

THE LAST DAYS OF REALITY

Mark Pesce

ON 1 MAY, readers of the *Australian* woke up to a story that seemed more appropriate for the pages of a dystopian sci-fi novel: ‘Facebook is using sophisticated algorithms to identify and exploit Australians as young as 14, by allowing advertisers to target them at their most vulnerable, including when they feel “worthless” and “insecure”, secret internal documents reveal.’ A 23-page Facebook document seen by the *Australian* marked ‘Confidential: Internal only’ and dated 2017, outlines how the social network can target ‘moments when young people need a confidence boost’ in pinpoint detail.

By monitoring posts, pictures, interactions and internet activity in real time, Facebook can work out when young people feel ‘stressed’, ‘defeated’, ‘overwhelmed’, ‘anxious’, ‘nervous’, ‘stupid’, ‘silly’, ‘useless’ and a ‘failure’, the document states. That evening, the panel on Ten’s nightly current affairs show *The Project* reflected a national mood of revulsion and anger at the exploitation and targeting of a population that was both unworldly and emotionally vulnerable. It might continuously proclaim its mission to ‘help everyone share’, but Facebook had been revealed in a new and inimical light.

Facebook quickly issued a press release: the social media giant would work to ‘understand the process failure and improve our oversight’. The relief was temporary. It soon came to light that the original document had been written by two of Facebook’s senior Australian executives. It seemed deliberate emotional monitoring and exploitation might be central to the entire rationale of Facebook. ‘In its statement to the *Australian*, Facebook refused to disclose if the practice exists elsewhere.’

Facebook stuck to its denials. In a follow-up article the next day the *Australian* reported that Facebook ‘tried to downplay the significance of the disclosure by saying it did not use personal information about children to target ads despite the data’s inclusion in a detailed sales pitch’.

At the time of writing, little more has been reported. Presumably everything is going on at Facebook as before this revelation, with no indication that any of its business practices have changed. Never transparent about how it applies the information supplied by its two billion monthly users, Facebook had been caught red-handed: exploiting the weak spots of teenagers in their moments of greatest vulnerability, watching, waiting



then delivering a targeted message at the moment of maximum impact. In another context, this could be read as something akin to torture or brainwashing. For Facebook, it was just business as usual.

The revelation forces us to confront some unpleasant thoughts about how the world works in 2017, and where things appear to be headed. As problematic as Facebook has become, it represents only one component of a much broader shift into a new human connectivity that is both omnipresent (consider the smartphone) and hypermediated—passing through and massaged by layer upon layer of machinery carefully hidden from view. The upshot is that it's becoming increasingly difficult to determine what in our interactions is simply human and what is machine-generated. It is becoming difficult to know what is real.

Before the agents of this new unreality finish this first phase of their work and then disappear completely from view to complete it, we have a brief opportunity to identify and catalogue the processes shaping our drift to a new world in which reality is both relative and carefully constructed by others, for their ends. Any catalogue must include at least these four items:

- the monetisation of propaganda as 'fake news';
- the use of machine learning to develop user profiles accurately measuring and modelling our emotional states;
- the rise of neuromarketing, targeting highly tailored messages that nudge us to act in ways serving the ends of others;
- a new technology, 'augmented reality', which will push us to sever all links with the evidence of our senses.

These four strands meet in a tangled nexus of technology and capital, each amplifying the others almost beyond imagining, warping the fabric of reality, offering

attractions so alluring many will find it difficult to resist, framing an emerging world that can only be termed 'post-real'. As we blithely approach the last days of reality, technological progress engulfs every way of knowing, accelerating into epistemological crisis. The real world is about to disappear. It all begins with fake news.

1. Facebook and the monetisation of propaganda

We'll tell you any shit you want to hear. We deal in illusions, man. None of it is true! But you people sit there day after day, night after night, all ages, colors, creeds. We're all you know. You're beginning to believe the illusions we're spinning here.

—Paddy Chayefsky, *Network*, 1976

In the closing days of the 2016 US presidential campaign another story—this time from BuzzFeed—seemed ripped from the pages of the same sci-fi novel. In the Republic of Macedonia, a network of teenagers published the very worst conspiracy theories about Hillary Clinton to a willing audience of Donald Trump's Facebook supporters. Each then shared these salacious claims across their own networks, who shared them even more widely, in a 'viral' assault on the Clinton campaign.

Over the past year, the Macedonian town of Veles (population 45,000) has experienced a digital gold rush as locals launched at least 140 US politics websites. These sites have American-sounding domain names such as World-Politicus.com, TrumpVision365.com, USConservativeToday.com, Donald-TrumpNews.co, and USADailyPolitics.com. They almost all publish aggressively pro-Trump content aimed at conservatives and Trump supporters in the US.

The young Macedonians who run these sites say they don't care about Donald Trump ... Several teens and young men who run these sites told BuzzFeed News that they learned the best way to generate traffic is to get their politics stories to spread on Facebook—and the best way to generate shares on Facebook is to publish sensationalist and often false content that caters to Trump supporters.

Those Macedonian teenagers had no particular political stake in the American election; they were in it for the money. The stories they found (or invented) would earn them hundreds to thousands of dollars as each—with the Google ads they carried—multiplied through Facebook.

The fake news stories floated past as jetsam on Facebook's 'newsfeed', that continuous stream of shared content drawn from a user's Facebook's contacts, a stream generated by everything everyone else posts or shares. A decade ago that newsfeed had a raw, unfiltered quality, the notion that everyone was doing everything, but as Facebook has matured it has engaged increasingly opaque 'algorithms' to curate (or censor) the newsfeed, producing something that feels much more comfortable and familiar.

This seems like a useful feature to have, but the taming of the newsfeed comes with a consequence: Facebook's billions of users compose their world view from what flows through their feeds. Consider the number of people on public transport—or any public place—staring into their smartphones, reviewing their feeds, marvelling at the doings of their friends, reading articles posted by family members, sharing video clips or the latest celebrity outrages. It's an activity now so routine we ignore its omnipresence.

Curating that newsfeed shapes what Facebook's users learn about the world.

Some of that content is controlled by the user's 'likes', but a larger part is derived from Facebook's deep analysis of a user's behaviour. Facebook uses 'cookies' (invisible bits of data hidden within a user's web browser) to track the behaviour of its users even when they're not on the Facebook site—and even when they're not users of Facebook. Facebook knows where its users spend time on the web, and how much time they spend there. All of that allows Facebook to tailor a newsfeed to echo the interests of each user. There's no magic to it, beyond endless surveillance.

Yet it *feels* magic, and that feeling keeps billions glued to their feeds. As Facebook felt its way towards a potent relationship between curating and 'stickiness'—how long users spend on the site per visit, per day, and per month—its approach changed, from filtering the torrent of updates towards a newsfeed that produced a positive reaction in the user, bringing them back again and again.

In 2014 an academic paper revealed that Facebook not only understood the power of the newsfeed to shape a user's mood, it had authorised a research experiment designed to manipulate users' moods by changing the emotional bias of their newsfeeds. As reported in the *Guardian*, Facebook worked with researchers from Cornell University, filtering the newsfeeds of nearly 700,000 Facebook users. One test reduced the amount of 'positive emotional content' users received from their Facebook friends. This resulted in fewer posts of positive emotional content from those users. Conversely, a second test reduced the amount of 'negative emotional content' users saw, and those users posted less negative emotional content themselves.

Facebook (and the researchers) characterised this as a phenomenon of 'emotional contagion'—downer friends can bring you down. Everyone else, however, noted that this experiment in emotional manipulation

occurred without the consent of these hundreds of thousands of users. Facebook had emotionally manipulated its users, and those users never knew—would never have known, had Facebook not published its findings.

Facebook claimed it was only an experiment and that routine newsfeed manipulations were not planned. Yet in a clear violation of research ethics Facebook had allowed researchers freely to experiment on its users, while withholding any knowledge of that experiment from those users. Why would they want to hide this experiment from the public? Most likely because those researchers were empirically demonstrating that Facebook had constructed a machine to manipulate mood, on a planetary scale.

Facebook claimed the experiment had somehow been approved without its knowledge, and—despite being found in possession of pervasive machinery to manipulate mood—promised not to use that power for good or ill, a series of denials and retractions that looks curiously similar to statements made by Facebook after the May 2017 revelations of its emotional manipulation of vulnerable teenagers.

What is clear is that Facebook has the power to sway the moods of billions of users. Feed people a steady diet of playful puppy videos and they're likely to be in a happier mood than people fed images of war. Over the last two years, that capacity to manage mood has been monetised through the sharing of fake news and political feeds atuned to reader preference: you can also make people happy by confirming their biases.

We all like to believe we're in the right, and when we get some sign from the universe at large that we are correct, we feel better about ourselves. That's how the curated newsfeed became wedded to the world of profitable propaganda. Those Macedonian

teenagers provided the raw fodder to help Hillary's Haters feel fully justified in their feelings, so those stories travelled far, quickly finding their way into the newsfeeds of those people Facebook knew were most likely to react positively to that kind of story in their newsfeed, earning profits for the posters, who then went looking for more salacious material, posting more and earning more.

This feedback loop of curating newsfeeds, fake news and confirmation bias makes it nearly impossible to break out of an accelerating cycle, creating a 'reality' that has very little to do with the facts on the ground. Someone who thinks Hillary is horrible will tend to see more of that in their Facebook feed, while someone who thinks Trump is horrible will see more of that. That's a serious problem in itself—leading to a collapse of consensus in the political sphere—and one now amplified by the direct economic interest Facebook has in providing a sticky newsfeed that leans towards a user's confirmation bias.

That has created a 'reality trap' for Facebook: if it stops confirming user biases, its users will drift away. If a newsfeed angers users with 'fake news' contradicting a user's world view, those users will pay less attention to Facebook—meaning Facebook earns less revenue from its advertisers. This locks Facebook and its users into an accelerating cycle of untruth.

Facebook has given users what they want in order to keep them from looking away. Getting what they want makes those users less tolerant of anything they don't agree with. If Facebook suddenly stopped giving users what they wanted, they would see a huge drop-off in usage. Now Facebook has to do everything in its power to ensure its users get exactly the newsfeed that will keep them happy—and thus sticky. A new technology—machine learning—gives Facebook the tool it needs to meet that challenge.

2. Machine learning and surveillance of the self

Last year, it was still quite human-like when it played, but this year it became a God.

—Ke Jie, May 2017

A few weeks after Facebook had been outed as teen manipulators, ‘AlphaGo’, a program written by a Google research team based in Britain, took on Ke Jie, a 19-year-old Chinese ‘9-dan’ (grand master) of Go, the ancient pebble game thought to be the most difficult and subtle of all board games. AlphaGo defeated humanity’s best player in all three matches, an event so disturbing that Chinese officials blocked live video coverage of the man-versus-machine contest at the last minute. AlphaGo had beaten its first 9-dan player in March 2016, when the program was barely 18 months old. Now it was indomitable. How did a computer program get to be so good at something so hard so quickly?

The last several years have seen huge leaps in ‘machine learning’, the capacity for computer programs to train themselves to get better at achieving their goals. AlphaGo began with little more than a basic knowledge of the rules of Go. Researchers added simple programming to help AlphaGo weigh the value of possible strategies as it played a match. Then AlphaGo played a range of human players, almost always losing to them. Yet AlphaGo learned from each of these losses. It began to avoid errors it had made previously, falling less frequently into the many possible mistakes, pitfalls and traps of Go.

Although AlphaGo could perform tens of billions of computations a second, it took thousands of games, played against very average human players, before it could beat a human reliably. Machine learning takes time. At that point Google’s researchers set AlphaGo against itself. AlphaGo knew

enough Go to play a passable game, so setting it against itself allowed AlphaGo, across hundreds of thousands of matches, to deepen its strategic sense of the game. The program took on a 2-dan player in October 2015 and won those matches before defeating that first 9-dan player in March 2016. All along, AlphaGo had been learning, from itself and from the players it had defeated. It has already played more games of Go than any human in history, using every game to learn how to be a better player in the next match.

The techniques employed by AlphaGo (and numerous similar projects) are becoming widespread across the entire tech sector. Cars equipped with the right sensors and computers can ‘observe’ a human driver over a period of time, eventually learning how to drive autonomously. The billions of words fed into Apple’s Siri, Google’s Assistant and Amazon’s Alexa teach them how to parse the subtleties of human language and follow our shades of meaning.

Facebook uses machine learning to predict what makes its users respond positively. In 2016 Facebook announced an effort to turn the entire platform into a gigantic machine learning system, with billions of ‘agents’ (each individual interactive instance) learning from the interactions of every one of its users, all the time. Just as AlphaGo gets better by playing human players, Facebook will get better at serving up just what makes its users stay stuck to the site. Watching and learning from every user interaction—every click and every scroll and every link followed to an outside website—Facebook builds a ‘profile’ to model the behaviour and interests of its users.

Once constructed, this profile can be queried on how a user might respond to a particular post or article. In essence, Facebook has constructed a simulation of each of its users through surveillance then uses that simulation to filter away the content that might drive users off the site.

Where the profile gets it wrong—and user engagement with Facebook observably declines—it learns from its mistake, just like AlphaGo, and will avoid repeating that behaviour. Conversely, where the agent sees its decisions lead to increased stickiness, it remembers that too, using what it has learned to finetune a user’s newsfeed.

Billions of these agents, continuously learning from every click and scroll (because Facebook uses data ‘cookies’, remember, to track its users across every website they visit anywhere on the web), gives Facebook a depth of insight into its users the East German Stasi would have envied. In a recent interview, design thinker Andy Polaine put it succinctly: ‘These systems, they observe everything. They *see* the patterns. They know me better than I know myself.’

Facebook hides these patterns from its users, and uses them without compunction. It doesn’t make money by liberating its users from any neurotic or destructive patterns of behaviour it observes, records and models. Instead, Facebook simply monetises them.

3. The persuadables

Goodness comes from within. Goodness is chosen. When a man cannot choose, he ceases to be a man.

—Anthony Burgess, *A Clockwork Orange*, 1962

Stand in front of a maths classroom, lecturing students on how numerous studies have shown that women just don’t have the brains for maths—then test them. You’ll find that the women in the class will measurably underperform compared to a class that hasn’t had that lecture. Never mind that the research is fabricated. Over the last few years researchers have shown that much of what anyone believes at any given moment can be charted as a response to some stimulus in their informational environment.

People appear to be far more open to persuasion than they’d like to believe: the right trigger, at the right moment, will change the way we behave.

Where these triggers occur randomly—catching sight of a starving child on the telly, for example—they’re just a normal part of a person’s mental and emotional range. But increasing evidence exists that these triggers have been utilised in the political sphere, via Facebook, and targeted to change behaviour in specific ways. The growing web of evidence—reported in detail in a May 2017 article in the *Observer*—centres on a firm named Cambridge Analytica. Founded by a hedge fund billionaire Robert Mercer, a world-class computer scientist, Cambridge Analytica harvests every bit of data publicly available about a voter, cross-references it to generate a profile of that voter, then—depending on the specifics of that profile—purchases Facebook advertising targeted at the voter and designed to ‘trigger’ that voter into making a desired voting choice.

Does the Cambridge Analytica method work? Robert Mercer has two friends who have been high-profile clients: former UKIP leader Nigel Farage and former Trump adviser Steve Bannon, who is also a former vice-president of the firm. Could it be that the against-the-odds successes of Brexit and Trump are owed not only to an anti-establishment revolt, but also to Cambridge Analytica’s capacities to drive ‘persuadable’ voters to their desired outcomes?

As we have seen, Facebook has been caught out using ‘emotional surveillance’, pushing vulnerable teenagers to purchase products to improve their self-image. Cambridge Analytica uses a variation on that same technique, targeting voters at their most persuadable moment, delivering (via their Facebook newsfeed) exactly the stimulus needed to get them to tick the box for Trump or Brexit.

The idea underlying Cambridge Analytica's methods isn't new. Its roots can be traced back to 'psychological operations' perfected by intelligence agencies in the first half of the twentieth century: spread a few lies about political opponents, plant a few fake documents, then lead the public to the obvious conclusion. That blunt weapon, targeted to reach whole populations, has become far more refined in a data-driven culture that builds profiles so specific and complete—knowing us better than we know ourselves—they offer a singular opportunity to 'move the needle' on a single individual. Cambridge Analytica delivers just what an individual needs to hear to help them 'make up their mind': and then vote the way the targeter intends.

Fake news provides the fodder for this targeting. A post revealing a candidate in a very negative light, targeted to the Facebook page of a voter identified as diffident, helps a voter decide to stay at home on polling day—a behaviour known to be a key factor in both the 2016 Brexit vote and the 2016 US presidential election. Make the idea of voting for a candidate odious and the opposing candidate wins by default.

In Australia the Liberal Party has engaged in conversations with Cambridge Analytica ahead of the next federal election. Cambridge Analytica are keen to set up shop, and while Australia's mandatory voting laws mean voters can't be encouraged to vote or kept from the polls, swing voters in marginal seats can be intensely profiled and targeted with highly tailored individual messaging designed to stimulate them with whatever message will suit the Liberal campaign. This profiling will drive provocative articles into the Facebook newsfeeds of persuadable voters. Although the Coalition appears to be entering the next federal election from a weak position, it could well be that Australia, like Britain and the United States, will see another 'unpredicted' result.

All of this brings to the surface age-old questions about free will and human agency. Are people really this flexible in their beliefs, and so susceptible to persuasion? We find comfort in the idea that we hold core beliefs that guide us in our passage through the world. Many *need* to believe this, rejecting the haunting image of 1984's Winston Smith, deep in the bowels of the Ministry of Love, tortured into seeing, at last, that $2 + 2 = 5$.

But what about the pleasurable authoritarianism of *Brave New World*? While George Orwell articulated the Zeitgeist of a crude and violent twentieth century, it's looking more as though Aldous Huxley took the longer view. In Huxley's dystopian fiction, a wholly planned society provided every distraction to keep its advanced industrial civilisation ticking along seamlessly, a lulled compliance aided by sex and 'soma'. That vision resurfaced in perhaps the most subversive film of the twentieth century, Paddy Chayefsky's *Network*: 'All necessities provided. All anxieties tranquilised. All boredom amused.' Perhaps now it is increasingly becoming the twenty-first-century version of reality.

4. Augmented reality and the Pollyanna machine

We believe that mixed reality is the platform that will define the world for generations to come, that it will change humanity, and make humanity better.

—Graeme Devine of Magic Leap

A fortnight before the *Australian* published its exposé on Facebook's offer to allow advertisers to manipulate vulnerable teenagers, Mark Zuckerberg took the stage at F8, Facebook's annual technical conference. After a few pleasantries about the state of the business (in rude health), Zuckerberg dived into a detailed demonstration of a technology he considers core to Facebook's future:

augmented reality. Where virtual reality is all about leaving this world behind in search of a fully synthetic universe, augmented reality blends the synthetic world and the real world into a seamless whole, crafting a 'technologically supported hallucination' that blurs the boundaries between reality and fantasy.

Augmented reality goes back nearly 50 years, with origins in the flight control systems for jet fighter pilots. With pilots overwhelmed by critical information, and under pressure to make life-or-death decisions, augmented reality crunched that sea of data into easily understood graphics, helping them make better, faster decisions. A few years ago Google reignited a dormant field when it introduced 'Glass': eyewear for generation cyborg, with a see-through projection lens that pasted data direct from the Googleplex over the outside world.

Both of these early versions of augmented reality require gear that covers the eyes. But you don't need all that kit. You can simply look through the screen of your smartphone. Millions got their first taste of augmented reality in July 2016, when *Pokemon Go* became the most popular mobile game in history, installed on hundreds of millions of mobile devices. *Pokemon Go* cleverly situated its virtual creatures in real-world landscapes, turning the idea of an electronic game inside out, using the real world as the game's setting and enticing players into multi-kilometre marches as they hunted their next *Pokemon*.

The wild popularity (and profit) of *Pokemon Go* transformed perceptions of augmented reality from a stodgy also-ran into the hottest tech novelty around. On the F8 stage, Zuckerberg walked through several demonstrations of capabilities Facebook is already pushing out to its two billion mobile users, changes that will reshape the way Facebook users see the world.

Starting with the view of the world passing through your smartphone camera,

Facebook's apps 'augment' the image, passing that along to the smartphone's display. Why? There are many uses, but one thing augmented reality does is convert the real world into a perfect (and selectively invisible) canvas for graffiti. Zuckerberg proudly showed a demo of a message that had been 'carved' into a wooden tabletop by someone who had passed that way before, a carving visible only through the smartphone.

It all seems innocuous enough, little more than just another toy for smartphones already crowded with such diversions, but Zuckerberg's final demo pointed to a bigger vision, and one fraught with problems. Enthusing about augmented reality as a technology that allows art to be put anywhere, he showed a demo of a large, white exterior wall of a building onto which an artist had (virtually) painted a large artwork. Zuckerberg showed that with augmented reality, every visible surface in the world could be drawn upon—rewritten—by Facebook.

Adding a little art to brighten an otherwise dull wall seems like an unalloyed good, but only if one completely ignores bad actors. What if that blank canvas gets painted with hate speech? What if, perchance, the homes of 'undesirables' are singled out with graffiti that only bad actors can see? What happens when every gathering place for any oppressed community gets invisibly 'tagged'? In short, what happens when bad actors use Facebook's augmented reality to amplify their own capacity to act badly?

But that's Zuckerberg: he seems to believe his creations will only be used to bring out the best in people. He seems to believe his gigantic sharing network would never be used to incite mob violence. Just as he seems to claim that Facebook's capacity to collect and profile the moods of its users should never be monetised—but, given that presentation unearthed by the *Australian*, Facebook tells a different story to advertisers.

Yet Facebook comes late to augmented reality. Google's ARCore has similar aims, while Apple touted its own efforts, known as ARKit—several years in the making—at its own developer conference a few weeks after F8. In September, when Apple's iOS 11 operating system was released, half a billion iPhones suddenly acquired sophisticated capacities in augmented reality. So Facebook is far from the only tech giant enabling a wholesale rewriting of the world, nor the only one to face these questions of intention and capacity. But, uniquely among these companies, Facebook has constructed a business model dependent on manipulating the mood of its users.

Facebook's augmented reality will need to make its users feel good about themselves and the world they live in. That's the essential bargain between Facebook and its users, and it means Facebook's augmented reality will 'accentuate the positive' as its core feature, echoing the same design decisions made for its newsfeed, but using them now subtly to rewrite observable reality.

All of this seems to be of little consequence when that's only a view through the screen of a smartphone, but augmented reality has been rapidly evolving into a form that looks much like a normal pair of spectacles, completely covering the eyes. Zuckerberg said as much at the front of his F8 presentation: 'Over the next 10 years, the form factor's just going to keep on getting smaller and smaller, and eventually we're going to have what looks like normal-looking glasses that can do both virtual and augmented reality.'

Those reality-bending spectacles are probably not as far away as Zuckerberg claims. Super-secretive Magic Leap, a Florida technology company, with two billion dollars in investment and led by Google and Chinese internet giant Alibaba, is arguably the best-funded start-up in history. Both investor firms are betting big on Magic Leap's stated goal of getting augmented reality spectacles

to market within the next few years. When they do come to market—likely hyped from all corners as the greatest tech innovation since the smartphone—those spectacles will offer the capacity to generate a thoroughly curated view of the world. The newsfeed will leap off the smartphone screen, plastering itself across the seen world, as Facebook works hard to keep users seeing what they want to see.

By virtue of the way they operate, augmented reality systems must simultaneously act as very sophisticated surveillance systems. In order to add or remove information about the world, these systems must scan that world continuously, creating a very valuable stream of data about the places people go and the things that catch their attention. As always, Facebook will be watching this and learning from users' passage through the world, feeding that data into their machine-learning profiles, and improving the capacity of those profiles to generate a view onto an ideal world. All of the pieces are now in place: we will soon be able to say goodbye to reality.

5. All together now/o brave new world

In the light of what we have recently learned about animal behavior in general, and human behavior in particular, it has become clear that control through the punishment of undesirable behavior is less effective, in the long run, than control through the reinforcement of desirable behavior by rewards, and that government through terror works on the whole less well than government through the non-violent manipulation of the environment and of the thoughts and feelings of individual men, women and children.

—Aldous Huxley, *Brave New World Revisited*, 1958

We live in a connected world, carrying smartphones with us nearly constantly, the

better to remain connected. Each of these smartphones has its own, invisible tethers to machines managing this connectivity, machines recording our passage through the world, analysing our activities then passing that information along. Just as humans share on Facebook, machines share what they know about us. This sharing shapes our interactions with these machines, and those interactions shape our view of the world. What these machines learn about us they feed back to us, creating a circle that feels smooth and seamless, because those machines, quiet and invisible, rarely merit our consideration.

That intangibility—how can you touch the ‘cloud’?—makes all of this feel a bit unreal, as though more about the play of ideas than the embodiment of those ideas in software, directing machines that increasingly shape what we see and what we know. These machines—connecting us, observing us, profiling us, learning from us, targeting us, and more and more shaping what we see—are perhaps the closest things to us; more intimate than any partner, more attentive than any child, patient and constant.

Hidden from view, these machines do the bidding of their owners, offering ‘users’ a few crumbs from a very rich table, hoarding the valuable insights garnered from this continual intimacy for themselves, and for their unseen clients. Although the web began, nearly 30 years ago, as a way to build peer-to-peer relationships, it has steadily evolved towards being a powerful technology of control. Consider that the monetisation of propaganda allows people to lie profitably; machine learning allows profitable lying to be targeted to a profile; and neuromarketing allows the lie to be delivered at just the right moment for maximum effect.

All of that seems wholly negative—until the qualifying effect of augmented reality is added to the mix, a technology that promises

to shape completely what we see. What we see will not be what we want to see for ourselves, but what the machines and those who have paid for the services of those machines want us to see.

This is not some hypothetical future. This is the present, now set for a tremendous acceleration via augmented reality. Just as the smartphone has allowed Facebook to gather information from users continuously—instead of only when they sat in front of a desktop browser—augmented reality means that Facebook and its clients will have the power continuously to adjust what users see, moment to moment. This will become their window on to the world. As that happens the individual perception of reality—as we have known it throughout history—will face its last days.

Each of these technologies in isolation has proven incredibly potent. Their unification—and the first three are already working together quite profitably—creates the most powerful apparatus for social control in human history. We have no precedent for this and no way to manage that transformation. To paraphrase Orwell: those who control reality control the present. Those who control the present control the future.

6. What is to be done?

A man without privacy is a man without dignity; the fear that Big Brother is watching and listening threatens the freedom of the individual no less than the prison bars.

—Sir Zelman Cowen, *The Private Man*,
1969 Boyer Lectures

Although so much of this story centres on Facebook, how it uses the data it gathers to advance its own ends and the ends of its clients, it is not about Facebook alone. This is a tale about a new data-driven culture and the exploitation of influence.

Facebook found itself in the right place at the right time. The smartphone, rising in parallel with the social media giant, provided Facebook with the perfect listening device to gather the data required to build the detailed user profiles needed to monetise influence. Facebook did whatever it could to stay ahead of its rivals in social networking. It made its site ‘sticky’—designing triggers that lure users into returning tens or hundreds of times a day—by tuning the newsfeed to a perfect reflection of a user’s needs and wants. Influence meant survival for Facebook.

Facebook’s user profiles, those billions of increasingly accurate simulations of human beings, can be interrogated, targeted, bought and sold to anyone who fronts up with the cash, because they’re not people: they are the best recording we know how to make of what a person is, and how a person behaves. As we move further into the age of machine learning, these profiles will improve dramatically. These machines already know us better than we do and very soon they will be able to simulate us with great accuracy.

It seems, yet again, as though all of this has veered into dystopian science fiction. There is a competition for influence. Now that we know it can be done, a lot of resources are being directed towards making these tools more effective: and there lies the real problem, one that extends far beyond Facebook. Even if Facebook vowed to retire all the tools it has created to exploit influence through profiling its users, even if it managed to keep that promise (one suspects its shareholders would balk at a decision that would diminish profits), it appears inevitable that some other actor will come along and fashion this conjunction of technologies for their own ends. Some of those actors will be commercial enterprises—firms such as Cambridge Analytica and Palantir (which provides similar profiling tools to the security establishment)—but the greater number will be nation-states, governments

seeking to maintain influence over their citizens.

In 2016 the Chinese government announced a ‘Social Credit System’ wherein each citizen will be given a rating that is an effective and public proxy for social worthiness, that is, their loyalty to the state. China already carefully controls the posts made on Weibo (a ‘microblogging’ site similar to Twitter), frequently censoring posts the government considers inappropriate and banning users who criticise government policy. Paired with this forthcoming rating scheme, it’s easy to imagine a near future where the Weibo feed comes principally from users with the ‘highest’ ratings. That rating system operates as a form of pre-censorship, making it difficult for ideas at variance with state orthodoxy to spread.

Governments throughout the world have taken careful note of the exploitation of influence: recording the public interactions of their citizens—what they read and post and share—to build profiles. Look no further than the Australian government’s thirst for your telecommunications ‘metadata’. They’ll use machine learning techniques to improve the accuracy of those profiles, then



“Just imagine everyone in the audience is fully clothed.”

target individuals with messages designed to keep them moving in sync with the aims of the state. All of this will happen behind the scenes—as has already been the case in Britain and the United States—cloaking authoritarian processes, hiding any hint of the persistent manipulation of reality.

Cambridge Analytica may want to scare you into voting for their candidate, Turkey and Russia might seek to keep restive oppositions in line, but all of these are merely the final, flickering instances of an architecture of power lifted straight from Orwell—‘a boot stamping on a human face, forever’. In the future, none of this will be experienced as coercive: the next stages of weaponised influence will feel good. Facebook does not coerce its users; it learns what they like and gives them more of that. That’s the fulcrum of this tilt headlong into a post-real era.

The future of power looks like an endless series of amusing cat videos, a universe cleverly edited by profiling, machine learning, targeting and augmented reality, fashioning a particular world view in which we will all comfortably rest. That’s already the case for billions of Facebook users, a lesson widely noted by those in power, carefully studied and soon to be widely copied. Facebook has been the beta test for a broad assault on all reality. As these techniques become universal, with the world now listening to us, then adapting to our wants and whims, while subtly shaping us to its ends, we lose our moorings and become entirely post-real.

7. Standover job

You gotta join Facebook. It’s going to be huge.

—Jimmy Wales, founder of Wikipedia,
May 2007

Facebook’s increasingly pervasive presence has seen a progressive routing of all social

discourse through the platform. In 2017 it’s just expected that communities of every flavour—chihuahua fanciers, neo-pagans, drag queens and almost anything else imaginable—have their own corners of the social network, where they connect and share the things important to them. This desire to ‘find the others’ created the initial move towards social networking platforms, and helps people feel less alone in an increasingly atomised world.

As more people joined Facebook, they shared more, about more things, making the site more interesting to more people, who also joined Facebook. This accelerating pile-on, known as a ‘network effect’, means that today Facebook dominates social discourse. Conversely, it has grown difficult to create a connected community outside Facebook. It’s where everyone is.

Almost eight years ago, Dr danah boyd, who had studied how communities of teenagers used social networks to ‘find the others’, building networks of support, wrote a prescient essay about Facebook’s growing power. ‘Facebook is a utility,’ she declared. ‘Utilities get regulated’:

Your gut reaction might be to tell me that Facebook is not a utility. You’re wrong. People’s language reflects that people are depending on Facebook just like they depended on the Internet a decade ago ... Don’t forget: we spent how many years being told that the Internet wasn’t a utility, wasn’t a necessity ... now we’re spending what kind of money trying to get universal broadband out there ...

And here’s where we get to the meat of why Facebook being a utility matters. Utilities get regulated ... We can argue about whether or not regulation makes things cheaper or more expensive, but we can’t argue about whether or not regulators are involved with utilities:

they are always watching them because they matter to the people.

According to an article in the *Information*, former Trump adviser Steve Bannon came to a similar conclusion:

Bannon's basic argument, as he has outlined it to people who've spoken with him, is that Facebook and Google have become effectively a necessity in contemporary life. Indeed, there may be something about an online social network or a search engine that lends itself to becoming a natural monopoly, much like a cable company, a water and sewer system, or a railroad.

What seemed in 2010 thinking from the edge has become obvious at the close of 2017. Facebook has become so central to twenty-first-century social discourse that it has become the de facto commons. The collision between public speech and private ownership means that Facebook has the capacity (and arguably, the legal right) to censor any speech it deems offensive. If you don't like it, Facebook implicitly says, you are free to go elsewhere. But there is no longer an elsewhere. The internet and Facebook have become synonymous in the minds and browsing habits of billions.

In July the German government's Federal Cartel Office (a rough equivalent to Australia's ACCC) launched an investigation into Facebook's monopolistic practices. According to an article in the *Independent*:

In the eyes of the Cartel Office, Facebook is 'extorting' information from its users, said Frederik Wiemer, a lawyer at Heuking Kühn Lueer Wojtek in Hamburg. 'Whoever doesn't agree to the data use, gets locked out of the social network community,' he said. 'The fear of social isolation is exploited

to get access to the complete surfing activities of users.'

German regulators believe Facebook uses its 'my way or the highway' attitude about data gathering as the price of admission to strong-arm users into surrendering all the data Facebook needs to profile and target them. The message is: give us all your data or risk being cut off from everyone you know and everything you love to do. It's a standover job using a very real threat of social shunning.

Having repeatedly abused its monopoly position as data gatherer and profile builder, Facebook appears headed for regulation; first in Germany, then across the European Union, and finally—perhaps years later—in the United States and Australia. Regulation will determine how much data can be gathered, how it can be used, on sold and shared with users. Regulation can bring transparency to purposefully opaque processes. Regulation can tame the manipulative beast of Facebook, though one suspects Facebook will learn how to adapt and profit from those regulations. It's good to be a monopoly, and even better when regulation supports monopoly status. Just ask the big four banks. Yet there's another way this could unfold: Facebook could cut a deal with the various national governments, a bargain that offers social stability in return for a permanent monopoly. Network cultures theorist John Robb spells it out on his *Global Guerillas* blog:

The success of [Facebook's] advertising platform will be based on the ability of Facebook to avoid intrusive government regulation. To accomplish that, Facebook will develop services it can provide governments to better secure, control, and manage their citizens in a volatile global environment. In exchange for these services, Facebook will avoid regulations that will limit its ability to make money.

Regulating Facebook enshrines its position as *the* data-gathering and profile-building organisation, while keeping it plugged into and responsive to the needs of national powers. Before anyone takes steps that would cement Facebook in our social lives for the foreseeable future, it may be better to consider how this situation arose, and whether—given what we now know—there might be an opportunity to do things differently.

8. Doing a solid

We risk being the first people in history to have been able to make their illusions so vivid, so persuasive, so ‘realistic’ that they can live in them.

—Daniel J. Boorstin, *The Image: A Guide to Pseudo-Events in America*, 1961

In its earliest days the web felt messy, chaotic, disorganised and altogether human. The first web browsers, such as Netscape Navigator, came with tools to help people compose their own websites. A million flowers bloomed, as people discovered they could share their opinions and interests, free from the ‘gatekeepers’ of publishing. Rapidly mushrooming from nothing at all to too much there, it quickly became almost impossible to find the interesting bits, until Google made finding things as easy as typing into a search box.

The web began as a lot of little things, not one or a few big things. Sir Tim Berners-Lee, inventor of the web, wanted to make it easy to tie lots of things together, linking pages, preserving their differences in a way that made all those differences immaterial. Those few big things such as Facebook developed the capacity to track users everywhere, build profiles and targeted, sticky newsfeeds, keeping users within their ‘walled gardens’. Without those developing profiles, Facebook would lose its ability to influence its users and centralise the web.

Berners-Lee, angry that his decentralised vision has been thwarted, has set to work on a new invention, one that seeks to restore the power of the profile to the web’s billions of users. Solid—the name of his project—promises to return profile data to the users who create it. Instead of Facebook collecting the list of sites you visit, people you connect with and things you like, Solid provides the capacity for users to expand and manage their own profiles. A user can decide if Facebook gets to use profile data, which data, if any, and for how long. A user can decide to store their profile data with Facebook or on their own smartphone. Solid brings transparency, choice and control to processes that have disappeared from view.

The Solid approach would starve Facebook of the profile data it needs to make itself irresistible to its users, so it’s likely to resist a move to ‘redcentralise’ the web and to respond with even more bright, shiny things to keep users entranced and glued to the site. Berners-Lee takes a longer view, citing the world before the web, when a few giant companies—such as AOL and Microsoft—controlled the online universe. ‘You can make the walled garden very very sweet,’ he says, ‘but the jungle outside is always more appealing in the long term.’

If he’s right and Solid succeeds, liberating user profile data from the companies that mine it to manipulate moods, buying habits, and elections, the future features less Facebook, but more manipulation. The power of the profile—the core of this weaponisation of influence—is here to stay. Anyone who wants power over another now knows to use these profiles. That can’t be stopped. However, we can treat profiles with the respect due to such powerful material.

Solid provides a foundation for a reimagination of the web, offering the opportunity of a path not yet taken, a possibility for a transformation in values, relationships, and economic models. The power of the

post-real can belong to us: Berners-Lee will ensure we have that choice. Two thousand four hundred years ago, Socrates commanded, 'Know thyself.' Berners-Lee makes a different request: 'Own thyself.' Establish control over the data that you create, make sure you are the sole owner of that data.

9. President Zuck

In June Mark Zuckerberg hit the road, going on a 'listening tour', learning from his millions of American users how to make Facebook an even better tool for sharing and communication. Rather than start such a tour in the most populous states on either coast, Zuckerberg headed right for the centre: the state of Iowa, famed for its moderate politics and its first-in-the-nation presidential primary.

Zuck for president? He'll be old enough in 2020, just past the constitutionally mandated minimum of 35. Already well known, the fifth-richest person in the world could fund a presidential campaign for a fraction of one

year's growth in his net worth. Via his shares he has personal control over the best tool yet fashioned to sway minds.

Zuckerberg downplays all of this, telling reporters he's simply trying to make Facebook better by listening to its users ... in Iowa. Combined with a recent declaration changing his status from professed atheist to believer—a necessity for any serious political contender—it begins to look as though Zuck protests too much.

He knows he's built Facebook into a machine that makes money from the exploitation of influence, a machine others already use to sway elections. Zuckerberg could be the next US president, if he wants it. With money and almost unimaginable influence, in the last days of reality that prize—and much else besides—is his for the taking. •

Mark Pesce is an inventor, writer, entrepreneur, educator and broadcaster with 35 years experience working in technology. He holds honorary appointments at the University of Sydney and UTS.