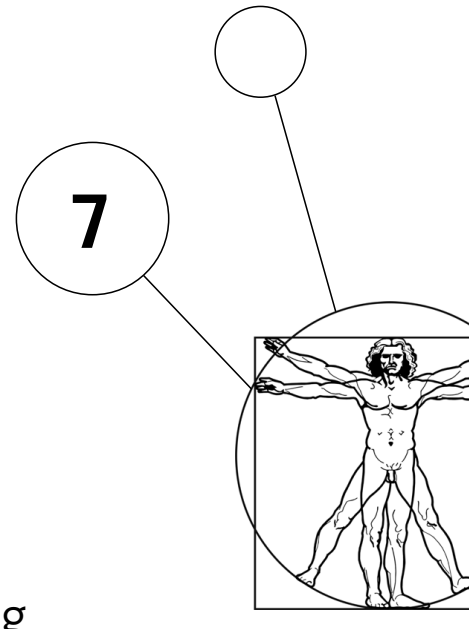

The Pox Is Spreading, Venice Is Sinking

How the age we're in magnifies systemic dangers and makes it harder to see them coming

The downside of linking up

In 1494, Ludovico Sforza, who for thirteen years had controlled the Duchy of Milan from behind the throne as regent for his too young nephew, stepped forward and seized the mantle of duke for himself. Alfonso II, the king of Naples (who himself had a claim to Milan), challenged the move and threatened to unseat this usurper. But Sforza was not without allies. Charles VIII, the powerful king of France, had a claim to the throne of Naples, and Sforza convinced Charles that now was the moment to press it.

Charles raised an army of 25,000–30,000 men – including 8,000 mercenaries drawn from all over Europe – to wage his Italian war.[1] Passing unimpeded through Milan, his men slashed and burned their way down the rest of the Italian peninsula and captured Naples in February 1495.

But Charles had miscalculated. His show of brutality and overwhelming force, which brought swift victory, also accomplished what no lesser threat could: it united the states of Italy and their allies in fear against a foreign foe. The Pope convened a Holy League to drive the French from Italy, and Venice, Spain, England, the Holy Roman Empire (very roughly, present-day Germany) and even Milan signed up. (Sforza had begun to fear that Charles would betray him and swallow his own duchy, too.)

On 6 July 1495, the massed armies of the Holy League met Charles VIII on a rain-soaked field at the Battle of Fornovo to decide Italy's fate – for the moment. The League had numbers on their side; Charles had the terrain. In less than two hours of fighting, the French lost 1,000 men; the League

lost twice that many. Both sides claimed victory; neither pressed the other to continue. Charles retreated back to France and his army scattered.[2]

But he left something behind.

Italian battlefield doctors noticed the disease first. It was unlike any affliction they had ever seen or that had ever been recorded, going back to the time of the Roman emperor Marcus Aurelius and the medical treatises of his physician, Galen.

Plague was well known – its symptoms terrible, its consequences mercifully swift. The victims spat blood, and in three days they were dead. But this was something new, something crueller. It debilitated its victims in horrid ways and left them to linger through months, even years, of flesh-eating filth and disgust. A medical chronicler of the time described sufferers' bodies 'covered with acorn-sized boils that emitted a foul, dark green pus'.[3] These boils often appeared first on the sex organs. Those who survived the disease's first stages developed 'tumors the size of bread rolls and ulcers that progressively, but simply dissolved skin'.[4]

It also spread quickly, thanks to the 'Noah's Ark nature' of both sides' armies.[5] Troops from both sides of Fornovo returned home, and mercenaries dispersed. By summer's end, the mysterious disease had terrorized towns throughout Italy, France, Germany and Switzerland. By the following year, it had hit Holland and Greece; the year after that, England and Scotland. Within four years of its first appearance, it had already touched the whole of Europe. Within another five years, except for some still-isolated exceptions, it was global.[6]

From its beginning, it was less fatal to Europe than the Black Death. To the rest of the world, it was certainly less devastating than some of the other diseases spread by European sailors, like the smallpox that decimated the Americas. As populations adapted to the new disease in their midst, its most repulsive symptoms subsided, and it eased into the lingering, chronic venereal disease we know today as syphilis.

Too complex to unravel, too concentrated to keep safe

Part II showed how human entanglement and development help genius to flourish, in both its individual and collective forms. Now we turn to the second consequence of these forces: flourishing *risk*.

We all need to develop a heightened concern for a particular kind of risk. It is not the direct danger that confronts us daily – things like getting hit by a car or being robbed. We're all well aware of such specific hazards already. Rather, it is the risks we don't see – the kind that creep slowly below our threshold of observation, and then shock all of us together. These 'butterfly defects' are widely felt but hard for anyone to see coming, because their causes are far removed from our day-to-day experiences and concerns. Such risks are not specific; they're *systemic*.†

Systemic risks flourish in the present age because the same connective and developmental trends that awaken genius also exacerbate the two conditions under which such risks breed: complexity and concentration.

Complexity demons

Most of us are well aware of the rising complexity of the present age: we see the evidence in our own lives. Flip back to the figures and graphs throughout Part I, and this heightened complexity is clear to see: in shifting patterns of global air travel; in the rising variety and number of cross-border financial investments; or in the growth of internet infrastructure. Meanwhile, developmental forces amplify this complexity by raising the volume of traffic that flows across these many, diverse connections, and by adding new nodes – be they new cities, universities, industrial zones, ports, power stations, labs, conferences or journals.

We've seen some of the benefits that complexity can bring. It increases the number and variety of good things that can touch us and that we can reach out to, and it's a major catalyst of creativity and idea generation.

From a risk perspective, too, complexity can be a good thing. The greater variety and volume of connections and flows create redundancy, of which the internet is contemporary life's best example. When one link goes down, its traffic reroutes almost instantaneously to alternatives, so that our end-user experience is not often interrupted. Complexity breeds benefits.

But it also presents a problem. The more complex our interactions become, the harder it is for us to see relationships of cause and effect. We develop cognitive 'blind spots' in our vision of the events around us. See Figure 7.1. How can we make good decisions, when we can't see their consequences?

† Ian Goldin and Mike Mariathasan explore the phenomenon and its consequences in their 2014 book, *The Butterfly Defect* from Princeton University Press.

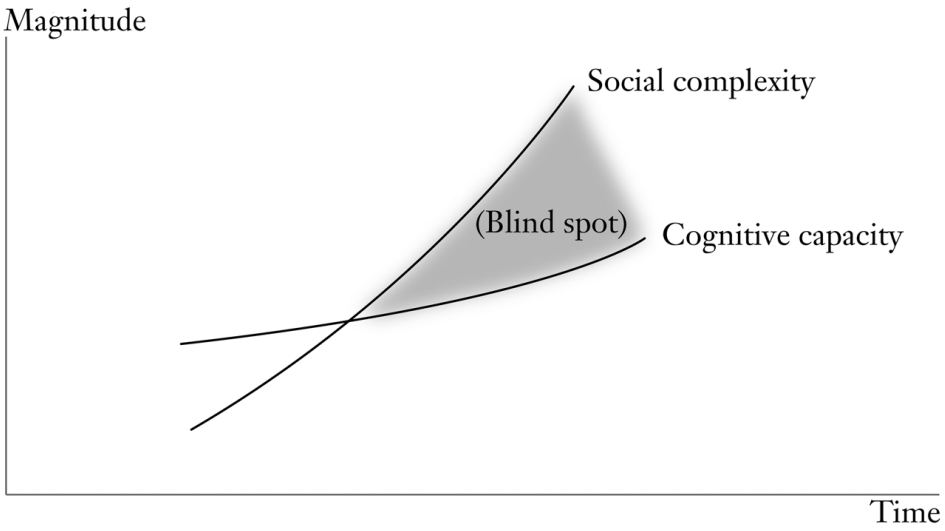

Figure 7.1 We develop blind spots when complexity rises faster than our understanding.

Complexity was a big part of how syphilis struck the last Renaissance, and why it struck as hard as it did.

The spread of disease has always been among the first unintended consequences of population exchange; syphilis was just a terrifying case in point. According to the most widely accepted theory today on the origin of syphilis, its emergence in Europe just three years after Columbus discovered the New World was no coincidence.† Likely, his sailors carried it back to Europe with them.[7] Just as the new sea link between Europe and the Americas injected the latter population with a lethal cocktail of European bugs – smallpox, typhus, measles, influenza, bubonic plague, cholera, malaria, tuberculosis, mumps, yellow fever and others – so, too, it exposed Eurasia to a new plague against which local populations had no immunity. Likewise, the disease's swift propagation throughout Europe and Asia is best explained by the variety and volumes of new economic linkages in Europe and across the sea, over which goods, people and livestock flowed in ever-growing numbers.

The complexity of the disease's origin and spread also left people in the Renaissance at a loss as to how to respond. The state of medical knowledge

† The theory of the American origin of syphilis is the most widely accepted, but has not been conclusively proven.

at the time meant that the causes of most diseases were beyond their understanding, even though disease was a familiar danger and society had worked out some effective responses. During the previous century, the Black Death and its many aftershocks had taught them to recognize plague, quarantine its victims and avoid the settlements it struck until the disease disappeared. Later bouts of plague scarred Europe far less, because such protocols were in place. But syphilis, so suddenly, was in too many places to be avoided and at every level of society, from the peasantry to the papacy. And it didn't disappear. It lingered like its victims, a chronic blight that couldn't be avoided, and to which society had to, somehow, adapt.

The absence of any clear understanding of cause and effect left a cognitive gap that society stuffed full of stereotype, superstition and ideological agendas. The Italians called it the French disease, since obviously Charles VIII had brought it with his army. The French called it the Neapolitan (Naples) disease, since it had been unknown to France until her soldiers returned home from that Italian kingdom. The Holy Roman Emperor Maximilian I saw the disease as God's divine retribution against man for his sins. What else could explain the sudden universality of a completely novel affliction, one that let its victims linger in agony and shame to contemplate their misdeeds? The most popular explanation was that God was punishing people for sins of the flesh. 'God has raised up new diseases against debauchery', in the words of John Calvin (1509–1564).[8] People attached great significance to the fact that the disease often appeared first on the sex organs, and that soldiers and prostitutes (the two professions most associated with sexual licence and moral disorder) ranked among the first and most frequent victims. 'A night with Venus, a lifetime with Mercury' went the saying, referencing the mercury salves that became the common treatment.

Concentration dilemmas

Concentration is a less obvious consequence of human entanglement and development, but just as destabilizing. How developmental forces produce concentrations is straightforward. Worldwide gains in health, wealth and education together yield a much bigger population that places far larger demands on existing social infrastructure, services, natural resources and the environment. Consider: humanity has been burning fossil fuels for millennia with seemingly little consequence. But now that a two-billion-member-strong global middle class has emerged, all wanting to drive cars, take flights, and run heaters and air conditioners, suddenly the collective weight of our energy consumption has to be reckoned with.

Concentrations also increase as the number and variety of our connections go up. At first glance, this seems counter-intuitive: if the variety of links across which goods, services, people and ideas can flow suddenly increases, shouldn't that diffuse human activity?

Yes and no. While it's true that more connections open more choices for where goods, services, people and ideas might flow, these things do not flow randomly. Instead, they flow to wherever we deem most desirable. Governments concentrate public infrastructure, and businesses concentrate their operations, wherever is deemed to be most efficient. Immigrants and job seekers concentrate wherever work is perceived to be plentiful and the quality of life is seen to be good. Industries concentrate wherever the supply of supportive talent, ideas and capital is abundant.

Connectivity presents choice. When many people make similar choices, concentrations result. Concentrations are not only geographical, but also conceptual and behavioural – from the standardization of managerial preparation in the form of MBA programmes, to the homogenization of crops and farming practices in today's agribusinesses, to the global harmonization of regulations governing banking and trade.

Like complexity, concentrations carry positives and problems. Part II focused on the former. Concentrations bring wealth, ideas, genius and fragments of capability up to a critical mass that catalyses creative achievements.

We've already hinted at its problems. Concentrations put stress on the supply of infrastructure, on resources, even on the sociability and goodwill that help us live together in peace. The higher the stress, the greater the chance of failure. All else being equal, concentrations also mean that when something fails, the costs are likely to be higher and the consequences more severe for more people. Imagine two identical solar flares knocking out the internet planet-wide, one in 1990, the other today. The first would have frustrated some military researchers and the physics community. The latter would be a global catastrophe. (This is more than a thought experiment. In July 2012, a coronal mass ejection (CME) from the sun tore directly through the earth's orbital path, missing our planet by only one week. The US National Academy of Sciences estimated that, had the CME hit, the damages to planetary electrical systems might have exceeded $2 trillion – more than forty times the cost of the costliest hurricane in US history (Katrina).[9])

When syphilis struck the last Renaissance, one of the reasons it hit so fast and so hard was that people had begun to concentrate more in towns and

cities. Urbanization compressed once self-sufficient families and villages into dense units of mutual interdependence. These crowded urban settings offered ideal conditions for the spread of disease. Rooms were overcrowded and public sanitation was stretched. Townspeople lived side by side with domesticated horses, pigs, chickens and their excrement. Itinerants were common; contact with new people and newly arrived travellers was frequent. Promiscuity was on the rise.[10] Poor environmental luck in the 1490s – flooding rivers, exceptionally cold winters – further weakened already vulnerable communities.[11]

Today, overcrowded and overstretched urban places are still a breeding ground for nature's killers.

New poxes

Disease has always been a fixture of human society. Flu epidemics affect an estimated 5–15 per cent of the global population annually, causing severe illness in three to five million people and death in one-quarter to one-half million people, every year.[12]

What is new is the emergence over the past few decades of fast-moving *pandemics* – viruses that spread easily among humans and can infect patients *worldwide* within a short time after their initial outbreak.[13] Already in the twenty-first century, humanity has confronted several such pandemic threats.

SARS

The first proof of these new fast-moving threats to global public health was Severe Acute Respiratory Syndrome (SARS). The first severe infectious disease to emerge in the twenty-first century, we've mostly forgotten about it by now, but in 2003 it suddenly seized the attention of the entire world, and 'pose[d] a serious threat to global health security, the livelihood of populations, the functioning of health systems, and the stability and growth of economies',[14] in the words of Carl Urbani, the World Health Organization (WHO) physician who first identified the disease (and died from it).

A viral infection with flu-like symptoms, SARS first appeared in November 2002, in Guangdong, China. Virologists now believe it transferred from

bats to humans, either directly or via live animal markets in that crowded corner of Asia. By itself, the inter-species hop was unremarkable, and at any other point in human history might have prompted nothing more than a local epidemic. Such was the thinking in Guangdong, when Liu Jianlun, a 64-year-old doctor who had been fighting on the front lines of the outbreak, travelled to nearby Hong Kong and checked into a popular business hotel. Casual contact was enough to cause infection. The hotel's roster of international guests shared elevators, dining halls and other surfaces with Liu. When they flew home – to Canada, Singapore and Vietnam – they carried the SARS virus with them. Within just four months, SARS had spread to every continent except Antarctica. It was the fastest viral propagation in human history.

Luckily, it was far from the most dangerous. By July 2003, some nine months after its initial detection, 8,300 cases (including 775 deaths) had been reported in thirty countries. Once the pathogen had been identified, the same networks that had enabled its spread supported a swift, globally coordinated response under the auspices of the WHO. Aggressive quarantine measures halted the spread of SARS and contained what could otherwise have been a global catastrophe. Part of this success was thanks to planning. The Global Public Health Intelligence Network, an early warning system, alerted the WHO to the new viral threat when it was still in its pre-pandemic stages. Part of it was luck. The first epicentres of the pandemic – Hong Kong, Toronto, Singapore – all enjoyed advanced and robust public health systems able to mobilize and enforce mass quarantine protocols. Had SARS struck less developed centres earlier – say, Lagos or Kinshasa – its ease of transmission could have produced a far worse scenario. In hindsight, SARS was 'the pandemic that didn't occur'.[15]

Still, its costs were significant. While fewer lives were lost than in the average annual flu season, the mortality rate was far higher – almost 10 per cent – and many survivors developed long-term respiratory complications. The scare, plus quarantines, cost the global economy at least $40 billion in 2003. Tourism and travel industries were especially hard hit, but the epidemic also impacted several large manufacturing zones where workers became infected.[16] The disease has still not been completely eradicated, and as yet no vaccine exists that is safe for human use.[17]

Ebola

More recent proof of our rising vulnerability to disease is Ebola. If SARS is the pandemic that didn't occur, Ebola is the epidemic that did.

In December 2013, in a remote mountain village in Guinea, West Africa, near the border with Liberia and Sierra Leone, a two-year-old boy named Emile fell suddenly, terribly ill. It began as a severe fever, and within a few days progressed to severe diarrhoea and bloody vomiting. Within one to two weeks, he was dead from loss of blood and fluids. Following local traditions, his family washed his body and mourners journeyed from surrounding villages to embrace it. Shortly thereafter, several of them began to exhibit the same symptoms.[18]

SARS causes respiratory problems, and occasionally death – usually among older people with a prior condition. Ebola causes massive haemorrhaging throughout the body and historically kills 50 per cent to 90 per cent of the healthy people it infects. Laboratories that research it give it the highest risk classification, biosafety level 4, on a par with anthrax and smallpox. Worldwide, fewer than 100 facilities meet the protocols to handle it safely, which include multiple airlocks and showers, ultraviolet irradiation chambers, air and water decontamination systems, and biohazard suits with independent air supplies for personnel.

West Africa had never seen Ebola before.[19] Since its discovery in 1976, it was almost exclusively a blight on rural Central and East Africa, principally Gabon, Uganda, Sudan and the Democratic Republic of Congo – all some 3,000–5,000 km away. That fact, plus austerity-era budget cuts at the WHO totalling almost $1 billion in 2010–2011, helps explain why no one detected the outbreak until it had spread.[20] By the time the WHO first announced the Guinean outbreak in March 2014, Ebola had already enveloped neighbouring Liberia and Sierra Leone and killed sixty people. In June 2014, with over 750 cases and 467 deaths recorded, it became the worst Ebola outbreak ever.[21] By mid-2015, after 28,000 reported cases and over 11,300 deaths, the human toll had become twenty times higher than all previous outbreaks put together.[22] (The true toll, including unreported cases, may be two to three times higher again.) Meanwhile, the economic losses inflicted upon this already desperately poor region – losses that the World Bank put at up to $4 billion – could eventually kill more people than the virus itself.[23]

If, like SARS, Ebola had struck places with robust public health systems, it could have been contained quickly. For all its deadliness, Ebola is not that easy to catch if one takes proper precautions. A person with measles infects on average eighteen other people; a person with Ebola infects on average fewer than two, in part because the disease kills its carriers quickly, but also because it doesn't spread well through the air. Body or fluid contact, which may include droplets transferred across a shared surface or handling an

infected corpse, is needed. This means that, as long as a community has the capacity to identify and quarantine the sick, to trace and quarantine their recent contacts, and to properly equip front-line medical responders, any outbreak can be swiftly bottled. (The effective containment of Ebola in Nigeria illustrates that with the necessary will and resources it can be contained, even in a relatively poor country.)

Instead, it struck where public health systems are practically non-existent. The three hardest-hit countries – Sierra Leone, Guinea and Liberia – rank among the world's poorest, and all have been the site of recent military coups or civil wars. Sierra Leone's public institutions were just beginning to recover from a long civil war (1991–2002); wounds from Liberia's war (1989–2003) were even deeper and were still being patched over by foreign peacekeepers. Spain, which suffered one local case of Ebola when a nurse became infected after helping two patients from West Africa, spends nearly $3,000 per capita on its health system. Sierra Leone spends $96; Guinea spends $32.[24] The United States, which suffered a handful of cases among returning aid workers, has 245 doctors per 100,000 people.[25] Liberia had fifty doctors, total, in a country of 4.3 million people at the start of the crisis, and several died in the outbreak's early stages.[26] Low levels of education in the remote regions where the virus emerged made the task of these feeble public health systems even harder. Although attitudes eventually changed, at the beginning of the outbreak villagers distrusted officialdom; they mobbed (and, on at least one occasion, killed) health workers who attempted to isolate victims,[27] and persisted with a combination of traditional healing and funeral practices, which encouraged families to bring the sick and dying into churches to be touched by the congregation and faith healers. Sierra Leone's first reported case of Ebola was a traditional healer. The next several cases were the villagers who washed her body.

From the initial outbreak through its rapid spread and eventual containment, the West African Ebola epidemic was a distinctly new kind of public health emergency. How Ebola migrated some 3,000 km from any previous appearance is not yet known. It could be related to recent growth in intra-African trade, which between 1995 and 2012 grew over 10 per cent per year in nominal terms.[28] It could be related to massive refugee flows out of the Democratic Republic of Congo (although most of those flows have been to the Congo's near neighbours). It could be that climate change is altering the habitat of the fruit bats that are the virus's most likely natural reservoir. Population pressures are bringing villagers into closer contact with each other and with the animals that inhabit the forests and provide bush meat.

How the scale of the outbreak exploded *is* known: for the first time ever, Ebola escaped the countryside and reached cities – in large part thanks to the growing links between the two. That is also why the developed world belatedly stepped in with money, health workers and soldiers to help halt its advance – by 2015, about $2.9 billion in combined pledges.[29] Soldiers from the United States and the United Kingdom built makeshift hospitals. Mobile technologies were brought in to help trace victims' recent movements and contacts. Rapid DNA sequencing technology was deployed to map the virus genome in hundreds of patients and get a complete picture of the origin and variety of strains active in the outbreak. (Five of the scientists involved in the initial sequencing effort in Sierra Leone died.[30]) Major pharmaceutical corporations and government labs began to fast-track vaccines through human trials.

A rural epidemic in a faraway place is a tragedy. A lethal virus, normally locked up under the strictest safety protocols known, roaming free in large cities that are linked to the world by air and seaports, is a global security threat. 'Never before in recorded history has a biosafety level 4 pathogen infected so many people so quickly, over such a broad geographical area, for so long', the WHO stated in September 2014.[31] If the virus had spread to other regions with weak health systems, such as developing parts of Asia (with which West Africa has been aggressively boosting trade), a worldwide catastrophe could have followed.

Contagion?

Looking ahead, a worrying pandemic threat on the horizon is H5N1 (bird flu).

Like SARS, H5N1 originated in Southeast Asia, in or near Hong Kong, when the new virus crossed the species barrier from poultry to humans around 1997. H5N1 demonstrates that not only human but also animal populations are now one globally connected pathogen pool. We are all partly to blame. Although infected birds might migrate long distances, the overlap between when they migrate and when they're infectious to other birds is slight; without our help, H5N1 would have remained a local epidemic among Southeast Asian bird populations.[32] Instead, via commercial trade in live birds and animals, H5N1 is today endemic in avian populations worldwide. It has killed tens of millions of birds and forced the culling of hundreds of millions more.

Scientists and public health authorities are watching H5N1 very, very closely. Since 2003, some 600 human cases have been reported in fifteen

countries (generally caused by close and prolonged contact with live infected birds).[33] Studies show that H5N1 is becoming progressively more pathogenic, can survive longer and is able to infect an expanding range of animals, including pigs, cats and dogs.[34] More than 60 per cent of human H5N1 patients have died, a rate that ranks H5N1 with Ebola and other lethal pathogens. Unlike Ebola, it spreads easily (among birds), like the common flu. Although it has not yet developed a reliable human-to-human transmission mechanism, lab researchers have already figured out the mutations that would give it that capability. Nothing but good luck prevents nature from duplicating this result.

The day we all hear of reliable human-to-human airborne transmission of H5N1, the world will come to a sharp stop. Would you get on a plane, knowing that a highly infectious virus with a 60 per cent kill rate might be on board with you? In reality, you won't face this choice: protocols already in place will shut borders and halt air traffic globally. Pandemic models show that a human H5N1 virus could easily surpass the Black Death of 1348–1350 as humanity's deadliest known extinction event. Depending on how early the pandemic is detected and how quickly a vaccine can be deployed, epidemiologists estimate that an H5N1 outbreak could infect up to one billion people and directly cause up to 150 million deaths.† Panic, riots, looting – in general, a fear-fuelled breakdown of social order – could push the death toll even higher. Survivors would suffer a global economic depression with losses in the trillions.

Other pandemic threats

The above three pandemic threats have so far grabbed the most headlines, but they are far from the only ones. In the past two decades, disease experts have identified more than thirty new or resurgent pathogens amidst human populations, including hepatitis C, Zika virus, cholera, malaria and the plague, and nature invents new ones each day.[35] In early 2013, a new bird flu, H7N9, was discovered in China; it causes severe respiratory problems and has been fatal in one-third of confirmed human cases. Unlike

† In 2013, GlaxoSmithKline announced the first FDA-approved vaccine for a strain of H5N1. The first batch may arrive by 2017. Assuming it proves effective against a future pandemic strain, three months would elapse before mass quantities of the vaccine could be made available. Existing plans aim to produce one billion doses within the first year of the virus's initial identification.

H5N1, birds infected with H7N9 can carry the virus without any signs of illness, making it much harder to detect.[36]

Another pandemic at least as terrible as any of the above is HIV/AIDS. To date, it has killed almost forty million people – more than the total population of Canada.[37] It was first identified by US public health authorities in 1981; at the time, perhaps 200,000 people worldwide lived with the disease. By the mid-1980s, that number had soared to three million; by 1990, to eight million; by 2000, to over forty million.[38] Still today, thirty-five million people are living with the disease; every year it kills over 1.5 million people and infects two million more.[39] But it does not threaten all humanity equally. More than two-thirds of those now living with HIV/AIDS are in Africa.[40] While persistent public health and education campaigns, plus antiretroviral therapies, have helped to contain the disease in the developed world, it continues to punish populations that are too poor to afford the necessary drugs or too slow to adapt their sexual practices to the reality of sexually transmitted diseases.

In its reach, persistence and consequences, HIV/AIDS is nothing less than a pandemic. That it no longer makes headlines in the developed world hints at another consequence of flourishing systemic risk: inequality. Renaissance moments magnify the differences in populations' capacities to prepare for, endure and recover from more frequent, more powerful shocks. (We take up this topic in the next chapter.)

Future health demons and dilemmas

So far, humanity has demonstrated that we're up to these new global health challenges. The global system for dealing with disease is among the most developed and effective of all international coordination efforts. The WHO and the national health authorities of its member nations have, for the most part, contained pandemic threats since the Second World War. Newer, more focused agencies like UNAIDS are proving effective at mobilizing global responses to specific blights. Yes, diseases propagate faster and farther than ever before, but so do our detection and response efforts.

On the other hand, the challenges are getting tougher. The complexities are growing. Recent simulations have shown that a contagious airborne pathogen (like H5N1) carried into any major airport, on any continent, would be global within three days at most.[41] If the infected individual took just two plane journeys prior to a public health quarantine, more than five billion people (75 per cent of humanity) would need to be vaccinated to

prevent a global pandemic. After three flights, global vaccination would be required.[42]

The concentration dilemmas are getting thornier. It is not a question of if, but when, a pandemic will strike a major political, financial or industrial centre and force its complete (albeit temporary) isolation from all physical flows in the global system – with hard-to-predict consequences for infrastructure services like energy and IT. But how can any large business avoid locating critical units in places like London, New York or China's Pearl River Delta?

The near-universal use of antibiotics and antimicrobials across the emerging global middle class, in everything from hospitals to cattle herds, is hastening nature's development of resistant superbugs. And our growing connectedness is spreading them worldwide. One superbug, MRSA, has already become a persistent nuisance (and occasionally, a serious threat) in hospitals and nursing homes everywhere. Another, a species of *Escherichia coli* discovered on pig farms in China in November 2015, is resistant to colistin – a powerful antibiotic that pig farmers mixed into feed to keep their herds healthy, but also humanity's 'weapon of last resort' for when all other known antibiotics fail. If this resistance spreads to other bacteria, once-simple infections could become untreatable.[43] For the health security of all, we must sharply reduce our antibiotic use (especially in animals) until replacements can be developed. But who can afford to leave themselves exposed?

In the developing world, cities are booming from a combination of migration and lower infant mortality. Cities offer better jobs, schooling, health and other services, and more opportunities than the countryside. They are also overcrowded and dirty. Human and animal populations are pressed side by side; the water supply is overstressed and easily contaminated. Such conditions were ground zero for SARS and H5N1, and will spawn many future pandemic threats. Those pathogens will threaten us all, and poor cities will be least equipped to confront them. But whether we are born among the rural poor, or born into an urban favela or slum, or born into an advanced economy that's questing for growth opportunities, how can we deny the lure of these emerging centres?

We have been tested many times, and for the most part, we have prevailed. But it is a biological certainty that pathogens will relentlessly assault our increasingly packed and interconnected populations and seek to turn our global shipping and transport infrastructure against us. Nature never gives up.

Merchants of destruction

Crisis then

Part I charted the transformation of Renaissance finance. New continental and intercontinental trade linkages, plus the lure of rapidly expanding coastal economies, shifted the centre of financial activity from the Mediterranean to the Atlantic. New financial instruments and markets were created to supply capital and insurance for increasingly complicated and costly mercantile activities. IOUs were deregulated so that they were no longer limited to the two parties who wrote them. Now, they could be freely bought and sold among third parties in secondary markets. That innovation gave rise to a continental money market, hundreds of times larger in value than the physical market it underwrote. The best-reputed merchant houses could write IOUs to the bourse on the strength of their name alone, with no goods to back them at all. In many ways, they could issue their own money.

Some sovereigns found this pot of easy money irresistible, and they approached merchant houses to raise vast sums on their behalf to finance their wars and ambitions. In one of the more famous episodes of the day, Charles I, king of Spain, borrowed some 850,000 gold florins from the Fuggers of Augsburg in 1519 to bribe the Electors of the Holy Roman Empire and get himself elected emperor.[44]

All of us today have some measure of wisdom about such situations. If we could, we might go back 500 years and warn Renaissance investors of what was likely to happen next. Unfortunately, they had to find out for themselves. The sixteenth-century Fuggers were so eager to lend out (at very high interest rates) to Europe's kings the large sums they could raise off the bourse, they discounted the possibility that even a sovereign might default. Charles I gained his sought-after office, but reneged on most of his debts. The Fuggers, and the myriad smaller investors across Europe who had bought Fugger bills, took heavy losses.

Crisis now

We, too, learned this lesson the hard way, during the 2007–2008 financial crisis and its prolonged aftermath. The hindsight-informed analysis has been repeated many times: American and European bankers played round after round of a highly lucrative game – lending cash to consumers and

homebuyers, moving the debts and risk off their books through securitization and credit derivatives, then lending again – until households were drowning in borrowing costs and the balance sheets of big financial institutions were awash with hundreds of billions in bad debts that would never be paid off. The integration of emerging economies into global capital markets amplified the risk: China was generating lots of cash, saw few investment opportunities in Asia after the 1997 Asian financial crisis, and so ploughed it into the US economy by buying US government debt. The foreign cash infusion helped dampen domestic interest rates and kept the game going longer.

We did not clearly understand how fragile these activities made the global financial system, and we broke it. Households who took out big mortgages in the belief that housing prices would never go down turned out to be wrong. Quantitative analysts or 'quants' who thought they had packaged iffy debts cleverly enough to reduce their risk and maintain their return turned out to be wrong. The 'insurance' (credit derivatives) that firms purchased to hedge against default turned out to be inadequate. The consequences were stupendously bad. By 2009, the financial crisis had already tallied up losses of $4.1 trillion, across every market in the world.[45] Roughly fifty million people lost their jobs worldwide; among those who managed to stay employed, a quarter-billion fell into the ranks of the 'working poor'.[46] In Africa, it is estimated that 30,000–50,000 children died – from starvation – as a direct result of the global economic downturn that followed.[47]

By now, this story has been told so many times – in interviews, editorials, books, and Hollywood-produced documentaries and dramas – that the main lessons easily get blurred in the back-and-forth blame-casting. But viewed through the Renaissance lens, the take-aways come back into focus.[48]

Complexity limits our foresight

The first lesson is how rising complexity makes risks within the financial system harder to see. Colliding connective and developmental forces produced a global financial system that was suddenly far bigger and more complex than just two decades before, which both made the danger harder to see and spread the danger more widely – to everyone.

Looking back, the dangers of rising complexity were obvious. At a systemic level, the balance sheets of the world's countries, institutions and individual investors became fatter and more interconnected (see Part I). At the

product level, financial instruments became more complex, largely thanks to the introduction of progressively more powerful computers into the portfolio-building process. Most large institutional investors, such as pension funds, were prohibited from investing in mortgages and other consumer loans. Entrusted with the accumulated retirement savings of whole industries, pension funds were only allowed to invest in assets that the major credit rating agencies (Standard & Poor's, Moody's and Fitch) deemed safe, and consumer debts were too small and risky to merit the agencies' attention. But what if sophisticated computer algorithms could construct big bundles from thousands of individual mortgages and debts – some more likely to default, some less – such that each resulting bundle met the agencies' thresholds for size and quality? This, essentially, is what mortgage lenders did: a process called securitization.[49] By securitizing mortgages and immediately selling them on to institutional investors, rather than holding those loans and default risks on their own books, mortgage lenders lost the incentive to scrutinize would-be homeowners' ability to repay. Mortgage quality declined, but neither the rating agencies nor the institutions that purchased these complicated products had the analytical powers or motivation to untangle this truth. Some funds insured themselves against the risk of default by buying insurance (called credit default swaps (CDSs)), but this practice only spread the danger to new sectors. Insurers were equally in the dark about the underlying mortgage risk, and pension funds were now *additionally* in the dark about the (in) solvency of the insurance companies whose swaps they bought.† Like a pandemic pathogen, toxic debts originated in a small backwater (subprime mortgage lending) and spread quickly through intertwined balance sheets to threaten the global financial system.[50]

From the top down and bottom up, the financial sector's tangled complexity muddled the vision of those standing in its midst. Neither private- nor public-sector actors saw the accumulating danger. As a Bloomberg columnist observed in 2008, '[The CEO of Bear Stearns] plays bridge, and [the CEO of Merrill Lynch] golfs while their firms collapse, not because they don't care their firms are collapsing, but because they don't know that their firms are collapsing'.[51] In its 2007 Global Financial Stability Report, the International Monetary Fund concluded that 'weakness has been contained

† AIG, the largest such insurer, famously received over $180 billion in federal and state government money to bail it out of insurance claims that would have bankrupted it several times over.

to certain portions of the subprime market, and is not likely to pose a serious systemic risk. Stress tests conducted by investment banks show that ... most investors with exposure to subprime mortgages through securitized structures will not face losses'.[52]

Their concept of risk was linear; it ended at the margins of their own balance sheet, and they either did not see or did not take seriously the big picture that emerged when everyone's balance sheets were laid side by side.

Concentrations weaken our resilience

The second lesson is about how rising concentrations make failures in the financial system more likely to occur. Leading up to the financial crisis, concentrations were increasing at every level.

At the firm level, capital and resources were concentrated into the new securitized mortgage and debt products. At the turn of the century, these products were niche offerings; by the outbreak of the crisis, they had become the second-largest class of asset-backed securities sold in the United States each year. Subprime mortgages were first.[53]

Industry concentration was also on the rise. In the United States between 1990 and 2008, the market share of the top three banks quadrupled from 10 per cent to 40 per cent. In the UK in 2008, the top three banks owned 80 per cent of the market (up from 50 per cent in 1997).[54] The phrase 'too big to fail' entered public discourse to describe these behemoths. Their executives knew their respective governments would never let them go bust – the ensuing chaos would be too great. Their investment discipline weakened – a phenomenon economists aptly call 'moral hazard'. The biggest financial institutions began to take excessive risks, knowing that should things go seriously awry, taxpayers would bail them out. And, indeed, we did.

Concentration also rose at the level of whole economies, as booming financial sectors loomed ever larger in the total economic mix. In the UK, between 1990 and the start of the crisis, the size of the financial sector grew from less than 6 per cent to almost 10 per cent of total GDP, and to over one-fifth of London's economic output.[55] More perilous was Iceland's situation. At the dawn of the new millennium, Iceland, with just over 300,000 citizens, was a small fisheries economy. By 2008, after radical deregulation of the island's tiny finance industry turned Iceland into a haven for European investment, Iceland's banks had racked up $75 billion in debt – one-quarter million dollars for every man, woman and child in the country. When the financial crisis hit, Iceland's kroner plummeted, and a once

simple but stable economy was crippled by the spiralling cost of servicing its foreign debts. Relative to the size of the host economy, it was, and remains, the biggest banking collapse in history.[56] Unemployment soared from zero to 10 per cent, and pensions were wiped out. The IMF bailed Iceland out – on the condition that Icelanders repay UK and Dutch investors a combined 6 per cent of total Icelandic GDP each year from 2017 to 2023. Iceland's financial sector became so big relative to the real economy that when it collapsed, so did the country.

Iceland also highlighted a rising regulatory concentration, as jurisdictions around the globe adopted a common policy of deregulation towards their domestic finance industries – a condition that Andy Haldane, in 2009 the executive director for financial stability at the Bank of England, described as a 'monoculture' that 'became, like plants, animals and oceans before it, less disease-resistant'.[57]

Each of these concentrations posed a genuine dilemma. Each one asked us to trade off legitimate private goals against poorly understood public dangers. What politician could afford to go against the deregulatory trend, when capital seemed so mobile, and loosening credit made voters feel so good? What financial firm could afford to stay out of a new market, when those entering it were profiting so highly? What person would not be tempted by the prospect of buying a house with little or no money down and building equity just by watching its value grow? All of which begs the question: Who, then, was to blame?

The financial crisis showed how difficult these dilemmas can be. Even if the risk of collapse had been more widely understood, it's not clear that we would have acted to prevent it.

Have we learned our lesson? Or will history repeat itself – again?

Taking nothing for granted

Infrastructure is vital. It is, literally, the *structure* that lies *below* (infra) contemporary life, atop which we build our economies, corporations, cities, families and individual life plans. Infrastructure includes the transportation networks over which raw materials, goods and services, people and ideas move; the systems that supply energy, food and water to our populations; the communication channels that handle everything from

remote monitoring of the electricity grid to broadcasting our *Pinterests*; and more.

This infrastructure is under threat, and that by itself is nothing new. Around the world, only a fortunate minority is served by public systems that reliably meet the demands placed upon them. For the global majority, however, the inadequacies of infrastructure are acute and make themselves painfully present in daily life. The same connective and developmental forces that boost health, wealth and populations are multiplying the demands upon lagging and aging infrastructure. Public belt-tightening in the wake of the financial crisis only exacerbates this strain, which is most acute in those areas most crucial to sustaining contemporary life: energy, water and food.

The World Economic Forum puts overall infrastructure investment needs at $100 trillion globally over the next twenty years.[58] It's a rich-world problem. The American Society of Civil Engineers gives current US infrastructure an overall grade of D+. The country's rail and bridges are 'mediocre'; roads, drinking water and waste management systems are 'poor'; levees and waterways score somewhere between 'below standard' and 'unfit for purpose'. Just to put American infrastructure 'in good repair' would take $3.6 trillion of public money through to 2020. Current spending levels ($2 trillion to 2020) will slow US urban decline, but won't help her cities seize new opportunities to flourish.[59]

It's also (and more urgently so) a poor-world problem. By 2020, the developing world will need to double the aggregate $800–900 billion it currently spends on infrastructure each year to meet rapidly rising demands.[60] India, for example, suffers persistent and increasingly severe electricity shortages. In urban centres, the power switches off an average of three hours per day outside the monsoon season, and seventeen hours per day in it, and roughly 40 per cent of the rural population is entirely without electricity. In July 2012, the largest blackout in history left over 600 million people – 9 per cent of the planet's population – without power for more than two days.

Such crises are urgent and immense, but they are well understood. Our responses may be inadequate, but at least we know how to respond: in 2014 alone, a half-dozen new multilateral infrastructure funds and facilities went up. However, the same complexities and concentrations that plague other social systems today also threaten infrastructure, and these risks are *not* well understood. Like all risks to basic social systems, they threaten severe consequences. But unlike traditional infrastructure risks, we do not have a clear idea how we should respond to them, and wealth and advancement offer little defence. The rich and the poor alike are made vulnerable.

Venice is sinking

This is the lesson that Venice learned during the Renaissance. Since the eleventh century, Venice had been the richest and most successful economy in Western Europe; by 1500, its citizens were the richest in the world on a per capita basis.[61] While most economies on the continent were occupied with primary industries – logging trees, raising cattle and grain, or mining useful stuff out of the ground – the Venetian economy was shockingly modern, dominated by trade and trade-related services. It was, in essence, an 'offshore bonded warehouse' (not unlike Singapore today) whose main resources were its location and its large pool of sophisticated merchants who understood supply and demand, consumer choice, on-time delivery and the importance of a supportive tax, legal and currency environment.[62] With its large merchant fleets, Venice was the world's leading naval power, with an undisputed monopoly over most Mediterranean trade, and the first European power to interact seriously and continuously with Islam.

Although Venice also produced fine glassware, silk, paper and other fine craft goods, the majority of Venice's wealth was invested in the trade of spices, which it imported into Europe via hundreds of middlemen and settlements along the Silk Roads and the Indian Ocean. See Figure 7.2.

Figure 7.2 The bulk of Venice's wealth was dependent on the world's longest supply chain.

Source: Greg Prickman (2008). 'The Atlas of Early Printing: Trade Routes.' University of Iowa Libraries. *Retrieved from atlas.lib.uiowa.edu; plus authors' analysis.*

Today, pepper is an optional condiment we sprinkle in our food according to taste. But in an era without refrigeration, pepper, saffron and other spices made the difference between palatable and unpalatable meat. As the continent developed and demand for spices grew, so too did the profits of the Venetian spice trade and the importance of its Mediterranean transport infrastructure to the city-state's economy. A pair of external shocks exposed that same infrastructure as the republic's greatest vulnerability.

The first was the Ottoman advance into the Mediterranean. In 1453, the Turks captured Constantinople, which had stood as an unassailable bulwark against Islamic expansion into European waters for centuries. A series of naval incursions into the Mediterranean followed, culminating with the Battle of Zonchio in 1499. It was the largest naval battle in history up to that time, comprising over 350 ships and 55,000 men, and Venice's failure to win it marked the decisive shift from Venetian to Turkish dominance over the eastern Mediterranean.

From an Ottoman perspective, these events demonstrated the positive impact of new connections and economic development upon their own empire's reach and resources. The Ottomans married gunpowder from China with Hungarian-designed cannons to help assail Constantinople's walls. They harnessed European naval technology to build galleys – bigger, faster and, for the first time, cannon-wielding – to wrest control of the Mediterranean from the world's most experienced naval fleet. They adapted, rather than dismantled, Venice's trade practices and outposts, in order to help fund their empire-building. And they followed up each war with negotiations that reopened trade and exchange between East and West (at revised prices, of course).

The Turks could be negotiated with; geography could not. In the long run, therefore, the second shock was more devastating. In 1499, word reached Venice: three Portuguese *ships* had been sighted at the spice markets of India. Vasco da Gama had found his sea route. 'At the receipt of this news, the whole city ... was dumbfounded and the wisest thought it was the worst news ever heard', recorded Girolama Priuli, a Venetian senator and banker.[63] Venice's merchants knew instantly what it meant: Portugal would buy in bulk, ship direct, and cut out the hundreds of small middlemen and exorbitant taxes that lined the overland route Venice had constructed and maintained over the past seven centuries. In a single masterstroke, a competitor had rendered Venice's Mediterranean spice infrastructure obsolete.

Some consequences were immediate. The following year, prices in Venice plummeted against the expectation of much better terms on the Atlantic coast. Many German spice buyers decamped and moved their business to Lisbon. Eventually, the decline of the overland spice trade proved less precipitous than first feared, because the sea route offered its own dangers: storms, piracy and hostile settlements up and down Africa's West and East coasts. But it was no less inevitable. Cities and regions along the overland trade routes that had flourished as centres of commerce and culture since antiquity – Baghdad, Beirut, Cairo, Damascus, the Black and Red Seas – declined into relative backwaters. Venice herself did not so much decline as start to fall behind. Her economic lords struggled long and bravely, trying in turn to pivot towards shipbuilding, manufacturing and agriculture, but their city-state was too poorly located to compete against emerging ocean-going empires.

Venetians, like so many others, thought their long boom would last forever. They were unable to see the disruptive, non-linear shocks coming – the rise of Ottoman power and the discovery of more and better trade routes – to which their own accumulating trade success had made them so vulnerable. Venice prepared against neither, succumbed to both and dragged a once-thriving transcontinental network down with it.

Thailand is flooding

We have likewise accumulated concentrations in present-day supply chains and infrastructure that make us vulnerable to sudden, hard-to-foresee shocks. Despite its intuitive promise to help diversify risk, 'globalization' has reduced the variety of much of our investment and activity, as we all independently arrive at similar conclusions about how and where to minimize cost, maximize efficiency or achieve other common objectives.[64] It is only recently that we have begun to appreciate how our private pursuits of similar outcomes have increased our collective vulnerability to shock events.

As the Venetians demonstrated 500 years ago, supply chains are especially likely to accumulate concentrations. Private business moves fast in response to profit-making motives. In 1990, Thailand had modest electronics and automotive industries.[65] By 2010, spurred on by Thailand's accession to the WTO in 1995, such manufactures accounted for 35 per cent of Thai GDP and 20 per cent of employment.[66] More than 40 per cent of global hard disk drive (HDD) assembly work and a large chunk of Japanese auto parts manufacturing had relocated to the Bangkok river valley.[67] Why? Because it

offered low-cost labour, preferential government policies and convenient access to nearby Asian hubs. The more businesses that followed this logic into the valley, the stronger the lure of the logic became – which is why when violent tropical storms in late 2011 flooded the region, the consequences – which were deep in Thailand – were widespread. Direct flood damage and losses totalled $40 billion and temporarily forced over two million people out of work.[68] Indirectly, the shutdown of auto part exports forced Nissan and Toyota to halt or delay production in Malaysia, Vietnam, Pakistan, the Philippines, the United States and Canada.[69] The loss of electronics manufacturing capacity caused consumer prices to spike worldwide and hit tech stocks hard on Japan's Nikkei and New York's NASDAQ. With almost half the world's hard drive production underwater, global production of personal computers stalled. Across the Pacific in Santa Clara, California, Intel lost $1 billion of revenue in 2011's fourth quarter.[70] Thailand's HDD manufacturing recovered, but only temporarily. The 2011 floods helped hasten the tech industry's transition to the solid-state drives made by Thailand's neighbours, and since 2013 Thai HDD exports have been falling.[71]

A similar story unfolded when Iceland's Eyjafjallajökull volcano erupted in spring 2010, spewing a cloud of ash over Western Europe that shut down, for six days, all three major airport hubs that connect Europe to the rest of the world – London Heathrow, Frankfurt and Paris Charles de Gaulle. Almost 100,000 flights were cancelled, and the ensuing chaos – ranging from cancelled organ transplants in European hospitals to rotted flowers and fruit in Kenyan and Zambian warehouses – cost the world economy an estimated $5 billion.[72]

New social complexity and concentrations did nothing to raise the risk of a volcanic eruption actually occurring, but they raised its costs once it did occur.

Other big infrastructure failures have been more of our own making. In August 2003, the worst power failure in North American history hit the northeastern United States and Canada. It cast more than fifty million people into darkness for more than thirty hours, at a cost of some $6–10 billion.[73] Until that moment, few in government or the utilities had believed that a single outage on that scale was even possible. But US power consumption had jumped almost 30 per cent in a decade, not least due to the lighting up of the internet. Deregulation and privatization, begun in the early 1990s, had increased the number of parties plugged into the grid from hundreds to thousands. Emerging smart grid devices alongside aging

power stations had complicated control systems. And greater use of renewable generation (which stops and starts according to the vagaries of sunshine and wind) had complicated load-balancing on the grid. Not surprisingly, a joint US–Canada task force concluded in the aftermath that the top two causes of the blackout were 'inadequate system understanding' and 'inadequate situational awareness'.[74]

Clearly, these episodes from our recent past have begun to sensitize us to systemic infrastructure risks. All told, in the first fifteen years of the twenty-first century, natural disasters taught some $2.5 trillion worth of lessons worldwide.[75] It is just as well we've had the chance to learn, because many other disasters loom. Throughout our global transportation networks, there exists much more variety, but far fewer ways to deliver it. Today, just the top thirty airports in the world encounter over 40 per cent of all international passengers and handle over two-thirds of all international freight.[76] The top ten seaports in the world touch fully 50 per cent of the global economy's container traffic.[77] The Strait of Malacca, the main lane between the Indian and Pacific Oceans, through which one-quarter of the world's traded goods and commodities now pass, is only 2.8 km wide at its narrowest point. The Gulf of Aden, connecting the Mediterranean Sea to the Indian Ocean via the Suez Canal, is equally significant, and almost as tight. A diminishing set of global platforms manufactures and delivers a startlingly large share of most everything we care about.

This is especially true of the internet, an entirely new source of systemic risk for the twenty-first century. The internet is so useful that we now use it for everything, and therein lies the danger. From a user perspective, the internet is an invisible field that connects us everywhere. But this user experience has a physical dimension – data centres and fibre-optic cables that concentrate our connectivity in dangerous ways. The same strait and canal choke-points through which the world's ships must pass are also the best routes along which to lay undersea cables. Every year in Africa, the networks of whole countries go dark when one of these threads is cut by a passing ship's anchor – or by an axe. In 2013, the Egyptian coastguard caught three men in a fishing boat who had pulled up, and were trying to hack through a Suez cable deliberately. Their target, the SEA-ME-WE 4 cable, is one of the principal data links between Europe, Africa and Asia. Cutting it would have choked networks on all three continents. It's also possible that some government agencies now possess 'kill switch' capabilities that can effectively shut down a target country's internet connections, by disabling essential servers, for example. That may explain why North Korea's internet

briefly went dark in November 2014, only days after the US government had accused the country of hacking Sony Pictures in retaliation for a recently released satirical film about the North Korean regime.

Meanwhile, the complexity of network hardware and software is making it steadily harder to maintain reliable services. Google, Microsoft, Amazon and Facebook operate around one million servers each, and downtime – due to glitches, natural disasters, equipment failure and human error – is a costly problem. Unexpected data loss and downtime cost businesses as much as $1.7 trillion in 2014, according to one global industry survey.[78] As we become more dependent on the internet, for example through wider adoption of cloud services, those costs will escalate.[79] And the exploitation of so-called 'zero-day' vulnerabilities – unknown bugs buried deep inside the code of widely distributed software or operating systems – threatens to interrupt services deliberately. Often, these bugs are fixed only after hackers have made use of them. In September 2014, a wave of attacks known as ShellShock exploited a core vulnerability in Mac and Linux operating systems to run malicious code on millions of computers. The bug had gone unnoticed for twenty years. Another zero-day vulnerability uncovered in November 2014, called Unicorn, had been present in every release of Microsoft Internet Explorer going back to 1995.[80]

The complexity of internet networks allows attacks like zero-day exploits to be performed with near-perfect anonymity. The most frequent kind of attack, distributed denial of service (DDoS), arranges to send dummy data requests to a victim's server from thousands of hijacked computers simultaneously, so that legitimate users can't get their own requests through. The internet was originally designed for sharing, not security, and perpetrators can hide in the open amidst the unwitting crowds they convene. Even when perpetrators are discovered – often overseas somewhere – limits of jurisdiction make it hard to bring them to justice.

Rare twenty years ago, today cybercrimes are ubiquitous: in our email, on the web, in social media, on mobile devices and on private networks. Their rapid spread has been enabled by maturing online marketplaces that supply cybercrime labour and tools, and that fence stolen goods. It is no longer a question of *whether* you will become a victim of cybercrime, but *when*. These crimes injure us personally, through the theft and ransom of identities, login information, webcam videos or Snapchat photos. They also use us to injure others, by making us unwitting accomplices in spam, phishing and email attacks, or by using our computers as web servers for malware and child pornography. And as more smart devices, from

appliances to automobiles to the locks on our house, connect to the 'Internet of Things', the range of injuries that cybercriminals can cause us will only widen. In July 2015, some 1.4 million Jeeps were recalled when researchers proved they could exploit a bug to hack, and crash, the vehicles remotely over the internet.[81]

Cybercrime also steals intellectual property and other secrets from institutions. In 2014, roughly one-half of small businesses, two-thirds of medium-sized companies and four-fifths of large enterprises worldwide were specifically targeted by a cybercrime.[82] Keith Alexander, the director of the US National Security Agency until 2014, described cyberespionage activity as the 'greatest transfer of wealth in history'.[83] In the United States alone, where half of all cyberattacks originate and are committed, corporate losses from cyberespionage may range from $300 billion to $400 billion per year.[84] These attacks also harm customers and clients, by exposing their personal data and making them more vulnerable to identity theft. A 2014 survey found that 43 per cent of US firms, including the largest online service providers, retailers and banks, had suffered a data breach in the past year.[85] A hack at JPMorgan Chase & Co stole the bank records of seventy-six million households and seven million small businesses.[86] Public data networks are also at risk. Back in April 2007, Estonia, one of the world's earliest adopters of paperless government and internet banking, was brought to a sudden stop when its banks, telecom companies, media outlets and government ministries were all hit by simultaneous DDoS attacks. More recently, in mid-2015, personnel records of 21.5 million current and former employees of the US government, including 5.6 million fingerprint images, were stolen when the Office of Personnel Management was hacked – possibly by a foreign government aiming to recruit informants or identify spies.[87] Other highly sophisticated malware initiatives, likely state-sponsored, have likewise penetrated embassies, research institutes and other sensitive targets of governments around the world.[88]

The rising scale of critical infrastructure connected to the internet – including defence, chemical, food, transportation, nuclear, water, financial, energy and other systems – means that not just cybercrime, but cyberwarfare is now possible. As of 2016, two major cyberattacks causing physical infrastructure damage have been publicly confirmed. In 2010, the Stuxnet worm sabotaged Iran's uranium enrichment infrastructure by infecting control systems and causing the uranium centrifuges to tear themselves apart.[89] (A similar worm had been aimed at North Korea's facilities, but

failed to reach its target because of the country's extreme isolation.[90]) And in 2014, a German steel mill suffered 'massive damage' after cyberattackers gained access to the plant's control systems and caused critical components to fail.[91] Many more such strikes are being attempted. The US Department of Homeland Security reported 245 serious incidents, the majority in energy and key manufacturing sectors, in 2014, ranging from unauthorized access, to malware infections, to data theft (which may serve as reconnaissance for future attacks).[92]

Gone are the days when we could take any of our infrastructure for granted.

Nature as infrastructure

That includes nature. Natural infrastructure, like climate, is the clearest example both of how risks have changed since the last Renaissance and of how the lessons learned 500 years ago are relevant to us.

What's changed is the sheer scale of human activity, amplified by a population that is seventeen times larger and a level of technology that consumes far more energy per person.[93] A half-millennium ago, the forces of nature seemed unconnected to human activities. Humanity could alter the landscape by agriculture and forestry, for example, but mostly nature was taken as a given – a power beyond our influence, let alone control. Today, that's no longer the case. There's no longer a clear-cut division between human-made and natural disasters, because the scale of human activity is sufficient to measurably affect planetary habitats, species diversity, weather, temperature, atmosphere, even sea level.

The lesson – that connective and developmental forces generate challenges of complexity and concentration – is equally obvious. Take climate change, for example. Humanity's relationship with the earth's climate has become one of the most complex phenomena within all science. Natural factors that go into understanding this relationship include, among many others: solar cycles, the earth's orbital variations, atmospheric and ocean currents, cycles of carbon absorption and release by plant and animal life, and the absorptive capacities of different planetary surfaces. Onto such phenomena we next need to overlay the impact of human society: our accelerating production of greenhouse gases, land use, ozone depletion, agriculture, deforestation and so on. Then the hard work begins, to discover the interactions, feedbacks and non-linear tipping points within and between these two giant sets of variables. Cause and effect are very difficult to see. That's half of the reason why it's hard to mobilize strong public action to mitigate

climate change – even when we suffer its consequences, like changing weather patterns, extreme hurricanes and other 'natural' disasters.

The second half of the reason is that climate change presents the paradigmatic concentration dilemma. It is a wholly unintended by-product of human initiative, adventurism, exploration, of connecting and cooperating – in short, of so much activity that we deem to be good. What are we to do, if in our individual pursuit of socially affirmed goals we contribute to an accumulation of carbon pollution that, as seems more and more certain, presents an existential threat?

Risk is flourishing, and the rising complexity and concentrations within our social systems are to blame.

These two factors present different challenges. With complexity, the hardest part of solving a problem is seeing it. If we could see the cause–effect relations at work, we might protect ourselves with some mix of managerial and technocratic solutions – but we can't, so we don't.

Complexity strains our cognition; concentrations strain our judgement. Concentrations are the collective consequence of all our individual choices – choices guided by free will, by ambition, by our duty to loved ones. What do we do when our private actions increase the risk of collective shocks we never intended? Even when we can see the shock coming, there's no easy answer.

We can't avoid these strains. They are the other side of the tangled, rapidly developing age we've been born into, and they permeate our lives – including, as the next chapter shows, our relationships with one another.