PART II

The panopticon of the soul

The transparent conception of the Open Society

8. The book of life

Christian eschatology holds that on the day of the Last Judgement, every person will face a final reckoning of their life's actions. A description of the consequences of the judgement, if one's life is found to be wanting, appears in the book of Revelation:

> And I saw the dead, great and small, standing before the throne; and the books were opened: and another book was opened, which is the book of life: and the dead were judged out of those things which were written in the books, according to their works.
>
> And whosoever was not found written in the book of life was cast into the lake of fire. (Revelation 20:12, 15)

The author of the Book of Life is, of course, God; omniscience is useful for fully chronicling each person's life. But even if you are not omniscient, you can still have a pretty good go at the job. Some people have recorded their own lives to such an extent that they could be drafting their entries in the Book of Life. Of these, the most famous is arguably Samuel Pepys, who from 1660 to 1669 wrote a diary of 1.25 million words. Pepys's diary provides a remarkable insight into his personal life. He describes at great length not only mundane aspects of his daily life but also his extramarital affairs. He also noted how he buried a wheel of Parmesan cheese in his garden to save it from the Great Fire of London.

Pepys might be the most famous diarist, but he falls short of being the most prolific. His diary of 1.25 million words is more than an order of magnitude shorter than the 17 million words written by the obscure poet Arthur Crew Inman,[1] and just barely 5% of that of the journalist Edward Robb Ellis, who managed 22 million words. The most extreme diarist on record is Reverend Robert Shields, who generated 37.5 million words of phenomenally boring prose by logging his life in five-minute intervals.[2]

One problem faced by extreme diarists is what Bertrand Russell called the Tristram Shandy paradox: the more assiduously they chronicle every detail,

How to cite this book part:

Alexander, J. McKenzie (2024) *The Open Society as an Enemy: A critique of how free societies turned against themselves*, London: LSE Press, pp. 101–171.
https://doi.org/10.31389/lsepress.ose.c.

the further behind they fall in the task.[3] Tristram Shandy was an extreme instance of this problem, requiring one year of labour for every day recorded. Karl Ove Knausgaard improves upon this in his 3,600-page work, *My Struggle*, but even he, on occasion, needs 80 pages to describe a single evening (Kachka 2014). But suppose we could automate the recording of our lives so that our history is written in real-time. Would we want to do that? How would that change our experience of living if we knew that everything we did was not only recorded but accessible to anyone?

Some people have tried to live that way. In 1996, Jennifer Ringley began broadcasting her life from her college dorm using a webcam (British Broadcasting Corporation 2016). In doing so, she became the first person to live life on the internet in real-time. Or at least that portion of her life falling within the line of sight of her webcam. (She added three more webcams a couple of years later when she moved.) The largely uncensored images gave viewers a rare insight into the life of a complete stranger. She broadcast her life in this way for seven years, finally unplugging the webcams in 2003.

Professor Morris Villarroel started logging his life in 2010 but, unlike Ringley, shows no sign of giving up. During most of his waking hours, Villarroel wears a chest-mounted camera that takes a picture every thirty seconds, logging around 1,200 photographs a day. Many photos feature the steering wheel of his car. In an interview with Hooper (2016), Villarroel explains he was inspired to begin this project after turning forty. "I was looking back on my life", he said, "and wondering what did I have to show? I wanted for the next 40 years to have a greater sense of what I had actually done during those years."

Technology such as the internet, social media, and the smartphone enable us to record and share our thoughts and activities to an extent few anticipated. Most of us don't engage in extreme lifecasting like Ringley or Villarroel, but many who use social media are willing to disclose a surprising amount of information about themselves. Shortly after founding Facebook in his dorm room at Harvard, Mark Zuckerberg remarked on this phenomenon to a friend:

> Zuckerberg: Yeah, so if you ever need info about anyone at Harvard,
> Zuckerberg: Just ask.
> Zuckerberg: I have over 4,000 emails, pictures, addresses, SNS.
> [Friend]: What? How'd you manage that one?
> Zuckerberg: People just submitted it.
> Zuckerberg: I don't know why.
> Zuckerberg: They "trust me".
> Zuckerberg: Dumb fucks.

Reporting on this for *Business Insider*, Carlson (2010) included the following contextualising remarks from Facebook's COO, Sheryl Sandberg: "Mark really does believe very much in transparency and the vision of an open society and

open world, and so he wants to push people that way [...] He hopes you'll get more open, and he's kind of happy to help you get there."

And so we turn to the transparent conception of the Open Society. Transparency is an important tool for eliminating corruption, reducing inequality, and ensuring that meritocratic aims are pursued fairly. Transparency facilitates accountability by allowing us to hold individuals responsible for their decisions. Transparency is also identified as a business virtue, providing instrumental benefits. A 2013 survey by TINYPulse (Kruse 2013) found that the most important factor in determining whether employees were happy was transparency of management, with a correlation coefficient of 0.93. Transparency was also found to make good business sense by improving customer satisfaction; in one report, Buell *et al.* (2014) found that in restaurants where customers and cooks were visible to each other, customer satisfaction increased by 17% and speed of service went up by 13%.

Yet, alongside all these virtues of the transparent Open Society, are the concomitant vices. We are social beings and, as such can feel considerable pressure to alter our behaviour and conform. In one famous experiment, Solomon Asch showed how group pressure could cause an individual to accept an explicitly incorrect group judgement about something as blatantly objective as which of three lines was longest.[4] How does living in a world in which every aspect of a person's life can be examined by others affect one's autonomy and capacity for personal growth, development, and expression?

All this raises a great many questions. How much transparency do we want in society? Is a society where the book of life is available to all a good thing? How does a transparent society rebalance the distribution of power between individuals, corporations, and the state? Does radical transparency help or hinder personal growth? What affect does radical transparency have on personal freedom and other individual liberties? Let us now turn to these questions.

9. Unwanted inferences

One problem with the transparent conception of the Open Society is that living generates *a lot* of information. In the wrong hands that information can be manipulated to reveal things about a person they might rather keep private. Given this, it's important for two reasons to think about the conditions under which that information is disclosed. First, when we think about transparency, we typically imagine the voluntary disclosure of information, such as what you say to a pollster, your doctor, or your tax accountant. In practice, much of the information generated by living is disclosed because it's impossible to avoid disclosure, or because the actions required to avoid disclosure are excessively onerous. Second, there's a problem with the metaphor; the term "transparency" suggests a symmetric relation, just like how light travels through a pane of glass in both directions. But much of the information flow concerning people lacks this symmetry. Information flows into a corporation, but that corporation tightly controls what information flows in the reverse direction. We'll consider both of these aspects in this chapter.

The involuntary nature of informational disclosure matters because it happens all the time. If you walk down the street listening to a podcast, other people can see your expression as you react to the audio. If you try to keep a poker face so as to avoid broadcasting your emotions, your poker face and emotional neutrality is on display. Walking into a shop signals that you think that that shop is (probably) worth going into, and so expresses an implicit endorsement of the establishment. Since the political leanings of newspapers are fairly well known, reading a newspaper in public suggests to observers your likely political affiliation. And so on.

Most of the time, we don't worry about this because these small amounts of information that are released seem arbitrary and insubstantial. Who cares if you are seen to go into a Starbucks with a copy of the *New York Times*, or a Wetherspoon's with a copy of the *Daily Mail*? The problem with this line of thought is that it assumes the value of any piece of information is restricted to its explicit propositional content. However, when a piece of information is considered in the wider social context in which it is embedded, those pieces of information, when aggregated, can reveal much more than we suspect.

Involuntary disclosures become all the more revealing when we consider the possibility of aggregating the drip-feed of personal information over time. Consider the act of shopping at your favourite superstore. The list of items that appear on the receipt at checkout doesn't disappear once you leave. It's useful for stores to know what items were purchased together and who purchased

them. One common way of collecting this data is by encouraging people to use store loyalty cards that offer small financial incentives for their continued use. Those loyalty cards provide one way of associating a purchase history with an individual. And even if a person doesn't have a loyalty card (or doesn't use one consistently), the purchase history can still be collected if a known credit card is used to make the payment.

In principle, even using cash would not necessarily enable a person to avoid accruing a purchase history associated with them. Face recognition technology could easily be installed at any checkout register to automatically identify people even if they paid with cash. It is easy to imagine companies choosing to pool information collected this way to develop a more fully-rounded profile of their customers.[1] As a proof of concept of something not all that different, in 2017, Alibaba, the world's largest retailer, demonstrated a "smile to pay" service at a fast food joint in Hangzhou (Russell 2017). People did have to sign up for the service, so it was voluntary in this case, but accurate facial recognition technology could make anonymous purchases essentially impossible in the future.

This matters because detailed consumer histories can reveal surprising facts about people. Consider the following example involving the American retailer Target. The problem Target wanted to solve was how to increase the number of people shopping regularly in their stores. It turns out that most people's shopping behaviour is pretty habitual; they do their weekly grocery shop at the same supermarket, they shop for clothes at certain retail outlets, and they shop for household goods at others. Target, like Wal-Mart, is a superstore that offers all of these under one roof. In principle, people could do most, if not all, of their shopping there, but people often don't because their habits are already set and resistant to change.

But there are certain rare events that *do* allow a person's behaviour to change. These events typically involve a substantial life-altering event that completely shatters the old routine, such as having a baby. In many parts of America, new births are officially registered and these data are available to the public. The problem with that, from Target's point of view, is that getting the information then is too late to be helpful – the new parents would be bombarded with advertising from a bunch of companies. The challenge was to try to identify potential new parents earlier. As Andrew Pole, a statistician who worked for Target, told the *New York Times*:

> We knew that if we could identify them in their second trimester, there's a good chance we could capture them for years. As soon as we get them buying diapers from us, they're going to start buying everything else too. If you're rushing through the store, looking for bottles, and you pass orange juice, you'll grab a carton. Oh, and there's that new DVD I want. Soon, you'll be buying cereal and paper towels from us, and keep coming back. (Duhigg 2012)

By extensively trawling through the available data, some of it collected at Target, some of it purchased from other sources, Pole was able to design a measure that calculated the probability that a woman was pregnant, given what Target knew. This was possible because there are certain changes to a woman's behaviour that are likely to occur when she becomes pregnant. For example, women in the first trimester tend to stock up on certain vitamin supplements.

Pole's measure was surprisingly accurate. In his *New York Times* interview, he told the story of a man who walked into a Target store in Minneapolis and insisted on speaking with the manager. The man said: "My daughter got this in the mail! She's still in high school, and you're sending her coupons for baby clothes and cribs? Are you trying to encourage her to get pregnant?" The manager apologised profusely and, a few days later, called him at home to apologise again. The man said: "I had a talk with my daughter. It turns out there's been some activities in my house I haven't been completely aware of. She's due in August."

Perhaps the most remarkable thing about Pole's measure is that it didn't require massive amounts of data – it only needed to look at a bundle of 25 select products. With that small set of data, Target was able to issue custom direct advertising that, indirectly, notified the woman's father that she was pregnant before she told him.

Such highly specific mail advertising is legal, but is it moral? On one hand, you might argue *yes*, since the measure only placed women in a certain reference class. Simply being part of a reference class doesn't determine anything, the argument goes. No harm was done. If we want to assign a probability to a single outcome (i.e., the probability of a *particular* woman being pregnant), we start by putting the event into the narrowest reference class we can using statistical data and then take into account other non-statistical information which can further influence our judgements about that probability. This non-statistical information can easily be more important than the information about the reference class. The British economist and philosopher John Maynard Keynes specifically warned about this back in 1921:

> Bernoulli's second axiom, that in reckoning a probability we must take everything into account, is easily forgotten in these cases of statistical probabilities. The statistical result is so attractive in its definiteness that it leads us to forget the more vague though more important considerations which may be, in a given particular case, within our knowledge. To a stranger the probability that I shall send a letter to the post unstamped may be derived from the statistics of the Post Office; for me those figures would have but the slightest bearing upon the question. (Keynes 1921, p. 322)

On the other hand, you might argue *no*, it isn't moral for Target to engage in this kind of advertising. Although any pregnant woman who chooses to

have a baby will eventually have to make that known (either when it becomes obvious that she is pregnant or when she gives birth), the exact moment she chooses to disclose that information, and to whom, is her right to decide and hers alone.[2] According to this line of thinking, Target's advertising committed a moral wrong by undermining the woman's control over this decision. How is her control undermined? A person seeing a brochure advertising baby clothes and baby furniture might well assume – as the father did – that there is a reason the woman is receiving such notices. That could cause uncomfortable conversations to be initiated between people. Depending on the nature of the relationship, the woman might be forced to admit something she did not want to admit. A woman may have very good reasons for wanting to keep her pregnancy secret; she may have an abusive partner she intends to leave and does not want to have his child.

Generalising the Target example, the worry about involuntarily disclosed information is that such information, when aggregated, can be used to make inferences about personal characteristics which the individual may not want to share and which, in some cases, are legally protected (Solove 2011, p. 27). In the UK, the Equality Act 2010 identifies nine "protected characteristics": age, disability, gender reassignment, race, religion, sex, sexual orientation, marriage and civil partnership, and pregnancy and maternity status. We have seen how even a not-particularly-sophisticated algorithm can make a pretty good guess about whether a woman is pregnant; it's not hard to imagine that other protected characteristics could be similarly unmasked. Protected characteristics are protected because, historically, they were frequently used as grounds for discrimination.[3] If they can be uncovered, it creates the possibility of covert discrimination under the cover of plausible deniability: someone identified as belonging, or likely belonging, to a certain reference class might just never make it to the short list. We know that this happens with respect to people's *names*: Bertrand and Mullainathan (2004) found in an experiment that when fake resumes were sent in response to help-wanted ads in Boston and Chicago, "white names [received] 50% more callbacks for interviews." More information isn't always better.

Although it might strike us as a bit creepy for a business to be able to determine from a distance when someone is pregnant, the resulting invasion of privacy might not strike us as deeply problematic, even if, all things considered, it was wrong. Had Target used the result of their inference differently, our moral qualms might have been assuaged. We might expect such self-policing on behalf of companies to take place, for it is in the self-interest of business to try and avoid aggravating and harming the consumer. The matter becomes much more problematic when we consider what could happen when the state becomes involved.

There are a number of countries where sex outside of marriage is a crime. In Saudi Arabia, flogging is a common punishment for sex outside marriage. In Iran, the punishment is 100 lashes and, occasionally, stoning to death (United

Kingdom Home Office 2016). There are many such examples. Target's pregnancy metric, developed in order to steer more customers through its doors, could instead be used as a tool of the state for identifying potential out-of-wedlock fornicators. Whereas a doctor might choose not to report a pregnant woman to the morality police out of sympathy for her plight, a computer algorithm has no such reservations.

If pregnancy can be identified by trawling through data, what about one's sexuality? According to the 2017 report from the International Lesbian, Gay, Bisexual, Trans and Intersex Association, there are 72 countries in the world that have criminalised same-sex relationships (Duncan 2017). Of those, Iran, Sudan, Saudi Arabia, Yemen, and parts of Somalia and Nigeria allow homosexuality to be punished by death. (Syria and Iraq have *de facto* death penalties for homosexuality, although they are enforced by agents other than the state.)

Estimates on the percentage of people who are gay vary. In 1948, Alfred Kinsey reported in his book, *Sexual Behavior in the Human Male,* that approximately 10% of men were gay. A later report by Janus and Janus (1993) provided a similar estimate, suggesting that 9% of men and 5% of women could be considered homosexual.[4] In an open-ended survey of Americans in 2002, Gallup found that these estimates were far too low. According to Gallup, 21% of men and 22% of women are homosexual (Robison 2002).

In a country where homosexuality is criminalised, people will either repress their sexuality or go to considerable lengths to hide it. Given the general base rates of homosexuality that social surveys have found, governments of countries where homosexuality is banned have good reason to suspect that there's still a fair bit of homosexual activity taking place. It's too expensive for an oppressive authoritarian regime to monitor *everyone* to catch homosexuals, but what if cues could be found in involuntarily disclosed data that enabled estimates of the chance that someone was homosexual? Target was able to identify pregnant women; can big data perhaps reveal cues to identify homosexual men and women? If yes, that would allow oppressive regimes to identify potential suspects and engage in cost-effective, targeted police surveillance. Lest you think this is pure speculation, Wang and Kosinski (2018) claimed that a neural network, trained from a database of 35,326 images, could distinguish between homosexual and heterosexual men 91% of the time (83% for women) when given five images of a person, exceeding the ability of human judges. That said, whether the methods used in this study could be rolled out more widely while retaining the purported accuracy has been questioned. *The Economist* wryly noted that the inclusion of dating site pictures in the training set were "likely to be particularly revealing of sexual orientation" (The Economist 2017a).

Even if we set aside concerns about how aggregated information could be used for the direct oppression of individuals, there are more subtle uses that we need to be wary of. For instance, women who buy birth control could be targeted by pro-life or pro-choice pressure groups. Temperance movements could do something similar to people who buy alcohol. Individuals who have recently lost their job due to a corporate relocation overseas could be targeted

on social media with political advertisements for an anti-immigration candidate; or, more worryingly, they could be targeted by fake news campaigns attempting to skew their understanding of the social and political landscape at a time when they are psychologically vulnerable. People who live in areas affected by certain demographic or economic changes could be identified as being more susceptible to extreme ideologies (whether it be white nationalist, radical Islamist, or other extremist groups), with a stream of notifications and stories fed to them on social media. Even if there's relatively little harm done by a company such as Target attempting to shift a few more goods by smartly targeting individuals, a real concern lies with possible political uses of the information that is available in the Open Society. We'll return to this topic in Chapter 14.

The ability to identify patterns in people's involuntarily disclosed data is not necessarily a bad thing. Many people *like* receiving recommendations from Amazon and Netflix. And I suspect many would find it acceptable for stores to automatically report to the police purchase histories which suggested harmful intent, such as buying items that, when combined, could be used to make improvised explosive devices.

Yet, when so much information is available, the risk of false positives greatly increases. Chapatti flour, a key ingredient in some homemade explosives, is also a staple of Indian cuisine. Fertilisers and acetone can be used to make bombs, as well as feed plants and remove nail varnish. One news story that went viral shortly after the terror attack on the Boston marathon by the Tsarnaev brothers involved six police officers allegedly visiting a family's house because of an accidental combination of innocent Google searches. Initial reports suggested that the wife's search for a new pressure cooker, her husband's search for a new backpack, and her son's Googling for news about the Boston attacks triggered the visit (Catalano 2013). (It later transpired that what *actually* prompted the visit was the husband searching for "pressure cooker bombs" and "backpacks" on his office computer at his former employer.) Even though the story turned out to be not quite what it seemed to be, the concern it calls attention to is a valid one. No system is perfect, and so false positives will always occur.

This is closely related to a problem Daniel Solove calls *distortion*: data collected about a person fails to represent the whole individual. Consequently, inferences made about a person can yield a skewed understanding of the true situation. Solove provides a nice example:

> Suppose government officials learn that a person has bought a number of books on how to manufacture methamphetamine. That information makes them suspect that he's building a meth lab. What is missing from the records is the full story: The person is writing a novel about a character who makes meth. When

> he bought the books, he didn't consider how suspicious the purchase might appear to government officials, and his records didn't reveal the reason for the purchases. (Solove 2011, p. 28)

Even if you aren't planning on writing the next *Breaking Bad*, this point should still give one pause. Each of us interacts with many people over the course of our life, with little control over how these interactions might be interpreted in the future. Here's an example from personal experience: the father of a boy I used to play with as a child turned out to be a notorious serial killer. As a lecturer at LSE I happened to teach Saif al-Gaddafi in a moral and political philosophy seminar. Neither of those have affected my life, aside from giving me some memorable examples to illustrate the "six degrees of separation" thesis (Milgram 1967). However, other people have not been so lucky. People have been banned from travelling to America simply because a student they taught in class turned out to be a terrorist. Unwanted inferences have the power to radically alter lives.

10. Lifting the veil

In the previous chapter, we considered the problem of unwanted inferences made about *ourselves* drawn from information that we have disclosed, generally involuntarily. Let's now consider a different way informational transparency can generate problems: cases where someone *wants* something to be known about them, because it works to their personal advantage, but doing so creates negative externalities for *other* people. Insurance premiums provide a nice illustration. The key idea underlying the insurance business is nothing more than risk pooling. Suppose that house fires in a certain area occur 0.01% of the time, and each house is worth £500,000. If you could persuade 10,000 people to pay £5 a month for protection, your little insurance scheme has an expected profit of £100,000 a year, and each person can sleep easy knowing that if their house burns down they don't have to find £500,000 to rebuild.

All of that sounds good, except people are remarkably sensitive to perceptions they are getting a raw deal. Someone who lives in a thatched-roof house has a much higher natural level of fire risk than someone who lives in an igloo. Why should they both pay the same amount for fire protection?

This came to a head in the EU regarding differential pricing for car insurance (Sinner and Neligan 2011). For whatever reason (some evolutionary psychologists would suggest that men and women have deeply ingrained natural differences towards risk), women are statistically better drivers than men, having fewer collisions and making fewer claims. A recent study by Laiou *et al.* (2016) found that women, despite being 51% of the EU population, account for only 24% of all road fatalities. Given this, if car insurers carved up the risk pools taking gender into account, women would stand to benefit by having lower insurance premiums than men. For a number of years, this was done. Women could purchase car insurance from companies that only sold to women drivers, and they paid lower rates than men. However, in March 2011, the European Court of Justice issued a ruling, known as the EU Gender Directive, that made it illegal for insurers to take gender into account when calculating insurance premiums. As a result, the car insurance premiums for women rose approximately 30%. (Interestingly age is a protected characteristic that is still legal to take into account for car insurance.)

Big data creates the possibility of generating extremely finely constructed risk pools for highly personalised insurance. Doing so means that insurance premiums can differ greatly from one person to another, with one person's benefit coming at the expense of another. When should this be allowed?

Car insurance provides an interesting case study, both regarding the issue about risk pools and how the technology used to adjust risk pools bleeds into other areas of moral concern. Telematics is a form of personalised insurance which uses data trackers installed in one's car to collect information about a person's driving habits. The data trackers can measure how long someone has been driving, how quickly they accelerate, how hard they brake, how tight they turn, and the location of the car. The selling point of telematics is that it allows drivers to not only *claim* that they are a safe driver but also to *prove* that they are a safe driver. Since safe drivers have reduced risk of accidents, this could allow a person to pay less for their car insurance. GPS devices fitted to cars also, in principle, would allow the automobile to be tracked if stolen, increasing the chances of recovery.

It's worth noting what else that information can be used for. Different areas of cities have different rates of vandalism and theft. If you park your car in an area with a higher than usual crime rate, the insurance company could be notified of the increased risk and correspondingly adjust your premium upwards. If you leave your car in your garage at home and don't drive it at all, the rate can be adjusted downwards. If you drive your car late at night or early in the morning when the risk of accident is higher, the rate can be adjusted upwards. Real-time adjustment of the premium sounds advantageous, in that a person can be billed according to the exact level of risk incurred.

Yet there's a hidden injustice lurking here. Crime rates tend to be higher in poorer areas. If you live in a high-crime area and are poor, telematics can result in you paying a higher insurance premium than a rich person who can afford to live in a safer area. If you are poor, you are also more likely to work unsocial hours, which means that you will drive your car during those hours identified as time of increased risk, and hence pay more. The personalisation of insurance which telematics makes possible could very easily result in the better-off benefiting at the expense of the worse-off. That doesn't seem fair.

Furthermore, telematics enables greater state surveillance. The GPS information collected would allow speeding tickets to be filed automatically, without any need for you to be caught by a speed camera or a police officer with a radar gun. That's not necessarily a bad thing; it arguably would make the roads safer. But automating the policing of speed limits in this manner could penalise people for innocent mistakes. You could be penalised if you accidentally speed while driving in a new area and are unfamiliar with the route, or if you briefly exceed the speed limit to avoid a dangerous situation, and so on. Although it's true that breaking the speed limit is against the law, we need to distinguish between intentional violations and accidental excesses. This could be accommodated by saying that a ticket would be issued only if the person was driving more than a certain amount over the limit. Yet it's worth bearing in mind that governments like revenue, especially easy revenue. One constant bone of contention in the UK has been the perceived egregious levels of fines for parking violations, on top of how expensive parking is in the first place.[1] Once a government realises how easy it is to collect money through automated

fines, the temptation to reduce the threshold above which a fine is issued would be great. Furthermore, suppose you are incorrectly accused of speeding, based on an error in the telematic data. What kind of audit trail would exist to allow you to not only appeal but to have some chance of proving your innocence?

In addition, there are privacy concerns. Who would have access to the data showing where you drove, and when? Would the insurance company be able to sell that information to other interested parties? Would that information be accessible by the state? It's easy to see that making such information available has both benefits and disadvantages. Certain crimes would become easier to solve. Fans of the first season of the podcast *Serial* will recall that a significant amount of time was spent trying to determine just where Jay and Adnan were at certain times.[2] Mandatory GPS trackers in cars would, in cases like that, be a useful source of information. However, it does raise a further question of how to confirm who was driving the car or travelling in it. We are now confronted with a new challenge whereby advocating for one form of transparency creates problems for which the solution is even more transparency. How do you confirm who was driving or travelling in the car? Why not use facial recognition? Maybe this could also be a way to reduce car theft. If the car can be operated only when a registered, recognised driver is behind the wheel, car thefts might be reduced. (But we might then wonder if the number of short-term kidnap events would increase.)

From the point of view of issues closer to home, parents would surely appreciate being able to check whether their children were driving safely and going where they said they were. (This latter question is already answerable with smartphone apps such as "Find my Friends".) However, the possibility of such parental verification sits ill at ease with the goal of encouraging children to become independent and autonomous individuals. Part of the reason why Americans have a longstanding love of the automobile is the sense of freedom that comes with it. Being able to exercise that freedom is part of personal development. After all, if you are old enough to assume the responsibility of driving, shouldn't you also be trusted to do what you say you are going to do?

Trust is an interesting attitude. There's an old Russian proverb, "Trust, but verify." It became Ronald Reagan's signature phrase as president and governed his approach towards negotiating arms control agreements with the Soviet Union. It has also become a popular phrase regarding parental strategies for raising children and managing other personal relationships both within families and within the workplace.

Trust is a relational concept; it either refers to *the belief in* (the noun form) or *believing in* (the verb form) the reliability, truth, or ability of a person or thing. When we think of how the transparent conception of the Open Society enables a "trust, but verify" approach, we need to be aware of two very different instances in which that approach can be used. In the first instance, it can be effective; in the second, "trust, but verify" is potentially harmful.

An organisation such as an insurance company may use "trust, but verify" because what they primarily care about is the *outcome*: Are you driving

safely? With arms control, what we care about is whether the number of nuclear weapons has been reduced or whether centrifuges capable of enriching uranium have been mothballed. As a means of enforcing outcomes, "trust, but verify" can be very effective.

However, in cases where we primarily care about a *relationship* between two people, "trust, but verify" is potentially harmful. If A trusts B, then that means that A believes in the reliability, truthfulness or ability of B. Yet the act of verifying suggests that A, in fact, does not fully believe in the reliability, truthfulness or ability of B; if A did, why would A need to verify? Consider a micro-manager who always verifies that their personal assistant puts stamps on the post. In what sense, if any, is that behaviour compatible with the state of *believing* that the personal assistant was reliable? Verification is associated with *mistrust* rather than trust. That's the whole *reason* why Reagan adopted his signature phrase with the Soviets.

One problem is that much of life involves interactions where we care *both* about the outcome and the relationship. In these instances, we need to understand the pros- and cons- of "trust, but verify". With respect to insurance companies, or your bank, we may be much less concerned about the relationship since few people have a close connection to their insurer or bank. However, in parent-child relationships, "trust, but verify" needs to be used sparingly so as to preserve the feelings of mutual respect and growing recognition of independence that happens as the child ages. The question of how to manage "trust, but verify" with respect to the state is a matter of great importance. An important component of state legitimacy in democratic societies is whether people trust the state. We can't all be present everywhere to monitor each vote in each election – much less be present in every legislative debate. If the state continuously employs "trust, but verify" with its citizens, will that signal of mistrust be reciprocated, eroding trust in the government and people's belief in the legitimacy of the state?

It may seem that we have moved quite far afield from the original question of what kind of information is acceptable to use in creating risk pools, but we really haven't. When it comes to informational transparency, the Open Society presents a multi-dimensional problem. One cannot answer questions about what information is acceptable to use for certain purposes without also addressing who may access that information and what kinds of controls and protections exist. Information collected for innocent purposes can just as well be used for nefarious purposes. This leads us towards two other classes of issues Solove identifies – problems of *exclusion* (do people know how their information is being used?) and *secondary use* (will information obtained for one purpose be used for another purpose without their consent?). Underlying all of these are questions of trust: to whom will we entrust our information and how can we ensure that our trust is not misplaced?

Let us revisit the question of personal information use in the Open Society, but from the point of view of health policy. Like car insurance, medical insurance references a number of individual factors that are under a person's

control: how often they exercise, their diet, and whether they smoke or drink. However, a person's health also depends on factors that they have no control over, for instance, their genetic inheritance. Both of these interact to produce the person's health state, along with a chance component.

Genetic testing used to be something out of science fiction; today, you can order a home testing kit from a company such as 23andMe for £149. It's important to distinguish between genetic *testing* and genetic *sequencing*. Genetic testing looks at specific sites on the genome in order to identify which variants of those genes a person has. This allows us to say if the person carries the gene for Tay-Sachs disease, cystic fibrosis, beta thalassemia, and other disorders. It can also track genetic associations associated with less important traits such as earwax type, finger length ratio, your ability to perceive bitter tastes, and the photic sneeze reflex. Genetic sequencing (also known as "whole genome sequencing"), in contrast, identifies the arrangement of nucleotide base pairs which constitute each person's unique genetic inheritance. (Unless you have an identical twin.)

If genetic testing seems comparatively cheap, the decline in the cost of genetic sequencing over the past two decades is nothing less than astonishing. According to the National Human Genome Research Institute, the Human Genome Project (HGP) invested approximately $2.7 *billion* on the technology and research activities associated with the HGP. The actual cost of sequencing the first human genome in 2003, according to their estimates, was about $300 million, depending on how you do the accounting (Wetterstrand 2024).

Since then, the decline in the cost of genetic sequencing has decreased faster than Moore's law. If you consider Moore's law that computing power doubles every two years and assume that the cost of manufacturing is constant, then that would mean that the cost of a given unit of computing power halves every two years. Figure 10.1 shows a plot tracking the decline in cost for sequencing an entire genome since 2001, contrasting it with a hypothetical decline in cost corresponding to Moore's Law. The cost for genetic sequencing roughly kept pace with Moore's law until 2008, when the original method of sequencing (Sanger sequencing) was replaced with the aptly named "Next-Gen" sequencing methods. By 2015, the cost of sequencing a whole genome was approximately $1,500. According to Jay Flatley, the CEO of Illumina, "[G]etting to a $500 genome is technologically possible, and we think certainly there's potential far beyond that" (Tirrell 2015).

Sequencing an individual genome is priced at a level that makes it feasible for a developed economy to start doing that for every one of its citizens. In 2015, slightly fewer than 4 million children were born in the United States, a figure relatively constant year-on-year (Martin *et al.* 2015). Assuming a cost of genetic sequencing of $1,500 per person, it would cost $6 billion to sequence each baby born in 2015. That sounds like a lot of money until you compare it to the overall budget. In 2015, the total expenditure by the US government was approximately $3.69 *trillion*. Genome sequencing for each new citizen would have required only 0.163% of the yearly budget. To put *that* in perspective,

Figure 10.1: The decline in the cost of sequencing an individual human genome

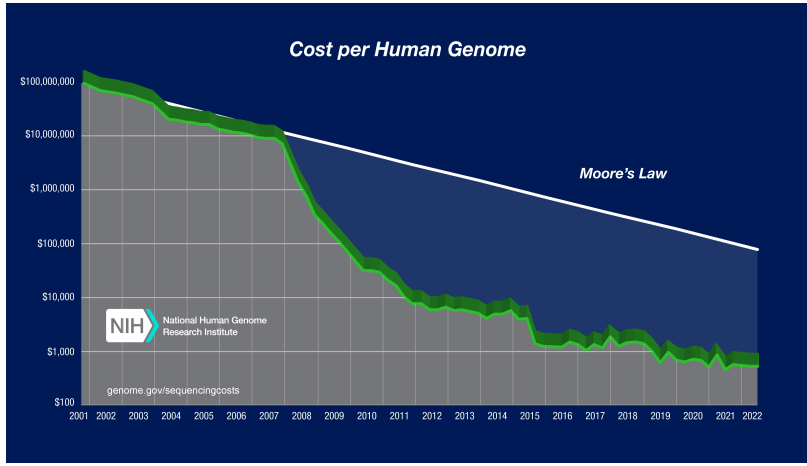

Source: Wetterstrand (2024). Courtesy: National Human Genome Research Institute, Public Domain (https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data).

the median US income in 2015 was $55,775 (Posey 2016), of which 0.163% is $90.91, or slightly shy of the price of 19 Big Macs (The Economist 2015). Future technological improvements will only make mass sequencing of the entire population cheaper.

Consequently, the following scenario is a real possibility. Medical insurance companies could require individuals to submit a genetic profile (partial or complete) so that they can be grouped into an appropriate risk pool. This information could be used along with an individual's aggregated consumer profile to determine whether someone has taken reasonable precautions to avoid triggering certain dispositions towards illnesses they might naturally have due to their genetic endowment. For example, someone with a disposition for type 2 diabetes, which can be prevented or delayed by following a healthy diet, could be determined to have not taken suitable precautions given their consumer profile. A person's consumer profile is an imperfect guide to their lifestyle; simply because someone bought eight cans of lager does not mean that they drank all or any of them. But years of data would provide enough evidence to significantly shift the balance of probabilities of whether a person maintained a healthy lifestyle.

From a philosophical point of view, one problem raised by all of the information which will become available about a person is how we (or society) assigns responsibility for what happens to them. This is a particularly acute problem in the case of healthcare, where demand is potentially limitless but

resources scarce. How should we treat people in need of a liver or lung transplant where the need for that transplant is due to lifestyle choices? Should the position of a person on a waiting list be determined solely on the basis of need, regardless of the cause, or should there be a "merit" element to it, where a person partially responsible for their health condition is demoted to a lower place on the list? Should a person who drives their car into a high-crime area be subjected to higher premiums (or a higher deductible if something untoward should happen)? Similar scenarios are easily envisioned. Should the generosity of unemployment benefits be conditional on the work history of a person? The possibilities multiply.

Underlying these issues is a question concerning the aspects of one's life over which one has control. One intuition many people have is that they need less protection from things that result from factors over which they have control. For example, why should society worry about how a person's information is used if they ultimately have control over how that information is released? If you don't want expensive car insurance, then buy the right kind of car, drive safely, and keep a clean record. When it comes to healthcare, then exercise, eat healthily, and don't drink too much. Read the fine print, look out for your own interests, and take care of yourself.

But what do we mean by "control"? Let's begin with the following definition by Dennett (1984): "*A controls B* if and only if the relation between A and B is such that A can *drive* B into whichever of B's normal range of states A wants B to be in." The idea behind this definition is that A has "control" over B exactly when there is robustness under counterfactual variations of the desires of A: if A were to change her mind about the state she wanted B to be in, A would still be able to drive B into that state.

There are a couple of features to note regarding this concept of control. First, there is an implicit dependence on context. This is clear from an example Dennett gives involving a person using a "wireless radio 'remote control' system" (how quaint that description sounds, today!) to control a model aeroplane. If the radio system loses power, then the person no longer has control over the model aeroplane. The concept of control, then, does not just involve a mere relationship between the person and the aeroplane, but a *three-way* relationship between the person, the aeroplane, and the environment.[3]

A second feature is that, in many cases, whether a person has control over something depends on the resources at their disposal, which can be brought to bear on the control problem. Call this the *budget* of the person.[4] For example, at present, I have some control over where I live; if I had a higher salary or more investments, I would have more control over where I live because then the set of possibilities would be larger. Although it's natural to think of a person's budget as referring to the wealth they have at their disposal, it also includes resources such as free time and social capital.

Putting these two aspects of control together, a third feature immediately emerges: control isn't necessarily preserved under aggregation. Suppose that a person A has three areas of her life which matter to her: B_1, B_2, and B_3. It can

be the case that A can control B_1 *or* B_2 *or* B_3 *in isolation* (so, in a sense, A has control over everything which matters to her), yet at the same time, A cannot control *the collection of* B_1 *and* B_2 *and* B_3. Why? When A controls B_1, it might either alter the environment in which A operates, or it might involve expenditure of some of A's budget. That then changes the conditions under which A approaches the control problem of either B_2 or B_3. A's ability to drive B_2 or B_3 into a particular desired state prior to driving B_1 into a certain state, might now be impaired. This also shows that when we think about multiple control problems faced by an agent, there can be *path dependence*.

The flip response is to say that this is nothing more than the basic fact that "you can't do everything". However, this misses the deeper point that when faced with issues such as how Target uses people's data to make inferences about their life, or how insurance companies use information to calculate rates, or what your genome reveals about your tendency to develop diabetes or dementia later in life, it's all too easy for people, organisations, or the government to view those particular cases *in isolation* and say that worries such as those I've expressed aren't a big issue because if a person feels uncomfortable they can avoid it. (That is, they can exercise control over the relevant area of their life.) But life is an *aggregate* control problem, where people are worried about controlling the *collective*. Our limited resources often mean we have to make cost-benefit calculations and relinquish control over part of our life simply because other things matter to us more. Yet the fact we *acquiesce* to something doesn't mean that we *accept* it.

We've covered a lot of ground in this chapter, but the underlying theme that has been driving the discussion is the following: a lot of the time people accept a great deal of transparency regarding their personal information because of the alleged benefits it gives them. But this is a devil's bargain for two reasons. First, social practices that require the disclosure of personal information because it benefits one category of person can create negative externalities for people from other categories. (This was illustrated in the discussion of using telematic data for car insurance.) Second, the alleged benefits of social practices that require the disclosure of personal information should be treated with scepticism, given the possibility of unexpected negative side-effects. (This was illustrated in the discussion of using genomic testing for health insurance.)

Both of these points matter because, although they raise concerns about how voluntarily disclosed personal information is used, they are generally not seen as providing decisive reasons against, in principle, rolling out social practices that use such information. Why is that? Because, in almost all cases, it can be argued that if someone is concerned about how their personal information might be used (for or against them) in particular instances, they have the ability to do something about it. As I said, people are thought to have *control* over the relevant aspect of their life, and could do something about it, if the way in which their personal information was used (or misused) was a cause of concern. However, as I then argued, even if that is true on a case-by-case basis,

it probably isn't true for a person's life *in aggregate*, because control isn't necessarily preserved under aggregation. And this establishes that the argument for why use of voluntarily disclosed personal information is not so bad (because people can always do something about it if it concerns them) is actually based on a false assumption.

The mere process of living generates a lot of information, much more than we can ever really comprehend. Is it a bad thing that, in the contemporary world, we are implicitly broadcasting all the time whether we like it or not? (Or, for that matter, whether we choose to or not?) This is the topic to which we turn in the next chapter.

11. Letting it all hang out

In the second quarter of 2017, Facebook had 2 billion active users worldwide. At the same time, Instagram boasted 800 million active users, Twitter (now X) 330 million, Pinterest 175 million, LinkedIn 467 million, and the messaging service WhatsApp over 1.3 billion users. By the fourth quarter of 2022, Facebook had added almost another billion active users, growing to 2.96 billion (Statista 2023a). Instagram had 1.28 billion users at the end of 2022, Twitter 368 million (Statista 2022a), and WhatsApp 2.24 billion (Statista 2022b). With the exception of Twitter, the major social media networks saw nothing but growth during that five-year period. In China, the largest social network, WeChat, had over 1.3 billion users at the end of 2022 (Statista 2023c), with 586 million active users on Weibo, the Chinese equivalent of Twitter (Statista 2023b). That's a lot of people broadcasting information about themselves.

The extraordinary growth of social media over the past decade is one of the greatest uncontrolled experiments we have conducted on ourselves. We are a social species, but the physical environment in which humans evolved is so vastly different from the environment on social media that it's not at all obvious that our natural dispositions, developed for face-to-face interactions in small groups, translate well to the new setting.[1] Here are a number of important differences: the relative *anonymity* of interactions, the *speed* at which communication can take place, the *scale* on which interactions can take place, and the temporal *endurance* of information thus generated. Let's work through these four differences in turn.

Consider anonymity. What is different from an interaction on social media and an ordinary interaction in the real world? Suppose I am at a dinner party with a group of friends and I make an assertion about politics. Although every person at the dinner party will have heard the assertion, my act of making the assertion ceases to exist once I have finished speaking. The only remaining trace of it will be a representation in the memory of my friends. If the assertion was controversial and my friends decide to talk to others about it once the dinner party is over, all they can do is *report* what I said as they remember it. The original speech act cannot be shared or reproduced. In some cases – if I say something particularly memorable or shocking – they might remember my assertion verbatim, but for most communications such exact recall is rare. Hence, if my friends decide to share what I said, what they communicate to others can only be a paraphrased version of their understanding of what I meant.

How do we understand what a speaker means? In a famous article, the philosopher of language H. Paul Grice offered the following analysis of attributing meaning to a speaker A who makes an utterance x:

> A must intend to induce by x a belief in an audience, and he must also intend his utterance to be recognized as so intended. But these intentions are not independent; the recognition is intended by A to play its part in inducing the belief, and if it does not do so something will have gone wrong with the fulfillment of A's intentions. Moreover, A's intending that the recognition should play this part implies, I think, that he assumes that there is some chance that it will in fact play this part, that he does not regard it as a foregone conclusion that the belief will be induced in the audience whether or not the intention behind the utterance is recognized. (Grice 1957, pp. 383–84)

That is, in order to say that a person A means p by x, three conditions have to hold. First, A intends to get the audience to believe p by means of x. Second, A also intends the audience to recognise that A *wants* them to come to believe p by x. Third, the audience does, at least in part, come to believe p because they recognise this intention on A's behalf.

If we accept Grice's account of speaker meaning, a number of differences between a social media interaction and an ordinary interaction are revealed. First, in an ordinary interactive context, such as a dinner party or a board meeting, the identity of the audience is fixed. Given this, the speaker then has the ability to select his or her words with the particular audience in mind, knowing what forms of communication are judged as acceptable by that audience. Second, the intention of the speaker, in cases of ambiguity, is frequently disambiguated by context or the shared background knowledge the audience possesses of the speaker. Both of these are important in ensuring effective communication, and both help to prevent radical misinterpretation of what the speaker says.

The difficulty arises when people post some x on social media, and x becomes detached from the speaker, the intended audience, and the context of its generation. The x can also be trivially reproduced in new contexts and to new audiences. Since the audience is no longer fixed, the ability to disambiguate or clarify the intention of the speaker by appealing to context or the speaker's background is lost.

In one sense, this is nothing new; in many respects, it is similar to the traditional publishing model where an author generates a text and releases it into the world. Yet there are some important dissimilarities. First, when we consider the traditional production of texts for public consumption, the authors or artists were aware of the fact that they were relinquishing control over who will encounter their work. Second, when people generated texts for publication or

presentation in the past, they were *intending* to communicate to a wide audience; they *wanted* to broadcast their message beyond a small set of people. Both of these factors would focus the mind on the form of expression used. Social media, by contrast, can encourage an illusive understanding of the nature of the audience, encouraging the adoption of an informal and unguarded tone. One may believe that one is just communicating with one's friends or followers. Yet the ability of social media to reach an audience very different than the one a person intended to address can have phenomenally damaging consequences, in two different ways.

First, national security services routinely trawl social media sites as part of their intelligence operations. This worked to the disadvantage of Leigh Van Bryan, who in 2012 tweeted a friend to say, "Free this week, for quick gossip/prep before I go and destroy America." Stripped of any context about the intention of the speaker, one faces the interpretive problem that "destroy" has a number of meanings, some innocuous, some not. We can speak of a sports team destroying their opponent (meaning, to defeat them in a match), a student destroying a test (meaning, to do well on it), a person destroying a hamburger (meaning, to eat in an uninhibited manner), or a person getting destroyed at a party (meaning, to become very intoxicated). Unfortunately, "destroy" also means "to demolish". When Van Bryan and his friend Emily Bunting arrived at Los Angeles International Airport, they were detained by the US Department for Homeland Security and questioned for five hours before being sent back to the UK (BBC 2012).

Second, posts on social media are easily transmitted, thereby changing the membership of the audience after the time of the original utterance. Since the expressed message is unchanged, each new recipient will view it as the speaker making the same utterance again, and will attempt to extract the speaker's meaning. This is problematic. Following Grice, although I may attempt to induce the belief p in audience A by uttering x, it does not follow that I would attempt to induce the same belief p in audience A' by making the same utterance x – even if every member of A is also part of A'. A form of communication, acceptable to audience A because they know the speaker and thus understand the intended interpretation of x, may not be acceptable to the wider audience A' because they lack the necessary information to know the speaker's intent. The anonymity of the speaker to the new audience thus creates a blank slate regarding the new audience's understanding of the speaker's intention, only constrained by a literal reading of the text. Statements involving irony or humour are thus particularly open to being misconstrued.

Let us now add to the mix how issues of both *speed* and *scale* affect the interaction. Jon Ronson, in *So You've Been Publicly Shamed*, discusses the role played by social media in the phenomenon of public shaming. Shame has played a part in all societies as a form of punishment and social control, but it has recently taken on a new, sinister, arbitrary, and disproportionate character when combined with social media.

In December 2013, Justine Sacco was travelling from the US to South Africa and she decided to pass time by tweeting sarcastic jokes and commentary about her experiences along the way. A typical example was her observation from Heathrow: "Chili – cucumber sandwiches – bad teeth. Back in London!" Where things went awry is when she decided to send one last tweet before boarding the plane to South Africa. She tweeted: "Going to Africa. Hope I don't get AIDS. Just kidding. I'm white!"

There are a lot of ways a person can read that last tweet. A charitable reading would interpret it as a send-up of white privilege, written by someone pretending to be an ugly American, expressing racist stereotypes. In doing so, it could also be seen as a historical allusion to a mistake made in the US when the AIDS epidemic first emerged there. (This might be an overly charitable reading.) How so? Because the AIDS epidemic in America was initially written off by many as a "gay cancer" (Kerr 2020). It was only when AIDS made inroads into the heterosexual population that it began to be taken seriously. Reading these historical allusions into the subtext of Sacco's message would have been a charitable interpretation.

When Justine landed in South Africa, she found the internet had not been charitable. In the space of 11 hours, her tweet had moved from being seen by her 170 followers to being the *top worldwide trend* on Twitter. Some of the commentary quoted by Ronson is disturbing for the sheer glee some people took in holding her feet to the fire:

> 'In light of @JustineSacco disgusting racist tweet, I'm donating to @care today' and, 'How did @JustineSacco get a PR job?! Her level of racist ignorance belongs on Fox News. #AIDS can affect anyone!' and, 'No words for that horribly disgusting, racist as fuck tweet from Justine Sacco. I am beyond horrified' and, 'I'm an IAC employee and I don't want @JustineSacco doing any communications on our behalf ever again. Ever' and, 'Everyone go report this cunt @JustineSacco' and, from her employers, IAC, 'This is an outrageous, offensive comment. Employee in question currently unreachable on an intl flight' and, 'Fascinated by the @JustineSacco train wreck. It's global and she's apparently *still on the plane*' and, 'All I want for Christmas is to see @JustineSacco's face when her plane lands and she checks her inbox/voicemail' and, 'Oh man, @JustineSacco is going to have the most painful phone-turning-on moment ever when her plane lands.' (Ronson 2015, p. 65)

Justine was fired from her job.

There are a couple of important things to note about this. First, interpreting the tweet as a "horribly disgusting, racist as fuck tweet" is making a claim about speaker meaning. During the 11 hours that Justine was on the plane, the

rapid expansion of the audience from her original pool of 170 followers to an audience of *millions*, the vast majority of who did not know her, meant that the speaker's intention was largely unknown. In the absence of any knowledge of the speaker or further contextual understanding, people were free to attribute to her whatever intention they wished due to the absence of evidence. And, once the tweet went viral, crowd dynamics and herding behaviour took over.

Second, it's worth remembering that the vast majority of us *never* address an audience of several million people. Because we never expect our words to be heard, read, or listened to by an audience of millions, we don't usually choose our words with that in mind. Most of us have, at one time or another, made a comment which, taken out of context and broadcast to several million people, could result in a fate similar to Justine's. But most of us haven't suffered that kind of public scrutiny of our carelessly chosen words, and most of us will continue to be lucky in the future. In this way, public shaming seems both arbitrary and disproportionate.

The important lesson to be drawn from this, is how the phenomenon of public shaming is a predictable consequence of how speaker meaning can be misattributed when the anonymity of the speaker is combined with a rapid scaling up of the size of the audience. Figuring out what a person intends is difficult, and we rely on a lot of contextual knowledge, along with background knowledge of the speaker, to do so. And, most importantly, in our day-to-day social interactions, if a friend says something rude or unkind or unpleasant, but is otherwise thought to be a good person, we look to environmental or external factors to explain the behaviour, rather than appealing to character flaws. Context, then, is crucial.

Yet Twitter almost seems to be designed to *encourage* people to take things out of context. There's only so much you can say, after all, in 140 (later, 280) characters.[2] Twitter isn't the only form of social media susceptible to the charge; comments on Facebook and photographs on Instagram or Snapchat can similarly mislead. The danger in all of these instances is that we experience the world and understand our actions through the omnipresent background context provided by our own internal narrative. Because *we* typically understand what we intend to mean by what we say and do, it's not normal for us to step back and reflect on how our statements and actions could be misinterpreted. However, when the background context exists only within our head, that gets lost over the internet.

What this chapter illustrates is yet another way that the radical transparency offered by the Open Society can be problematic. Digital communication via social media differs sufficiently from ordinary methods of face-to-face communication that our informal understanding of how to determine a speaker's meaning (e.g., Grice's theory) cannot simply be carried over without modification. This means that it is very easy for miscommunication to occur,

with unusually harmful consequences for persons involved given the four factors that distinguish communication via social media from other forms (i.e., anonymity, speed, scale, and endurance).

It is important to note that the concerns raised here about miscommunication are ones that occur in the context where I am explicitly restricting attention to cases where people are engaged in *good faith* communication. Good faith communication already presents a number of problems regarding interpretation of meaning because of the possible uses of humour and irony (and other rhetorical devices), which are well-known to present risks of being misunderstood. *Bad faith* communication, in contrast, involves people deliberately attempting to mislead or spread misinformation. The problem of bad faith communication is, in some sense, not new. There have always been liars and hucksters and people willing to say anything you want to hear in order to get you to do something for them. What the internet and social media enable is for this bad faith communication to occur on a scale which was previously unimaginable. One interesting question this presents is, how do we decide who to trust? As that question is, I believe, deeply connected to the re-emergence of tribalism and the rejection of the communitarian conception of the Open Society, that investigation will be postponed until Chapter 24 of Part IV.

12. Don't you forget about me

In *Funes the Memorious*, Jorge Luis Borges tells the story of Ireneo Funes, who, after injuring his head in a fall from a horse, develops the ability to remember everything that happens to him. Such an ability at first blush sounds desirable. Think about how much time you would save! Passing the written test at the local department of motor vehicles would be a cinch, you'd be the most sought-after team member for the next pub quiz, and filing your taxes would be trivial.

Except, in these fantasies, there's one crucial assumption: we only remember the things we *want* to remember. We envision ourselves like a modern-day John von Neumann, who purportedly could recall every book he'd ever read (Goldstine 1980). We don't envision our memories polluted with all the *faux pas* we have ever made, all the rejections we have experienced, and all the suffering of loved ones in their last moments. When those memories are included, Funes' gift no longer seems quite as attractive.

A select group of people actually do have the ability to remember everything that happens to them. The first person to be diagnosed with Highly Superior Autobiographical Memory (HSAM) was Jill Price, who emailed Dr James McGaugh, the director of UC Irvine's Center for the Neurobiology of Learning and Memory, on 8 June 2000 to complain about a problem with her memory. The problem? She remembered too much. From her point of view, the memories inserted themselves into her consciousness in a manner beyond her control and often against her will (McRobbie 2017).

Since McGaugh's discovery of HSAM, around 60 people worldwide are now thought to have the condition. One common reaction, although not universally shared, is that HSAM is both a blessing and curse. In an interview with the BBC, Nicole Donohue, who has HSAM, said: "It can be very hard to forget embarrassing moments. You feel the same emotions – it is just as raw, just as fresh […] You can't turn off that stream of memories, no matter how hard you try." Another person with the condition, Nima Veiseh, agreed: "It is like having these open wounds – they are just a part of you" (Robson 2016). Yet, at the same time, Nima put forward an interesting theory about how having HSAM has made him a better person: "Some say 'forgive and forget', but since forgetting is a luxury I don't have, I need to learn to genuinely forgive. Not just others, but myself as well."

The idea that forgetting is advantageous has been around for a while. Back in 1890, William James wrote in *The Principles of Psychology* that "If we remembered everything, we should on most occasions be as ill off as if we remembered nothing." The virtues of forgetting have perhaps been underappreciated, though, since we generally tend to undervalue that which we experience

in excess. Furthermore, the virtues of having a good memory have been noted for ages; Socrates, in *The Phaedrus*, argued that the invention of writing was bad because "this discovery of yours will create forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves."

If forgetting is a virtue, what are we to make of an Open Society in which our personal information will be always available, instantly discoverable, and impossible to forget? Removing information from the internet has become the modern-day analogue of a Sisyphean task. Nima's challenge of learning how to live when forgetting is not an option is one which all of us increasingly face. And although the "right to be forgotten" has recently begun to appear in legislation of the EU and elsewhere,[1] whether there is, in fact, a right to be forgotten is a contested issue. In the US, some have argued that the right to be forgotten sits uneasily, or outright conflicts, with the freedom of speech guaranteed by the First Amendment.

What will it be like to live in a world where nothing can be forgotten? If, as Nima suggests, the inability to forget requires the ability to genuinely forgive, are we capable, as a society, to be so forgiving? Let's broach that question from the point of view of one society: the US. There is good reason to think that, in some sense, the US has become a less forgiving society over the past few decades. Truth in sentencing laws, which began to appear in the mid-80s, restricted chances for prisoners to be paroled for good behaviour. Increasing numbers of people were locked up as part of the "war on drugs," even for non-violent crimes. The penal code became a lot tougher, as well. The first "three strikes and you're out" law was passed in California on 7 March 1994. If a felon with two prior convictions for serious or violent felonies was found guilty of a third offence, the law required that they be sentenced for a period between 25 years and life. Three strikes laws are popular with voters, as they give the appearance of being tough on crime; some version of the three strikes laws can be found in 28 states. However, these laws also generate remarkable injustices. In 1995, Curtis Wilkerson, a 33 year-old Californian man who had prior convictions from 1981 (he had served as a lookout in a number of robberies when he was 19), was caught attempting to shoplift a pair of white tube socks valued at $2.50. He was sentenced to 25 years to life. A number of other examples of disproportionate sentencing exist.[2]

In addition to longer prison sentences and the growing elimination of judicial discretion on sentencing (that is, the greater use of "mandatory minimums"), there are further punishments people experience even if they aren't locked up for life. A criminal record in the US hampers one's ability to start again due to restrictions on the kinds of work former convicts are able to get. The *Economist* noted that a number of jobs in the US require a licence of some sort, and licences are typically denied to convicted felons. Being a firefighter in California, for example, requires certification as an emergency medical technician, which not many felons achieve. Any licence which includes a "good

moral character provision" rules out a person with a felony conviction. Oklahoma's restrictions are particularly onerous: "state licensing boards completely banned convicted felons from almost 40 professions ranging from asbestos-abatement contractor to embalmer, and from landscape architect and podiatrist to wrecker, a job which usually entails removing debris from building sites" (The Economist 2017b).

Social policy towards people convicted of crimes provides one perspective on the *zeitgeist* of the US. It suggests people aren't particularly willing to forget and forgive the sins of others. How does that attitude manifest itself when we consider the permanent record of people's past left on social media? Remember, a behaviour doesn't have to be *illegal* for people to disapprove and act differently towards a person. It turns out that even though social media tends to involve the personal aspects of one's life, that can have negative consequences for people in a variety of ways. Let us call this phenomenon "the policing of the personal". It can take several forms: organisations and individuals using available information about a person's private life to penalise them in ways they wouldn't be inclined to if that information wasn't available, and – more worryingly – people recognising this possibility and *internalising* these expectations and altering their behaviour in ways which they would not otherwise.

In 2013, *The New York Times* reported a Kaplan survey of 381 college admissions officers. Over 30% admitted to having looked up applicants on Facebook or other forms of social media. Of those, 30% said that they had found negative information which affected the individual's application. The results of Facebook searches by admissions officers have led to applications being rejected, even though it was almost certain that other individuals who were admitted had engaged in similar behaviour, but just didn't leave a detectable record. In response to this phenomenon, high school students are now frequently advised to clean up their social media profiles by removing posts related to alcohol and to use personal email addresses that are appropriate for a general audience.

Things don't get appreciably better once you leave university and go to work. *Time* magazine reported in 2015 that over 50% of companies check the social media profiles of job applicants (Kumar 2015). Once you make it over the various hurdles and secure a job, the trawling doesn't end there. A 2012 report from the IT research firm Gartner predicted that the proportion of corporations monitoring the behaviour of employees on social media will rise to 60% by 2015. Such monitoring, as you might expect, occasionally results in the business discovering that their employees aren't always saying positive things about their employer. A number of people have been fired for "inappropriate" remarks made on social media.

The problem here is that whereas it's one thing to hold someone accountable for behaviour that violates the law or clearly stated guidelines, such as rules regarding procurement, what counts as an "inappropriate" remark is a judgement call. Inappropriate to whom, and for what reason? In 2013, Stephanie Bon learned that the new chief executive of Lloyds Banking Group could earn up to £8.3 million in their first year of employment. That evening she posted

the following remark on Facebook: "New boss gets £4,000 an hour. I get £7 an hour. That's fair." The next day she was fired. In 2017, Juli Briskman was fired from her job as a marketing and communications specialist for Akima, a federal contractor, after using a photograph of herself flipping off Trump's motorcade as her profile picture on Twitter and Facebook (Walters 2017). Should an off-hand remark or gesture result in a person losing their job, simply because their employer doesn't like it? Bon's comment could be viewed as a critique of social inequality and pay differences, especially in light of the financial crisis. Briskman's gesture, made while cycling and off work, is legally protected speech, even if a bit rude.

The cases have important differences in their legal contexts because Bon was working in the UK whereas Briskman was working in the US, but they both illustrate the phenomenon of a person experiencing retribution for expressing legal speech in their personal life. The difficulty is that social media allows speech within one's personal sphere to intersect with the public sphere in ways which "normal" speech does not. Given that social norms have changed and opting-out of social media is increasingly less feasible, we need to ask if it is acceptable for corporations and organisations to use their power over us to restrict personal expression simply because they dislike it. (It's worth noting that these corporations and organisations have no problem using social media to push *their* own interests,[3] so it's not like they have a principled objection against their brand appearing in social media, generally.) The economic power exercised by companies and organisations is possible because most of us need to work to feed ourselves and afford a place to live; and this economic necessity exists because we were born into a pre-existing economic system at a particular place in the structure over which we had no control. (Think of this as an economic variant of the birthright lottery, from Chapter 2.) Is it right that something I have no control over explicitly restricts my freedom of speech and ability to express myself? The rich and powerful have many advantages which the rest of us do not. Should those advantages be expanded to include not just fancy hotels and fine wines but also the ability to express themselves with impunity?[4]

One line of argument some people advance is that the freedom of speech does not mean the freedom to speak without consequences. In the case of Juli Briskman, she was employed in Virginia which has "at-will" employment. States with such laws allow private-sector employers to fire people at any time, for any reason. We need to ask whether such broad laws are fit for purpose. One consequence of "at-will" employment is that it leads to the second aspect of policing the personal – the internalisation of the expectations of others. When a company can fire a person for any reason (e.g., when a social media mob demands it), any online speech act becomes fraught with danger. Any speech, no matter how innocuous on the surface, will be run through a content filter which not only asks "Is this what I really think? Does this express the point I am trying to make?" but, in addition, "How will my employer respond? What will someone who knows nothing about me think?"

It is true that freedom of speech does not mean the freedom to be a jerk or, worse, without consequences. But it *does* mean the freedom to express religious or political beliefs (among others) in many contexts with protection against certain sorts of reprisals. The issue at stake is whether one can speak freely *within* the limits of socially acceptable discourse without fear of reprisal.[5] A right which cannot be exercised without fear of reprisal undermines the very idea of that being a *right* in the first place. (Here, it is important to distinguish between fear of reprisal due to people violating the law and fear of reprisal due to people acting entirely within the law.)

One worry about the transparent Open Society and its inescapable memory is that it does not seem to fit well with how many people actually evaluate the character of individuals. If we consider the phenomenon of public shaming, the unforgiving attitude demonstrated by the criminal justice system in the US, and the uncompromising attitude of many employers (and admissions officers) towards social media posts, it seems that the way many people evaluate others shares a lot in common with the following biblical quote from Matthew:

> Every good tree bringeth forth good fruit; but a corrupt tree bringeth forth evil fruit. A good tree cannot bring forth evil fruit, neither can a corrupt tree bring forth good fruit. Every tree that bringeth not forth good fruit is hewn down, and cast into the fire. Wherefore, by their fruits ye shall know them. (Matthew 7:17–20, KJV)

Why are so many so willing to engage in public shaming, deny criminals rehabilitation opportunities, and reject or fire people based on social media revelations? Because of the underlying assumption that people have stable and binary characters; small samples of a person's behaviour are treated as revealing the whole. In addition, people commit the fundamental attribution error (Ross 1977), where they explain the behaviour of *others* via character traits downplaying environmental influences. Thus, when we see something we dislike, we tend to explain it as that person being a bad person instead of some other cause. People falling for the fundamental attribution error would interpret Justine Sacco's tweet as indicative of racist beliefs rather than a bad joke tweeted unthinkingly.

This matters because people are complex aggregates of both good and bad, subject to environmental influences more than we think (Doris 2002). Good people can do bad things, and *vice versa*. Hitler was a vegetarian who loved dogs. Gandhi, in addition to liberating India from colonial rule, tested his ability to resist sexual temptation by sleeping naked with his teenage grandniece (Connellan 2010). Virginia Woolf has been accused of anti-Semitism, and Pablo Picasso and J. D. Salinger were misogynists. When it comes to famous individuals, we seem better at separating judgement of their achievements, both good and bad, from their personal attributes. Yet when it comes to

ordinary individuals, maintaining that division is harder. This is another challenge presented by the transparent conception of the Open Society: the omnipresent availability of our past providing a voyeuristic window into a person's life, in a world unwilling or incapable of accepting the complexity of people.

13. Returning to the past

Some have claimed that the loss of privacy in a transparent society is nothing more than a return to the natural state of humanity. In small hunter-gatherer societies – the original position of *homo sapiens* – there was no "right to privacy". In a small society, everyone knew what everyone else was doing. There was more emphasis on the community and less on the individual. Jared Diamond calls attention to the contrast quite vividly:

> [Privacy is] an unusual concept by the standards of world cultures, most of which provide little individual privacy and don't consider it a desireable ideal. Instead, common traditional living arrangements consist of an extended family inside a single dwelling, or a group of huts or shelters around a single clearing, or a whole band sleeping in one communal shelter. Unthinkably to most modern Americans, even sex between a couple traditionally goes on with a minimum of privacy. The couple's hammock or mat is visible to other couples, and the couple's young children may be sharing the same mat but are merely expected to close their eyes. (Diamond 2012, p. 224)

As an aside, it's worth noting that this anthropological fact challenges Hannah Arendt's claim that "from the beginning of history to our own time it has always been the bodily part of human existence that needed to be hidden in privacy, all things connected with the necessity of the life process itself" (Arendt 1959, p. 72). Privacy has not nearly been so central to the human condition.

Yet even if life in small-scale societies is typically lived under the watchful eyes of others, we should not assume that people don't find that grating and, hence, that the modern conception of privacy is not something worth safeguarding. In his ethnography of the life of the Kragur on Kairiru island in Papua New Guinea, Smith argues:

> I believe even villagers find their small-scale social world a strain at times. Families sometimes build second houses near their gardens […] not only to save time when there is so much work to be done, but also, some told me, so that they can occasionally escape from the noise of dogs and children, the frequent public gatherings, and the lack of privacy of life in the closely packed village. (Smith 1994, pp. 60–61)

This suggests a more nuanced understanding of the value of privacy in traditional societies than Diamond suggests.

If the Open Society threatens to undermine people's ability to live private lives, is that necessarily a bad thing? After all, what is the value of privacy? Josh Cohen, author of *The Private Life*, offers one answer to this question: "Privacy, precisely because it ensures we're never fully known to others or to ourselves, provides a shelter for imaginative freedom, curiosity and self-reflection. So to defend the private self is to defend the very possibility of creative and meaningful life" (Preston 2014).

How does privacy provide shelter for "imaginative freedom, curiosity and self-reflection"? It doesn't seem obvious that privacy is *necessary* to pursue those ends. Imagine an artists' commune where everyone lives openly in an environment of tolerance, mutual respect, and intellectual freedom. For the sake of argument, let's assume the commune was also founded in perpetuity by a large foundation grant so that it was financially self-sufficient, removing the need for members to worry about economics. Such a place would seem to be one where "imaginative freedom, curiosity and self-reflection" could co-exist with an absence of privacy.

I venture that Cohen's claim about the value of privacy relies on the empirically contingent fact that people, in social settings, can be pressurised to behave in ways contrary to how they would otherwise, in the absence of others. The Asch experiments, mentioned in Chapter 8, provide one vivid example. More generally, the phenomenon we are discussing is people's willingness to conform to a *social norm*. Part of the value of privacy, then, is its ability to protect people's autonomy by reducing the influence of others.

Bicchieri (2005) introduced an influential definition of a social norm. The following definition from her later work, *Norms in the Wild*, is an updated version:

> A *social norm* is a rule of behaviour such that individuals prefer to conform to it on condition that they believe that (a) most people in their reference network conform to it (empirical expectation), and (b) that most people in their reference network believe they ought to conform to it (normative expectation). (Bicchieri 2017, p. 35)

From this definition, one problem with Cohen's defence of the value of privacy becomes apparent. If individuals *prefer* to conform to a particular rule, given that the empirical and normative expectations are met, how can we say a person's autonomy is undermined by their conforming to the norm? Being able to act in accordance with one's *preferences* seems, in part, a requirement of autonomy.

Perhaps what is at stake concerns the source of an individual's preference. A heroin addict's preference for heroin might have been acquired through a free

and autonomous choice of the person, yet at the same time, the preference created by the addiction prevents the addict from exercising autonomy. Similarly, a preference to conform to a particular norm can undermine a person's autonomy unless that preference to conform to a norm is the right kind of preference (i.e., not like that of the addict) and acquired in the right way (i.e., not through coercion). The difficulty with this line of defence is that it makes the value of privacy turn on a distinction between legitimately and illegitimately acquired preferences and how they are generated through one's social existence. If there is a value to privacy, and I believe that there is, it should admit a more direct and substantive defence.

A different line of argument can be found in John Stuart Mill's seminal work *On Liberty*:

> As it is useful that while mankind are imperfect there should be different opinions, so it is useful that there should be different experiments of living, that free scope should be given to varieties of character short of injury to others, and that the different modes of life should be proved practically. (Mill 1859)

Here, the idea is that privacy provides a degree of protection for persons to engage in "experiments of living". The difference between this defence and Cohen's defence is that, here, *society* is the beneficiary of a right to privacy; for Cohen, it was the individual. Another advantage of this defence of privacy is that it doesn't require an appeal to any *specific* values, such as "imaginative freedom, curiosity and self-reflection" or the pursuit of a "creative and meaningful life". If you didn't believe those were important, should you still value privacy? I think so; even philistines can value privacy.

There's another reason we should be sceptical of the claim that the transparent Open Society is nothing to worry about because it's just the natural state of humanity. The small bands and tribes that were the dominant form of social organisation for most of human existence were relatively *egalitarian*.[1] This matters because one fact about people is that our happiness not only depends on absolute outcomes (e.g., am I getting enough food?) but also on *relative* outcomes (e.g., how much is my opinion valued?). Life in a transparent Open Society may be unproblematic when the society is small and egalitarian because my life, compared to others, is about as good as everyone else's. However, when the basis of comparison in an Open Society expands to include millions or billions with great levels of inequality, the widespread knowledge of relative inequality can cause us to feel very differently about our lives than we would have otherwise.

One much-discussed example of how relative perceptions matter for people's perception of happiness is the Easterlin paradox. In 1974, Richard Easterlin argued that, despite the growth in income in the US over the past 50 years, there had been no corresponding increase in happiness (Easterlin 1974). Figure 13.1 illustrates this for both the US and five European countries for the

30-year period from 1973 to 2004. In both cases, happiness levels remained flat despite considerable increases in income in all countries. Why is this? Layard (2005, p. 45) explains the phenomenon as follows: "People are concerned about their relative income and not simply about its absolute level. They want to keep up with the Joneses or if possible to outdo them."

Easterlin's claim has not gone unchallenged. It has been suggested that a link between happiness and income does exist in developing countries which began with a low level of GDP per capita. Easterlin's original observations have been reconciled with this by theorising that once a country has become wealthy enough to satisfy the basic needs of its people, happiness starts to level off, with other life matters becoming more important. Clark *et al.* (2008, p. 96) note that, "It has been argued that once an individual rises above a poverty line or 'subsistence level,' the main source of increased well-being is not income but rather friends and a good family life." Stevenson and Wolfers (2008, p. 9) go further, disputing the existence of a satiation level at all: "new large-scale datasets covering many countries point to a clear, robust relationship between GDP per capita and average levels of subjective well-being in a country. Furthermore, we find no evidence that countries become satiated – the positive income-happiness relationship holds for both developed and developing nations." In response, Easterlin *et al.* (2010) rejects the Stevenson and Wolfers findings on the grounds that it rests "almost entirely on the short-term positive association between life satisfaction and GDP" in certain transition countries.

One concern is that much of this debate is taking place at too high a level of abstraction to be truly useful. Setting aside the question over whether any link between happiness and income is meaningful over a 30-plus year period, given all the cultural and political shifts that occur,[2] we might ask what happens if we shift the focus from wealth to types of consumption. In an interesting study, Hsee *et al.* (2009) investigate the connection between happiness and how it is generated by money, acquisition, and consumption. Intuitively, you might think that a person who receives a certain amount of money experiences happiness by reflecting on her expected consumption. However, Hsee *et al.* found that this isn't always the case.[3] They write:

> We posit that utility of money has two rather independent components: its value *per se* (monetary experience) and its consumption consequence (consumption experience). These two types of happiness obey different hedonic principles: Monetary experience depends on relative monetary value, whereas consumption experience depends on absolute consumption level. (Hsee *et al.* 2009, p. 400)

Something similar was found if one measured the happiness associated with the *acquisition* of a good, separate from its consumption. But it would be too quick to conclude that happiness produced by consumption only depends on

Figure 13.1: Two illustrations of the Easterlin paradox

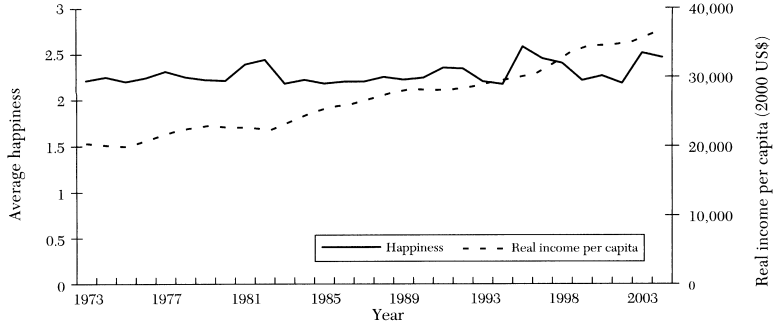

Notes: Happiness is the average reply to the question: "Taken all together, how would you say things are these days? Would you say that you are…?" Responses were coded as (3) Very Happy, (2) Pretty Happy, and (1) Not too Happy.

(a) Happiness and Real Income Per Capita in the United States, 1973–2004

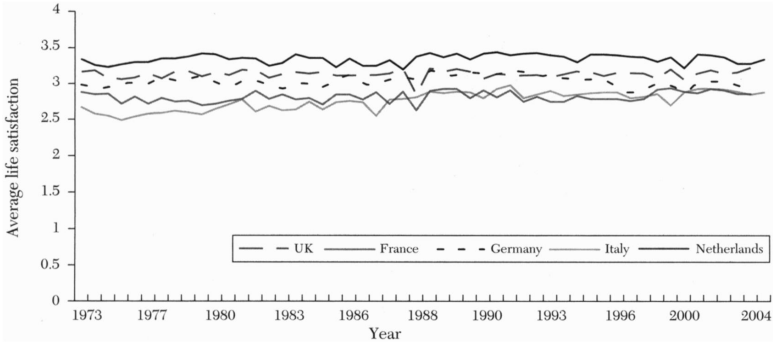

Notes: Happiness is the average reply to the question: "On the whole how satisfied are you with the life you lead?" Responses were coded as (4) Very Satisfied, (3) Fairly Satisfied, (2) Not Very Satisfied, and (1) Not at all Satisfied.

(b) Life Satisfaction in Five European Countries, 1973–2004

Source: Figures 1 and 2 in Clark et al. (2008). Copyright AEA, reproduced with permission from the authors.

absolute levels of the good being consumed; it also depends on the kind of good. Goods such as a warm house or a delicious cup of coffee will produce happiness corresponding to an absolute evaluation, whereas a good such as expensive jewellery will produce happiness according to relative evaluations.

The point is this: relative comparisons matter greatly for human happiness, regardless of whether we are talking about the receipt of money, the acquisition of a good, or even the *consumption* of certain types of goods (such as the wearing of jewellery). Constant exposure to the Rich Kids of Instagram can cause a person who would otherwise be content to become unhappy. This matters because the role played by relative comparisons is often lost in the narrative of capitalism, with the rhetoric of a rising tide lifting all boats. Reflecting on the US, Andy Warhol put the point well:

> What's great about this country is that America started the tradition where the richest consumers buy essentially the same things as the poorest. You can be watching TV and see Coca-Cola, and you know that the President drinks Coke, Liz Taylor drinks Coke, and just think, you can drink Coke, too. A Coke is a Coke and no amount of money can get you a better Coke than the one the bum on the corner is drinking. All the Cokes are the same and all the Cokes are good. Liz Taylor knows it, the President knows it, the bum knows it, and you know it. (Warhol 1975)

Warhol is right that all the Cokes are the same and all the Cokes are good, but you won't enjoy your Coke as much in your one-bedroom shotgun shack if you are looking at a billionaire drinking one on his megayacht. With the top 1% of global wealth holders now possessing 50.1% of all wealth, according to a 2017 Credit Suisse report, we have moved very far away indeed from the equality of the natural state of humankind.

Interestingly, this problem of perception gets worse as you get richer. Catherine Rampell argues that one reason why many rich people don't feel particularly rich is because of how *unequal* the wealth distribution is, even at the very top (Rampell 2011). Figure 13.2 plots the log of income for the various percentiles in the US for 2010. What's striking is how rapidly inequality *increases* towards the upper end of the plot. An income of $10,000 puts one just below the 10th percentile, but *quintupling* the income to $50,000 would move that person to just above the 54th percentile. However, when a person with an income of $250,000, right above the 96th percentile, increases their income to $1,250,000, they would just climb above the 99.5th percentile. Paul Krugman elucidates this well with the following metaphor: imagine society "as being something like a long street running up a hill, in which rising altitude goes along with rising income" (Krugman 2011). This street metaphor offers a psychological explanation for why the rich don't feel as happy as we might expect: people don't consider their overall position with respect to the entire

Figure 13.2: Natural log of income versus percentiles for 2010 US

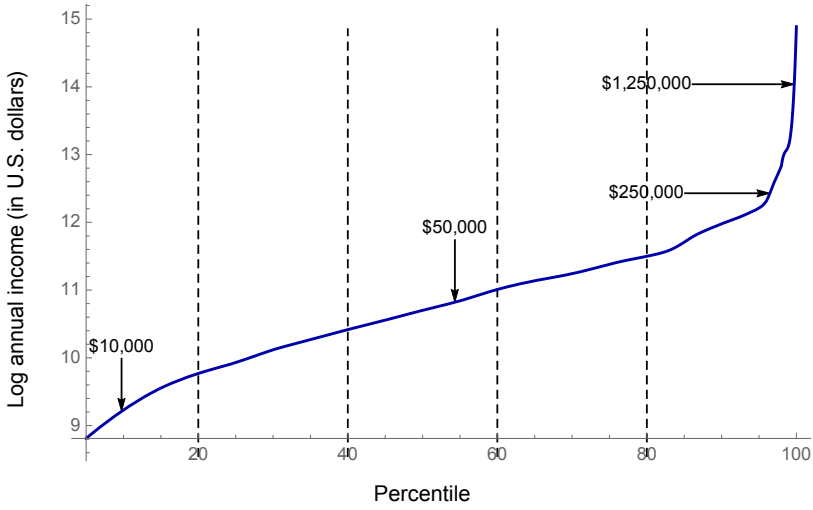

Source: created by the author using data from Rachel Johnson, Urban-Brookings Tax Policy Center Microsimulation Model (version 0509–7).
Notes: the idea to plot log income rather than just income versus percentiles is due to DeLong (2011).

street but only compare themselves to their neighbours on either *side*. For people in the lower percentiles, the difference in income (altitude) on either side is about the same: your worse-off neighbour differs from you about as much as you differ from your better-off neighbour. Yet when you look at people on the upper percentiles, that symmetry is lost; your worse-off neighbour differs from you much less than you differ from your better-off neighbour. If happiness is based on relative comparisons, this is a precarious position to be in; it will be perceived as much easier to slip down the ladder than to climb up because the distance between the rungs increases the higher you climb.

And so another problem with the transparent Open Society is that it can lead to greater levels of unhappiness because of the constant awareness of relative inequality, *regardless of what level one is at*. In addition to the decrease in happiness, there's another problem that's been identified – social anxiety. The constant exposure to news and images of other people's experiences on social media has given rise to a new term, *Fear of Missing Out* (FOMO). Przybylski *et al.* (2013) define FOMO as: "a pervasive apprehension that others might be having rewarding experiences from which one is absent." Although it hasn't yet made it into that great catalogue of mental illnesses published by the American Psychiatric Association, the DSM, Oxford Dictionaries Online added "FOMO" in 2013 (Dirda 2013).

There are a couple of reasons why the fear of missing out has become so salient. The first has to do with the immediacy and detail with which people can chronicle their experiences. Fifty years ago, photography was slow or expensive. You either sent film away to be developed, or you developed it yourself, or you used an instant camera. In 1972, the Polaroid SX-70 instant camera cost $180, with film costing $6.90 for a pack of ten pictures. In 2017 dollars, that would be $1,031 for the camera and $40 for ten pictures. Home videos were a rarity, and telephone calls were expensive. As a result, sharing details of experiences was difficult.

Now fast-forward to April 2023 (when I am writing this); the cheapest iPhone costs $429. That phone gives you virtually unlimited digital photographs and video, all of which can be shared immediately. Voice and video calls can still be expensive if you aren't on the right data plan (some things never change), but that's trivially avoidable with an internet connection and apps. As a result, people don't need to carefully select what experiences they share; anything they want to share can be shared.

This ease of sharing interacts with the second reason why FOMO is so salient: people don't generally want to broadcast all aspects of their life, but only the better parts, carefully curated. In a society where personal identity is a commodity and people are urged to develop their personal brand, people want to present themselves positively – although context-dependency can make this tricky. For example, photographs of someone having a fantastically hedonistic experience in Magaluf may not be well received by a potential future employer, as we have seen. But those photographs would, at least, show the person to be having a good time. It is much rarer to find someone sharing details of their misery in the moment without it being part of a larger narrative arc tending towards positivity.[4] There are the occasional ironic outliers, of course. "Sad desk lunch" does exist as a Tumblr site (Fisher 2015), but hasn't been updated since October 2015.

The ease of sharing experiences, along with a selection bias towards the positive, combines with a third cause of FOMO. When people engage with social media, they tend not to be in the midst of a positive experience of their own at the time of engagement. Instead, they are commuting, attempting to study in the library, in a lecture, at their desk at work, or bored at home. If they *are* engaging with social media while in the middle of an experience – at a club or at a concert – they are typically looking at experiences *different in kind* from the one they are having. In the first case, the contrast between one's banal existence with the positive experience seen will obviously be striking. In the second case, one might well wonder if the person's current experience, although nice enough, is still not quite as good as what the others are experiencing. These comparisons will not always be negative, of course; my point is that the comparisons will often be enough to plant a kernel of doubt – the fear of missing out.

What we have seen in this chapter and the previous one are some of the individual consequences of making so much information about our private lives

readily available. We focused on issues surrounding our inability to escape our past and carve out a new identity for ourselves (e.g., Chapter 12) and the psychological effects of being constantly exposed to information about *other* people (e.g., this chapter). The focus of these chapters has primarily been on how information about ourselves and other people can affect us in a variety of ways. However, another obvious consequence of making information about ourselves available is that people can *use* that information to *manipulate* us without us being aware of it. That is the topic of the next chapter.

14. We'll be watching you

Jeremy Bentham, the co-founder of utilitarianism with James Mill, is remembered for many things – the idea that the measure of right and wrong was "the greatest happiness of the greatest number", his work in economics and the law, his commitment to gender equality and reforming attitudes towards homosexuality, and the fact that his body was preserved after his death and put on permanent display in University College London.[1] Yet, perhaps the idea most closely associated with him is one related to an obsessive 20-year project on prison reform: the Panopticon. Bentham's Panopticon has become synonymous with surveillance, monitoring, and social control. The basic design is a building in which all prisoners are housed in cells observable by a single person (see Figure 14.1) without the prisoners *knowing* that they are being watched. This was meant to instil constant fear in the prisoners so that they had to behave all the time.

What we have seen is how the transparent Open Society creates a panopticon of the soul. Your desires are knowable by tracking what you buy or what you read and watch on the internet. GPS in mobile phones makes your movements knowable whenever you leave the house. The ability of any microphone-equipped device, whether it is a mobile phone or a virtual assistant, to be turned on silently means that, in principle, your personal conversations could be recorded and shared without your knowledge. And the explosion of social media means that people's beliefs and interests, hopes and fears are logged in real time. Add in the risk of public shaming, losing your job because your employer spots something they don't approve of (or of not being hired in the first place), all confounded by the indelible shadow of the past. In many ways, this is more troubling than Bentham's Panopticon because, there, only your *body* was being watched; now, people can get inside your head. The ability to pressure people to conform is great. But there is a greater problem.

The greater problem is that we know people's behaviour systematically fails to be rational in a number of ways. This was first revealed in a joint work by Daniel Kahneman and Amos Tversky in the early 1970s which led to Kahneman getting the Nobel Prize in Economics in 2002,[2] followed a few years later by Richard Thaler receiving the 2017 Nobel Prize in Economics for his work in behavioural economics. People's systematic deviation from rational behaviour goes well beyond framing effects, confirmation bias and nudging, extending into political behaviour. When people's choices are strongly influenced by things other than facts, the more you know about what someone believes and wants, the more power you have to influence that person's behaviour indirectly, without them knowing.[3]

Figure 14.1: Bentham's Panopticon

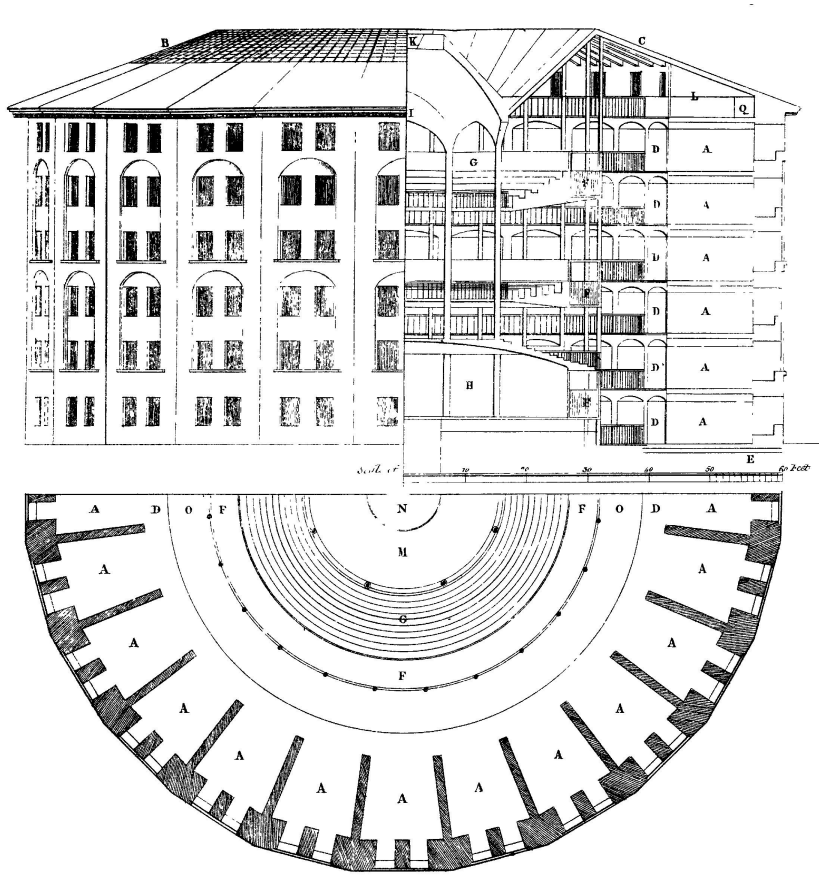

Source: Reveley (1791). Available under CC BY-SA 4.0 via Wikimedia Commons.

Democracies have long posed a problem in how to explain people's voting behaviour using the rational actor model. For starters, it's not clear why anyone would bother to vote. In large populations, the chances of any single voter being pivotal in an election are essentially zero. (That said, examples of single voter pivotality do exist – the ironically named Marcus "Landslide" Morton was elected governor of Massachusetts in 1839, receiving 51,034 of 102,066 votes.) Since voting costs time and effort, why would anyone turn out to vote if it probably won't make a difference? A lot of ink has been spilt trying to show, despite appearances to the contrary, that voting is rational.[4]

There's another assumption lurking in the background: that people are generally responsive to reasons when they decide for whom to vote. This is why democracies have held political debates where candidates discuss and argue about policy and why politicians have worried about their stance on certain issues and articulated an overarching coherent worldview to ground their positions. People, according to rational voter theory, listen to debates, weigh reasons for and against, and ultimately plump for the side they find most persuasive. Unfortunately, research suggests this model is far from the truth.

Westen (2008) conducted an experiment on partisan voters in the US, in the run-up to the 2004 election. Each voter was presented with six pairs of statements by John Kerry, six by George W. Bush, and six by "politically neutral" persons (e.g., Tom Hanks). The pairs of statements consisted of an initial assertion, followed by a second one which contradicted the first. After being presented with the conflicting information, voters were asked to rate on a scale of 1 to 4 the extent to which they thought the two statements were contradictory.[5] The point of the experiment was to identify how partisan voters, and their brains, reacted when exposed to such information.

The findings were what you might suspect. Partisan voters clearly identified the conflicts as contradictions for those on the *other* side of the political spectrum, but they were much less likely to see contradictions in statements made by people on *their* side. In addition, Westen found that "the neural circuits charged with regulation of emotional states seemed to recruit beliefs that eliminated the distress and conflict partisans had experienced when they confronted unpleasant realities." Interestingly, this was done with little activation of those parts of the brain that are involved in explicit reasoning. The most striking result was the following:

> Once partisans had found a way to reason to false conclusions, not only did neural circuits involved in negative emotions turn off, but circuits involved in positive emotions turned *on*. The partisan brain didn't seem satisfied in just feeling *better*. It worked overtime to feel *good*, activating reward circuits that give partisans a jolt of positive reinforcement for their biased reasoning. (Westen 2008)

So not only did partisan voters not see the contradictions made by their own side, they felt *good* when they didn't see them.

A lot of other research has reported on the less-than-rational aspects of how voters behave. Here's a small sample. Klor and Winter (2018) draw upon experimental and empirical evidence to argue for a "bandwagon effect" from polling data. That is, people who believe the election is close and are on the side with the slight majority are significantly more likely to turn out to vote.[6] Payne *et al.* (2010) suggest, based on survey data, that implicit racial prejudice might have affected the 2008 US presidential election, in that "even explicitly

rejected attitudes can influence important political decisions." They found that "Americans higher in implicit prejudice were less likely to vote for Obama, but not more likely to vote for McCain." (After Trump's 2016 election, we might also wonder whether people are much better at hiding explicit racial prejudice than we thought: see Stephens-Davidowitz 2017 for a discussion of how Google searches cast light on the seedy underbelly of the American psyche.)

The tendency to feel disgust also turns out to be a predictor of how people vote. Inbar *et al.* (2009) find that the tendency to feel disgust is correlated with people's political and moral attitudes; in particular, persons more likely to feel disgust are more likely to self-describe as conservative. Inbar *et al.* (2012) provide further evidence that the connection between being sensitive to disgust and political conservatism holds even when you control for other factors such as personality traits. Bloom (2014) found that when subjects are primed with associations of harm and the emotion of disgust, they evidenced greater moral conviction on political issues.

Finally, what are some factors connected with why voters prefer a candidate? Caprara and Zimbardo (2004) examined the self-reported personalities of voters and how they aligned with the public presentation of candidates amongst centre-right and centre-left candidates in Italy. They found a degree of affinity between the two: voters' personalities were more similar to the politicians they preferred than the politicians they didn't. Reflecting on this finding, they speculate:

> Either citizens' political preferences are in accord with their self-reported personality, with voting serving an expressive function with regard to self-perception, or that citizens assimilate their preferred candidates' personalities to their own. In both cases, personality characteristics that are reported and inferred may be critical to strengthening the bond between voters, parties and candidates. (Caprara and Zimbardo 2004, p. 586)

This is one reason why, in American elections, pollsters ask questions such as, "Which candidate would you most like to have a beer with?" Establishing an emotional connection with a candidate matters because, once someone has picked a candidate, that colours their interpretation of subsequent information.

Furthermore, when a person's behaviour doesn't conform to the rational actor model, knowing about things that *do* influence them allow for carefully crafted, targeted advertising to be aimed at them. Westen suggests:

> If you're running a campaign, you shouldn't worry about offending the 30 percent of the population whose brains can't process information from your side of the aisle unless their lives depend on it (e.g., after an attack on the US mainland). If you're a Republican, your focus should be on moving the 10 to 20 percent of

> the population with changeable minds to the right and bring your unbending 30 percent to the polls. (Westen 2008)

If you can identify people who are likely single-issue voters, you then saturate their social media with ads showing how your candidate is better on that issue. Eric Greitens, the Republican nominee for the governor of Missouri in 2016, released an advertisement consisting of nothing more than him firing a Gatling gun while a voice-over uttered a few platitudes. Such an ad wasn't designed to win over fence-sitters with subtle argumentation: it was a direct appeal to the gut of gun lovers. And one reason stories of Hillary Clinton's handling of official government emails while Secretary of State featured endlessly in 2016 was because, even though the FBI said there was nothing to it, repeated mentions of the emails eroded Clinton's perception of trustworthiness amongst fence-sitters. The thing about emotional judgements is that they don't need to be supported by *reasons*, much less *good* reasons. If getting a voter to feel a way towards a candidate is the primary thing that matters, then any method of forming that association is as good as another.[7] Given this, the real danger of fake news is that it can plant seeds of doubt in an undecided voter's mind, tipping them one way or another, and then legitimate news stories will be interpreted as providing *retrospective* reasons offered by a person to explain why they voted the way they did.

In addition, while elections are decided by who gets the most votes, we mustn't forget the electioneering tactic of trying to *suppress* the other side's vote. It doesn't matter if your candidate is less popular, all things considered, than the other person if enough of the *other* person's supporters stay at home.[8] Carefully targeted negative advertising seeks to depress voter turnout by lowering the enthusiasm of voters for candidates they might otherwise support. In Western democracies with low voter turnout, this can have important effects on the outcome and perceived legitimacy of the purported mandate.[9]

In 2017, reporters uncovered some of the methods pioneered by Cambridge Analytica to influence voter behaviour in the 2016 Brexit referendum, the 2016 presidential campaign in America, and elections in other countries around the world. These methods involve combining many of the sources of data we've discussed in this part of the book, and then targeting specific individuals. Cambridge Analytica acquired data from Facebook, constructed psychological profiles of individuals, and then combined that information with data purchased from other sources – "everything from magazine subscriptions to airline travel" (Cadwalladr 2017). These detailed psychological profiles allowed them to identify likely trigger issues for individuals: effective pressure points for nudging them in one direction or another. A former employee described their operations in 2013 as being similar to psychological warfare:

> That's what it is. Psyops. Psychological operations – the same methods the military use to effect mass sentiment change. It's

> what they mean by winning 'hearts and minds'. We were just doing it to win elections in the kind of developing countries that don't have many rules. (Cadwalladr 2017)

And that, of course, is the real problem. In the best of times, elections are a crude instrument for crafting future policy. Politicians aren't held accountable (at least, until the next election) for promising something they fail to deliver.[10] Voters project onto a candidate their beliefs about what that candidate stands for. A candidate becomes a Rorschach concept for each individual voter, and how a voter perceives a candidate is based partly on fact and partly on gut response. From a campaign manager's point of view, all that matters is whether you can get someone to prefer your candidate (or issue, if we are talking about a referendum) more than the alternatives. And that preference doesn't have to be based on true beliefs, because elections aren't about whether people have *good* reasons for their preferences. The transparent Open Society risks creating a world where we, the electorate, have our own information used against us, manipulating us like pawns in a political game.

Although I have focused on how our personal information can be used by others to manipulate us for explicitly *political* ends, the point is easily generalised. How many times have you had a conversation with a friend at a party and then, soon afterwards, seen an advertisement for the very thing you were speaking about? (See McNutt & Partners 2021 for a discussion of this phenomenon.) The point is that all of our mental attitudes – our beliefs, feelings, and desires – are vulnerable to being shaped by someone who has access to our information. Even when we think we are acting autonomously, making a decision based on evidence from the world around us, it doesn't hurt to keep a degree of critical awareness and ask *who benefits* from our choice? The fact that we *believe* we made our own decision is entirely compatible with that decision also having been shaped to further someone else's interest.

15. Concluding remarks

What, then, is the overall assessment of the transparent conception of the Open Society? The first point is that when faced with all these issues it is easy to feel a sense of despair at reversing the trend. The technological drive is towards ever more information transparency and more collection of data, creating more opportunities for our personal information to be misused. There is simply too much money at stake and too many powerful institutions (governmental or otherwise) with a vested interest in collecting as much information as possible about individuals by any means possible. The Stasi would have loved the information gathering methods which now exist.[1]

However, a second point is that our age of information transparency, in many other instances, provides social goods worth preserving. Many people want to know whether our pension funds are being invested in tobacco companies or armaments industries. We would like to see *all* of the trials performed by drug companies, not just the ones showing a drug performed better than placebo with sufficiently few side effects that it could be approved for the market. We want to know who funded certain studies or policy papers, who paid for political advertisements, and who gave what to politicians or judges. And it is important, sometimes, to know whether people with certain criminal records live in the neighbourhood or have applied for certain jobs. Any attempt to prevent potential abuses of informational transparency needs to be weighed against the potential benefits such transparency can provide. But one message which emerges from our investigation is that the current realisation of the transparent conception of the Open Society is too much in favour of making the lives of ordinary people transparent to companies and governments, and too little in favour of making the operations of companies and governments transparent to ordinary people.

Do we want to live in a panopticon of the soul? As we have seen, the abuse of personal information means that in countries with lax labour protections, individuals can be fired from their job for engaging in *legal behaviour* during *out-of-work hours* as a *private citizen*. And the way information was handled in the past allowed companies such as Cambridge Analytica to use personal information to craft fine-tuned campaign advertisements targeting hot-button issues likely to trigger an emotional response in the viewer, who would not know that they were being deliberately manipulated. When it is impossible to prevent our information from being collected, we need protection to stop our information being weaponized against us.

Striking the right balance will not be easy. It requires a public discussion about the kind of society we want to live in, and how the power conferred by the possession of information should be distributed across society. For example, should it be the case that social media companies can retain the information they collect indefinitely? In asking that question, let's set aside what the actual terms and conditions said when you agreed to use the platform: those terms and conditions were often written during the Wild West period of the internet, when services were developing faster than social norms or the law could adapt. Should it be the case that companies can access the information – anonymised or not – which people voluntarily disclose? It would be possible for social media companies to use end-to-end encryption so that *they* were unable to access the information stored on their servers even though users could. But that would prevent companies from monetising the information of its users, and it would no longer be possible for them to provide the service for free. Would advertisers be less willing to advertise on a service if they cannot aim highly targeted ads at users, given the already poor rate at which ad impressions are converted into useful outcomes?[2] Would users be willing to pay for it? If social media provides a valuable public service – the digital equivalent of a public square (albeit often a shouty, misanthropic one) – should that kind of service be brought into public ownership and regulated by an independent, politically neutral body?

If we want to rebalance transparency in the Open Society, what kind of regulation should be put in place, and what kind of trade-offs are we willing to accept regarding our privacy and the material benefits generated? One worry I have is that conversation is not being held. Another worry concerns the move from *regulation* to *enforcement*. Although it is all well and good to pass laws regarding how companies can use our information, laws provide no protection if they are not enforced. And even if they are enforced, the punishment for violation needs to be sufficiently severe so as to avoid companies simply pricing in fines as part of the "cost of doing business".[3]

But we have faced bleak situations in the past and have managed to turn things around. The economic advancements of the Industrial Revolution led to atrocious working conditions for many, yet laws were introduced, and the most egregious evils of that system were curtailed. Food safety standards were introduced, and environmental protection legislation was passed. These advances need continuous protection, but advances *were* made. And although many parts of the world still suffer from the excesses of exploitative capitalism, it is worth remembering the following statement from Popper's philosophical colleague: "no social tendency exists which could not be altered if the individuals concerned both wanted to alter it and possessed the appropriate information" (Watkins 1953). When faced with deepfake videos and troll factories using AI to carpet-bomb the internet with misinformation, there are options. We could, for example, choose as a society to have less technology in our lives. Why we will probably not choose that outcome is an interesting question, and one which I address in the final chapter of this book.

It could be argued that it is not possible to curtail the transparent conception of the Open Society simply because there is too much money to be made and too much power at stake. But I think that conclusion is wrong for two reasons. First, although economics is a hugely important determinant of social practice, it is not the *sole* determinant. A nice historical example to illustrate the point can be found in the abolition of slavery by the British empire.[4] At the time slavery was abolished, there were still vast sums of money to be made from the institution. The abolitionists were able to push through radical change even though this was against the economic interests of a great many people and cost the government of the time a lot of money.[5] The abolitionists, through effort and moral persuasion, eliminated a practice which had existed for several hundred years *despite* its continued profitability.[6] This shows that, in principle, it is possible to change social practices for the betterment of humanity even in the face of economic interests. If we were take a principled stand and resist the panopticon of the soul, the fact that there are enormous profits to be made by mining our information does not mean resistance is futile.

Second, I think there are good moral reasons to think that, even though there are vast sums of money to be made through the collection, analysis, and sale of personal information, this is not the kind of world we should want. I've outlined a number of reasons to think this in this part of the book, but let me mention one last reason to bring this chapter to a close. The reason involves a transformation in the way we think about human beings, and a human's ultimate source of value.

You might have thought it was bad enough when capitalism brought about the transformation of individuals into *consumers*. Instead of conceiving of people as agents with hopes, ambitions, goals, and projects they want to see realised, we shifted the focus to the products and resources used along the way. As long as people consumed, increasing economic activity and helping GDP grow, we paid less attention to the side effects of all that consumption. Data mining, unleashed upon the boundless quantities of information about individuals, runs the risk of transforming *consumers* into mere *commodities*. The second form of the Categorical Imperative resonates here: "Act in such a way that you treat humanity, whether in your own person or in the person of any other, never merely as a means to an end, but always at the same time as an end." Human beings should not *just* be commodities in an economic game, where our information is collected, analysed, and sold so that we can be manipulated. The fundamental values of democracy, the Enlightenment, and civil society risk being undermined when we allow our information to be turned against us, transforming us into a mere means to bring about another person's end, often at our own expense.

Notes to Part II: The panopticon of the soul

8. The book of life

[1] Published as *The Inman Diary: Volumes 1–2: A Public and Private Confession*, Harvard University Press (1990). I suggest an electronic version, if possible, as the hardback is 1,600 pages.

[2] The *New York Times* published an obituary shortly after his death in 2007 (Martin 2007). Those interested in reading the contents of his diary will be disappointed, as the terms and conditions of Shields' bequest to Washington State University include the requirement that no one be permitted to read the diary until 2057. However, if you search Google you can find excerpts, hence my claim about the "phenomenally boring prose".

[3] Bertrand Russell introduced the Tristram Shandy paradox in his 1903 book, *The Principles of Mathematics*. The original purpose of the paradox was to illustrate one of the antinomies of set theory: that an infinite collection can be put into a one-to-one correspondence with a proper subpart. According to the paradox, Tristram Shandy writes his diary very slowly, taking one year of his life to record the events of one day. However, if Shandy lives for an infinite length of time, he will complete his diary even though he always falls further behind. How is this possible? Shandy finishes the diary entry for the N^{th} day of his life at the end of the N^{th} year. If he were to fail to complete his diary, there must be some day M whose events are not adequately reported. But we can say with certainty when he finishes writing about the M^{th} day, and since he lives infinitely long each day gets recorded.

[4] In the original experiment, Asch (1951) put a single experimental subject in a group with seven confederates and asked each member of the group to match a given line with one of three other lines of unequal length. When the confederates stated their beliefs about the match, they were all *unanimously wrong*, with the margin of error ranging between half an inch (1.27cm) and one and three-quarters of an inch (4.44cm), in various trials. The error was thus clearly, *obviously*, recognisable. The point of the experiment was to see what people did when they were in a *minority of one* pitted against a *unanimous majority*. Asch found that, while some subjects defied the majority, one-third changed their judgements to the majority view in one-half or more of the trials. These results were confirmed in a later paper (Asch 1956).

9. Unwanted inferences

[1] Face recognition technology is already used throughout the world for security purposes and to increase ease-of-entry in public venues. Football stadiums are using it to prevent known hooligans from entering matches. On a more entertaining note, Cohen (2023) reported that James Dolan, the owner of Madison Square Garden, uses face recognition technology to identify lawyers representing people suing him to prevent them from entering the venue.

[2] Although, on this point, recall our discussion of Nozick's objection to the All-Affected principle in Chapter 4.

[3] For example, my mother was denied a job as a schoolteacher in northern Minnesota in the late 1960s because the official policy was that they didn't hire married women.

[4] One point of imprecision with this statistic is how to define what it means to be homosexual. In their survey, Janus and Janus (1993) found that 22% of men and 17% of women reported having had at least one homosexual experience. However, of this subset, only 39% of men and 27% of women said that they "frequently" had homosexual experiences or had ones which were "ongoing". These latter estimates are what generate the 9% and 5% estimates, respectively.

10. Lifting the veil

[1] To provide some context, there is a worry that when a society cracks down hard on minor, unintended violations of the law by generally law-abiding citizens while, at the same time, allowing criminals to get away with egregious violations, public attitudes towards the law and the police change. For example, in the UK, "from 2020, not a single personal, vehicle or bike theft was solved by police in between half and two-thirds of the 30,100 neighbourhoods in England and Wales" (Hymas 2023). At the same time, UK local councils "issued 19,631 parking fines per day in 2022, up 12 per cent from 2021 and a surge of 36 per cent from 2020" Saunders (2023).

[2] Modern smartphones, with inbuilt GPS, give a totally different solution to this problem.

[3] Dennett does talk about the environment in a later example involving a person attempting to control an automated fuel refinery by intervening in the environment where the refinery operates by depressing the aviation fuel market. But this is better read as an illustration of how the relation of control can, in some instances, be transitive. An agent A, by controlling B, who controls C, can thus be said to control C. However, the passage isn't presented as an example of how the control relationship has implicit contextual dependence.

[4] One could treat the budget as part of the environment. Yet, I wish to distinguish them. I take environmental dependency to refer to features outside of A or B that matter *at a particular moment in time* for A to control B, such as whether a remote has working batteries or not. The budget, on the other hand, involves resources that can be used or exchanged to make it possible for A to control B in the future.

11. Letting it all hang out

[1] In a famous study, Dunbar (1992) investigated a correlation between the size of the neocortex and the size of the social group in primates. The basic idea is

that maintaining stable social relations is cognitively expensive, which places an upper limit on the size of a social group. By examining data from 38 genera of social primates, Dunbar proposed that if humans follow the same trend, we can typically maintain around 150 stable relationships.

[2] In 2023, Twitter expanded the character limit to 4,000 if you were a US-based Twitter Blue subscriber.

12. Don't you forget about me

[1] It is closely related to the "right to erasure", and appears in Recitals 65 and 66 of the European Union's General Data Protection Regulation (GDPR). Recital 65, on the right of erasure, specifies that "a data subject should have the right to have his or her personal data erased and no longer processed where the personal data are no longer necessary in relation to the purposes for which they are collected or otherwise processed, where a data subject has withdrawn his or her consent or objects to the processing of personal data concerning him or her, or where the processing of his or her personal data does not otherwise comply with this Regulation." Recital 66, on the right to be forgotten, states: "To strengthen the right to be forgotten in the online environment, the right to erasure should also be extended in such a way that a controller who has made the personal data public should be obliged to inform the controllers which are processing such personal data to erase any links to, or copies or replications of those personal data."

[2] In recent years, California has attempted to address the most egregious injustices generated by the "three strikes" law. In 2012, Proposition 36 was passed, which reformed the law by requiring the third strike to be a serious or violent felony rather than any crime.

[3] In 2014, the UK's Advertising Standards Authority criticised Mondelez, the owner of Oreo brand biscuits, for paying vloggers in the UK to promote Oreos without labelling the videos as an advertisement. This shows that companies are clearly willing to exploit loopholes in regulation regarding social media when it's to their advantage.

[4] In the UK, the wealthy have a greater ability to exercise control over the Open Society than the rest of us through the use of injunctions to prevent the media reporting stories which they dislike. Furthermore, there has been increased use in what are known as "super-injunctions" that prevent the media from disclosing the fact that a wealthy individual has obtained an injunction preventing news organisations from reporting a story.

[5] Here we encounter, once again, the importance of Popper's observation that the Open Society should tolerate everything except intolerance. Firing someone from their job simply because they expressed a reasonable point is a high order of intolerance. Such intolerance deliberately cultivates a climate of fear, preventing discourse required as part of the Enlightenment conception of the Open Society, which we consider in Part III.

13. Returning to the past

[1] Service (1962) introduced four categories of human social organisation: bands, tribes, chiefdoms, and states. The smallest form of organisation, the band, typically consisted of about 10–30 individuals, often one or two extended families. The next form, the tribe, consisted of about 200 individuals. Each stage in the transition from bands to states is accompanied by a corresponding increase in size, political centralisation, and social stratification. The first state came into existence only around 3,400 BC, in the Fertile Crescent (see Diamond 2012, p. 12).

[2] For example, during the 30-year period from 1973 to 2004, US society experienced (a) the legalisation of abortion with the Roe v. Wade decision in 1973; (b) the end of the Vietnam war in 1975; (c) a significant push by the women's movement attempting to ratify the Equal Rights Amendment (which failed in 1982); (d) the growing conservative backlash against the counterculture movement of the 60s; (e) longstanding problems concerning inflation; (f) Reagan's election, followed by a recession; (g) the end of the Cold War and the collapse of the Soviet Union; (h) the first Gulf War; (i) the bursting of the the dot-com bubble; and (j) the 9/11 attacks and subsequent "war on terror". (That's just a subset of the major events.) Why focus on just GDP?

[3] In one experiment, they divided a pool of subjects into two groups: a "poor" group and a "rich" group, where members of the poor group received coupons worth either one or two points, and members of the rich group received coupons worth either five or ten points. (Notice that this divides each group into two: you could be a poor member of the poor group, with a coupon valued at only one point, or a rich member of the poor group, with a coupon valued at two points.) The coupon could be exchanged for a 100ml glass of milk, with the number of teaspoons of milk powder used to make the glass of milk corresponding to the number of points on the coupon. Subjects were asked to assess their happiness first upon receipt of the coupon, and then after drinking the glass of milk. What Hsee *et al.* found was that the happiness reported after receipt of the coupon didn't depend on whether a person was a member of the rich or poor group – it just depended on their relative positioning *within* that group. However, when people reported their happiness after consumption, the happiness level depended just on absolute properties of the milk.

[4] I want to distinguish between different kinds of negative valence and the forms it can take. People do post about family funerals or experiences of surviving abusive relationships or dysfunctional families. My point is that these stories often belong to a larger positive narrative of survival, endurance, and overcoming. Or they can illustrate involvement in individual or collective acts of protest against injustice. A person engaging in self-harm may post photos of the act, but that is generally viewed as a call for help. The act I am claiming to be rare is the sharing, by a person, of the kind of ordinary, day-to-day negative experiences which leave one feeling overwhelmed, out of control, alienated, disquieted, unhappy or helpless, without a concomitant attempt to seek

redress. From the point of view of managing your personal brand, failure, depression, anxiety, alienation, and anomie make poor copy.

14. We'll be watching you

[1] *Most* of it is still available for viewing by the public. The one exception is his head, which didn't respond to the embalming nearly as well as the rest of his body and went cheesy shortly after his death. Bentham's head was thus replaced by a wax replica. If you search the internet, you can find images of the current state of Bentham's preserved head, but I don't recommend it.

[2] Amos Tversky died in 1996 and the Nobel Prize cannot be received posthumously.

[3] Much has been written recently about the ethics and appropriateness of using "nudges" to shape people's behaviour (see Thaler and Sunstein 2009). While I think there are some cases where nudging is not ethical, my concerns here are with much more invasive and troubling applications.

[4] The main contending theories invoke additional concerns of the voter that make voting instrumentally rational despite the low probability of being pivotal. For example, if you believe that you have a *duty* to vote, then you will make an effort to vote because it is rational to do your duty (Mackie 2014). Alternatively, there is the expressivist theory of voting, which says that a person casts a vote in order to signal that they have certain values or belong to a certain group (Brennan and Lomasky 1993). This second point relates to issues we'll cover in Part IV.

[5] Yes, I know; two statements can be contradictory *to a degree* in the same way a woman can be a little bit pregnant. The statements weren't *literal* contradictions. They conflicted in that their natural interpretation, which required filling in some hidden assumptions, would be contradictory, yet they admitted other interpretations which reconciled the conflict.

[6] The explanation they give is that in such circumstances people overestimate their chances of being pivotal. One might also wonder if there is some regret aversion at work, too; imagine how bad you would feel in a close election where your side was ahead, but you *didn't* vote, only to have the other side squeak through to victory.

[7] In Chapter 16 of Part III, I'll explore this point further, looking at how laws protecting freedom of speech can be exploited by the unscrupulous for this purpose.

[8] Insidious methods of voter suppression include putting bureaucratic hurdles in place to discourage certain classes of people from voting, or denying people the right to vote due to prior convictions. As an example of the former method, Hebert and Lang (2016) report: "In North Carolina, the legislature requested racial data on the use of electoral mechanisms, then restricted all

those disproportionately used by blacks, such as early voting, same-day registration and out-of-precinct voting. Absentee ballots, disproportionately used by white voters, were exempted from the voter ID requirement […] The documents acceptable for proving voters' identity in North Carolina were the ones disproportionately held by whites, such as driver's licenses, US passports, and veteran and military IDs, and the ones that were left out were the ones often held by poor minority voters, such as student IDs, government employee IDs and public assistance IDs. The Texas voter ID law was designed the same way: There, officials accepted concealed-weapon licenses but not student or state employee IDs." An example of the second method is what is known as felony disenfranchisement. Lai and Lee (2016) note that 10% of the voting population in Florida, a major swing state, is disenfranchised due to prior felony convictions. However, my interest here lies with less blatant methods of attempting to influence voter turnout.

[9] The UK Brexit referendum may provide an illustration of this. The Brexit verdict was delivered by a 52% to 48% win (BBC News 2016), or 17,410,742 votes to 16,141,241. The population of the UK in 2016 was 65.4 million people, of which 46,501,241 were in the electorate. Of those eligible to vote, 72.2% participated, or approximately 33.6 million. The number of people who *could* have voted, but didn't, was more than 25 times the winning margin of Vote Leave.

[10] I will return to this point in Chapter 29.

15. Concluding remarks

[1] Space constraints have prevented me from covering more than a mere fraction of what is possible and what has been done. For example, consider Pegasus, the software originally designed by the NSO Group to track terrorists and members of organised criminal networks. Pegasus allowed governments to circumvent the security protections that were built into the software of smart phones, allowing the user's emails, text messages, photos, appointments, to be harvested. Pegasus even allowed governments to record the keystrokes of users *while using applications that provided encrypted communication*.

There's no doubt that this power gave police an advantage in the fight against crime. Pegasus led to the arrest of the Mexican drug lord El Chapo, and was used to break up a number of child-abuse rings (Kitroeff and Bergman 2023). However, interest in Pegasus spread well beyond what one might consider its reasonable use. Pegasus has been implicated in the assassination of Jamal Kashoggi. It was proven to have been installed on the phone of his wife, Hanan Elatr, months before his murder. In addition, the FBI purchased the Pegasus software "for product testing and evaluation" purposes *after* the software was blacklisted by the Biden administration (Mazzetti and Bergman 2023). Although there is, at the time I write this, no evidence that the FBI has used Pegasus against the US population, the FBI has given demonstrations of other

software, known as Phantom, that could hack American phones. A brochure obtained by the *New York Times* advertised Phantom to US law enforcement claiming that it could "turn your target's smartphone into an intelligence gold mine" (Levenson 2022).

[2] Measuring the efficacy of online advertising is tricky, but two standard key performance indicators are used – the click-through-rate (CTR), and the click-conversion rate (CCR). Because people do not click on an advertisement every time they see it, you need to first measure, on average, how many times you need to show an advertisement to people before someone eventually clicks on it. The CTR, a percentage, is defined as follows:

$$\frac{\text{Total number of ad clicks}}{\text{Total number of ad impressions}} \times 100 = \text{CTR}.$$

The CTR varies across social media platforms and the kind of ads we are talking about. For instance, the average CTR in 2023 for Facebook was 0.9% (Gardner 2023). So just over one hundred ad impressions were required in order to get a single user to click.

But not every click by a user results in a positive outcome for the company that placed the ad, be it a sale, a new contract, or whatever. (This is known as a "conversion".) This is where the CCR comes in. The CCR attempts to measure how many clicks are required in order to generate an outcome the company wants. It is defined as:

$$\frac{\text{Total number of ad coversions}}{\text{Total number of ad clicks}} \times 100 = \text{CCR}.$$

Again, taking Facebook as an example, the average CCR across industries is 9.31%. That means that just under one in every ten people who click on an ad takes further action to the company's benefit. So even *with* highly specific targeted advertising, it takes about 1,000 impressions in order to get one positive result.

[3] We will revisit some of these questions in Chapter 29.

[4] The abolition of slavery in the British empire took several decades. The slave *trade* was officially abolished by the Slave Trade Act in 1807, but existing slaves were not freed in colonies of the British empire until years later. The Slavery Abolition Act 1833 paved the way for the gradual elimination of slavery in most parts of the British Empire, with exceptions for territories controlled by the East India Company, Ceylon, and St. Helena. The 1833 Act freed slaves younger than six years of age, but converted slaves older than six to apprentices, with the apprenticeships ending in two stages: one lot ending on 1 August 1838 and the second lot ending on 1 August 1840. In addition, the 1833 Act set aside twenty million pounds (equivalent to nearly £2.2 billion in 2023) for the financial compensation of slave-owners. Slavery was not abolished in areas

controlled by the East India Company until the Indian Slavery Act 1843 was passed. However, *even then* slavery was replaced by other forms of indentured labour (see Dingwaney 1985).

[5] It has been argued (see Williams 1944) that the institution of slavery was already in decline as a result of industrialisation, and the abolitionist movement merely hastened its end. However, the modern consensus is that this was far from the case. Drescher (2010) argued that abolition of slavery actually occurred at a time when it was still very much economically viable, and that ending the slave trade imposed great expense on the government. A detailed discussion of this topic can be found in Macaskill (2022) (see, in particular, the section "The Contingency of Abolition").

[6] Although it is worth noting that the practice of slavery persisted longer than most people know: Mauritania was the last country to officially abolish slavery, in 1981.

References

Arendt, Hannah. (1959). *The Human Condition*. Anchor Books.

Asch, S. E. (1951). "Effects of Group Pressure Upon the Modification and Distortion of Judgments." In Guetzkow, Harold (ed.), *Groups, Leadership and Men: Research in Human Relations*, 177–190. Rutgers University Press.

Asch, Solomon E. (1956). "Studies of independence and conformity: I. A minority of one against a unanimous majority." *Psychological Monographs: General and Applied* 70 (9):1–70. https://doi.org/10.1037/h0093718

BBC. (8 March 2012). "Caution on Twitter urged as tourists barred from US." https://www.bbc.co.uk/news/technology-16810312

BBC News. (2016). "EU Referendum Results." https://www.bbc.co.uk/news/politics/eu_referendum/results

Bertrand, Marianne and Sendhil Mullainathan. (2004). "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *The American Economic Review* 94 (4):991–1013. https://doi.org/10.1257/0002828042002561

Bicchieri, Cristina. (2005). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press. https://doi.org/10.1017/CBO9780511616037

Bicchieri, Cristina. (2017). *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190622046.001.0001

Bloom, Pazit Ben-Nun. (2014). "Disgust, Harm, and Morality in Politics." *Political Psychology* 35 (4):495–513. https://doi.org/10.1111/pops.12053

Brennan, Geoffrey and Loren Lomasky. (1993). *Democracy and Decision: The Pure Theory of Electoral Preference*. Cambridge University Press. https://doi.org/10.1017/CBO9781139173544

British Broadcasting Corporation. (18 October 2016). "Jennicam: The first woman to stream her life on the internet." https://www.bbc.co.uk/news/magazine-37681006

Buell, Ryan W., Tami Kim, and Tsay Chia-Jung. (2014). "Cooks Make Tastier Food When They Can See Their Customers." *Harvard Business Review* 92 (11):34–35.

Cadwalladr, Carole. (May 7 2017). "The great British Brexit robbery: how our democracy was hijacked." *The Guardian*. https://www.theguardian.com/technology/2017/may/07/the-great-british-brexit-robbery-hijacked-democracy

Caprara, Gian Vittorio and Philip G. Zimbardo. (2004). "Personalizing Politics: A Congruency Model of Political Preference." *American Psychologist* 59 (7):581–594. https://doi.org/10.1037/0003-066X.59.7.581

Carlson, Nicholas. (13 May 2010). "Well, These New Zuckerberg IMs Won't Help Facebook's Privacy Problems." *Business Insider*. https://www.businessinsider.com/well-these-new-zuckerberg-ims-wont-help-facebooks-privacy-problems-2010-5

Catalano, Michele. (1 August 2013). "My family's Google searching got us a visit from counterterrorism police." *The Guardian*. https://www.theguardian.com/commentisfree/2013/aug/01/government-tracking-google-searches

Clark, Andrew E., Paul Frijters, and Michael A. Shields. (2008). "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *Journal of Economic Literature* 46 (1):95–144. https://doi.org/10.1257/jel.46.1.95

Cohen, Andrew. (16 February 2023). "Facial recognition in sports: The biometrics technology shaping ticketing, payments, security, more." *Sports Business Journal*. https://www.sportsbusinessjournal.com/Journal/Issues/2023/01/09/Technology/facial-recognition-technology-sports-stadiums.aspx

Cohen, Josh. (2013). *The Private Life: Why We Remain in the Dark*. Berkeley: Counterpoint.

Connellan, Michael. (27 January 2010). "Women suffer from Gandhi's legacy." *The Guardian*. https://www.theguardian.com/commentisfree/2010/jan/27/mohandas-gandhi-women-india

DeLong, Brad. (12 January 2011). "On the Richness of the Rich Once Again." https://delong.typepad.com/sdj/2011/01/on-the-richness-of-the-rich-once-again.html

Dennett, Daniel C. (1984). *Elbow room: The varieties of free will worth wanting*. MIT Press.

Diamond, Jared. (2012). *The World Until Yesterday*. Penguin Books.

Dingwaney, Manjari. (1985). "Unredeemed Promises: The Law and Servitude." In Patnaik, Utsa and Manjari Dingwaney (eds.), *Chains of Servitude: Bondage and Slavery in India*, chapter 9, 283–347. Sangam Books.

Dirda, Michael. (28 August 2013). "Oxford Dictionaries adds 'twerk,' 'FOMO,' 'selfie,' and other words that make me vom." *The Washington Post*. https://www.washingtonpost.com/lifestyle/style/oxford-dictionaries-adds-twerk-fomo-selfie-and-other-words-that-make-me-vom/2013/08/28/678ddd48-102c-11e3-8cdd-bcdc09410972_story.html

Doris, John M. (2002). *Lack of Character: Personality and Moral Behavior*. Cambridge University Press. https://doi.org/10.1017/CBO9781139878364

Drescher, Seymour. (2010). *Econocide: British Slavery in the Era of Abolition*. The University of North Carolina Press, 2nd edition.

Duhigg, Charles. (16 February 2012). "How Companies Learn Your Secrets." *The New York Times Magazine*. https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

Dunbar, Robin IM. (1992). "Neocortex size as a constraint on group size in primates." *Journal of Human Evolution* 22 (6):469–493. https://doi.org/10.1016/0047-2484(92)90081-J

Duncan, Pamela. (27 July 2017). "Gay relationships are still criminalised in 72 countries, report finds." *The Guardian*. https://www.theguardian.com/world/2017/jul/27/gay-relationships-still-criminalised-countries-report

Easterlin, Richard A. (1974). "Does Economic Growth Improve the Human Lot? Some Empirical Evidence." In David, R. and M. Reder (eds.), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, 89–125. New York: Academic Press. https://doi.org/10.1016/B978-0-12-205050-3.50008-7

Easterlin, Richard A, Laura Angelescu McVey, Malgorzata Switek, Onnicha Sawangfa, and Jacqueline Smith Zweig. (2010). "The happiness—income paradox revisited." *Proceedings of the National Academy of Sciences* 107 (52):22463–22468. https://doi.org/10.1073/pnas.1015962107

Fisher, Kira. (13 October 2015). "Sad Desk Lunch." https://saddesklunch.com

Gardner, Madelyn. (17 February 2023). "Social Advertising Benchmarks for 2024." https://www.brafton.com/blog/social-media/social-advertising-benchmarks/

Goldstine, Herman. (1980). *The Computer from Pascal to von Neumann*. Princeton University Press.

Grice, H. Paul. (1957). "Meaning." *The philosophical review* 66 (3):377–388. https://doi.org/10.2307/2182440

Hebert, J. Gerald and Danielle Lang. (3 August 2016). "Courts are finally pointing out the racism behind voter ID laws." *The Washington Post*. https://www.washingtonpost.com/posteverything/wp/2016/08/03/courts-are-finally-pointing-out-the-racism-behind-voter-id-laws/

Hooper, Richard. (3 November 2016). "The man who records his entire life." *BBC*. https://www.bbc.co.uk/news/magazine-37631646

Hsee, Christopher K, Yang Yang, Naihe Li, and Luxi Shen. (2009). "Wealth, warmth, and well-being: Whether happiness is relative or absolute depends on whether it is about money, acquisition, or consumption." *Journal of Marketing Research* 46 (3):396–409. https://doi.org/10.1509/jmkr.46.3.396

Hymas, Charles. (25 June 2023). "Police solve no thefts at all in most neighbourhoods." *The Telegraph*.

https://www.telegraph.co.uk/news/2023/06/25/theft-decriminalised-two-thirds-neighbourhoods-police/

Inbar, Yoel, David Pizarro, Ravi Iyer, and Jonathan Haidt. (2012). "Disgust sensitivity, political conservatism, and voting." *Social Psychological and Personality Science* 3 (5):537–544. https://doi.org/10.1177/1948550611429024

Inbar, Yoel, David A. Pizarro, and Paul Bloom. (2009). "Conservatives are more easily disgusted than liberals." *Cognition and emotion* 23 (4):714–725. https://doi.org/10.1080/02699930802110007

Inman, Arthur Crew. (1985). *The Inman Diary: A Public and Private Confession*, volume 1 and 2. Harvard University Press.

James, William. (1890). *The Principles of Psychology*, volume 1. New York, NY: Holt.

Janus, Samuel S. and Cynthia L. Janus. (1993). *The Janus Report on Sexual Behavior*. John Wiley & Sons, Inc.

Kachka, Boris. (27 May 2014). "Why One Norwegian Author's 6-Volume, 3,600-Page Book Is on the Verge of Breaking Out in America." *New York Magazine*. https://www.vulture.com/2014/05/karl-ove-knausgaard-boyhood-profile.html

Kerr, Ted. (2 July 2020). "39 Years Later, The New York Times' 1981 'Gay Cancer' Story Continues to Distort Early AIDS History." *The Body*. https://www.thebody.com/article/new-york-times-1981-gay-cancer-story-distorts-aids-history

Keynes, John Maynard. (1921). *A Treatise on Probability*. London: Macmillan.

Kinsey, Alfred C., Wardel B. Pomeroy, and Clyde E. Martin. (1948). *Sexual Behavior in the Human Male*. Philadelphia and London: W. B. Saunders Company.

Kitroeff, Natalie and Ronen Bergman. (18 April 2023). "How Mexico Became the Biggest User of the World's Most Notorious Spy Tool." *The New York Times*. https://www.nytimes.com/2023/04/18/world/americas/pegasus-spyware-mexico.html

Klor, Esteban F. and Eyal Winter. (2018). "On Public Opinion Polls and Voters' Turnout." *Journal of Public Economic Theory* 20 (2):239–256. https://doi.org/10.1111/jpet.12274

Krugman, Paul. (12 January 2011). "Why Does Inequality Make the Rich Feel Poorer?" *The New York Times*. https://archive.nytimes.com/krugman.blogs.nytimes.com/2011/01/12/why-does-inequality-make-the-rich-feel-poorer/

Kruse, Kevin. (2 December 2013). "Transparency Eats Culture For Lunch." *Forbes*.

https://www.forbes.com/sites/kevinkruse/2013/12/02/transparency-eats-culture-for-lunch/

Kumar, S. (22 May 2015). "Why Monitoring Employees' Social Media Is a Bad Idea." *Time*. https://time.com/3894276/social-media-monitoring-work/

Lai, K. K. Rebecca and Jasmine C. Lee. (October 6 2016). "Why 10% of Florida Adults Can't Vote: How Felony Convictions Affect Access to the Ballot." *The New York Times*. https://www.nytimes.com/interactive/2016/10/06/us/unequal-effect-of-laws-that-block-felons-from-voting.html

Laiou, Alexandra, Katerina Folla, George Yannis, Robert Bauer, Klaus Machata, Christian Brandstaetter, Pete Thomas, and Alan Kirk. (2016). "Comparative analysis of road accidents by gender in Europe." *Injury Prevention* 22 (Suppl 2):A305–A305. https://doi.org/10.1136/injuryprev-2016-042156.856

Layard, Richard. (2005). *Happiness: Lessons from a New Science*. London: Penguin.

Levenson, Michael. (28 January 2022). "F.B.I. Secretly Bought Israeli Spyware and Explored Hacking U.S. Phones." *The New York Times*. https://www.nytimes.com/2022/01/28/world/middleeast/israel-pegasus-spyware.html

Macaskill, William. (2022). *What We Owe The Future: A Million-Year View*. Oneworld Publishing.

Mackie, Gerry. (2014). "Why It's Rational to Vote." In López-Guerra, Claudio and Julia Maskivker (eds.), *Rationality, Democracy, and Justice: The Legacy of Jon Elster*, 21–49. Cambridge University Press. https://doi.org/10.1017/CBO9781107588165.005

Martin, Douglas. (29 October 2007). "Robert Shields, Wordy Diarist, Dies at 89." *The New York Times*. https://www.nytimes.com/2007/10/29/us/29shields.html

Martin, Joyce A., Brady E. Hamilton, Michelle J. K. Osterman, Anne K. Driscoll, and T. J. Mathews. (2015). "Births: Final Data for 2015." *National Vital Statistics Reports* 66 (1):1–69.

Mazzetti, Mark and Ronen Bergman. (2 April 2023). "A Front Company and a Fake Identity: How the U.S. Came to Use Spyware It Was Trying to Kill." *The New York Times*. https://www.nytimes.com/2023/04/02/us/politics/nso-contract-us-spy.html

McNutt & Partners. (25 January 2021). "Why We See Digital Ads After Talking About Something." https://www.mcnuttpartners.com/why-we-see-digital-ads-after-talking-about-something/

McRobbie, Linda Rodriguez. (8 February 2017). "Total recall: the people who never forget." *The Guardian*.

https://www.theguardian.com/science/2017/feb/08/total-recall-the-people-who-never-forget

Milgram, Stanley. (1967). "The small world problem." *Psychology Today* 2 (1):61–67. https://doi.org/10.1037/e400002009-005

Mill, John Stuart. (1859). *On Liberty*. London: John W. Parker and Son, West Strand.

Payne, B Keith, Jon A Krosnick, Josh Pasek, Yphtach Lelkes, Omair Akhtar, and Trevor Tompson. (2010). "Implicit and explicit prejudice in the 2008 American presidential election." *Journal of Experimental Social Psychology* 46 (2):367–374. https://doi.org/10.1016/j.jesp.2009.11.001

Posey, Kirby G. (2016). "Household Income: 2015 — American Community Survey Briefs." Technical report, United States Census Bureau. https://www.census.gov/content/dam/Census/library/publications/2016/acs/acsbr15-02.pdf

Preston, Alex. (3 August 2014). "The death of privacy." *The Guardian*. https://www.theguardian.com/world/2014/aug/03/internet-death-privacy-google-facebook-alex-preston

Przybylski, Andrew K., Kou Murayama, Cody R. DeHaan, and Valerie Gladwell. (2013). "Motivational, emotional, and behavioral correlates of fear of missing out." *Computers in Human Behavior* 29 (4):1841–1848. https://doi.org/10.1016/j.chb.2013.02.014

Rampell, Catherine. (11 January 2011). "Why So Many Rich People Don't Feel Very Rich." *The New York Times*. https://archive.nytimes.com/economix.blogs.nytimes.com/2011/01/11/why-so-many-rich-people-dont-feel-very-rich/

Robison, Jennifer. (8 October 2002). "What Percentage of the Population Is Gay?" *Gallup*. https://news.gallup.com/poll/6961/what-percentage-population-gay.aspx

Robson, David. (26 January 2016). "The blessing and curse of the people who never forget." *BBC Future*. https://www.bbc.com/future/article/20160125-the-blessing-and-curse-of-the-people-who-never-forget

Ronson, Jon. (2015). *So You've Been Publicly Shamed*. Riverhead Books.

Ross, Lee. (1977). "The intuitive psychologist and his shortcomings: Distortions in the attribution process." *Advances in Experimental Social Psychology* 10:173–220. https://doi.org/10.1016/S0065-2601(08)60357-3

Russell, Bertrand. (1937). *The Principles of Mathematics*. London: George Allen & Unwin, Ltd., 2nd edition.

Russell, Jon. (4 September 2017). "Alibaba debuts 'smile to pay' facial recognition payments at KFC in China." *TechCrunch*. https://techcrunch.com/2017/09/03/alibaba-debuts-smile-to-pay/

Saunders, Hayley. (12 January 2023). "Councils issue nearly 20,000 parking fines per day." *DirectLineGroup*. https://www.directlinegroup.co.uk/en/news/brand-news/2023/12012023.html

Service, Elman R. (1962). *Primitive Social Organization: an Evolutionary Perspective*. Random House, New York.

Sinner, Michele and Myles Neligan. (1 March 2011). "EU court bans insurers from pricing on gender." *Reuters*. https://www.reuters.com/article/uk-europe-insurance/eu-court-bans-insurers-from-pricing-on-gender-idUKTRE7201O720110301/

Smith, Michael French. (1994). *Hard Times on Kairiru Island: Poverty, Development, and Morality in a Papua New Guinea Village*. University of Hawaii Press: Honolulu. https://doi.org/10.1515/9780824843267

Solove, Daniel J. (2011). *Nothing to Hide: The False Tradeoff between Privacy and Security*. Yale University Press.

Statista. (December 2022a). "Number of Twitter users worldwide from 2019 to 2022." https://www.statista.com/statistics/303681/twitter-users-worldwide/

Statista. (May 2022b). "Number of unique WhatsApp mobile users worldwide from January 2020 to June 2022." https://www.statista.com/statistics/1306022/whatsapp-global-unique-users/

Statista. (February 2023a). "Number of monthly active Facebook users worldwide as of 4th quarter 2022." https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/

Statista. (March 2023b). "Number of monthly active users of Weibo Corporation from 1st quarter of 2014 to 4th quarter of 2022." https://www.statista.com/statistics/795303/china-mau-of-sina-weibo/

Statista. (March 2023c). "Number of monthly active WeChat users from 2nd quarter 2011 to 4th quarter 2022." https://www.statista.com/statistics/255778/number-of-active-wechat-messenger-accounts/

Stephens-Davidowitz, Seth. (2017). *Everybody Lies: What the Internet Can Tell Us About Who We Really Are*. Bloomsbury Publishing.

Stevenson, Betsey and Justin Wolfers. (2008). "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox." *Brookings Papers on Economic Activity* 2008:1–87. https://doi.org/10.1353/eca.0.0001

Thaler, Richard H. and Cass R. Sunstein. (2009). *Nudge: Improving Decisions about Health, Wealth and Happiness*. Penguin.

The Economist. (16 July 2015). "A few dollars less." *The Economist*. https://www.economist.com/finance-and-economics/2015/07/16/a-few-dollars-less

The Economist. (2017)a. "Advances in AI are used to spot signs of sexuality." *The Economist*.

The Economist. (26 October 2017)b. "Preventing ex-convicts from working is silly." *The Economist*. https://www.economist.com/united-states/2017/10/26/preventing-ex-convicts-from-working-is-silly

Tirrell, Meg. (14 December 2015). "Unlocking my genome: Was it worth it?" *CNBC*. https://www.cnbc.com/2015/12/10/unlocking-my-genome-was-it-worth-it.html

United Kingdom Home Office. (November 2016). "Country Policy and Information Note — Iran: Adulterers."

Walters, Joanna. (7 November 2017). "Woman who gave Donald Trump the middle finger fired from her job." *The Guardian*. https://www.theguardian.com/us-news/2017/nov/06/woman-trump-middle-finger-fired-juli-briskman

Wang, Yilun and Michal Kosinski. (2018). "Deep Neural Networks Are More Accurate Than Humans at Detecting Sexual Orientation From Facial Images." *Journal of Personality and Social Psychology* 114 (2):246–257. https://doi.org/10.1037/pspa0000098

Warhol, Andy. (1975). *The Philosophy of Andy Warhol: From A to B & Back Again*. Harcourt Brace Jovanovich.

Watkins, J. W. N. (May 1953). "Ideal Types and Historical Explanation." *The British Journal for the Philosophy of Science* 3 (9):22–43. https://doi.org/10.1093/bjps/III.9.22

Westen, Drew. (2008). *The Political Brain: The Role of Emotion in Deciding the Fate of the Nation*. Public Affairs.

Wetterstrand, Kris A. (1 February 2024). "DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP)." National Human Genome Research Institute. https://www.genome.gov/sequencingcostsdata

Williams, Eric. (1944). *Capitalism and Slavery*. Chapel Hill: The University of North Carolina Press.