep people of all political views agreed on some central shared view of reality – but does it also risk letting bullshit pass unchecked?

This reminder of restrictions on the media from law and from regulation should also serve as an antidote to politicians making accusations about how outlets use anonymous sources. US newspapers haven't opted to sign up to (or even create) industry-wide regulators, but they have internal ombudsmen and other safeguards, and they remain subject to libel laws and other restrictions. This should cast attacks from Donald Trump against the media using 'anonymous sources' in a new light; the President has suggested that news outlets should only use on-the-record sources in their stories.<sup>11</sup> Many government staff have no permission to talk to the media without explicit permission, so the only way people can hear a view other than the administration's official line is if off-the-record sources are allowed. Donald Trump himself has used anonymous sources to claim that Obama faked his birth certificate and applied to university as a foreign student, that 'Clinton's hacked emails' led to Iran's execution of a nuclear scientist, and more.<sup>12</sup> Anonymous sources do risk making it possible for politicians and outlets to invent 'news' from nowhere – but this is less true for mainstream outlets than for others. If a UK publication runs a big story based on an anonymous source, it is saying it's confident its anonymous source would – if

necessary – back up their claims in court or in front of the regulator. There are reasons to be sceptical of anonymous sources, but for most mainstream outlets, they're not wholly invented when they're making major claims – though former tabloid reporters have admitted fabricating quotes on stories where it was unlikely to lead to lawsuits.<sup>13</sup>

And then there's climate change, and how the media covers it – an issue with the ability to reshape the world, leading to tens of millions of extra deaths by the end of the century,<sup>14</sup> and which would require major action by government, business and citizens to tackle. In terms of expert opinion, more than 97 per cent of published climate scientists agree climate change is real and is caused by the activities of humans.<sup>15</sup> The picture is very different in politics, where the issue has become a cause of controversy between the religious right, pro-business right and liberals, especially in the USA.

The media is more used to covering issues through debate and argument, and many flagship news programmes rely on this format to produce segments on issues. This means even well-intentioned programmes looking to produce good-quality news can be left covering climate change badly, as a 2014 parliamentary report into coverage of the issue found.<sup>16</sup> 'The producers of the recent Today Programme piece on the new IPCC report tried, we are told, more than a dozen qualified climate scientists willing to give an opposing view but could not find a single one (a hint, perhaps, that there is indeed a scientific consensus on global warming),' said the report, quoting Professor Steve Jones, author of a report for the BBC Trust on climate change coverage. 'Instead, they gave equal time to a well-known expert and to [an] Australian retired geologist with no background in the field: in my view a classic [case] of "false balance".'

Another issue across newspapers and broadcasters is that news desks aren't really set up to cover something like climate change, a serious issue which develops slowly and unsurprisingly, as James Randerson – then an environmental journalist at *The Guardian* – told the committee, speaking of the 'tendency for news desks to like things that are new and surprising and favour the underdog. A general issue with science reporting is that mavericks tend to get more coverage than perhaps they deserve.'<sup>17</sup>

Science journalists have been covering topics where false or fringe views challenge established and well-evidenced mainstream consensus for far longer than most political journalists. The media were among the most enthusiastic boosters of the anti-vaccination movement, but were also instrumental in exposing the scientific fraud at its core. Some media outlets enthusiastically boost climate conspiracy theories while others battle to cover the issue fairly and accurately while also reflecting the political rows and debates around it. If political discourse is indeed facing a new high tide of bullshit, talking to science reporters and editors might serve as a good way to find what will and won't work – they've been in the trenches of this battle for years.

The instinctive response of many US outlets to the rise of Trump, the fall of trust in their output, and the polarisation of US discourse seems to be promising to double-down, committing to doing even more of the type of reporting they've always been doing. News outlets have reported how Trump and his White House spokesman told an average of around four untruths a day in the early days of the administration.<sup>18</sup> And yet most of the battles on the ground between White House reporters and the administration are to keep open access to the briefings in which they're told these untruths and half-truths. The fight is to maintain the illusion of business-as-usual, a status quo which arguably no longer exists.<sup>19</sup>

It's clearly a temptation for some in the mainstream to use the rise of Trump as an I-told-you-so, and some have taken the bait. 'This is what happens when we stop paying for quality journalism,' says Asha Dornfest in a Medium post warning that 'authoritarian governments choke off access to information. By abandoning quality news sources, we've already done much of that work ourselves.'<sup>20</sup> Dornfest tells her readers to stop getting news from Facebook, start paying for local and national newspaper subscriptions, and pay for public radio. Her calls to action may be helpful, but they're also exactly what mainstream journalists would like to hear: if the public are divided and believing in bullshit, it's their own fault for not reading us any more. Sure, why not. But how many Trump supporters are going to read that post and return to elite media on bended knee?

Others have been more questioning: veteran *Washington Post* report Dan Balz, who's had more than 1,000 front-page stories, warned against the rise of polling-driven and model-driven news – where every headline

was based on a shift in a forecast in elaborate models aimed to predict the election. ‘We have all come to rely on data which turned out not to be specifically accurate, whether it was polling, predictive modelling, almost anything you looked at in one way or another broke down in the end,’ he told CNN.<sup>21</sup> ‘I think we have to re-examine a lot of the ways we go about doing journalism.’

Others are blunter still. Margaret Sullivan, the *Washington Post* columnist and respected former public editor of the *New York Times*, warned that Trump’s presidency would be ‘hellish’ for journalists, that he would be a ‘gaslighter in chief’ who would ‘relentlessly manipulate’ the media and ‘punish journalists for doing their jobs’. Her conclusion was less business-as-usual than some of her colleagues. ‘Journalists are in for the fight of their lives. And they are going to have to be better than ever before, just to do their jobs,’ she said.<sup>22</sup> ‘They will need to work together, be prepared for legal persecution, toughen up for punishing attacks and figure out new ways to uncover and present the truth.’

This hope that journalists may unite to respond to a new kind of President may prove as futile as hoping the public will spontaneously regain their trust in traditional journalism. In a somewhat despairing post, Russian journalist Alexey Kovalev warned his ‘doomed colleagues in the American media’ that they ‘can’t hurt this man with facts and reason’, urged them not to expect any camaraderie as some outlets would always go along with Trump’s conferences, and advised them to ‘expect a lot of sycophancy and softballs’ from colleagues towards the President.<sup>23</sup>

We have to hope Kovalev’s counsel is one of despair, and too far on the negative side – but other responses are evidently leading too far towards business-as-usual. While giving up won’t help tackle the rise of bullshit, we do have to acknowledge that while messages like ‘let’s keep on doing our jobs, with more energy and determination than ever’ might be in journalists’ comfort zones, there’s no reason to believe they will work either. So far, the principal weapon against bullshit in the mainstream press has been the rise of the fact-checker. In our next and final section, we’re going to look at how they’ve fared – and then look into what else might help get us through the post-truth era.