

NIALL FERGUSON

The Square and the Tower

*Networks, Hierarchies and the Struggle
for Global Power*



ALLEN LANE

an imprint of

PENGUIN BOOKS

ALLEN LANE

UK | USA | Canada | Ireland | Australia
India | New Zealand | South Africa

Allen Lane is part of the Penguin Random House group of companies
whose addresses can be found at global.penguinrandomhouse.com



Penguin
Random House
UK

First published 2017
001

Copyright © Niall Ferguson, 2017


The moral right of the author has been asserted

Set in 10.5/14 pt Sabon LT Std
Typeset by Jouve (UK), Milton Keynes
Printed in Great Britain by Clays Ltd, St Ives plc

A CIP catalogue record for this book is available from the British Library

ISBN: 978-0-241-29046-0

www.greenpenguin.co.uk



Penguin Random House is committed to a sustainable future for our business, our readers and our planet. This book is made from Forest Stewardship Council® certified paper

Our Networked Age

Networks, it seems, are everywhere today. In the first week of 2017, the *New York Times* ran 136 stories in which the word ‘network’ appeared. Just over a third of the stories were about television networks, twelve were about computer networks and ten about various kinds of political network, but there were also stories about transport networks, financial networks, terrorist networks, health-care networks – not to mention social, educational, criminal, telephone, radio, electricity and intelligence networks. To read all this is to behold a world, as the cliché has it, ‘where everything is connected’. Some networks link militants together, others connect medics, still others are between automated teller machines. There is a cancer network, a jihadi network, an orca network. Some networks – too often referred to as ‘vast’¹ – are international, while others are regional; some are ethereal, others underground. There are networks of corruption, networks of tunnels, networks of espionage; there is even a tennis match-fixing network. Network attackers battle network defenders. And all of this is breathlessly covered by terrestrial, cable and satellite networks.

In *Bleak House* it was fog that was ubiquitous. Today it is networks that are, to borrow from Dickens, up the river and down the river. ‘The alternative to networking is to fail,’ we read in the *Harvard Business Review*.² ‘A key reason why women lag behind in leadership,’ the same journal asserts, ‘is that they are less likely to have extensive networks to support and promote them as potential leaders.’³ Another *HBR* article shows that ‘mutual fund portfolio managers placed larger concentrated bets on companies to which they were connected through an education network’, and that those

investments performed better* than average.⁴ However, not everyone would infer from this that the ‘old boy’ network is a benign force, worthy of emulation by old girls. In finance, some ‘expert networks’ have been revealed to be channels for insider trading or interest rate rigging.⁵ Networks have also been blamed for the global financial crisis of 2008: specifically, the increasingly complex network that turned the world’s banks into a global transmission and amplification system for losses on US subprime mortgages.⁶ The world described by Sandra Navidi in *Superhubs* may seem glamorous to some. In her words, a ‘select few’ – she names just twenty individuals – ‘preside over the most exclusive and powerful asset: a unique network of personal relationships that spans the globe’. These relationships are forged and maintained in an even smaller number of institutions: the Massachusetts Institute of Technology, Goldman Sachs, the World Economic Forum, three philanthropic entities, among them the Clinton Global Initiative, and the Four Seasons restaurant in New York.⁷ Yet one of the core messages of Donald J. Trump’s successful election campaign in 2016 was that these were the very ‘global special interests’ that stood behind the ‘failed and corrupt political establishment’ personified by Hillary Clinton, the candidate he defeated.⁸

No account of the 2016 US presidential election will be complete without a discussion of the roles played by media networks, from Fox News to Facebook to Twitter, the victorious candidate’s network of choice.† One of many ironies of the election was that Trump’s network-driven campaign directed so much of its fire at Clinton’s elite network – a network to which Trump himself had once belonged, as the Clintons’ presence at his third wedding attested. Just a few years before the election, an entity called ‘The Trump Network’ – set up in 2009 to sell products like vitamin supplements with Trump’s endorsement – had gone bankrupt. Had Trump lost the election, he would have launched Trump TV as a television network. One of the

* The return was 21 per cent when both portfolio manager and chief executive went to the same university and took the same degree with some chronological overlap, compared with 13 per cent when there was no such connection.

† At the time of writing, Donald J. Trump has 33.8 million followers on Twitter. He himself follows just forty-five individuals or institutions.

many reasons why he did not lose was that Russia's intelligence network did its utmost to damage his rival's reputation, using the website WikiLeaks and the television network RT as its principal instruments. In the words of a partly unclassified report by the US intelligence agencies, 'the Russian President Vladimir Putin ordered an influence campaign in 2016' that was intended to 'denigrate Secretary Clinton, and harm her electability and potential presidency', reflecting the Kremlin's 'clear preference' for Trump. In July 2015, according to the report, 'Russian intelligence gained access to Democratic National Committee (DNC) networks and maintained that access until at least June 2016', systematically publishing the emails it obtained through WikiLeaks. At the same time, 'Russia's state-run propaganda machine – comprised of its domestic media apparatus, outlets targeting global audiences such as RT and Sputnik, and a network of quasi-government trolls – contributed to the influence campaign by serving as a platform for Kremlin messaging to Russian and international audiences.'⁹

Another reason Trump won, however, was that the Islamist terrorist network known as Islamic State carried out multiple attacks in the twelve months before the election, including two in the United States (in San Bernardino and Orlando). These attacks enhanced the appeal of Trump's pledges to 'expose', 'strip out' and 'remove one by one . . . the support networks for Radical Islam in this country', and 'totally dismantle Iran's global terror network'.¹⁰

We live, in short, in 'the network age'.¹¹ Joshua Ramo has called it 'the Age of Network Power'.¹² Adrienne Lafrance prefers 'the Age of Entanglement'.¹³ Parag Khanna even proposes a new discipline – 'Connectography' – to map 'the Global Network Revolution'.¹⁴ 'The network society', according to Manuel Castells, 'represents a qualitative change in the human experience'.¹⁵ Networks are transforming the public sphere and with it democracy itself.¹⁶ But for better or for worse? 'Current network technology . . . truly favours the citizens,' write Google's Jared Cohen and Eric Schmidt. 'Never before have so many people been connected through an instantly responsive network', with truly 'game-changing' implications for politics everywhere.¹⁷ An alternative view is that global corporations such as Google are systematically achieving 'structural domination' by exploiting networks to

erode national sovereignty and the collectivist politics that it makes possible.¹⁸

The same question can be asked of the effect of networks on the international system: for better or for worse? For Anne-Marie Slaughter, it makes sense to reconfigure global politics by combining the traditional ‘chessboard’ of inter-state diplomacy with the new ‘web . . . of networks’, exploiting the advantages of the latter (such as transparency, adaptability and scalability).¹⁹ The stateswomen of the future, she argues, will be ‘web actors wielding power and exercising leadership alongside governments’ with ‘strategies of connection’.²⁰ Parag Khanna looks forward with relish to a ‘supply-chain world’ in which global corporations, megacities, ‘aerotropolises’ and ‘regional commonwealths’ engage in an endless but essentially peaceful ‘tug-of-war’ for economic advantage that resembles ‘a massive multiplayer game’.²¹ Yet it seems doubtful – not only to Joshua Ramo, but also to his mentor Henry Kissinger – that such tendencies are likely to enhance global stability. ‘The pervasiveness of networked communications in the social, financial, industrial, and military sectors,’ Kissinger has written:

has . . . revolutionized vulnerabilities. Outpacing most rules and regulations (and indeed the technical comprehension of many regulators), it has, in some respects, created the state of nature . . . the escape from which, according to Hobbes, provided the motivating force for creating a political order . . . [A]symmetry and a kind of congenital world disorder are built into relations between cyber powers both in diplomacy and in strategy . . . Absent articulation of some rules of international conduct, a crisis will arise from the inner dynamics of the system.²²

If the ‘first world cyberwar’ has already begun, as some have claimed, then it is a war between networks.²³

The most alarming prospect of all is that a single global network will ultimately render *Homo sapiens* redundant and then extinct. In *Homo Deus*, Yuval Harari argues that the age of large-scale ‘mass cooperation networks’ based on written language, money, culture and ideology – products of carbon-based human neural networks – is giving way to a new era of silicon-based computer networks based on

algorithms. In that network, we shall quickly find ourselves about as important to the algorithms as animals currently are to us. Disconnection from the network will come to mean death for the individual, as the network will be maintaining our health around the clock. But connection will ultimately mean extinction for the species: 'The yardsticks that we ourselves have enshrined will condemn us to join the mammoths and Chinese river dolphins in oblivion.'²⁴ On the basis of Harari's bleak assessment of the human past, these would seem to be our just deserts.²⁵

This book is about the past more than it is about the future; or, to be precise, it is a book that seeks to learn about the future mainly by studying the past, rather than engaging in flights of fancy or the casual projection forward of recent trends. There are those (not least in Silicon Valley) who doubt that history has much to teach them at a time of such rapid technological innovation.²⁶ Indeed, much of the debate I have just summarized presupposes that social networks are a new phenomenon and that there is something unprecedented about their present-day ubiquity. This is wrong. Even as we talk incessantly about them, the reality is that most of us have only a very limited understanding of how networks function, and almost no knowledge of where they came from. We largely overlook how widespread they are in the natural world, what a key role they have played in our evolution as a species, and how integral a part of the human past they have been. As a result, we tend to underestimate the importance of networks in the past, and to assume erroneously that history can have nothing to teach us on this subject.

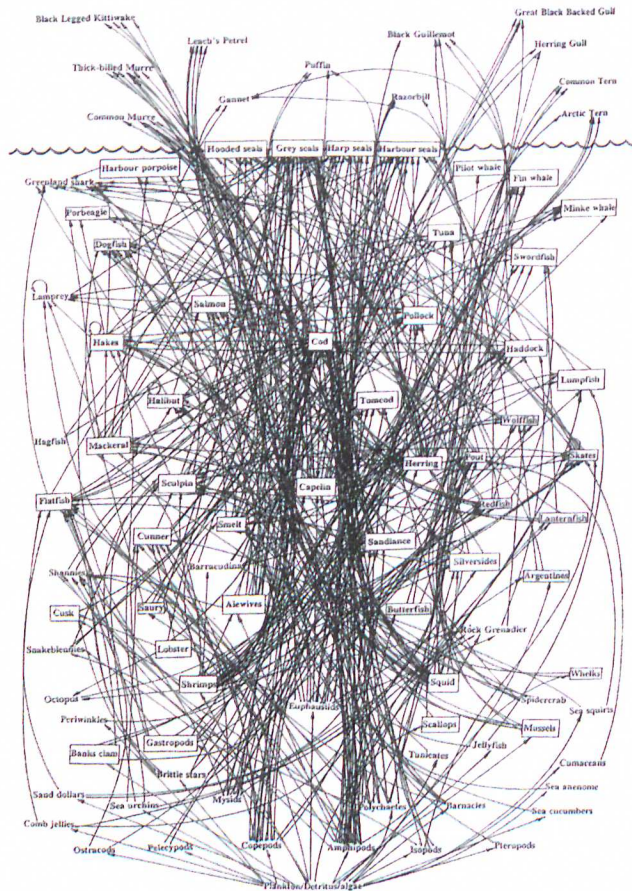
To be sure, there have never been such large networks as we see in the world today. Nor have the flows of information – or, for that matter, disease – ever been so rapid. But scale and speed are not everything. We shall never make sense of the vast, swift networks of our own time – in particular, we shall have no inkling whether the network age will be joyously emancipatory or hideously anarchic – if we do not study the smaller, slower networks of the past. For these, too, were ubiquitous. And sometimes they were very powerful indeed.

3

Networks, Networks Everywhere

The natural world is to a bewildering extent made up of ‘optimized, space-filling, branching networks’, in the words of the physicist Geoffrey West, from the human circulatory system to a colony of ants, all of which have evolved to distribute energy and materials between macroscopic reservoirs and microscopic sites over an astonishing twenty-seven orders of magnitude. The animal circulatory, respiratory, renal, neural systems are all natural networks. So are plant vascular systems and the microtubial and mitochondrial networks inside cells.¹ The brain of the nematode worm *Caenorhabditis elegans* is the only neural network to have been comprehensively mapped, but more complex brains will in due course be given the same treatment.² From worms’ brains to food chains (or ‘food webs’), modern biology finds networks at all levels of life on earth.³ The sequencing of the genome has revealed a ‘gene regulatory network’ in which ‘nodes are genes and links are chains of reactions’.⁴ The delta of a river is a network, too: your school atlas mapped those. Tumours form networks.

Some problems can only be resolved by network analysis. Scientists seeking to explain the massive algal blooms that afflicted the San Francisco Bay in 1999 had to map the network of marine life before they could identify the true cause. A similar mapping of neural networks was necessary to establish that the hippocampus is where the human memory resides.⁵ The speed with which an infectious disease spreads has as much to do with the network structure of the exposed population as with the virulence of the disease itself, as an epidemic amongst teenagers in Rockdale County, Georgia, made clear twenty years ago.⁶ The existence of a few highly connected hubs causes the



2. A partial food web for the 'Scotian Shelf' in the north-west Atlantic. Arrows go from the prey species to the predator species.

spread of the disease to increase exponentially after an initial phase of slow growth.⁷ Put differently, if the 'basic reproduction number' (how many other people are newly infected by a typical infected individual) is above one, then a disease becomes endemic; if it is below one, it tends to die out. But that basic reproduction number is determined as much by the structure of the network it infects as by the innate infectiousness of the disease.⁸ Network structures can also condition the speed and accuracy with which a disease is diagnosed.⁹

In prehistory, *Homo sapiens* evolved as a cooperative ape, with a

unique ability to network – to communicate and to act collectively – that set us apart from all other animals. In the words of the evolutionary biologist Joseph Henrich, we are not simply bigger-brained, less hairy chimpanzees; the secret of our success as a species ‘resides . . . in the *collective brains* of our communities’.¹⁰ Unlike chimpanzees, we learn socially, by teaching and sharing. According to the evolutionary anthropologist Robin Dunbar, our larger brain, with its more developed neocortex, evolved to enable us to function in relatively large social groups of around 150 (compared with around fifty for chimpanzees).¹¹ Indeed, our species should really be known as *Homo dictyous* (‘network man’) because – to quote the sociologists Nicholas Christakis and James Fowler – ‘our brains seem to have been built for social networks’.¹² The term coined by the ethnographer Edwin Hutchins was ‘distributed cognition’. Our early ancestors were ‘obligate collaborative foragers’, who became interdependent on each other for food, shelter and warmth.¹³ It is likely that the development of spoken language, as well as the associated advances in brain capacity and structure, was part of this same process, evolving out of ape-like practices such as grooming.¹⁴ The same could also be said of practices such as art, dance and ritual.¹⁵ In the words of the historians William H. McNeill and J. R. McNeill, the first ‘worldwide web’ in fact emerged around 12,000 years ago. Man, with his unrivalled neural network, was born *to network*.

Social networks, then, are the structures that human beings naturally form, beginning with knowledge itself and the various forms of representation we use to communicate it, as well of course as the family trees to which we all necessarily belong, even if only some of us possess detailed genealogical knowledge. Networks include the patterns of settlement, migration and miscegenation that have distributed our species across the world’s surface, as well as the myriad cults and crazes we periodically produce with minimal premeditation and leadership. As we shall see, social networks come in all shapes and sizes, from exclusive secret societies to open-source movements. Some have a spontaneous, self-organizing character; others are more systematic and structured. All that has happened – beginning with the invention of written language – is that new technologies have facilitated our innate, ancient urge to network.

Yet there is a puzzle. For most of recorded history, hierarchies dominated networks in their scope and scale. Men and women were mostly organized into hierarchical structures, with power concentrated at the very top in the hands of a chief, lord, king or emperor. By contrast, the average individual's network was stunted in its scale. The typical peasant – and that word roughly describes what most human beings were for most of recorded history – was stuck in a tiny cluster called a family, inside a slightly larger cluster called a village, with almost no links to the wider world. This was how most human beings lived as recently as a hundred years ago. Even today, the inhabitants of Indian villages are, at best, connected in a 'social quilt . . . a union of small cliques where each clique is just large enough to sustain cooperation by all of its members and where the cliques are laced together'.¹⁶ A key role in such isolated communities is played by the 'diffusion-central' individuals commonly known as gossips.¹⁷

So oppressive were traditional small-scale networks that some individuals preferred to retreat into complete isolation. Robert Burns's song 'Naeboddy' celebrates self-reliance as a kind of defiant disconnection:

I hae a wife o' my ain,
I'll partake wi' naeboddy;
I'll tak Cuckold frae nane,
I'll gie Cuckold to naeboddy.

I hae a penny to spend,
There, thanks to naeboddy;
I hae naething to lend,
I'll borrow frae naeboddy.

I am naeboddy's lord,
I'll be slave to naeboddy;
I hae a gude braid sword,
I'll take dunts frae naeboddy.

I'll be merry and free,
I'll be sad for naeboddy;
Naeboddy cares for me,
I care for naeboddy.


From the Lone Ranger to the High Plains Drifter, such insular individuals have been recurrent heroes of Western cinema. In the Coen brothers' film *Blood Simple* (1984), the narrator inhabits a world of unbridled, brutal individualism. 'Go ahead, complain,' he says, 'tell your problems to your neighbour, ask for help – and watch him fly. Now in Russia, they got it mapped out so that everyone pulls for everyone else – that's the theory, anyway. But what I know about is Texas. And down here . . . you're on your own.'¹⁸

Nevertheless, such rampant individualism is the exception, not the rule. As John Donne memorably put it in his 'Devotions upon Emergent Occasions':

No man is an island, entire of itself; every man is a piece of the continent, a part of the main. If a clod be washed away by the sea, Europe is the less, as well as if a promontory were, as well as if a manor of thy friend's or of thine own were. Any man's death diminishes me, because I am involved in mankind; and therefore *never send to know for whom the bell tolls; it tolls for thee.*

Man is indeed a social animal and the misanthrope is shunned as well as shunning. The puzzle is why and how we natural networkers have for so long been in thrall to vertically structured and rigidly institutionalized hierarchies.

The word *hierarchy* derives from ancient Greek – *ἱεραρχία* (*hierarchia*), literally the 'rule of a high priest' – and was first used to describe the heavenly orders of angels and, more generally, to characterize a



3. Google n-gram of the frequency of appearance of the words 'network' and 'hierarchy' in English-language publications between 1800 and 2000.

stratified order of spiritual or temporal governance. Up until the sixteenth century, by contrast, the word 'network' signified nothing more than a woven mesh made of interlaced thread. Occasionally, Shakespeare uses the words 'net' and 'web' metaphorically – Iago's plot against Othello is a 'net that shall enmesh them all' – but 'network' itself does not appear in any of his plays.¹⁹ Scientists in the seventeenth and eighteenth centuries discerned that there were networks in nature – from spiders' webs to the human circulatory system of veins and arteries – but it was not until the nineteenth century that the term began to be used more metaphorically, by geographers and engineers to describe waterways and railways, and by writers to characterize the relations between people. The poet Coleridge (1817) spoke of a 'net-work of property', the historian Freeman (1876) of a 'network of feudal tenures'.²⁰ Even so, until around 1880, books published in English were more likely to contain the word 'hierarchy' than the word 'network' (see figure 3). It is possible retrospectively to subject the political and social relationships depicted in Anthony Trollope's 1869 novel *Phineas Finn* to network analysis,²¹ but the word 'network' does not appear once in the text. Only in the later twentieth century did 'networks' begin to proliferate: first transport and electrical networks, then telephone and television networks, finally computer and online social networks. And not before 1980 was 'network' used as a verb to connote purposive, career-oriented socializing.

5

From Seven Bridges to Six Degrees


The formal study of networks dates back to the mid-eighteenth century, the heyday of the East Prussian city of Königsberg, home of the philosopher Immanuel Kant. Among the sights of Königsberg were the seven bridges across the Pregel River that connected the two riverbanks to the two islands in the middle of the river, as well as linking the islands (see figure 4). It was a familiar conundrum to natives of the city that it was impossible to take a walk that crossed all seven bridges just once, without re-crossing any of them.* The problem attracted the attention of the great Swiss-born mathematician Leonhard Euler, who in 1735 invented network theory to demonstrate formally why such a walk was impossible. In the simplified graph (see figure 5), there are four ‘nodes’, representing the two main banks of the river and the smaller and larger islands, and seven ‘edges’, representing the bridges that connected them. Formally, Euler demonstrated that the possibility of a path that traverses each edge only once must depend on the *degree* of the nodes (the number of edges touching each node). The graph must have either two nodes with an odd number of edges or none. As the graph of the Königsberg bridges has four such nodes (one with five edges, the others with three), it cannot have a Eulerian path. A walk that crossed each bridge just once would be possible only if one edge – the bridge connecting the two islands – were removed; then only two nodes would have an odd-numbered

* Disappointingly, Kant’s daily walk – so punctual that people were said to set their watches by it – did not include the seven bridges. According to the poet Heinrich Heine, he preferred to walk eight times up and down a tree-lined street, thereafter known as the Philosopher’s Way.

degree. Since Euler's time, the basic units of graph theory – which he originally called 'the geometry of position' – have been nodes (or vertices) and edges (or links).

Nineteenth-century scientists applied this framework to everything from cartography to electrical circuits to isomers of organic components.¹ That there might also be *social* networks certainly occurred to some of the great political thinkers of that age, notably John Stuart Mill, Auguste Comte and Alexis de Tocqueville – the last of whom grasped that the rich associational life of the early United States was crucial to the working of American democracy. However, none attempted to formalize this insight. The study of social networks may therefore be said to date from 1900, when the schoolteacher and amateur social scientist Johannes Delitsch published a matrix that mapped the friendships of the fifty-three boys he had taught in his 1880–81 class.² Delitsch identified a close relationship between the boys' social affinities and their academic ranking – which in those days was the basis of classroom seating. Somewhat similar work was done three decades later in New York, where the idiosyncratic Austrian-born but anti-Freudian psychiatrist Jacob Moreno used sociograms to study the relationships between the 'delinquent' girls in a reformatory school in Hudson, NY. His research – published in 1933 as *Who Shall Survive?* – showed that the surge in the number of runaway girls in 1932 was explicable in terms of the runaways' positions in the school's social network of 'attractions and repulsions', which were both racial and sexual (see plate 2). Here, Moreno proclaimed, were 'the social forces which dominate mankind'. The book was, he believed, 'a new bible, the bible for social conduct, for human societies'.³


Thirty years later, the linguist and bibliographer Eugene Garfield devised a similar graphical technique to visualize the history of scientific fields by creating a 'historiograph' of citations. Citation indices and 'impact factors' have since become standard measures of academic achievement in science. They are also a way of mapping the process of scientific innovation – revealing, for example, the 'invisible colleges' implied by networks of citation, which look very different from the actual colleges that employ most scientists.⁴ However, such metrics may simply show that scientists tend to cite the work of



4. Euler's figure 1 from his *Solutio problematis ad geometriam situs pertinentis* (1741). Those wishing to test Euler's theorem literally cannot do so, as two of the seven original bridges did not survive the bombing of the city in the Second World War and two others were demolished after the city became Soviet-controlled Kaliningrad.

like-minded scientists. As the old adage says, birds of a feather flock together. What is true of citations is true more generally. The chances are that, when two nodes are linked to a third node, they will also be linked to each other, because (in the words of the economist James E. Rauch) 'two people who know me are more likely to know each other than are two randomly selected people'.⁵ A triad, all three members of which are connected by positive sentiments, is said to be 'balanced' and exemplifies the idea that 'the friend of my friend is my friend'. Another triad, two members of which do not know each other, despite both knowing the third member, is sometimes called a 'forbidden triad'. (A variant, with two members that are amicable but one that is hostile, represents the uncomfortable case when 'the enemy of my friend is also my friend'.)⁶

'Homophily' – our tendency to gravitate towards people similar to us (also known as assortativity) – might therefore be regarded as the first law of social networks. Everett Rogers and Dilip Bhowmik were the first sociologists to suggest that homophily might be disadvantageous, in narrowing the range of an individual's milieu; there was, they suggested, an 'optimal heterophily'. Was homophily a kind of



5. Simplified graph of Euler's Königsberg bridge problem. Only by removing the edge in the middle (the bridge linking the two islands in figure 4) can one solve the problem.

self-segregation? In the 1970s, Wayne Zachary plotted the friendship network between members of a university karate club. It clearly revealed two distinct clusters within the club. Homophily can be based on shared status (ascribed characteristics such as race, ethnicity, sex or age, and acquired characteristics such as religion, education, occupation or behaviour patterns) or shared values, in so far as those can be distinguished from acquired traits.⁷ A familiar illustration is the tendency for American schoolchildren to self-segregate by race and ethnicity (see plate 3), though recent research suggests that this tendency varies significantly between racial groups.⁸

Can such graphs show us which individuals are important? It was not until the twentieth century that scholars and mathematicians formally defined importance as 'centrality'. The three most important measures of importance in formal network analysis are degree centrality, betweenness centrality and closeness centrality. Degree centrality – the number of edges radiating out from a specific node – captures what might be called sociability: the simple number of relationships an individual has to others. Formalized by the sociologist Linton Freeman in the late seventies, betweenness centrality measures the extent to which information passes through a particular node. Just as commuters, by individually seeking the shortest route to their destination, concentrate traffic in a few congested

intersections, so too people in a network often rely on key individuals to connect them to otherwise distant individuals or groups. The individuals with high betweenness centrality are not necessarily the people with the most connections, but the ones with the important connections. (In other words, it's not how many people you know that matters; it's who you know.) Finally, closeness centrality measures the average number of 'steps' it takes for a node to reach all other nodes and is often used to discover who has the best access to information, assuming that this is widely distributed.⁹ In their different ways, individuals in social networks with high degree centrality, betweenness centrality or closeness centrality act as 'hubs'.

The mid-twentieth century also saw important advances in how we understand a network's aggregate properties, which are often invisible from the perspective of any individual node. At the Massachusetts Institute of Technology, R. Duncan Luce and Albert Perry proposed the use of 'clustering' coefficients to measure how much a group of nodes are connected, with a clique being the extreme case in which each node is connected to all the others in the network. (Technically, the clustering coefficient is the proportion of social triads which are fully connected, meaning that each member of any trio is connected to the other two.) The 'density' of a network is a similar measure of interconnectedness.

The importance of such measures became apparent in 1967, when the social psychologist Stanley Milgram conducted a famous experiment. He sent out letters to randomly chosen residents of Wichita, Kansas, and Omaha, Nebraska. The recipients were asked to forward the letter directly to the intended final recipient – respectively, the wife of a Harvard divinity student and a stockbroker in Boston – if that person was known personally to them, or to forward it to someone they believed might know the final recipient, provided they knew that intermediary on a first-name basis; and also to send Milgram a 'tracer' postcard saying what they had done. In all, according to Milgram, forty-four of the 160 letters from Nebraska ultimately got through.¹⁰ (A more recent study suggests it was just twenty-one.)¹¹ The completed chains allowed Milgram to calculate the number of intermediaries required to get the letter to its target: on average five.¹² This finding had been anticipated by the Hungarian author Frigyes

Karinthy, in whose story 'Láncszemek' ('Chains', published in 1929) a character bets his companions that he can link himself to any individual on earth they choose to name through no more than five acquaintances, only one of which he has to know personally. It was also borne out by separate experiments by other researchers, notably the political scientist Ithiel de Sola Pool and the mathematician Manfred Kochen.

A network that connects two nodes via five intermediaries has six edges. The phrase 'six degrees of separation' was not coined until John Guare's 1990 play of that title, but it therefore had a long pre-history. Like the concept of a 'small world' (made famous by the Disneyland ride devised in 1964), or the more technical concept of closeness, it neatly summed up the mid-twentieth century's growing sense of interconnectedness. There have since been numerous variations on the theme: six degrees of Marlon Brando, six degrees of Monica Lewinsky, six degrees of Kevin Bacon (which became a board-game), six degrees of Lois Weisberg (the mother of one of Malcolm Gladwell's friends) and the academic equivalent, six degrees of the mathematician Paul Erdős, himself a pioneer of network theory, as we have seen.¹³ Recent research suggests the number is now closer to five than six, which suggests that technological change since the 1970s has perhaps been less transformative than is commonly supposed.¹⁴ For the directors of Fortune 1000 companies, however, it is 4.6.¹⁵ For Facebook users it was 3.74 in 2012,¹⁶ and just 3.57 in 2016.¹⁷

6

Weak Ties and Viral Ideas

What makes this kind of finding so intriguing is that we tend to think of our networks of friends as relatively small clusters or cliques of similar, like-minded people, isolated from other groups whose members have different affinities with one another. The fact that we are all nevertheless just six degrees away from Monica Lewinsky is explained by what the Stanford sociologist Mark Granovetter called, paradoxically, ‘the strength of weak ties’.¹ If all ties were like the strong, homophilic ones between us and our close friends, the world would necessarily be fragmented. But weaker ties – to the ‘acquaintances’ we do not so closely resemble – are the key to the ‘small world’ phenomenon. Granovetter’s initial focus was on the way people looking for jobs were helped more by acquaintances than by their close friends, but a later insight was that, in a society with relatively few weak ties, ‘new ideas will spread slowly, scientific endeavours will be handicapped, and subgroups separated by race, ethnicity, geography, or other characteristics will have difficulty reaching a *modus vivendi*’.² Weak ties, in other words, are the vital bridges between disparate clusters that would otherwise not be connected at all.³

Granovetter’s was a sociological observation, based on interviews and similar data, and was subject to refinement on the basis of field studies. These revealed, for example, that strong ties matter more to the poor than weak ties, suggesting that the tightly knit networks of the proletarian world might tend to perpetuate poverty.⁴ It was not until 1998 that the mathematicians Duncan Watts and Steven Strogatz demonstrated *why* a world characterized by homophilic clusters could simultaneously be a small world. Watts and Strogatz classified networks in terms of two relatively independent properties: the

average closeness centrality of each node and the network's general clustering coefficient. Beginning with a circular lattice in which each node was connected only to its first- and second-nearest neighbours, they showed that the random addition of just a few extra edges sufficed drastically to increase the closeness of all nodes, without significantly increasing the overall clustering coefficient.⁵ Watts had begun his work by studying the synchronized chirping of crickets, but the implications of his and Strogatz's findings for human populations were obvious. In Watts's words, 'the difference between a big- and a small-world graph can be a matter of only a few randomly required edges – a change that is effectively undetectable at the level of individual vertices . . . [T]he highly clustered nature of small-world graphs can lead to the intuition that a given disease is "far away" when, on the contrary, it is effectively very close.'⁶

For economists, too, advances in network science had important implications. Standard economics had imagined more or less undifferentiated markets populated by individual utility-maximizing agents with perfect information. The problem – resolved by the English economist Ronald Coase, who explained the importance of transaction costs* – was to explain why firms existed at all. (We are not all longshoremen, hired and paid by the day like Marlon Brando in *On the Waterfront*, because employing us regularly within firms can reduce the costs that arise when workers are hired on a daily basis.) But if markets were networks, with most people inhabiting more or less interconnected clusters, the economic world looked very different, not least because information flows were determined by the networks' structures.⁷ Many exchanges are not just one-off transactions in which price is a matter of supply and demand. Credit is a function of trust, which in turn is higher within a cluster of similar

* Coase argued in 'The Problem of Social Cost' (1960, 15) that 'in order to carry out a market transaction it is necessary to discover who it is that one wishes to deal with, to inform people that one wishes to deal and on what terms, to conduct negotiations leading up to a bargain, to draw up the contract, to undertake the inspection needed to make sure that the terms of the contract are being observed, and so on'. Organizations like firms and indeed states exist to lower or eliminate such transaction costs with, for example, standardized long-term employment contracts. Larger units can do this more efficiently, hence 'economies of scale'.


people (e.g. an immigrant community). This has implications not only for employment markets, the case studied by Granovetter.⁸ Closed networks of sellers can collude against the public and deter innovation. More open networks can promote innovation as new ideas reach the cluster thanks to the strength of weak ties.⁹ Such observations prompted the question of how exactly networks are formed in the first place.¹⁰

In practice, it seems clear how networks form. From Avner Greif's eleventh-century Maghribi traders in the Mediterranean¹¹ to the modern entrepreneurs and managers studied by Ronald Burt, scholars have produced a rich literature on the role of business networks in generating social capital¹² and promoting – or discouraging – innovation. In Burt's terminology, competition between individuals and firms is structured by networks, with 'structural holes' – the gaps between clusters, where weak ties are lacking – as 'entrepreneurial opportunities for information access, timing, referrals, and control'.¹³ Brokers – people who are able to 'bridge the holes' – are (or should be) 'rewarded for their integrative work' because their position makes them more likely to have creative ideas (or less likely to suffer from group-think). In innovative institutions, such brokers are always appreciated. However, in most contests between an innovator-broker and a network inclined towards 'closure' (i.e. insularity and homogeneity), the latter often prevails.¹⁴ This insight applies as much to academic philosophers as to the employees of an American electronics company.¹⁵

An entire subfield of 'organizational behaviour' now occupies a foundational place in most master's programmes in business administration. Among its recent findings are that managers are more likely to be networkers than non-managers;¹⁶ that a 'less hierarchical network may be better for producing solidarity and homogeneity in an organizational culture';¹⁷ and that brokers are more likely to succeed in spanning structural holes if they 'fit culturally into their organizational group', while those who are 'structurally embedded' fare better when they are 'culturally distinct'. In sum, 'assimilated brokers' and 'integrated nonconformists' tend to do better than their peers.¹⁸ Here, too, network theory offers insights that have utility beyond the typical corporate workplace satirized in Ricky Gervais's *The Office*.

After all, office networks are seldom very large. Yet network size matters because of Metcalfe's law – named after the Ethernet inventor, Robert Metcalfe – which (in its original form) stated that the value of a telecommunications network was proportional to the square of the number of connected compatible communicating devices. This is in fact true of networks generally: put simply, the greater the number of nodes in a network, the more valuable the network to the nodes collectively. As we shall see, this implies spectacular returns to very large, open-access networks and, conversely, limited returns to secret and/or exclusive networks. Yet even in the largest networks there are nodes that act as brokers or hubs.

The phrase 'to go viral' has become a tiresome cliché, the holy grail of advertisers and marketers.¹⁹ Nevertheless, network science offers the best way of understanding why some ideas can spread very rapidly. Ideas – and indeed emotional states and conditions such as obesity – can be transmitted through a social network, not unlike a contagious virus. However, ideas (or 'memes', to use the evolutionist's neologism) are generally less contagious than viruses. Biological and computer viruses typically carry out a 'broadcast search' across a network, as their goal is to spread themselves as far as possible, targeting every neighbour of every node they infect. We, by contrast, instinctively select the members of our network to whom we want to communicate an idea or from whom we are likely to accept one as credible.²⁰ An early contribution was the so-called 'two-step flow of communication model', associated with the sociologists Paul Lazarsfeld and Elihu Katz, who argued in the 1950s that ideas flowed from the media to the wider population via opinion 'leaders'.²¹ Other late-twentieth-century researchers sought to measure the speed at which news, rumours or innovations moved. More recent research has shown that even emotional states can be transmitted through a network.²² Though distinguishing between endogenous and exogenous network effects is far from easy,²³ the evidence of this kind of contagion is clear: 'Students with studious roommates become more studious. Diners sitting next to heavy eaters eat more food.'²⁴ However, according to Christakis and Fowler, we cannot transmit ideas and behaviours much beyond our friends' friends' friends (in other words, across just three degrees of separation). This is because the



6. The foundational concepts of network theory. Each dot in the graph is a node, each line an edge. The dot labelled 'Hub' has the highest degree centrality and betweenness centrality. The nodes labelled 'Cluster' have a higher density or local clustering coefficient than other parts of the graph.

transmission and reception of an idea or behaviour requires a stronger connection than the relaying of a letter (in the case of Milgram's experiment) or the communication that a certain employment opportunity exists. Merely knowing people is not the same as being able to influence them to study more or over-eat. Imitation is indeed the sincerest form of flattery, even when it is unconscious.

The key point, as with disease epidemics, is that network structure can be as important as the idea itself in determining the speed and extent of diffusion.²⁵ In the process of going viral, a key role is played by nodes that are not merely hubs or brokers but 'gate-keepers' – people who decide whether or not to pass information to their part of the network.²⁶ Their decision will be based partly on how they think that information will reflect back on them. Acceptance of an idea, in turn, can require it to be transmitted by more than one or two

sources. A complex cultural contagion, unlike a simple disease epidemic, first needs to attain a critical mass of early adopters with high degree centrality (relatively large numbers of influential friends).²⁷ In the words of Duncan Watts, the key to assessing the likelihood of a contagion-like cascade is 'to focus *not* on the stimulus itself but on the structure of the network the stimulus hits'.²⁸ This helps explain why, for every idea that goes viral, there are countless others that fizzle out in obscurity because they began with the wrong node, cluster or network.

7

Varieties of Network

If all social network structures were the same, we would inhabit a very different world. For example, a world in which nodes were randomly connected to one another – so that the numbers of edges per node were normally distributed along a bell curve – would have some ‘small world’ properties, but it would not be like our world.* That is because so many real-world networks follow Pareto-like distributions: that is, they have more nodes with a very large number of edges and more nodes with very few than would be the case in a random network. This is a version of what the sociologist Robert K. Merton called ‘the Matthew effect’, after the Gospel of St Matthew: ‘For unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken away even that which he hath.’† In science, success breeds success: to him who already has prizes, more prizes shall be given. Something similar can be seen in ‘the economics of superstars’.¹ In the same way, as many large networks expand, nodes gain new edges in proportion to the number that they already have (their degree or ‘fitness’). There is, in short, ‘preferential attachment’. We owe this insight to the physicists Albert-László Barabási and Réka Albert, who were the first to suggest that most real-world networks might follow a power law distribution or

* Random networks were first studied by the famously prolific and much cited mathematician Paul Erdős and Alfréd Rényi, one of his many co-authors. A random graph is constructed by placing n nodes on a plane, then joining pairs of them together at random until m edges are used. Nodes may be chosen more than once, or not at all.

† Matthew 25:28.

be ‘scale-free’.* As such networks evolve, a few nodes will become hubs with many more edges than other nodes.² Examples of such networks abound, ranging from the directorships of Fortune 1000 companies to citations in physics journals and links to and from web-pages.³ In Barabási’s words:

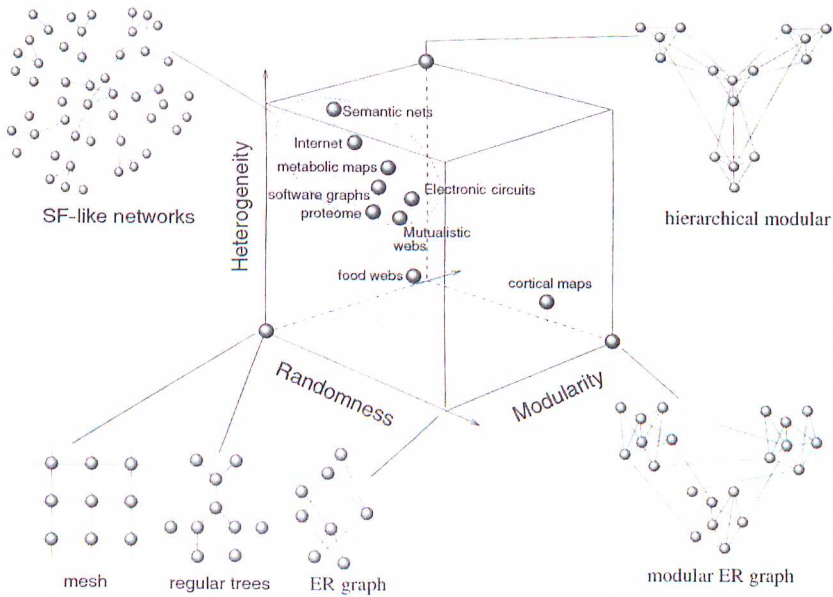
[T]here is a hierarchy of hubs that keep those networks together, a heavily connected node closely followed by several less connected ones, trailed by dozens of even smaller nodes. No central node sits in the middle of the spider web, controlling and monitoring every link and node. There is no single node whose removal could break the web. A scale-free network is a web without a spider.⁴

In the extreme case (the winner-takes-all model), the fittest node gets all or nearly all the links. More often, there is a ‘fit get rich’ pattern whereby ‘a heavily connected node [is] closely followed by several less connected ones, trailed by dozens of even smaller nodes’.⁵ Other intermediate network structures can also be found: for example, the friendship networks of American adolescents are neither random nor scale-free.⁶

In a random network, as Erdős and Rényi showed long ago, each node within the network has approximately the same number of links to other nodes. The best real-world example is the US national highway network, in which each major city has roughly the same number of highways connecting it to others. An example of a scale-free network is the US air traffic network, in which a large number of small airports are connected to medium-sized airports, which in turn connect to a few huge and busy hubs. Other networks are more highly centralized without necessarily being scale-free. One way of

* Distributions that follow a power law are said to have ‘fat tails’, as the relative likelihood of very high degree and very low degree are higher than if links were formed at random. Technically, the term ‘scale-free’ refers to the fact ‘that the relative frequency of nodes of degree d , compared to nodes of degree d' , is the same as the relative frequency of nodes with degree kd compared to nodes of degree kd' , when rescaling by an arbitrary factor $k > 0$.’ In a scale-free network there is no typical node, and yet the ‘scale’ of difference between nodes appears the same everywhere. Put differently, the scale-free world is characterized by fractal geometry: the town is a large family, the city a large town, and the kingdom a large city.

VARIETIES OF NETWORK




8. Varieties of network (SF: scale-free; ER: Erdős-Rényi, i.e. random).

We can now see clearly that, far from being the opposite of a network, a hierarchy is just a special kind of network. As figure 9 shows, edges in an idealized hierarchical network follow a regular pattern, like that of an upside-down tree (or the roots of a tree). To construct a hierarchical network, start with the top node, and add a certain number of subordinate nodes. To each subordinate node add the same number of subordinates again, and so on. The key is to always add nodes downwards, but never connect nodes laterally. Networks built in this way have special properties. For one, there are no cycles, that is, no path that leads from a node back to itself. There is only one path connecting any two nodes, which clarifies chains of command and communication. More importantly, the top node has the highest betweenness and closeness centrality – that is, the system is designed to maximize that node's ability both to access and to control information. As we shall see, few hierarchies achieve such a total control over information flows, though Stalin's Soviet Union came close. Most organizations are in practice only partially hierarchical, not

unlike the 'cooperative hierarchies' of the natural world.⁸ It may be helpful, nevertheless, to think of a pure hierarchy as in some sense 'anti-random', in that the promiscuous connectivity associated with networks – above all, clustering – is prohibited.

These varieties of network should not be regarded as static categories. Networks are rarely frozen in time. Large networks are complex systems which have 'emergent properties' – the tendency of novel structures, patterns and properties to manifest themselves in 'phase transitions' that are far from predictable. As we shall see, a seemingly random network can evolve with astounding speed into a hierarchy. The number of steps between the revolutionary crowd and the totalitarian state has more than once proved to be surprisingly few. By the same token, the seemingly rigid structures of a hierarchical order can



9. Hierarchy: a special kind of network. In the example shown here, the node at the top has the highest betweenness and closeness centrality. Other nodes can communicate with the majority of other nodes only through that one ruling hub.

disintegrate with astounding rapidity.⁹ This comes as no shock to the student of networks. We know now that the random addition of a very small number of new edges can radically reduce the average separation between the nodes. It would not require too many additional edges in figure 9 to destroy the ruling node's near monopoly on communication. This helps explain why emperors and kings throughout history have fretted about conspiracies. Cabals, camarillas, cells, cliques, coteries: all such terms have sinister connotations in the context of a monarchical court. Hierarchs have long been uneasily aware that fraternizing amongst subordinates can be the prelude to a palace coup.

8

When Networks Meet

The final conceptual challenge – and the most important one for the historian – is to consider how different networks interact with one another. The political scientist John Padgett and his co-authors have proposed a biochemical analogy, arguing that organizational innovation and invention are both the results of interaction between networks, which takes three basic forms: ‘transposition’, ‘refunctionality’ and ‘catalysis’.¹ In itself, a resilient social network will tend to resist change to its production rules and communication protocols. It is when a social network and its patterns are transposed from one context and refunctioned in another that innovation and even invention can occur.²

As we shall see, Padgett has used this insight to explain changes in the economic and social structure of Florence in the time of the Medici, when banking partnerships were incorporated into city politics. It clearly has a more general applicability, however. Networks are important not just as transmission mechanisms for new ideas, but as the sources of the new ideas themselves. Not all networks are likely to foster change; on the contrary, some dense and clustered networks have the tendency to resist it. But the point of contact between diverse networks may be the place to look for novelty.³ The question is what the nature of that point of contact is. Networks can meet and fuse amicably, but they may also attack one another, as happened (in an example to be discussed below) when Soviet intelligence successfully penetrated the elite networks of Cambridge students in the 1930s. In such contests, the outcome will be determined by the relative strengths and weaknesses of the rival networks. How adaptable and resilient are they? How vulnerable to a disruptive contagion? How

reliant on one or more ‘superhubs’, the destruction or capture of which would significantly reduce the stability of the whole network? Barabási and his colleagues simulated attacks on scale-free networks and found that they could withstand the loss of a significant fraction of nodes, and even of a single hub. But a targeted attack on multiple hubs could break the network up altogether.⁴ Even more dramatically, a scale-free network could quite easily fall victim to a contagious node-killing virus.⁵

But why would one network attack another, as opposed to peaceably linking to it? The answer is that most attacks on social networks are not initiated by other networks but are ordered or at least encouraged by hierarchical entities. The Russian meddling in the 2016 US election is a case in point: according to US intelligence, as noted above, it was authorized by President Putin, one of the most unabashed autocrats in the world, but it was directed not just at the Democratic National Committee but at the whole complex of US media networks. This illustrates the fundamental difference between networks and hierarchies. Because of their relatively decentralized structure, because of the way they combine clusters and weak links, and because they can adapt and evolve, networks tend to be more creative than hierarchies. Historically, as we shall see, innovations have tended to come from networks more than from hierarchies. The problem is that networks are not easily directed ‘towards a common objective . . . that requires concentration of resources in space and time within large organizations, like armies, bureaucracies, large factories, vertically organized corporations’.⁶ Networks may be spontaneously creative but they are not strategic. The Second World War could not have been won by a network, even if superior networks (of atomic scientists or cryptographers) played an important part in the Allied victory. Not only that, but networks are as capable of creating and spreading bad ideas as good ideas. In cases of social contagion or ‘cascades’ of ideas, networks can spread panics as readily as they can communicate the wisdom of crowds – crazes for burning witches as easily as harmless manias for photographs of cats.

True, today’s networks are better designed than the US electricity grid in the 1990s, which was so fragile that the failure of a single power line in western Oregon caused the tripping of hundreds more

lines and generators. Yet we know that even a robust network can tip over into dysfunction as it grows and evolves: the normal congestion and delays at US airports are a case in point, as airlines compete to service hubs, but end up clogging them.⁷ Quite aside from the Internet, there is little doubt that a targeted attack on US electrical and transport infrastructure would have devastatingly disruptive consequences. As Amy Zegart has said, the United States is simultaneously the most powerful and the most vulnerable actor in the cyberwar theatre. 'The cyber threats of tomorrow,' she has warned, 'could disable the cars that we drive, the airplanes that we fly, they could turn off power or water to cities across the country for days or weeks or longer, they could incapacitate our military or even turn our own weapons against us.'⁸ And yet the United States 'seems unwilling to acknowledge the basic facts about new cyber-technologies or our cyber-vulnerabilities, let alone take the necessary measures to attribute, deter and defend against future attacks'.⁹ The May 2017 epidemic, when the WannaCry 'ransomware' infected hundreds of thousands of computers in 150 countries, encrypting their hard drives and demanding payment in Bitcoin, exposed the vulnerability not only of European countries but also, ironically, of Russia, to criminal attacks.

The reality is that we find it very hard indeed to fathom the implications of the growth of networks in our own time. For every article extolling their positive effects in empowering the young and enlivening democracy – for example, in the Arab revolutions of 2010–12 – there is another warning of their negative effects in empowering dangerous forces – for example, political Islam. For every book prophesying a 'singularity' in which a 'global brain' or 'planetary superorganism' arises from the Internet,¹⁰ there is another foreseeing collapse and extinction.¹¹ Anne-Marie Slaughter expects that 'the United States and other powers will gradually find the golden mean of network power: not too concentrated and not too distributed', and looks forward to the emergence of 'a flatter, faster, more flexible system, one that operates at the level of citizens as well as states'.¹² Writing before 9/11, Graham Allison was relatively confident that the United States would have an inbuilt advantage in a world of global networks.¹³ Yet Joshua Ramo is far less optimistic. 'The simple, once-appealing idea that connection is liberation is wrong,' he writes. 'To connect now is

to be encased in a powerful and dynamic tension.’ The inability of old leaders to make sense of the Network Age is ‘the reason [their] legitimacy . . . is failing, the reason our grand strategy is incoherent, the reason our age really is revolutionary’. In his eyes, ‘the fundamental threat to American interests isn’t China or Al Qaeda or Iran. It is the evolution of the network itself.’¹⁴

In one respect only does there seem to be consensus: few futurologists expect established hierarchies – in particular, traditional political elites, but also long-established corporations – to fare very well in the future.¹⁵ Francis Fukuyama is unusual in arguing that hierarchy must ultimately prevail, in the sense that networks alone cannot provide a stable institutional framework for economic development or political order. Indeed, he argues, ‘hierarchical organization . . . may be the *only* way in which a low-trust society can be organized’.¹⁶ By contrast, the iconoclastic British political operative Dominic Cummings hypothesizes that the state of the future will need to function more like the human immune system or an ant colony than a traditional state – in other words, more like a network, with emergent properties and the capacity for self-organization, without plans or central coordination, relying instead on probabilistic experimentation, reinforcing success and discarding failure, achieving resilience partly through redundancy.¹⁷ This may be to underestimate both the resilience of the old hierarchies and the vulnerabilities of the new networks – not to mention their capacity to fuse to form even newer power structures, with capabilities potentially greater even than those of the last century’s totalitarian states.