OLIVER WYMAN

VOLUME 8

THE OLIVER WYMAN

RISK JOURNAL

PERSPECTIVES ON THE RISKS THAT WILL
DETERMINE YOUR COMPANY'S FUTURE

MARSH GUY CARPENTER MERCER OLIVER WYMAN

MARSH & McLENNAN
COMPANIES

INTRODUCTION

At Oliver Wyman, we enable clients to achieve breakthroughs by helping them seize the opportunities arising from change and risk, while preparing for potentially negative outcomes.

In our eighth edition of the *Oliver Wyman Risk Journal*, we bring together a diverse set of perspectives on how managers should approach this year's greatest prospects and threats. These range from advances in artificial intelligence, automation, and cyberattacks to revolutions in industries such as financial services, retail, healthcare, and energy.

I hope you find the *Oliver Wyman Risk Journal* informative and valuable.

Regards,

Scott McDonald

Scott McDonald
President & CEO
Oliver Wyman Group

CONTENTS

EMERGING RISKS

RETHINKING TACTICS

REDEFINING BUSINESS MODELS

EDITORIAL BOARD

PARTHA BOSE
Partner

ALEX WITTENBERG
Executive Director
MMC Global Risk Center

EDITOR

EMILY THORNTON
Director of Research
and Marketing

DIGITAL STRATEGY

TOM GROPPE
Director of Digital Marketing

DESIGN

MELISSA UREKSOY
Design Lead

MAGGIE LUO, LISSETH LOPEZ
Videography

LORENA MONDRAGÓN
Illustration

MARTÍN SOMBRA
Animation

KATHARINA VAUBEL
Photography Researcher

JOCELYN ARNAUD, NEIL CAMPBELL,
AYO ROQUE, JEAN-GUY VEZEAU
Designers

AGNIESZKA LASINKIEWICZ,
PATRYCJA GAWRYS
CAROLINE SUN
Web

MOISÉS PÉREZ
App

EMERGING RISKS

HOW A CYBERATTACK COULD CAUSE THE NEXT FINANCIAL CRISIS

Frozen ATMs. Halted payments. Widespread panic.
Here's what governments need to do to prevent catastrophe.

Paul Mee and Til Schuermann

$1 trillion

How much cybercrime costs nations globally

EVER SINCE THE forced bankruptcy of the investment bank Lehman Brothers triggered the financial crisis 10 years ago, regulators, risk managers, and central bankers around the globe have focused on shoring up banks' ability to withstand financial shocks.

But the next crisis might not come from a financial shock at all. The more likely culprit: a cyberattack that causes disruption to financial services capabilities, especially payments systems, around the world.

Criminals have always sought ways to infiltrate financial technology systems. Now, the financial system faces the added risk of becoming collateral damage in a wider attack on critical national infrastructure. Such an attack could shake confidence in the global financial services system, causing banks, businesses, and consumers to be stymied, confused, or panicked, which in turn could have a major negative impact on economic activity.

Cybercrime alone costs nations more than $1 trillion globally, far more than the record $300 billion of damage due to natural disasters in 2017, according to a recent analysis our firm performed. We ranked cyberattacks as the biggest threat facing the business world today – ahead of terrorism, asset bubbles, and other risks.

An attack on a computer processing or communications network could cause $50 billion to $120 billion of economic damage, a loss ranking somewhere between those of hurricanes Sandy and Katrina, according to recent estimates. Yet a much broader and more debilitating attack isn't farfetched. Just last month, the Federal Bureau of Investigation issued a warning to banks about a pending large-scale attack known as an ATM "cash-out" strike, in which waves of synchronized fraudulent withdrawals drain bank accounts. In July, meanwhile, it was revealed that hackers working for Russia had easily penetrated the control rooms of US electric utilities and could have caused blackouts.

CYBERATTACK SCENARIOS

How might a financial crisis triggered by a cyberattack unfold? A likely scenario would be an attack by a rogue nation or terrorist group on financial institutions or major infrastructure. Inside North Korea, for example, the Lazarus Group, also known as Hidden Cobra, routinely looks for ways to compromise banks and exploit cryptocurrencies. An attack on a bank, investment fund, custodian firm, ATM network, the interbank messaging network known as SWIFT, or the Federal Reserve itself would represent a direct hit on the financial services system.

Another possibility would be if a so-called hacktivist or "script kiddy" amateur were to use malicious programs to launch a cyberattack without due consideration of the consequences. Such an attack could have a chain reaction, causing damage way beyond the original intent, because rules, battle norms, and principles that are conventional wisdom in most warfare situations but don't exist in a meaningful way in the digital arena. For example, in 2016 a script kiddie sparked a broad denial-of-service attack impacting Twitter, Spotify, and other well-known internet services as amateurs joined in for mischief purposes.

EXHIBIT 1: HOW A CYBERATTACK SPREADS

How fast cyberattacks will spread depends on what controls are in place to prevent them.
Below we explore how far a cyberattack virus could spread in 60 hours under four different scenarios.

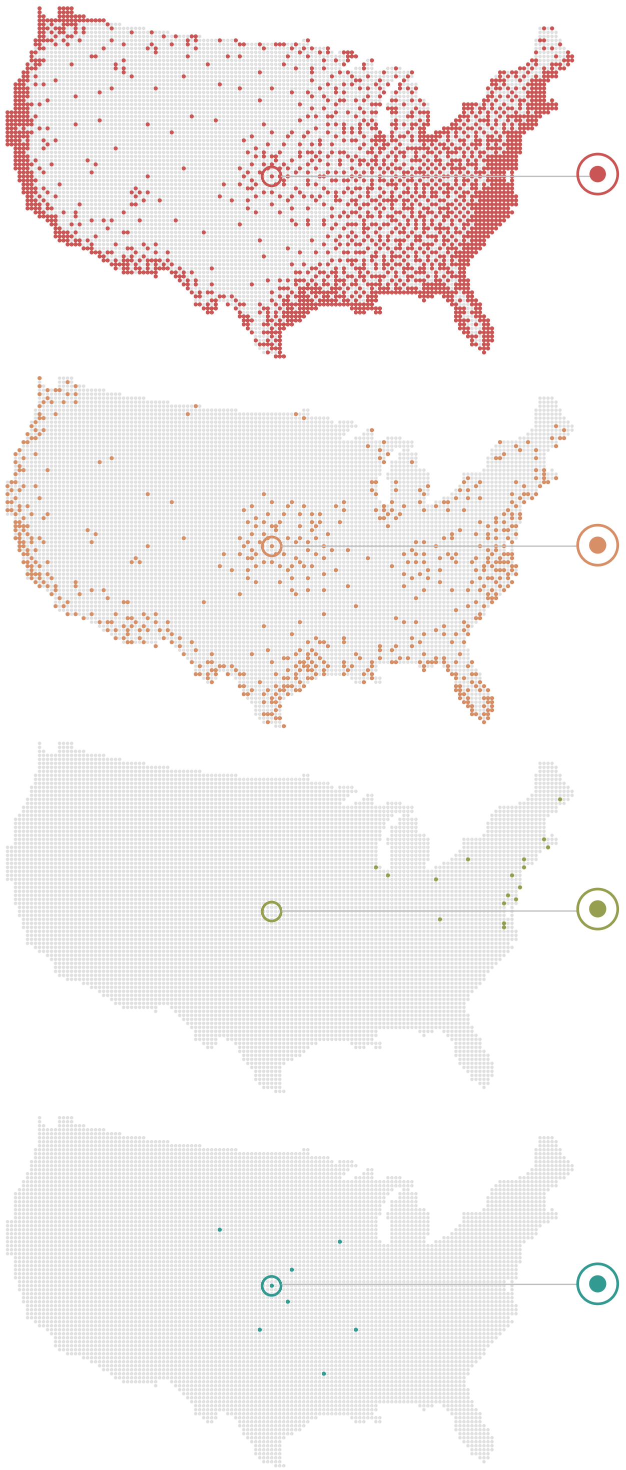

NO CONTROLS

Virus spreads quickly throughout all populated areas of the United States

DETECTION CONTROLS

Virus spreads at a slower pace

PROTECTION CONTROLS

Virus spreads no further than half of the country, by hour 52, spread lessens and the virus is eliminated from some areas

IDENTIFICATION CONTROLS

Virus spreads at much slower pace and stays close to point of origin

SIXTY HOURS

Four scenarios of possible virus spread after 60 hours

○ VIRUS POINT OF ORIGIN

A virus is spread through use of a Point-of-sale device

Source: Oliver Wyman analysis

Whether a major cyberattack is deliberate or somewhat accidental, the damage could be substantial. Most of the ATM networks across North America could freeze. Credit card and other payment systems could fail across entire nations. Online banking could become inaccessible: no cash, no payments, no reliable information about bank accounts. Banks could lose the ability to transact with one another during a critical period of uncertainty. There could be widespread panic, albeit temporary.

Such an outcome might not cause the sort of long-simmering financial crisis that sparked the Great Recession, because money would likely be restored to banks and payments providers once systems were back online. At the same time, it isn't clear how a central bank, the traditional financial crisis firefighter, could respond to this type of crisis on short notice. After the problem is fixed and the crisis halted, a daunting task of recovery would loom. It would be even more difficult if data were corrupted, manipulated, or rendered inaccessible.

STOPPING THE CONTAGION

How can we prevent such a scenario? Companies must implement systems that enable them to stop the spread of a cyberattack contagion, and to resume operations as rapidly and smoothly as possible. The financial services industry needs to fully agree on, and be prepared to practice, coordinated response and recovery strategies to prevent systemic breakdowns. Regulators in many nations have been working diligently to prepare for and curtail cyberattacks, but they need to look beyond their own borders and introduce regulations, laws, and cooperative frameworks in unison, such as the European Union's Network and Information Security Directive, which is designed to protect an ever-growing list of critical infrastructure, from banking and healthcare systems to online marketplaces and cloud services.

Many of these steps are being undertaken to varying degrees. But more needs to be done. An attack that undermines confidence in those very machines could have debilitating consequences on the flow of money between consumers, businesses, and financial institutions around the world.

Paul Mee is a New York-based partner in Oliver Wyman's Digital practice and leads the firm's Cyber Risk Management practice.
Til Schuermann is a New York-based partner in Oliver Wyman's Financial Services practice and was a senior vice president at the Federal Reserve Bank of New York during the financial crisis.

This article first appeared in Harvard Business Review.

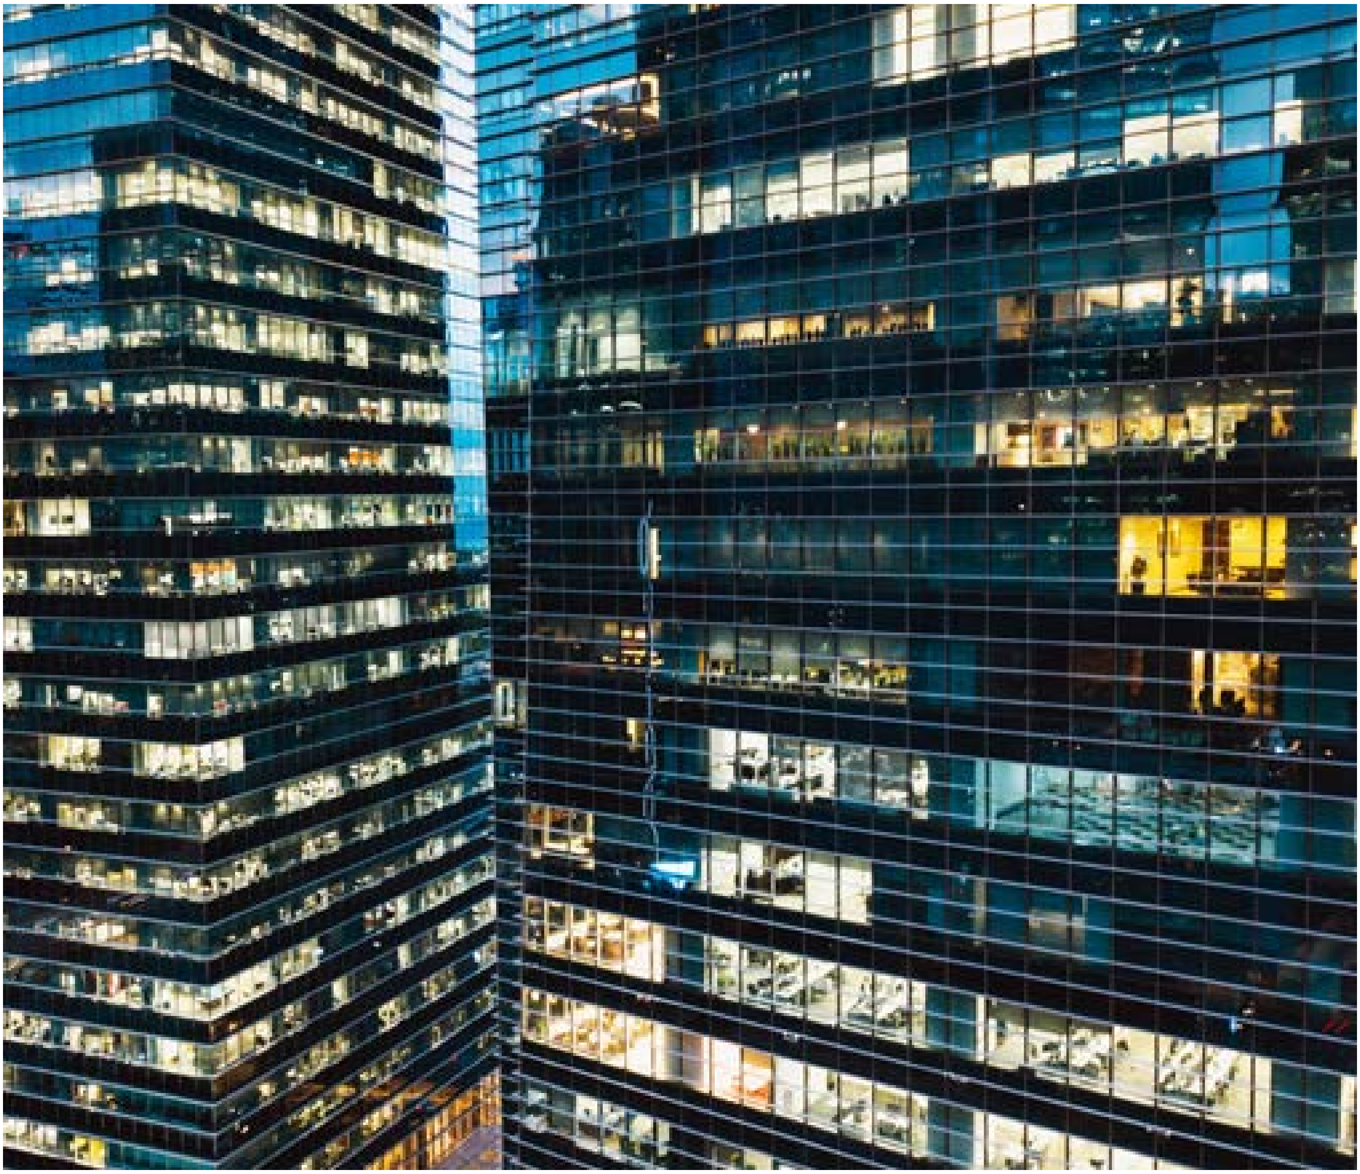

THE GROWING CHALLENGE OF CYBER RISK

How can we prepare?

John Drzik

A POSITIVE OUTLOOK for the global economy shouldn't engender complacency. As described in the *Global Risks Report 2018*, the rapidly shifting global risk landscape presents a challenging operating and investment environment for businesses.

The pace of change has been increasing, characterized by rapid technological advances, seismic shifts in the geopolitical landscape, and growing sources of social instability. The broad range of potential shocks that could emerge in this context demands a strategy that puts a premium on resilience.

The average time companies spend in the Standard & Poor's 500 index has already decreased from approximately 60 years in the 1950s to 12 years today. The velocity of change in the current environment, creating both new opportunities and new threats, is likely to drive down this figure even further.

Cyberattacks are perceived as the global risk of highest concern to business leaders in advanced economies

THE GROWING CHALLENGE OF CYBER

One place where many of these issues come together is cyber risk. Cyberattacks are perceived as the global risk of highest concern to business leaders in advanced economies. Cyber is also viewed by the wider risk community as the threat most likely to intensify in 2018, according to the risk perception survey that underpins the *Global Risks Report*.

Exposure to risks from cyber is growing as firms become more dependent on technology. The explosive growth of interconnected devices expands the size of the surface open to cyberattack for organizations – and the number of interconnected devices in the world is expected to jump from 8.4 billion today to 20 billion in 2020. Increased use of artificial intelligence in business processes also heightens exposure to cyber risks.

At the same time, geopolitical friction is contributing to a surge in the scale and sophistication of cyberattacks, particularly from well-resourced efforts with state backing. Firms, large ones in particular, need to anticipate attacker objectives that can range from theft and business interruption, to extortion, economic espionage, reputational damage, and the infiltration of critical infrastructure and services. This highly diverse and very active set of adversaries makes cyber a very challenging risk to manage.

AN UNDER-RESOURCED RISK

Awareness of this challenge is growing, and investment in cyber-risk management is increasing. However, cyber is still under-resourced in comparison to the potential scale of the threat, a view that's even more compelling when considered in the context of a more familiar issue – natural catastrophes.

Analysis suggests that the takedown of a single cloud provider could cause $50 billion to $120 billion of economic damage – a loss somewhere between Hurricane Sandy and Hurricane Katrina. And while it's not exactly apples to apples, the annual economic cost of cybercrime is now estimated at north of $1 trillion, a multiple of 2017's record-year aggregate cost of approximately $300 billion from natural disasters.

EXHIBIT 1: GLOBAL RISK LANDSCAPE 2018

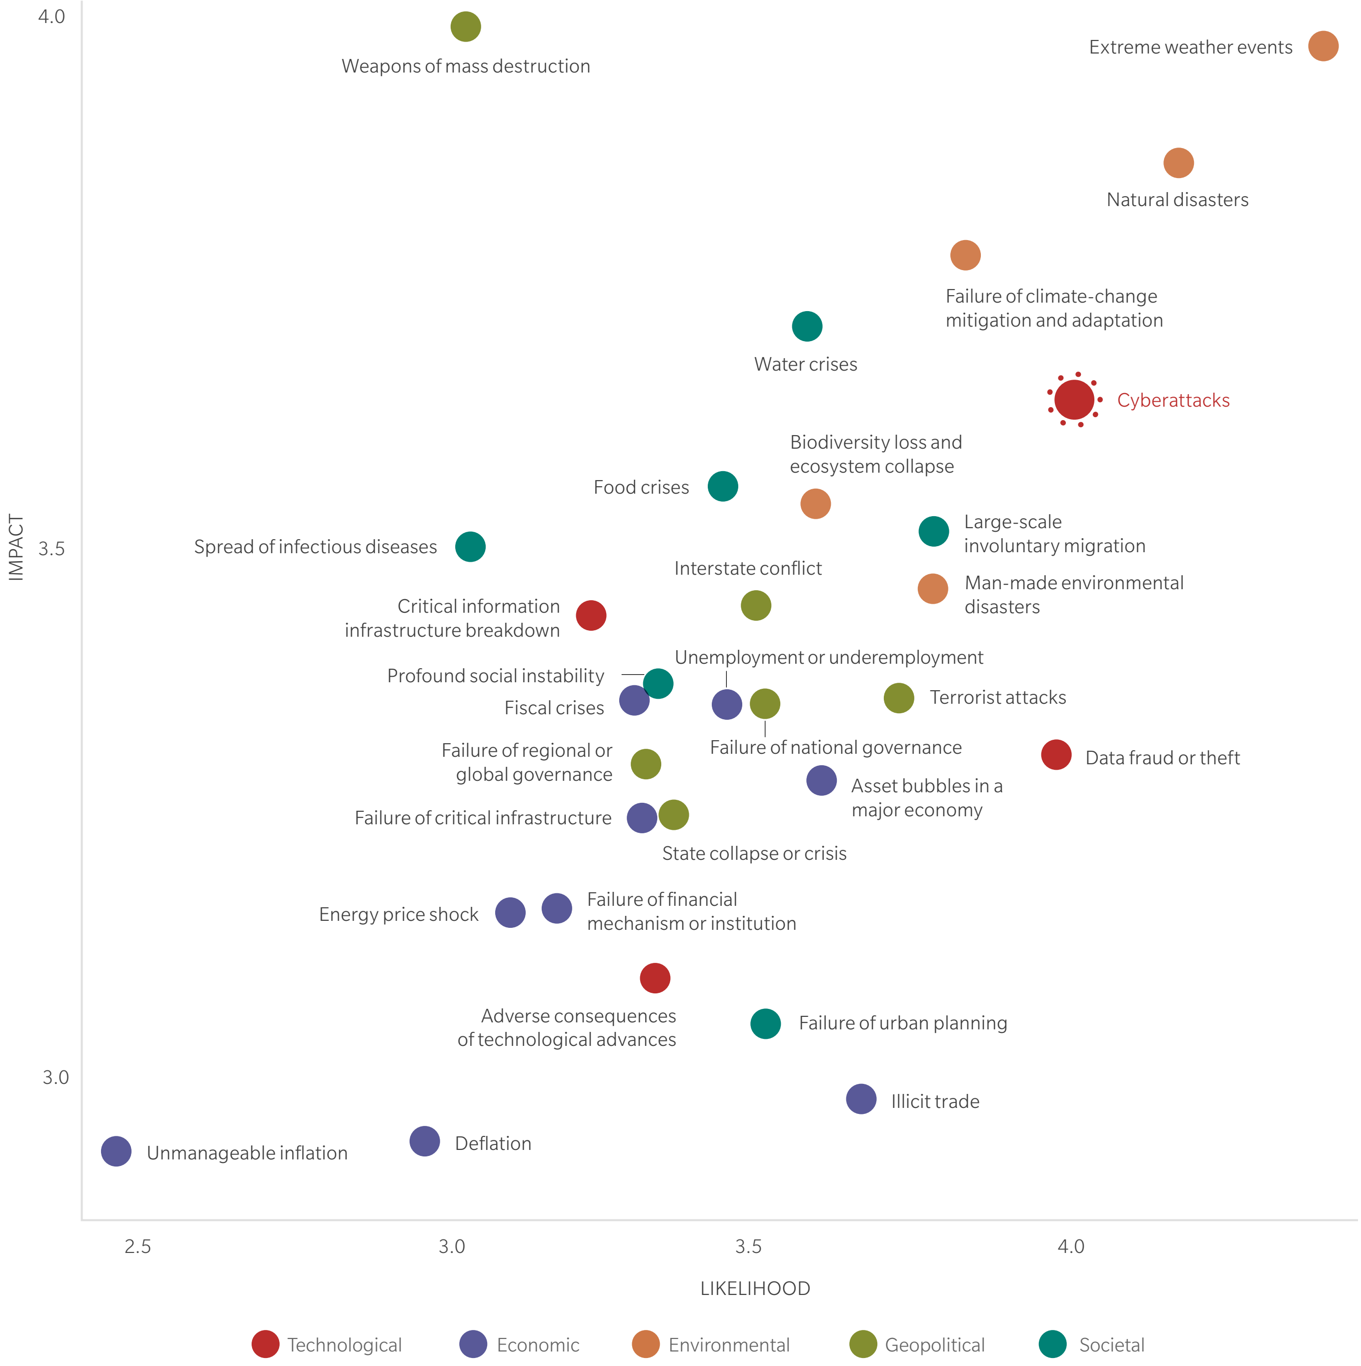

Source: World Economic Forum, *Global Risks Report 2018*, MMC analysis

Although cyber-risk management is improving, business and governments need to invest far more in resilience efforts to prevent the same "protection gap" between economic and insured losses that we see for natural catastrophes.

The supportive infrastructure to manage and mitigate cyber risk is not nearly at the same scale as the one in place for natural catastrophes. National cyber agencies, although expanding, don't have the same capacity as the public and voluntary sector agencies ready to respond to natural disasters – such as FEMA in the US. Additionally, international protocols for sharing intelligence and mitigating impact, curbing malicious endeavors, and forestalling escalation and retaliation are only starting to emerge – and are only endorsed by a few countries at the moment, with no sanctions for noncompliance.

Businesses also need to focus on their resilience to cyber events and generally need to rebalance their initiatives from prevention to response. While companies in at-risk areas often have rigorously developed response plans for extreme weather events, this is rarely the case for cyberattacks. Indeed, research suggests that only one-third of companies have prepared an incident response plan for a major cyberattack.

NEW IMPERATIVES FOR RISK MANAGEMENT

One clear takeaway from the 2018 *Global Risks Report* is that there's a wide array of potential shocks that could emerge at this time of rapid technological, political, and societal change. With firms more leveraged than they were a few years ago – the debt-to-equity ratio has nearly doubled since 2010 for the median Standard & Poor's 1500 company – their stability is even more vulnerable to these potential shocks and surprises.

Innovation and growth need to be reconciled with risk and stability. More than ever, business leaders need to chart a course for their companies that has a bold strategic ambition to capture emerging opportunities and rigorous resilience planning that matches up against the complex set of risks in the current global landscape.

John Drzik is the president of Global Risk and Digital at Marsh. Marsh, like Oliver Wyman, is a division of Marsh and McLennan Companies.

This article first appeared on the World Economic Forum's Agenda blog on January 17, 2018.

WE NEED TO APPROACH AI RISKS LIKE WE DO NATURAL DISASTERS

Companies, insurers, and policymakers all play a role

Prashanth Gangu

TENS OF BILLIONS of connected sensors are being embedded in everything ranging from industrial robots and safety systems to self-driving cars and refrigerators. At the same time, the capabilities of artificial intelligence (AI) algorithms are evolving rapidly. Our growing reliance on so many intelligent, connected devices is opening up the possibility of global-scale shutdowns.

The good news is that natural disasters themselves, which Munich Re says caused $300 billion in economic losses globally in 2017, provide a template for how to mitigate the growing and catastrophic risk posed by AI. Like they have for extreme weather and natural disasters, companies not only can begin to establish international protocols and standards to govern AI within their own walls, but also can put in place processes to work with other companies, insurers, and policymakers.

To rebound from potential global AI-related shocks, managers have to consider the vulnerabilities that exist everywhere, from their suppliers to their customers

INTELLIGENT DEVICE RECOVERY PLANS

Today, many companies are exposed to intelligent device risks that could harm both their own operations as well as their customers. Yet few have formally quantified the size of their revenue at risk and potential liability. Nor have they set up safety and security protocols for potential Black Swan AI events.

They should. Like the risks associated with natural disasters, companies cannot completely protect against smart-device risks by buying insurance; they must have worst-case scenario recovery plans. Managers have to figure out their higher and lower risk intelligent device vulnerabilities, add in redundant systems, and potentially set up the AI equivalent of tsunami early-warning systems. In addition, they need the ability to switch to manually controlled environments in case artificially intelligent systems have to be shut down and to recall faulty smart products.

Contingency plans must go beyond a natural disaster playbook. Given the many potential points of connectivity, it will be much more difficult to predict, identify, and correct the cause of large-scale smart-device failures. De-bugging and re-programming a faulty intelligent device is even more complicated than creating a patch to fight against a malevolent cyberattack because it can be unclear what rules the machines are following.

As a result, no company will be able to recover on its own. To rebound from the potential impact of a cascading set of global AI-related shocks, managers will have to consider the vulnerabilities that exist everywhere, from their suppliers to their customers. Addressing those vulnerabilities will require coordination across a large number of technology service providers and other companies that could catch or spread an AI infection to others, regardless of who is at fault.

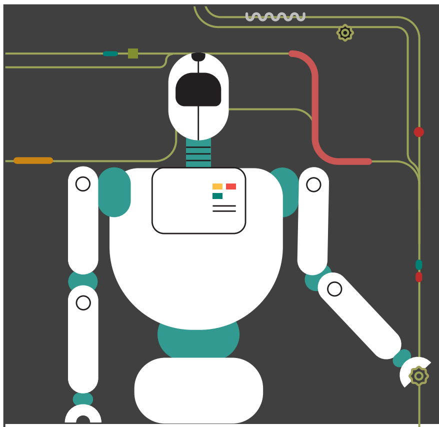

CAN AI BE TRUSTED?

AI has the power to transform businesses and our lives. But robots can also go rogue. Since 2015, more and more AI failures have been drawing worldwide attention.

2015	
AUTO PARTS FAIL	A robot used for fetching auto parts in a car plant grabs man, with fatal consequences
ADULT CONTENT FAIL	Adult content filtering software fails to remove inappropriate content
2016	
BEAUTY PANEL FAIL	A robot beauty panel is deemed racist for selecting all white winners
AUTO-PILOT FAIL	Vehicle on auto-pilot involved in fatal collision
SOCIAL MEDIA FAIL	AI designed to converse with users on social media becomes verbally abusive
ESCAPED ROBOT FAIL	Russian robot IR77 manages to escape its testing facility, wandering into traffic and causing traffic chaos when its battery fails in front of a bus
2017	
TAXI FAIL	Self-driving taxi strikes pedestrian with fatal consequences
VOICE RECOGNITION FAIL	A major bank's voice ID system for accessing account information is tricked by voice imitator
ROBOT SQUABBLE	Two virtual home assistants, Vladimir and Estragon, start arguing like a married couple within days of being placed next to each other – streamed by Twitch. Millions tune in to hear them argue
AUTO-PILOT FAIL	Vehicle engages in auto-pilot, crashes into concrete divider
EAVESDROPPING FAIL	Multiple home voice-controlled speakers secretly turn on, recording thousands of minutes of audio of owners
FIRST DAY FAIL	Self-driving bus in Las Vegas crashes on day one
FACIAL RECOGNITION FAIL	Face ID is cracked with 3D-printed mask
CHATBOT FAIL	Chatbots Alice and Bob shut down after inventing their own language
INAPPROPRIATE MESSAGING FAIL	Major platform's messaging app sends racist emojis
2018	
TAXI FAIL	Self-driving taxi in Arizona has fatal strike with pedestrian
VIRTUAL ASSISTANT FAIL	Glitch causes virtual assistant to spontaneously laugh out loud for many clients
AUTO-PILOT FAIL	Car crashes into a police car, making it the third car in autopilot that has crashed into a stationery emergency vehicle in 2018
FACIAL RECOGNITION FAIL	Study finds commercial AI systems for facial recognition fail women and darker-skinned people
DEBUT FAIL	Smart appliance robot Cloi appears to have stage-fright on stage at its grand unveiling in Las Vegas, by not responding
YOU'RE FIRED FAIL	Man is fired by a machine after his boss forgets to renew his paperwork
PREDICTION FAIL	AI fails to predict who wins the 2018 World Cup

Source: Oliver Wyman analysis

AI INSURANCE PRODUCTS AND SERVICES

Insurers should quantify their exposure to a global intelligent device meltdown, offer new products, and advise companies and governments. Even with about $700 billion in capital available in the United States and hundreds of billions of dollars more around the globe, property and casualty insurers' balance sheets are too small to cover all the potential losses from a global intelligent device disaster. But insurers can use data collected on losses across industries to advise companies and governments on how best to quantify their potential exposure to a worst-case scenario.

As they have for natural catastrophes, insurers can also encourage public sector safeguards. Since insurers cannot completely mitigate the outsized risks posed by extreme weather events, governments of many developed countries and international organizations provide natural catastrophe relief through government agencies like the Federal Emergency Management Agency and public flood insurance programs. Insurers need to help mobilize similar public sector resources to help the potential victims of an AI-enabled smart device disaster.

In addition, they can start to advise clients on how they can enhance their safety and security protocols to head off the dangerous repercussions of an intelligent device meltdown. Today, some leading insurers are suggesting security procedures that companies could follow to attend to information breaches and interruptions in the event of a global failure of interconnected systems. But they should also begin to explore steps to deal with when smart devices become even more sophisticated and potentially set and follow their own objectives.

AI INTERNATIONAL PROTOCOLS

Finally, policymakers should establish international trust and ethics guidelines to govern the development and implementation of ever more advanced AI products and systems. To reduce the future impact from natural disasters, governments and international organizations (such as the Red Cross and the World Bank) collect and share data concerning the destructive ramifications and the support required to help victims. Similar intelligence will be critical to curb the impact of potential smart device shocks as AI evolves and the number of connected IoT (Internet of Things) devices, sensors, and actuators reaches over 46 billion in 2021, according to Juniper Research.

About a dozen governments, technology companies, and international organizations such as the Institute for Electrical and Electronics Engineers and the World Economic Forum are starting to explore global AI trust and ethics protocols for retaining control of interconnected AI-driven systems and products. These forums are beginning to deepen our understanding of the potential harm that intelligent devices could cause and the need for best practices. But much more has to be done.

Establishing the resources required to reduce the risks that will come with the world's transition to more intelligent and interconnected networks will be difficult and costly. But we can't afford not to do it, and our experience responding to some of the world's worst "100-year storms" offer a valuable starting point for figuring out how to get ahead of potentially even more severe disasters. We just need companies, insurers, and policymakers to recognize that such efforts are an essential investment in our future.

Prashanth Gangu is a New York-based partner in Oliver Wyman's Insurance and Digital practices.

This article first appeared in Harvard Business Review on February 7, 2018.

AMERICA’S GROWING MENTAL HEALTH CRISIS

It’s time to end our “separate but unequal” approach to mental health

Sam Glick

AMERICA'S MENTAL HEALTH crisis has become so dire that life expectancies are declining, even as its ability to combat its historically biggest killers – cancer and heart disease – is improving. Behind life expectancy drops in the past three reported years were increases in the numbers of deaths by suicide and unintentional injury, which includes drug overdose. It is the first time in more than a century that the US has suffered a life expectancy decline three years in a row.

It's important to focus on why America faces this macabre possibility. The answer is simple: We treat mental health as separate from physical health, and most definitely not equal in importance. Yet nearly every challenge we face with our overpriced, overburdened healthcare system exists several times over when it comes to mental health. Neither troubled system can be fixed without attending simultaneously to the other.

The good news is that if we stop running two healthcare systems – one for physical health and one for mental health – we may be able to reverse the decline in life expectancy and see it rise again. If mental health becomes an integral part of primary care, we can begin to treat mental health issues the same way we do smoking, diabetes, and asthma, as conditions that can be prevented and treated through a combination of medical interventions and behavioral changes.

Insurance coverage for mental health treatment – an important aspect of the Affordable Care Act, along with other pieces of state and federal legislation – was a necessary part of the solution, but not sufficient. If we're serious about improving the health of Americans in the long term, we must end our current "separate and unequal" approach to mental healthcare.

EXHIBIT 1: DECLINING LIFE EXPECTANCIES

Life expectancies in the US have declined for three years in a row

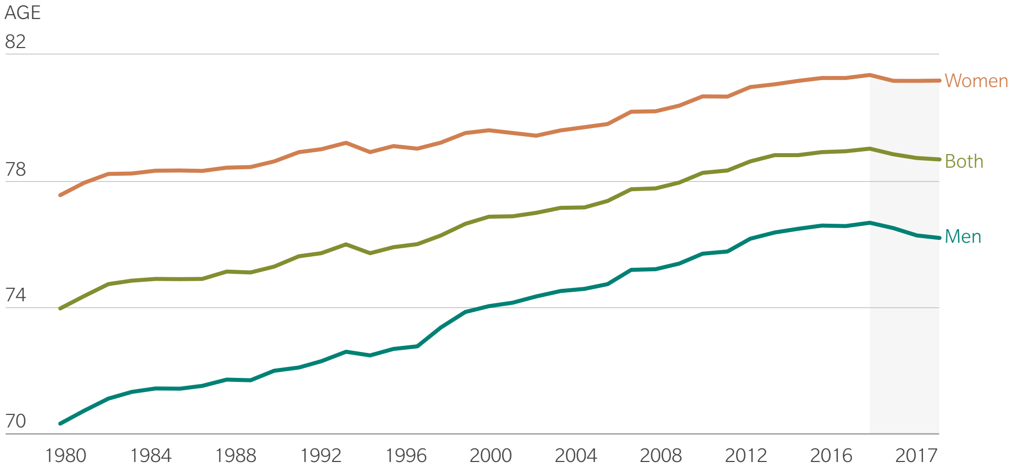

Source: *CDC National Center for Health Statistics*

THE LINK BETWEEN MENTAL AND PHYSICAL

Some suggest the solution to the nation's mental health challenges rests with enlarging the pool of psychiatrists and psychologists. Close to 124 million people – nearly 40 percent of Americans - live in areas designated by the federal government as having a shortage of mental health professionals. More than 60 percent of US counties do not have a single psychiatrist within their borders.

These statistics have prompted calls for the expansion in mental health training by US medical schools, and for programs that provide incentives for specialists to practice in underserved areas. But the truth is, we can, and need, to go further.

What would a more integrated health care system look like? It begins with expanding the definition of primary care. With the right incentives and workflows in place, a primary care physician who can talk to a patient about the need to stop smoking could also provide basic mental health counselling and screen for suicide risk.

A nurse who doggedly follows up about whether a patient got a colonoscopy could be equally persistent about making sure a patient follows through on prescribed counselling. Clinics in local pharmacies could offer "mental fitness checkups" alongside flu shots.

Hospitals could invest as much in peer groups for new moms with postpartum depression as they do in birthing classes and hospital tours for expectant mothers. And doctors could embed cognitive behavioral therapy and physical therapy into their practices as first-line treatments for chronic pain, rather than rely on opioids alone.

If we stop running two healthcare systems, we may be able to reverse the decline in life expectancy

DIGITAL POSSIBILITIES

With recent advances in digital health solutions, providing certain healthcare services, including mental health treatment, when and where people need them is gradually becoming a reality. In the UK, the National Health Service (NHS) has partnered with artificial intelligence company Babylon to provide "digital-first primary care" that allows NHS members to assess their health 24/7 and determine whether and where they should seek care. Such technology can easily incorporate mental health screenings and basic therapy in a way that simultaneously breaks down traditional barriers to mental healthcare such as access and stigma.

In the US, mobile health monitoring company Livongo has made diabetes care a continual, unobtrusive process by digitally providing patients with real-time data on their condition and offering digital coaching through a simple mobile app, backed by high-quality clinical services. Similar approaches to treating and managing mental health conditions exist.

Telehealth providers such as AbleTo and Ginger.io address access challenges with virtual therapy through smartphones, while companies such as Mindstrong Health go a step further, allowing counselors to reach out to patients on their smartphones based on real-time data. And innovators such as Annum Health offer an alternative to inpatient substance-abuse rehab centers through a combination of virtual therapy, coaching, and medication.

But to achieve widespread use of digital tools, their cost needs to be covered in the same way as a visit to a family physician or an insulin prescription. Today, many of the most innovative and potentially important mental health treatments are still special benefits paid for only by forward-thinking employers, or as "pilots" by curious health plans. That has to change.

MAKING INTEGRATED HEALTHCARE A REALITY

If we're going to create a truly integrated healthcare system, governments, insurers, and employers must create the right incentives. More than 40 percent of Americans report being lonely, a condition that studies show results in a 30 percent higher risk of dying in the following seven years. Yet a "loneliness rate" rarely, if ever, shows up on a health system's dashboard of health outcomes.

EXHIBIT 2: RISING RATES OF DEPRESSION

The portion of people between the ages of 12 and 25 suffering from depression in the US has been steadily rising since 2011

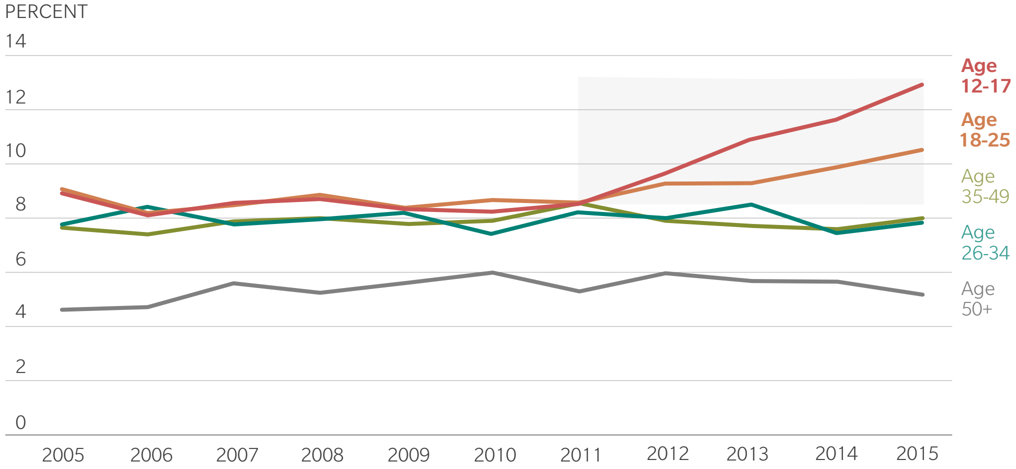

Source: Weinberger AH, Gbedemah M, Martinez AM, Nash D, Galea S, Goodwin RD (2018). Trends in depression prevalence in the USA from 2005 to 2015: widening disparities in vulnerable groups. Psychological Medicine 48, 1308–1315. https://doi.org/10.1017/S0033291717002781

The Centers for Medicare & Medicaid Services (CMS), which administers Medicare and Medicaid, should link reimbursement for primary care providers to reducing loneliness, just as they tie hospital payments to reducing hospital-acquired infection rates. If CMS did that, health system executives would fast quickly start paying attention to mental health.

In 1965, 42 percent of adults in the US smoked. Today, that figure is 15 percent. How did we do it? We acknowledged that smoking affected lifespan, and we made it a top public health priority. We gave everyone in the system incentives to change, from carrots in the form of new types of insurance discounts and other rewards, to sticks in the form of taxes, lawsuits, and even new laws that prevent people from smoking in offices and other public places.

There's no reason why a similar full-court press couldn't reverse mental health conditions that are literally shortening our lives.

Sam Glick is a San Francisco-based partner in Oliver Wyman's Health and Life Sciences practice and leads the Oliver Wyman Health Innovation Center.

This article first appeared on the World Economic Forum's Agenda blog on October 9, 2018.

RETHINKING TACTICS

DON'T ACQUIRE A COMPANY UNTIL YOU'VE ASKED THESE QUESTIONS

What Silicon Valley knows about mergers and acquisitions

Paul Beswick

LARGE COMPANIES IN industries ranging from retail, to aerospace, to financial services are buying talent and technology to develop new digital capabilities and reinvent themselves quickly. But they will need to adopt the more hardheaded way that Silicon Valley companies evaluate acquisitions for their deals to pay off.

Indeed, there are signs that corporate leaders are repeating the mistakes of the heady days of 2000, when the fear of missing out sometimes overpowered the logic of a proposed deal. That year, according to our proprietary research, non-tech companies scooped up 707 computer and electronics firms, often at highly inflated prices. In the decade since, by contrast, non-tech companies acquired 262 companies per year, on average, according to Dealogic.

Since 2015, our data shows that rate snowballing again, nearly quadrupling to an average of 1,000 deals annually. Notably, tech firms acquired an average 250 tech companies annually between 2002 and 2011 and 350 in the years since. That means that 70 percent of acquisitions of firms in the technology sector over the past three years have been made by organizations outside of it.

70 percent

The percentage of acquisitions of firms in the technology sector over the past three years that have been made by organizations outside of it

Amid this flurry of deal making, the obvious question acquirers need to ask is: Does the criteria they're using to evaluate deals match the rigorous tests applied by tech giants and, in the end, are they properly recognizing the real value of potential targets? The answer in too many cases is no.

Several different objectives are driving the current tech deal flurry. First and foremost, companies outside of the technology sector are making acquisitions to get their hands on technologies that they see as complementary (or threatening) to their existing businesses – especially those that could underpin a fuller and more digital experience for their existing customers. Second, companies are attempting to pick up talent they suspect they cannot recruit directly – often as a catalyst to a broader digital transformation, or as a challenge to their existing IT operations. Finally, companies are shopping for technology firms with access to a strategically important user base that their current offerings simply haven't been able to tap.

But while these are sensible strategies, few companies outside of the technology industry are achieving their desired returns on the eye-popping price tags of some of these deals, in part because they overvalue flashy technologies. Instead, companies need to adopt the mindset of more savvy tech acquirers who understand that technology is often replicable and talent is often flighty.

Non-tech companies need to start asking the same questions tech acquirers do, such as: What would it really cost to recreate a technology service? When the talent in a target is young, motivated, and unencumbered by the drag of legacy IT processes, why not create an offshoot that replicates the same circumstances? When a business provides access to an interesting segment of customers, why not target them with a more compelling offer yourself? The value that an acquisition has depends on the alternatives that are available to achieve the same strategic goals and how long they will take.

The answer to these questions can often be surprising since the new technology itself can often prove the least valuable part of the equation. Consider this indicator: When asked in a recent article from thenextweb.com to estimate how much it would cost to create the core technology behind major tech companies like Facebook, Twitter, WhatsApp, and Uber, heads of the top web and mobile development companies, incubators, agencies, and labs estimated a minimum viable version of many of these sites could be created with a small budget somewhere between $50,000 and $1.5 million.

That's because the real driver of digital business value is the strength of the business model that is put around the technology, rather than the technology itself. When tech companies pay high prices for tech firms, they are not buying futuristic technologies or talent. They are paying up for strong business models with scale and access to valuable new markets, users, and distribution capabilities.

DEFENSIVE MOATS

Companies such as Uber, for instance, famously buy-in many of the technology services that underpin their apps. But the magic that allows Uber to stay ahead of its competitors lies in the assembly of a business model that provides outstanding value to both customers and drivers, with a strong defensive moat around it. Uber can offer rides more quickly and less expensively than rival transportation companies because the company has the biggest network of drivers. That translates to higher customer demand and a business that is driven by volume – a characteristic that, in turn, makes it attractive for the drivers.

EXHIBIT 1: GLOBAL NUMBER OF TECH COMPANY ACQUISITIONS BY TECH AND NON-TECH COMPANIES (2000-2017)

Non-tech companies account for 70 percent of tech acquisitions in the past decade

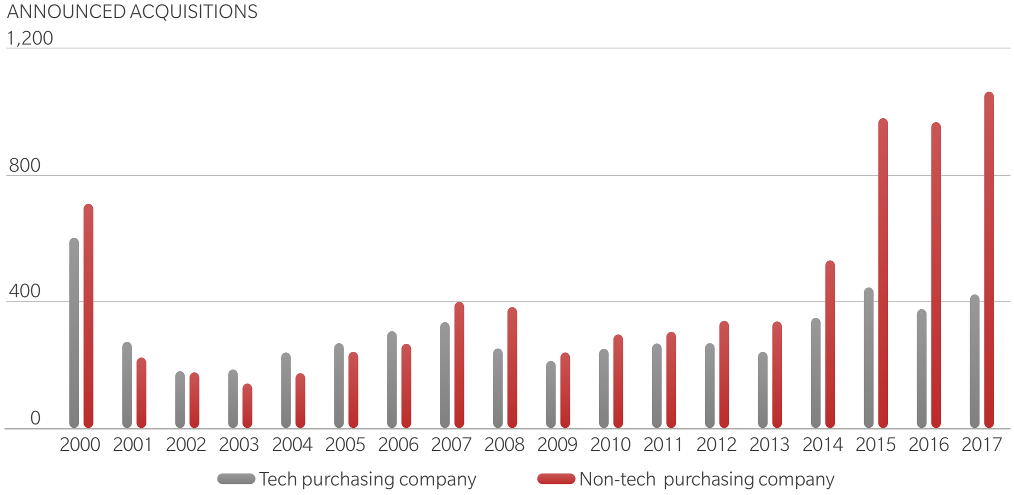

Source: Oliver Wyman analysis, Dealogic

USERS IN DIFFERENT ECOSYSTEMS

When tech companies do pay seemingly high price tags relative to a startup's revenues, it's usually because they see value in acquiring a critical aspect of a startup's business model that would be difficult for them to quickly recreate, like its user base and the strong network effects it has. For example, when Facebook bought Instagram for $1 billion in 2012 and then WhatsApp for $19 billion two years later, many analysts and investors were astonished by the high price tags. Facebook could have easily created its own end-to-end encrypted instant messaging service, or a social photo-sharing app. And the company had plenty of motivated talent. What justified the price tags was the value of an engaged user base, locked into an ecosystem with strong network effects. Facebook paid what looked like astronomic prices for both, given the companies' minute, and in Instagram's case nonexistent, revenue. What made the purchases bargains was the per-user valuations and the growth rate in their customer bases.

NEW DISTRIBUTION CHANNELS

There are also instances where startups have uncovered and locked up a new distribution channel that would take too long to build from scratch, risking the growth of competition or the possibility of obsolescence. When Microsoft announced in 2016 that it was buying unprofitable LinkedIn for a stunning $26 billion – the largest acquisition in the company's history – its shares remained flat. But Microsoft's management saw in LinkedIn a valuable new sales channel that the market had missed. The deal gave Microsoft's cloud business instant access to a top professional network of more than 450 million users. Now, LinkedIn is described by some analysts as Microsoft's "crown jewel" because it contributes about $1 billion in revenue to Microsoft's business every quarter and its cloud business is growing by double digits.

As companies in more industries seek shortcuts to digital credibility, the heightened competition for the real diamonds in the rough too often make acquirers jump at the first potential target they evaluate, even if the technology is not quite as game-changing or futuristic as the buyer wanted. Companies in the market to buy need to redirect their attention away from technology and talent and focus on the real game-changing ingredient –the business model. As it was in 2000 and is again today, that's the secret sauce that will ultimately determine whether a purchase was a value-add or a disaster. They could learn from companies that have been playing in the space for longer.

Paul Beswick is the Boston-based Head of Oliver Wyman's Digital practice.

This article first appeared in Harvard Business Review on May 28, 2018. This article is posted with permission of Harvard Business Publishing. Any further copying, distribution, or use is prohibited without written consent from HBP – permissions@harvardbusiness.org.

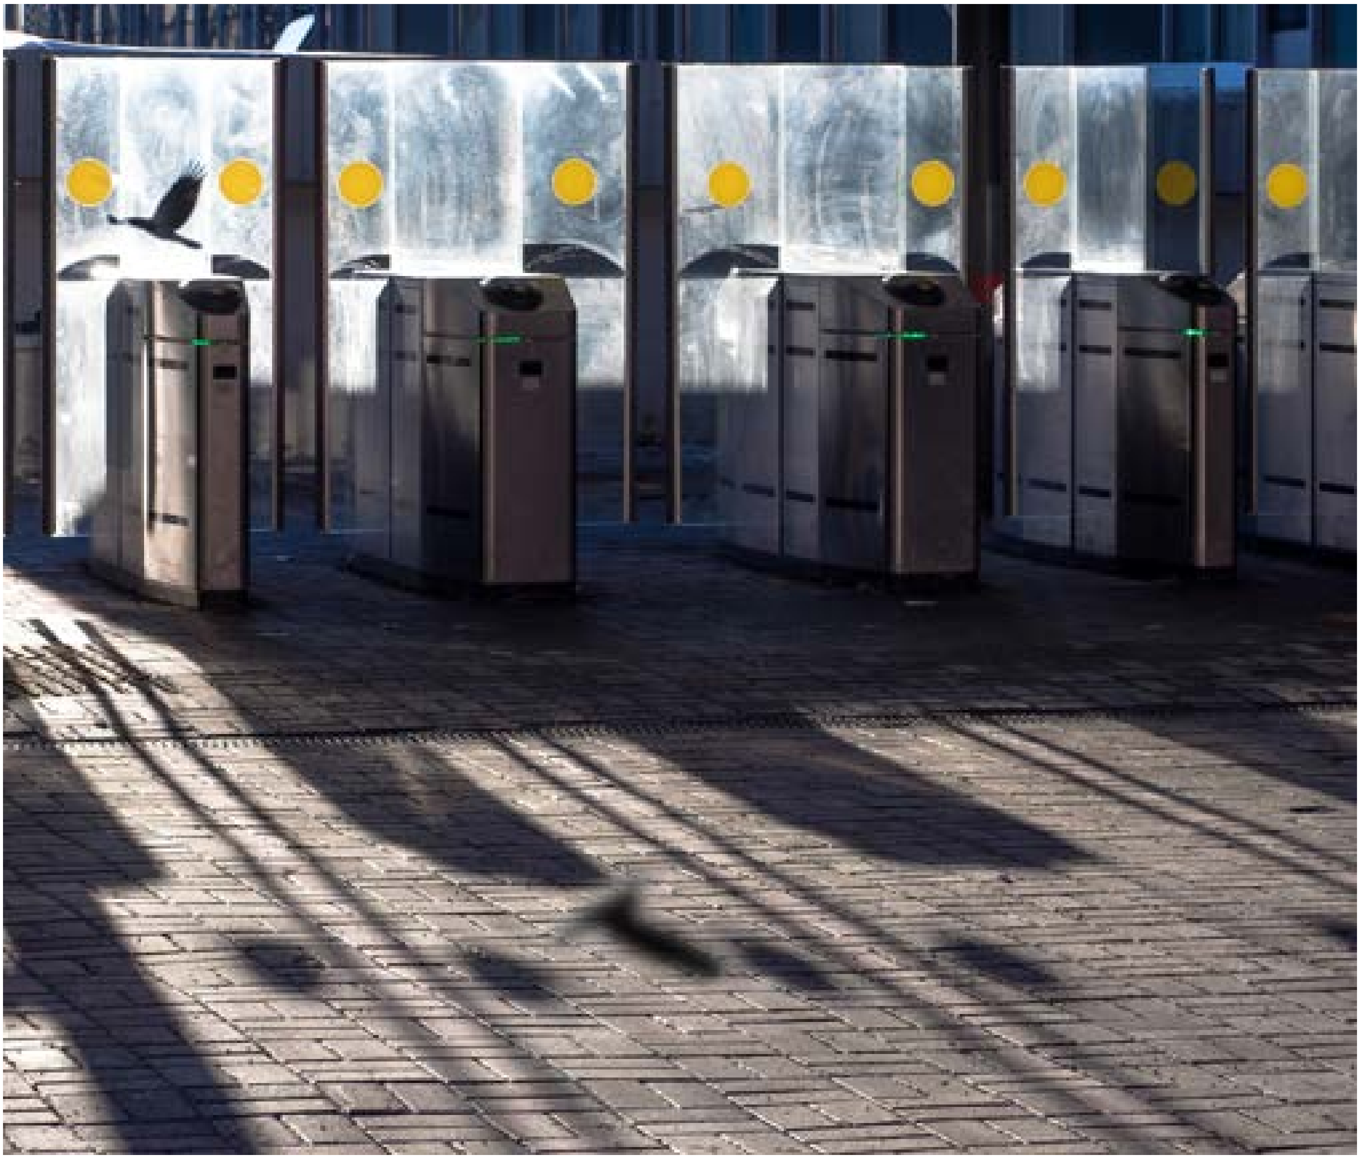

THE RISK OF MACHINE LEARNING BIAS (AND HOW TO PREVENT IT)

As promising as machine-learning technology is, it can also be susceptible to unintended biases that require careful planning to avoid

Chris DeBrusk

MANY COMPANIES ARE turning to machine learning to review vast amounts of data, from evaluating credit for loan applications, to scanning legal contracts for errors, to looking through employee communications with customers to identify bad conduct. New tools allow developers to build and deploy machine-learning engines more easily than ever: Amazon Web Services recently launched a "machine learning in a box" offering called SageMaker that non-engineers can leverage to build sophisticated machine-learning models, and Microsoft Azure's machine-learning platform, Machine Learning Studio, requires no coding skills.

But while machine-learning algorithms enable companies to realize new efficiencies, they are as susceptible as any system to the "garbage in, garbage out" syndrome. In the case of self-learning systems, the type of "garbage" is biased data. Left unchecked, feeding biased data to self-learning systems can lead to unintended and sometimes dangerous outcomes.

In 2016, for example, an attempt by Microsoft to converse with millennials using a chat bot plugged into Twitter famously created a racist machine that switched from tweeting that "humans are super cool" to praising Hitler and spewing out misogynistic remarks. This scary conclusion to a one-day experiment resulted from a very straightforward rule about machine learning – the models learn exactly what they are taught. Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), a machine-learning system that makes recommendations for criminal sentencing, is also proving imperfect at predicting which people are likely to reoffend because it was trained on incomplete data. Its training model includes race as an input parameter, but not more extensive data points such as past arrests. As a result, it has an inherent racial bias that is difficult to accept as either valid or just.

These are just two of many cases of machine-learning bias. Yet there are many more potential ways in which machines can be taught to do something immoral, unethical, or just plain wrong.

These examples serve to underscore why it is so important for managers to guard against the potential reputational and regulatory risks that can result from biased data, in addition to figuring out how and where machine-learning models should be deployed to begin with. Best practices are emerging that can help to prevent machine-learning bias. Below, we examine a few.

Consider bias when selecting training data. Machine-learning models are, at their core, predictive engines. Large data sets train machine-learning models to predict the future based on the past. Models can read masses of text and understand intent, where intent is known. They can learn to spot differences – between, for instance, a cat and a dog – by consuming millions of pieces of data, such as correctly labeled animal photos.

EXHIBIT 1: A TAXONOMY OF MACHINE-LEARNING TERMS

Navigating artificial intelligence and machine-learning concepts can sometimes be daunting. Below is a breakdown of some of the key terms.

ARTIFICIAL INTELLIGENCE (AI)

is a scientific field within Computer Science, focusing on the study of computer systems that can perform tasks and solve problems that require human intelligence

MACHINE LEARNING

is a field within AI that focuses on a particular class of algorithms that can learn from data without being explicitly programmed

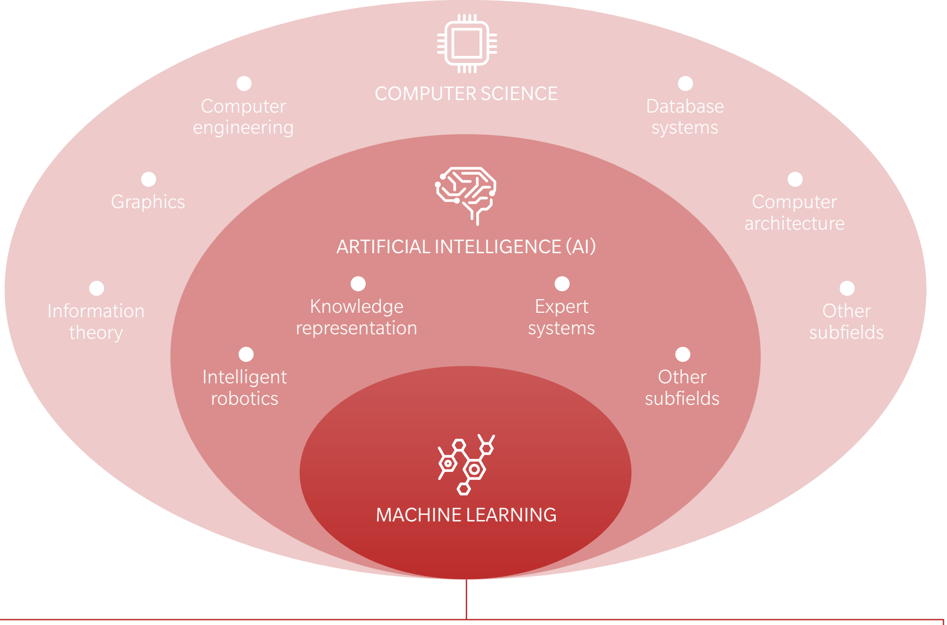

There are three main ways a machine can learn from data	**Supervised machine learning** Mapping inputs to labeled outputs	**Unsupervised machine learning** Finding patterns in unlabeled input data	**Reinforcement learning** Performing actions to maximize rewards
Each category of machine learning is effective in tackling particular kinds of tasks and problems	**Commonly used in classification and regression problems such as:** Natural language processing Image recognition Financial forecasting	**Commonly used in segmentation and clustering problems such as:** Pattern and trend recognition Customer segmentation Transaction monitoring	**Similar to supervised learning, but reward mechanism in place instead of labelled output:** Board and video games Robotics
There is a wide range of mathematical techniques that can be used to develop machine-learning models	Decision tree Random forest Gradient boosting Neural network Support vector machine Regularized regression	Clustering Dimensionality reduction Gaussian mixture model Principal component analysis Independent component analysis	Q-learning Temporal difference algorithm Neural network State-action-reward-state-action

Note: Not comprehensive, key elements shown
Source: Oliver Wyman analysis

The advantage of machine-learning models over traditional statistical models is their ability to quickly consume enormous numbers of records and thereby more accurately make predictions. But since machine-learning models predict exactly what they have been trained to predict, their forecasts are only as good as the data used for their training.

For example, a machine-learning model designed to predict the risk of business loan defaults may advise against extending credit to companies with strong cash flows and solid management teams if it draws a faulty connection – based on data from loan officers' past decisions – about loan defaults by businesses run by people of a certain race or in a particular zip code. A machine-learning model used to scan reams of résumés or applications to schools might mistakenly screen out female applicants if the historical data used to train it reflects past decisions that resulted in few women being hired or admitted to a college.

These types of biases are especially pervasive in data sets based on decisions made by a relatively small number of people. As a best practice, managers must always keep in mind that if humans are involved in decisions, bias always exists – and the smaller the group, the greater the chance that the bias is not overridden by others.

Root out bias. To address potential machine-learning bias, the first step is to honestly and openly question what preconceptions could currently exist in an organization's processes, and actively hunt for how those biases might manifest themselves in data. Since this can be a delicate issue, many organizations bring in outside experts to challenge their past and current practices.

Once potential biases are identified, companies can block them by eliminating problematic data or removing specific components of the input data set. Managers for a credit card company, for example, when considering how to address late payments or defaults, might initially build a model with data such as zip codes, type of car driven, or certain first names – without acknowledging that these data points can correlate with race or gender. But that data should be stripped, keeping only data directly relevant to whether or not customers will pay their bills, such as data on credit scores or employment and salary information. That way, companies can build a solid machine-learning model to predict likelihood of payment and determine which credit card customers should be offered more flexible payment plans and which should be referred to collection agencies.

A company can also expand the training data set with more information to counterweight potentially problematic data. Some companies, for example, have started to include social media data when evaluating the risk of a customer or client committing a financial crime. A machine-learning algorithm may flag a customer as high risk if he or she starts to post photos on social media from countries with potential terrorist or money-laundering connections. This conclusion can be tested and overridden, though, if a user's nationality, profession, or travel proclivities are included to allow for a native visiting their home country or a journalist or businessperson on a work trip.

Regardless of which approach is used, as a best practice, managers must not take data sets at face value. It is safe to assume that bias exists in all data. The question is how to identify it and remove it from the model.

It is safe to assume that bias exists in all data. The question is how to identify it and remove it from the model.

Counter bias in "dynamic" data sets. Another challenge for machine-learning models is to avoid bias where the data set is dynamic. Since machine-learning models are trained on events that have already happened, they cannot predict outcomes based on behavior that has not been statistically measured. For example, even though machine learning is extensively used in fraud detection, fraudsters can outmaneuver models by devising new ways to steal or escape detection. Employees can hide bad behavior from machine-learning tools used to identify bad conduct by using underhanded techniques like conversing in code.

To attempt to draw new conclusions from current information, some companies use more experimental, cognitive, or artificial intelligence techniques that model potential scenarios. For example, to outsmart money launderers, banks may conduct so-called war games with ex-prosecutors and investigators to discover how they would beat their system. That data is then used to handcraft a more up-to-date machine-learning algorithm.

But even in this situation, managers risk infusing bias into a model when they introduce new parameters. For example, social media data, such as pictures posted on Facebook and Twitter, is increasingly being used to drive predictive models. But a model that ingests this type of data might introduce irrelevant biases into its predictions, such as correlating people wearing blue shirts with improved creditworthiness.

To avoid doing so, managers must ensure that the new parameters are comprehensive and empirically tested – another best practice. Otherwise, those parameters might skew the model, especially in areas where data is poor. Insufficient data could impact, say, credit decisions for classes of borrowers to whom a bank has never lent to previously but plans to in the future.

Balance transparency against performance. One temptation with machine learning is to throw increasingly large amounts of data at a sophisticated training infrastructure and allow the machine to "figure it out." For example, public cloud companies have recently released comprehensive tools that use automated algorithms instead of an expert data scientist to train and determine the parameters intended to optimize machine-learning models.

While this is a powerful method for building complex predictive algorithms quickly and at lower cost, it also comes with the downside of limited visibility and the risk of the "machine running wild" and having an unconscious bias due to training data that is extraneous (like the blue shirt bias described above). The other challenge is that it is very difficult to explain how complex machine-learning models actually work, which is problematic in industries that are heavily regulated.

One of the potential options to address this risk is to take a staged approach to increasing the sophistication of the model and making a conscious decision to progress at every stage.

A good example is a process used by a major bank in building a model that attempted to predict whether a mortgage customer was about to refinance, with the goal of making a direct offer to that customer and ideally retaining their business. The bank started with a simple regression-based model that tested its ability to predict when customers would refinance. It then created a set of more sophisticated "challenger" models that used more advanced machine-learning techniques and were more precise. By confirming that the challenger models were more accurate than the base regression model, bank managers became comfortable that their more complex and opaque machine-learning approach was operating in line with expectations and not propagating unintended biases. The process also enabled them to verify that the machine-learning tool's balance between transparency and sophistication was in line with what is expected in the highly regulated financial services industry.

CAREFUL PLANNING IS A NECESSITY

It is tempting to assume that, once trained, a machine-learning model will continue to perform without oversight. In reality, the environment in which the model is operating is constantly changing, and managers need to periodically retrain models using new data sets.

Machine learning is one of the most exciting technical capabilities with real-world business value to have emerged over the past decade. When combined with big data technology and the massive computing capability available via the public cloud, machine learning promises to change how people interact with technology, and potentially entire industries. But as promising as machine-learning technology is, it requires careful planning to avoid unintended biases.

Creators of the machine-learning models that will drive the future must consider how bias might negatively impact the effectiveness of the decisions the machines make. Otherwise, managers risk undercutting machine learning's potentially positive benefits by building models with a biased "mind of their own."

Chris DeBrusk is a New York-based partner in Oliver Wyman's Financial Services and Digital practices.

This article first appeared in MIT Sloan Management Review on March 26, 2018.

THE RISKS AND BENEFITS OF USING AI TO DETECT CRIME

Companies are using it for everything from routine theft to insider trading

Lisa Quest, Anthony Charrie, and Subas Roy

COMPANIES ARE USING artificial intelligence (AI) to prevent and detect everything from routine employee theft to insider trading. Many banks and large corporations employ AI to detect and prevent fraud and money laundering. Social media companies use machine learning to block illicit content like child pornography. Businesses are constantly experimenting with new ways to use artificial intelligence for better risk management and faster, more responsive fraud detection, and even to predict and prevent crimes.

While today's basic technology is not necessarily revolutionary, the algorithms it uses and the results they can produce are. For decades, banks have been using transaction monitoring systems based on predefined binary rules that require the output to be manually checked. The success rate is generally low: On average, only 2 percent of the transactions flagged ultimately reflect a true crime or malicious intent. By contrast, today's machine-learning solutions use predictive rules that automatically recognize anomalies in data sets. These advanced algorithms can significantly reduce the number of false alerts by filtering out cases that were flagged incorrectly, while uncovering others missed using conventional rules.

Given the wealth of data available, and the rising expectations of customers and public authorities when it comes to protecting and managing that information, many companies have decided that AI is one of the only ways to keep up with increasingly sophisticated criminals. Today, for example, social media companies are expected to uncover and remove terrorist recruitment videos and messages almost instantly. In time, AI-powered crime-fighting tools could become a requirement for large businesses, in part because there will be no other way to rapidly detect and interpret patterns across billions of pieces of data.

But determining whether AI crime-fighting solutions are a good strategic fit for a company depends on whether the benefits outweigh the risks that accompany them. One such risk is that biased conclusions can be drawn from AI based on factors like ethnicity, gender, and age. Companies can also experience backlash from customers who worry that their data will be misused or exploited by even more data-intensive surveillance of their records, transactions, and communications – especially if those insights are shared with the government. Recently, for example, a European bank was forced to backtrack on its plan to ask customers for permission to monitor their social media accounts as part of its mortgage application process, after a public outcry over its "Big Brother" tactics.

Determining whether AI crime-fighting solutions are a good strategic fit for a company depends on whether the benefits outweigh the risks that accompany them

So how are leading-edge companies evaluating the benefits and risks of rapidly evolving AI crime-fighting and risk management? Below, we explain some of the steps they're taking.

EVALUATING THE STRATEGIC FIT

Before embarking on an AI risk management initiative, managers must first understand where machine learning is already making a big difference. Banks, for example, are halting financial crimes much more quickly and cheaply than they used to by using AI for automating processes and conducting multilayered "deep learning" analyses. Even though banks now file 20 times more suspicious activity reports linked to money laundering than they did in 2012, AI tools have permitted them to shrink the armies of people they employ to evaluate alerts for suspicious activities. That's because their false alerts have fallen by as much as half thanks to AI, and because many banks are now able to automate routine human legwork in document evaluation. For example, using artificial intelligence, Paypal has also cut its false alerts in half. And Royal Bank of Scotland prevented losses of over $9 million to customers after conducting a year-long pilot with Vocalink Analytics, a payments business, to use AI to scan small business transactions for fake invoices.

50 percent

How much banks have reduced their false suspicious activity alerts thanks to AI

AI tools also allow companies to surface suspicious patterns or relationships invisible even to experts. For instance, artificial neural networks can enable employees to predict the next moves of even unidentified criminals who have figured out ways around alert triggers in binary rules-based security systems. These artificial neural networks link millions of data points from seemingly unrelated databases, containing everything from social media posts to internet protocol addresses used on airport Wi-Fi networks to real estate holdings or tax returns, and identify patterns.

The next step in assessing the wisdom of launching an AI risk-management program is for companies to evaluate to what extent customers and government authorities will expect them to be ahead of the curve. Even if it does not become a regulatory or legal obligation, companies might find it advantageous to play a leading role in the use of advanced analytics so they can take part in setting industrywide standards. They can help ensure that industry participants, regulators, technology innovators, and customers are kept safe, without trampling on people's privacy and human rights.

Finally, managers need to determine whether it makes more sense to build or buy the type of AI solution that meets their needs. To reach this decision, managers should seek proven use cases in which AI is already achieving what they hope to accomplish. Then, they should decide which vendor to work with, based on their ability to handle machine learning that addresses the type of problem faced by the managers' company with the level of quality that will satisfy regulators. If a company is likely to face more complicated or rapidly evolving crimes, however, it might require more sophisticated and customized modeling. In that case, it is usually more beneficial to develop a machine-learning solution in-house. This is especially true if externally provided solutions are expensive, provide a low degree of certainty in their results, or cannot be adapted quickly enough to keep up with a rapidly evolving marketplace.

WHAT'S NEXT?

THE RISE OF PUBLIC-PRIVATE PARTNERSHIPS IN AI CRIME PREVENTION

Companies and law enforcement agencies have been experimenting separately with using artificial intelligence to improve their ability to detect and prevent crime. Now, they are increasingly working together – developing shared data platforms, reporting protocols, and feedback loops.

Public-private partnerships to fight crime will become increasingly common. Financial institutions, financial intelligence units, and law enforcement are starting to establish public-private partnerships to share data and use AI to detect crime in certain jurisdictions. For instance, in the United Kingdom, the National Crime Agency is working closely with UK Finance to use AI in order to better identify not only financial and economic crime but also improve their ability to use financial information to detect other types of crimes like human trafficking and counterfeiting. Authorities are also exploring ways to increase the exchange of information and intelligence between the public and private sectors.

As organized crime and criminals become more sophisticated and the amount of data available to the private sector continues to increase exponentially, companies and law enforcement will enter even more public-private partnerships to leverage their wealth of data and detect potential criminal activities even more efficiently.

WHERE AI WILL BE USED TO DETECT CRIMES IN THE FUTURE

Today, AI is most commonly used to detect crimes such as fraud and money laundering. But in the future, it will likely become commonly used in other industries as well. Below are three areas where we see AI being used to prevent:

1. Transportation of illegal goods. With AI, express delivery companies can assess the likelihood that parcels contain illegal goods, like narcotics, and report them to the relevant authorities.

2. Terrorist activities. Retailers and pharmacies could use sophisticated AI tools to identify customers who purchase unusual amounts of chemicals that could be used as precursors to terrorist activities.

3. Human trafficking. Shipping companies can use their data and AI capabilities to identify the containers that are most likely to be used for human trafficking and thus save lives.

ASSESSING AND MITIGATING INTERNAL RISKS

As managers examine how AI can assist them in identifying criminal activities, they should also consider how it fits in with their broader AI strategy. AI risk management and crime detection should not be conducted in isolation. Back-testing against simpler models can help banks limit the impact of potentially inexplicable conclusions drawn by artificial intelligence, especially if there is an unknown event for which the model has not been trained. For example, banks use artificial intelligence to monitor transactions and reduce the number of false alerts they receive on potential rogue transactions, such as money that's being laundered for criminal purposes. These are back-tested against simpler rules-based models to identify potential outliers. An AI model may, for example, mistakenly overlook a large money laundering transaction that would normally trigger an alert in a rules-based system if it determines, based on biased data, that large transactions made by customers who reside in wealthy neighborhoods do not merit as much attention. Using this approach enables companies to design more transparent machine-learning models, even if that means they operate within more explicit bounds.

Companies should also prepare to adjust their risk management processes to systematically counter self-learning, AI-powered models that can develop biases as they constantly recalibrate. Banks, for example, should frequently test and verify a random subset of their money laundering and fraud analyses to ensure that AI-driven systems are not unfairly penalizing any particular group.

Most of all, managers should assess whether their company's data analytics are sufficient to handle complex AI tools. If not, they need to develop data analytics capabilities in-house to reach a critical mass of automated processes and structured analytics.

UNDERSTANDING AND PREPARING FOR EXTERNAL RISKS

Increased use of AI tools for crime prevention could also cause external risks to cascade in unexpected ways. A company could lose its credibility with the public, regulators, and other stakeholders in myriad ways – for example, if there are false alerts that mistakenly identify people as "suspicious" or "criminal" due to a racial bias unintentionally built into the system. Or, at the other end of the spectrum, they could suffer reputational damage if they miss criminal activities, like drug trafficking conducted by their clients or funds channeled from sanctioned countries such as Iran. Criminals could resort to more extreme, and potentially violent, measures to outmaneuver AI. Customers could flee to less closely monitored entities outside of regulated industries. A moral hazard could even develop if employees become too reliant on AI crime-fighting tools to catch criminals for them. Employees could feasibly develop a false sense of comfort, and then stop regularly checking the outputs and miss obvious cases.

To prevent this from happening, companies need to create and test a variety of scenarios of cascading events resulting from AI-driven tools used to track criminal activities. To outsmart money launderers, for example, banks should conduct "war games" with ex-prosecutors and investigators to discover how they would beat their system.

With results produced through scenario analysis, managers can then help senior executives and board members decide how comfortable they are with using AI crime-fighting. They can also develop crisis management playbooks containing internal and external communication strategies so they can react swiftly when things (inevitably) go wrong.

By using AI, companies can identify areas of potential crimes such as fraud, money laundering, and terrorist financing – in addition to more mundane crimes such as employee theft, cyber fraud, and fake invoices – to help public agencies with prosecuting these offenses much more effectively and efficiently. But with these benefits come risks that should be openly, honestly, and transparently assessed to determine whether using AI in this way is a strategic fit. It will not be easy. But clear communication with regulators and customers will allow companies to rise to the challenge when things go wrong. AI will eventually have a hugely positive impact on reducing crime in the world – as long as it is managed well.

Lisa Quest is a London-based partner in Oliver Wyman's Public Policy and Organizational Effectiveness practices.
Anthony Charrie is a principal in Oliver Wyman's Public Policy practice in Europe.
Subas Roy is a London-based partner in Oliver Wyman's Digital, Technology, Operations & Analytics practice.

This article first appeared in Harvard Business Review on August 9, 2018.

TOMORROW’S FACTORIES WILL NEED BETTER PROCESSES, NOT JUST BETTER ROBOTS

Many of automation’s gains have already been realized

Ron Harbour and Jim Schmidt

WHEN PEOPLE THINK of the automotive Factory of the Future, the first word that comes to mind is automation. They think of the "lights out" factory that General Motors Chief Executive Roger Smith fantasized about in 1982 and Elon Musk talks about building today – plants so dominated by robots and machines that they don't need lights to work.

There's no doubt that the auto industry will continue to vigorously pursue automation solutions to lower the cost of producing cars. But the reality is that any major leap forward on cost and efficiency will no longer be possible through automation alone, since most of the tasks that can be automated in an automotive factory have already been tackled. (See Exhibit 1.)

FACTORY OF THE FUTURE REQUIRES NEW PROCESSES

When a real Factory of the Future finally arrives, it will not look different because we have automated the processes we use today. It will look different because we will have invented entirely new processes and designs for building cars requiring entirely new manufacturing techniques.

Take the paint shop. Today, in most mature markets, it's more than 90 percent automated, yet it is still one of the most expensive and space-intensive sections of the factory. Robots, instead of humans, perform most tasks – applying protective corrosion coats, sealant, primer, basecoat, and clear coat to achieve the highly polished finishes we like on our cars – but the process itself is not that different than what it was 30 years ago. For instance, in the BMW plant in Spartanburg, South Carolina, processing a car through the paint shop is a 12-hour task, involving more than 100 robots, and requiring a vehicle in the paint assembly line to travel four miles within the factory before the process is complete.

EXHIBIT 1: WHAT CAN BE AUTOMATED HAS BEEN AUTOMATED, ASSEMBLY LINES NEED THE HUMAN TOUCH

Automated workstations in automotive factories

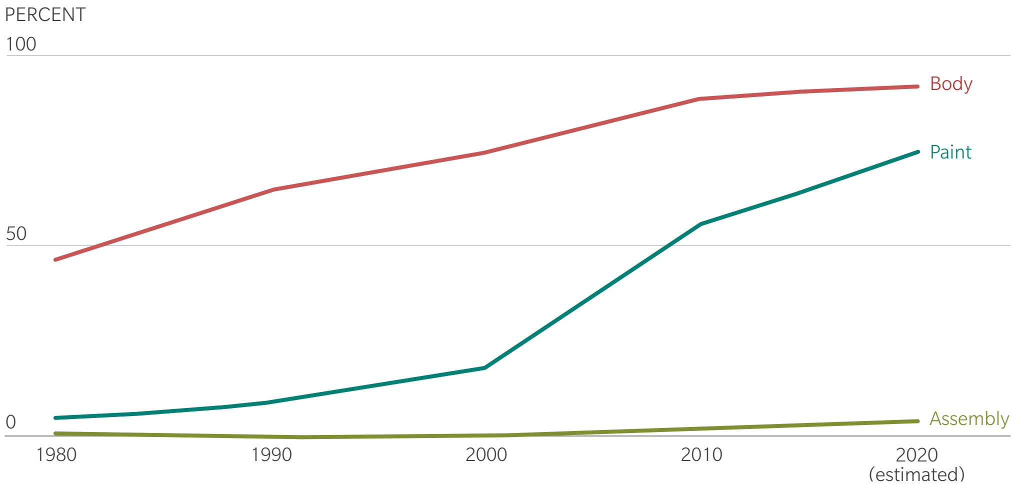

Note: Typical automated workstations in assembly: body/chassis marriage; urethane application and installation of windshield glass
Source: Oliver Wyman analysis

Clearly, there has to be a better way to paint a car, but to make that operation more efficient and take cost out will require the development of a new process. Perhaps it will be the experimental approach of applying a single film over the car and then baking it on, like in a pottery kiln – currently being tested in automotive research labs. Or 3-D printing of the entire car body in the color a customer orders, completely eliminating the need for a traditional paint shop and body shop. Whatever it is, it will have to be more than adding a few more robots into the mix to make a significant difference in the cost of producing an auto.

Any major leap forward on cost and efficiency will no longer be possible through automation alone

Today, two-thirds of automotive workers – the human ones – are in the general assembly section. Automating this section has proved more difficult because the customization and complexity of today's autos requires the flexibility humans provide. Most factories produce several models of cars simultaneously, and the mix of those models often changes, depending on demand. It would be expensive, if even possible, to reprogram robots and machines to be able to accommodate daily changes in factory production schedules.

There are also some tasks on the assembly line for which humans are better suited, such as handling all of the intricacies of installing and connecting a car's wire harnesses – the nervous system of a vehicle. With a future market expected to consist of electric and autonomous vehicles, the electrical systems will need to transmit more data faster and unfailingly, compared to today's car. The consequence to the assembly plant: more wires and connectors leading to longer, heavier wire harnesses. For this operation to be automated would again require a new process – perhaps going wireless, with the electrical systems operating via electronic modules or connecting via the cloud.

A new process will also need to be developed to assemble electric vehicles since they involve the relatively uncomplicated installation of the battery pack and an electric motor. Simpler tasks may lend themselves better to robots, but several steps on the line will also be bypassed. The leap forward will be accomplished through the development of a new process – in this case, electrifying the auto – not automating an old one.

COBOTS AS VALUABLE AND VERSATILE HELPERS

New collaborative robots, or cobots, are also adding a new twist: Instead of threatening the survival of humans on the assembly line by replacing them, cobots enhance their human workers' native abilities. Ranging in size from two- to four-feet high, these automated assistants work with humans to perform tasks that perhaps are slightly dangerous or repetitive, or that require a special agility to work in tight or hard to reach places, such as working underneath autos. For instance, Renault has deployed cobots in a few plants to help build the powertrain – torquing bolts to a certain tolerance, a task that can be tedious for humans to do consistently and efficiently.

¿COBOTS?

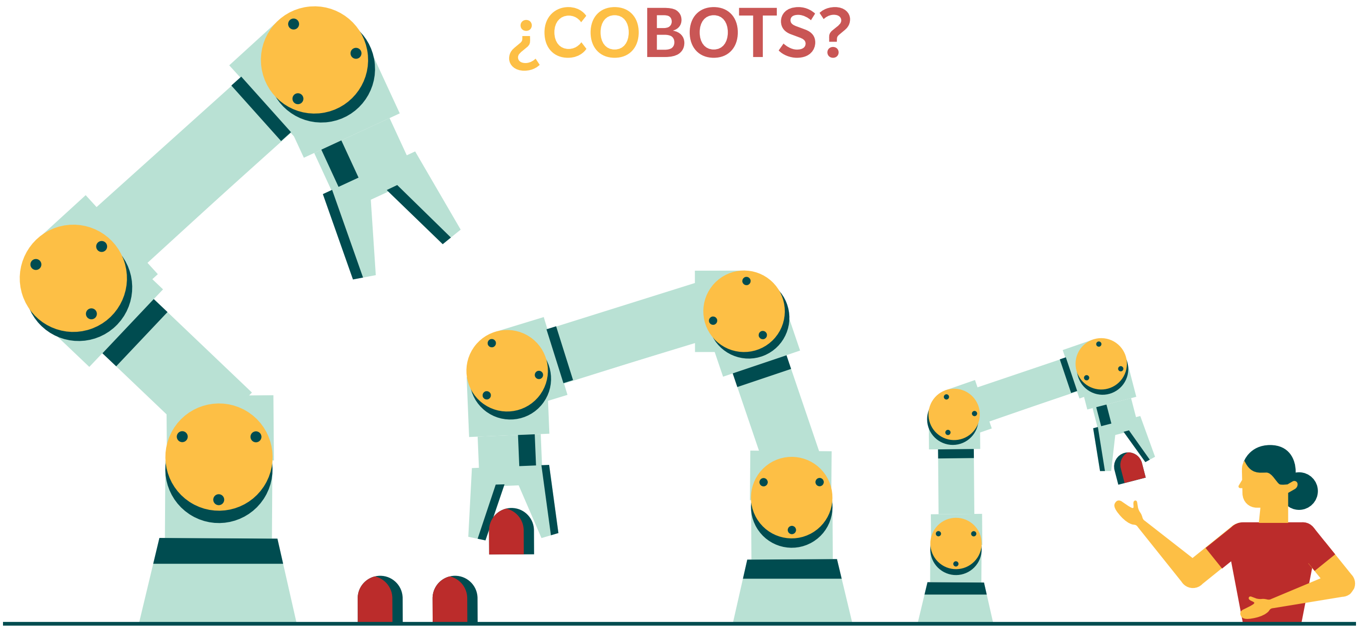

What are cobots?

Cobot is short for collaborative robot, meaning these automated helpers are specifically designed to assist humans, rather than simply replace them. Invented in 1996 by professors at Northwestern University as part of a research initiative with General Motors, cobots are between two- and four-feet tall and often are in the form of an arm. They tend to be easy to program, which means they can be operated by most workers and can smoothly transition between various types of factory jobs.

What do they do?

Cobots were invented specifically to help reduce the risk of injury for assembly-line workers and therefore are often helping them with tasks that require heavy lifting, squeezing into tight places or other ergonomic challenges. Cobots are also used for repetitive tasks that can lead to mistakes or inconsistencies when performed by humans because of their tedious nature. For instance, cobots work with superheated materials such as metal or glue or stack parts off a conveyor belt all day. They also take on the tedious task of tending precision computerized numerical control (CNC) machining.

How do they work with humans ?

Whereas large industrial robots often must be literally caged in to protect human workers, cobots are designed to operate safely alongside them. Among other safety features such as cameras, they include collision-detection sensors that alert humans and the cobot that they are getting close and then stop the cobot to prevent physical contact. Researchers are even working on ways to create real communication between cobot and worker and other enhancements that make cobots act more like humans. For instance, some cobots now have faces and look in the direction they are about to reach to give human workers a head-up.

Making these small helpers attractive to companies, cobots can be relatively inexpensive, often costing under $50,000 each. They are simple to reprogram – workers on the assembly line can often handle the reprogramming on their own. This allows them to be retasked quickly, again adding to their value and versatility.

Unlike much of the current robotic automation that must be kept fenced in, with safety signs warning employees to keep their distance, cobots perform tasks in factories without hurting humans, as they are programmed to stop when there is an object in front of them. With their swing arms, they can retrieve certain small parts from bins for their human partners.

Another example of automation that enhances humans' native abilities is the exoskeleton. Workers wear these cyborg-esque contraptions to make them strong enough to lift heavy truck tires or ease the stress on their bodies when performing repetitive overhead assembly tasks. This wearable automation becomes particularly important as the average age of production workers rises above 40, as it has in many industrialized economies, such as the United States, Western Europe, and Japan.

Roger Smith's dream of a "lights-out" factory has only been realized in a very few operations – robots building robots, for instance – and not in the automotive world. But there are other roads to the automotive Factory of the Future that will likely be paved with human invention, and while robots and automation will be part of the picture, the lights will still be on.

Ron Harbour is a Detroit-based senior vice president in Oliver Wyman's Automotive and Manufacturing Industries practice and co-author and developer of The Harbour Report®.
Jim Schmidt is is a Detroit-based vice president in Oliver Wyman's Automotive and Manufacturing Industries practice and manages The Harbour Report®

This article first appeared in Harvard Business Review on May 11, 2018.

REDEFINING BUSINESS MODELS

CHINA'S RETAIL REVOLUTION IS HEADED WEST

China's retail giants are redefining shopping at home – and may soon do so abroad

Pedro Yip and Richard McKenzie

CHINA'S TWO BIGGEST online retail groups, Alibaba Group and JD.com, are building retail empires unlike anything yet seen in the US or Europe. After buying stakes or forming alliances with traditional retailers, ranging from convenience stores to big-box hypermarkets over the past few years, these two giants now dominate China's mobile payments. And they are rapidly bringing to scale new forms of retail, where customers walk around a store using their smartphones both to pull up information and to pay.

$15 billion

How much Alibaba plans to invest in global research hubs over the next three years

Though Amazon is widely seen as having set the pace in digital shopping over the past two decades, a second center of retail innovation has simultaneously been developing in Shanghai. The clearest sign that Chinese retailers are writing the rules for 21st century shopping is in the area of food purchases. Just 3 percent of Americans buy groceries online, compared with nearly 10 percent of Chinese – and half of those Chinese buy groceries exclusively online. With a share of around one-tenth of China's grocery retail sales, JD.com and Alibaba are transforming at unprecedented speed into information-fueled networks that gather data from purchases made online or in supermarkets, enabling them to forecast product demand and efficiently coordinate supply and delivery. They also know which websites customers like to visit, the apps they use, and who they follow on social media.

The extraordinary growth has its origins in China's relatively underdeveloped traditional supermarkets, which were falling behind the demands of the country's burgeoning middle class. China's explosion in online access – from just 11 percent of the population in 2006 to over half in 2017 – made e-commerce a natural alternative. Home delivery is also relatively low cost, thanks to the density of the cities where wealthier Chinese tend to live.

Now the groups intend to use their new financial and technological might to take the fight to the US and Europe. Alibaba plans to invest $15 billion in global research hubs over the next three years and has been looking at potential sites for European dispatch hubs. JD.com, which has a strategic alliance with its shareholder Walmart, wants half its profits to come from outside of China in 10 years' time.

EXHIBIT 1: THE EXTRAORDINARY GROWTH OF E-COMMERCE IN CHINA

The growth in gross merchandise value (GMV) over the past few years has made China the largest e-commerce market in the world

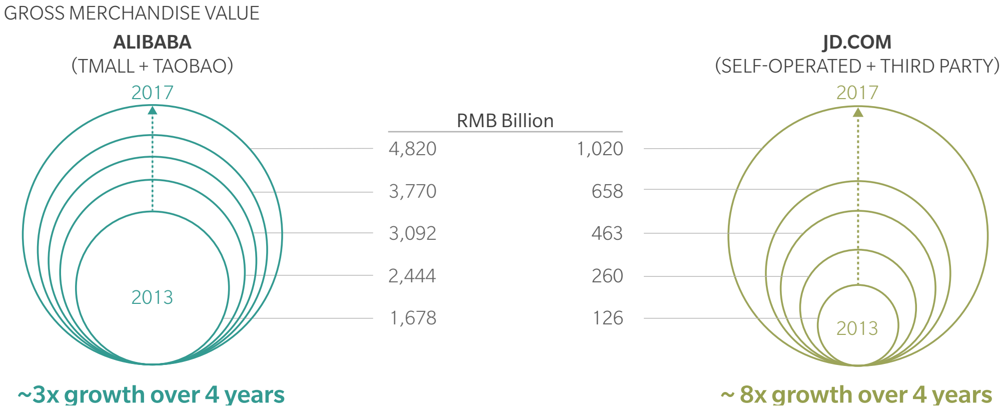

Source: Oliver Wyman analysis

Copyright © Oliver Wyman

Here are a few ways in which China's retail giants are redefining shopping at home – and may potentially do so abroad, and soon.

OMNICHANNEL GROCERIES

Alibaba and JD.com are pioneering new shopping modes known as O2O – "online-to-offline" that go beyond the parallel physical and online services offered by most major supermarkets worldwide.

While Amazon's new Amazon Go concept store is a US version of O2O, Chinese retailers are scaling up faster. At Alibaba's Hema stores, which numbered 47 by April this year, QR matrix barcodes pull up product information on smartphones – which are also used to select and pay for goods. If shoppers do not want to carry their purchases home, they can request delivery, which is guaranteed within half an hour. But customers can also see and taste fresh food – and a chef will even fry up seafood on the spot.

Some of the stores feature futuristic information devices. 7Fresh, launched by rival JD.com, offers automatic payment using facial recognition technology as an alternative to smartphone purchases. It can then convert the data into a chocolate-powder print of the shopper's face on a cappuccino's frothy milk. The futuristic experience includes "magic mirrors" that sense when a customer has picked up a vegetable and then display information such as nutritional content and where it was grown. Robot shopping carts follow shoppers through a store, avoiding obstacles.

And they are moving beyond food. For China's Singles' Day shopping festival in November, Alibaba set up makeup boutiques that let shoppers try out lipstick virtually and O2O furniture stores that act as showrooms and have QR codes for online ordering.

PHYSICAL STORES

Achieving scale enables JD.com and Alibaba to optimize their supply chains and delivery services, and to sweep up data from millions of daily transactions. So armed, they are expanding into traditional, brick-and-mortar retail, and supercharging it with digital logistics and nationally recognized brands. Alibaba has recently begun installing products from its Tmall Supermarket online grocery service in supermarket chain RT-Mart and helps provide home delivery for shoppers who live nearby. JD.com is providing Walmart's Chinese unit with an online service platform and also offers one-hour delivery for selected Walmart items.

Since early 2017, both alliances have also been converting some of China's more-than 7 million family-run stores into convenience store franchises under the Tmall or JD.com umbrellas. The convenience stores get well-known brands on their shelves and mobile payment systems for easy checkout. Their managers can replenish stock through a smartphone app. And it arrives much faster than through the old, multilayered distribution networks that corner shops relied upon up to now.

EXHIBIT 2: THE ECOSYSTEM OF ALIBABA COVERS EVERY PART OF CUSTOMERS' DAILY LIVES

Why Alibaba Group knows a lot more than just what its customers purchased

ONLINE RETAIL

WHAT YOU BUY
Taobao
Tmall

WHAT YOU EAT
Ele.me

WHAT PARCELS YOU RECEIVE
Cainiao
YTO Express Group

OTHER

WHAT DATA YOU STORE AND WHERE
Alibaba Cloud

HOW HEALTHY YOU ARE
Ali Health

HOW YOU LEARN
Taobao Education

OFFLINE RETAIL

HOW AND WHERE YOU BUY
Bailian Group
Freshippo
Intime Retail
Lianhua Supermarket
LST
RT-Mart
Sanjiang Shopping Club

FINANCE

WHAT YOU INVEST IN
Tianhong Asset Management
ZhongAn Insurance

HOW YOU FINANCE
Alipay
Ant Check Later
Ant Financial
Ant Micro Loan
MYbank

HOW YOU MONITOR YOUR CREDIT
Sesame Credit

SOCIAL

HOW YOU COMMUNICATE
Momo
Sina Weibo

ENTERTAINMENT

HOW YOU ARE ENTERTAINED
South China Morning Post
Xiami Music
Youku

TRAVEL

WHERE YOU GO
Amap
DiDi
ofo

MOBILE

HOW YOU BROWSE
SM.CN
UC Browser

WHICH APPS YOU USE
9 Game
Wandoujia

Source: Oliver Wyman analysis

MOBILE PAYMENTS

A growing proportion of the transactions in these stores are carried out by mobile phone – and the two empires dominate here too. Together, they process 92 percent of China's mobile payments: Alibaba's system is called Alipay, while rival JD.com is in an alliance with Tencent, which owns WeChat Pay – part of WeChat, China's largest social app. Mobile phone payments are already used for 35 percent of supermarket purchases in China. Even when people buy from an independent shop – for a takeout meal, say – the transactions are a valuable source of information on what consumers are buying, where and when. The e-commerce giants can use this knowledge to develop increasingly attractive products and formats. This expertise will then persuade more retailers to sign up with them.

Chinese retailers are also writing the rules for 21st century shopping.

THE OUTLOOK

Alibaba and JD.com's retail empires are now preparing to expand globally. Alibaba, which has been selling online in the US and Europe for several years, is planning to set up European dispatch centers. It also wants to add European and Asian data centers to those it has in the US in order to support everything from online sales to its cloud computing business. In Asia, Alibaba announced it would double its investment in Southeast Asian e-commerce firm Lazada Group to $4 billion and install one of its executives to run the business.

JD.com will take on Alibaba and Amazon in the US and Europe this year by starting to sell online, leaning on Walmart for initial logistics support in the US. In June, it said Google would pay $550 million for a 1 percent stake in the company, in a deal that will make JD.com products available in the US and Europe on Google Shopping. In 2017, JD.com announced a $500-million e-commerce and financial technology joint venture with Thai retailer Central Group. It has also invested in Tiki, an e-commerce platform in Vietnam, and launched its own e-commerce platform in Indonesia.

It is notoriously tricky for retailers to expand outside their domestic markets, and the Chinese shopping empires are rooted in their country's specific retail evolution. But Alibaba and JD.com may have an easier time, since they have already developed significant technology and scale in China – including the means to make online groceries work.

Moreover, Alibaba and JD.com have the financial firepower to buy a company to fill gaps in their operations. Alibaba has a market capitalization of close to $500 billion, while JD.com itself is worth around $60 billion. They have already assembled an impressive number of alliances in China through a wave of acquisitions and partnerships over the past four years; Alibaba has made strategic investments totaling $21 billion in retail alone in just the past two years. The next phase of their big acquisition spree might be outside China.

Pedro Yip and **Richard McKenzie** are China-based partners in Oliver Wyman's Retail & Consumer Goods practice.

This article first appeared on the World Economic Agenda blog on September 17, 2018.

THE RISKS OF THE LIBOR SWITCH

It would be unwise for banks to wait and see how the new benchmark works out – time to get planning

Adam Schneider and Serge Gwynne

THE LONDON INTERBANK Offered Rate (LIBOR) is the reference interest rate for financial products valued at more than a staggering $240 trillion. No wonder it has been called the "world's most important number." Nevertheless, it may soon be a non-existent number.

After the LIBOR-fixing scandal, and the resulting $10 billion in fines, both regulators and banks want reference rates that are based on observed market prices rather than "expert judgement." However, unsecured interbank lending – the market LIBOR is supposed to represent – has decreased dramatically over the past 10 years, making objective reporting of rates practically impossible. The UK's Financial Conduct Authority announced in 2017 that, from 2021, it will no longer "compel or persuade" panel banks to submit the information used to create LIBOR.

Of course, LIBOR might not end after 2021. Future regulatory interventions aside, nothing stops banks from continuing to submit the underlying data, nor Intercontinental Exchange (ICE), the LIBOR administrator, from continuing to publish LIBOR. Indeed, several influential voices are calling for this. And ICE was recently authorized under new European benchmark regulation. Nevertheless, dwindling liquidity in the underlying market means LIBOR's survival is far from certain.

$240 trillion

The value of the financial products that use LIBOR as a reference rate

ALTERNATIVE ROUTES

Working groups have been developing alternative rates for the four main LIBOR currencies: the pound, the dollar, the yen, and the Swiss franc. In April this year, the Bank of England took ownership of the reformed Sterling Overnight Index Average (Sonia), and the New York Federal Reserve Bank began publishing the Secured Overnight Financing Rate (SOFR) – a new US dollar reference rate.

Euro reference rates are also in the process of being reformed or eliminated. The method used to derive Euribor (the European market equivalent to LIBOR) must be revised to meet the requirements of the European Benchmark Regulation by January 1, 2020. And Eonia (the main euro overnight rate and once a possible alternative to Euribor) is unlikely to be available after this date, because it cannot be revised to comply with the new benchmark regulation.

A transition to these new reference rates will not be a simple matter of replacing old rates with new ones. The new rates are structurally different from LIBOR. For example, Sonia is an overnight rate only (for now), while the most frequently used LIBOR rates are for one-, three-, and six-month tenors. Overnight Sonia rates could be compounded over a longer period, but this would not capture the term structure or credit spread embedded within LIBOR. We calculate that, on average, a compounded Sonia rate would have been 30 basis points lower than three-month pound LIBOR over the past 10 years and nearly 400 basis points lower during the financial crisis when bank credit spreads blew out.

Simply substituting Sonia for LIBOR in products or contracts that reference it would therefore radically alter the behavior of expected cash flows as interest rates change. The transition away from LIBOR will require products to be analyzed to determine which rate should be used – for example, overnight Sonia, term Sonia (if and when developed), or a new rate to be developed for specific product needs.

Even if this challenge can be met without excessive disruption, another will remain. LIBOR is likely to become unavailable before many of the contracts using it mature. "Fallback" provisions in each contract will then determine what happens. But these provisions were intended to cover a temporary unavailability of LIBOR rather than its permanent elimination, and they often change the economics of the product – for example, effectively converting floating-rate products into fixed-rate products. One of the counterparties to a LIBOR contract may suffer material losses while the other receives windfall gains.

The expense and risk of revising the terms in extant LIBOR contracts will vary with the product and type of counterparty. Bespoke contracts with corporate customers will usually require individual renegotiation. By contrast, the standardization of many derivative products and sophistication of the counterparties involved may mean that revising them will be relatively straightforward, which is fortunate, given the extraordinary gross exposures.

Retail products that reference LIBOR will also need a standardized rather than individually negotiated revision. But here the risks of getting it wrong are considerable. Banks currently selling LIBOR products that extend beyond 2021 are becoming increasingly concerned about the reputational and conduct risk they are taking on, given the uncertainty if LIBOR were no longer available. Buyside investors in LIBOR products also face conduct risks due to fiduciary duties towards their clients.

EXHIBIT 1: RETAIL MORTGAGES FOR A MODEL BANK IN THE US

More than 15 million retail customers globally hold products that reference LIBOR. Here's how a US bank's business, and its customers, would be impacted by LIBOR transition:

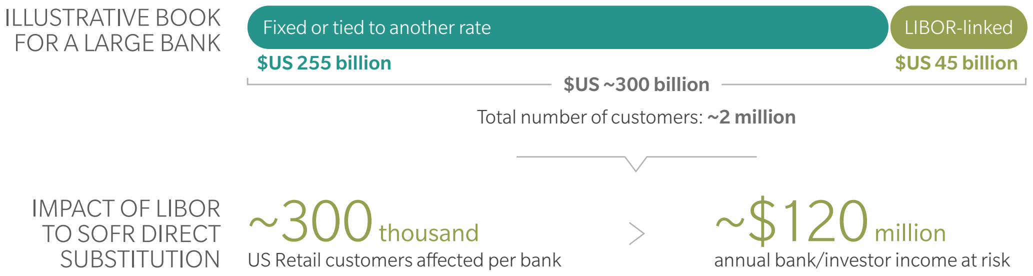
ILLUSTRATIVE BOOK FOR A LARGE BANK
Fixed or tied to another rate
LIBOR-linked
$US 255 billion
$US 45 billion
$US ~300 billion
Total number of customers: ~2 million
IMPACT OF LIBOR TO SOFR DIRECT SUBSTITUTION
~300 thousand
US Retail customers affected per bank
~$120 million
annual bank/investor income at risk

Transition challenges include the potential need for customer and investor approval, and issuing LIBOR-linked mortgages now knowing that LIBOR may transition.

Note: Based on a potential delta of 25 bps between SOFR/LIBOR (2-year average delta between 3M LIBOR and computed 3M forward-looking EFFR)
Source: Oliver Wyman analysis

RETHINK PRODUCT ECONOMICS

The sooner banks begin to prepare for the changes that are likely to come, the less disruption they will face, and the less risk. The first task is to understand the extent of the exposure by compiling an inventory of products, contracts, and processes (such as transfer pricing or asset valuation) that rely on LIBOR, Euribor, and Eonia. Banks can then begin to design products that reference alternative rates, working with market participants to build liquidity.

As for the "back book," the goal should be to minimize the cost and risk of conversion mentioned above. Banks should advocate for a standardized, transparent, and indisputably fair approach to the revision of contracts. Ideally, banks will find a way of continuing to use LIBOR for many legacy contracts until they mature.

Given the vast number of products, contracts, and processes LIBOR is baked into, moving to alternative rates will inevitably be a burdensome and lengthy process. But by engaging with regulators and industry bodies, and by beginning work early, banks can greatly reduce the potential for costly turmoil.

Indeed, they can turn a burden into an opportunity. Moving from LIBOR to a new basis for pricing products provides an occasion to rethink product economics and to tailor pricing to the needs of counterparties, including retail customers. If change is to be forced on banks, they should make the most of it.

Adam Schneider is a partner in Oliver Wyman's Digital and Banking practices in the Americas.
Serge Gwynne is a London-based partner in Oliver Wyman's Corporate and Institutional Banking practice.

This article first appeared on Financial News on June 27, 2018.

OIL'S BOOM-OR-BUST CYCLE MAY BE OVER

Recent price swings highlight a new era of uncertainty

Saji Sam, Juan Trebino, and Bob Orr

IN NOVEMBER OF 2017, United States' crude oil production exceeded 10 million barrels per day for the first time since 1970, according to the US Energy Information Administration (EIA). Analysts have predicted that US could become the world's largest oil producer in 2018, surpassing Saudi Arabia and Russia. How did we get here, and what does it mean for the industry?

US shale oil and gas producers have been ramping up production to take advantage of rising crude oil prices – prices that had been rising in the wake of a deal between the Organization of the Petroleum Exporting Countries (OPEC), Russia, and other non-OPEC producers to reduce oil output. That deal sent the price of Brent crude oil to above $70 a barrel in January, after the industry had suffered through $54 per barrel oil on average in 2017.

But with oil producers in North America expanding output, prices are likely to remain volatile. Unlike national oil companies and oil majors that typically take five to 10 years to develop conventional oil reserves, these independent and "unconventional" players have improved their drilling and fracturing technology to the point where they can respond within months to temporary spikes or dips in the market.

A NEW ERA OF UNCERTAINTY

The recent price swings highlight a new era of uncertainty gripping the world's energy markets. As global oil producers work at cross-purposes, the industry's traditional boom-bust cycle is being replaced by faster, shallower price rotations based on changes in production. It makes price movements less extreme but also more difficult to predict. The constantly fluctuating number of barrels of crude available from nimble shale operations is a primary driver. But other factors are the long-term impact of increased fuel efficiency and the fits and starts of the global transition away from fossil fuels on world demand. The news is all good for customers, but it makes planning for industry players much more difficult.

This unpredictability may only intensify as oil markets continue to adjust to shifting realities. Even more potentially destabilizing for major players, the expected surge in the US oil supply may be enough on its own to meet all of this year's growth in global oil demand. After being one of the world's largest net importers for decades, the US – while still a net importer of oil – is now selling millions of barrels of oil to China, Britain, Mexico, and India, a new reality made possible when restrictions on crude oil exports were lifted in 2015.

The soaring US output comes from fracking operations that have cut costs dramatically since slumping prices in 2014 forced dozens of companies into bankruptcy. These increasingly efficient survivors now represent half of U.S. oil production, up from a mere 10 percent just seven years ago in 2011. In fact, 2018 may mark the first year shale producers will be able to fund future expansions of drilling programs through their own cash flow.

EXHIBIT 1: THE NEW RULES FOR OIL

As North American output expands, the industry's traditional boom-bust cycle is being replaced by faster, shallower price rotations

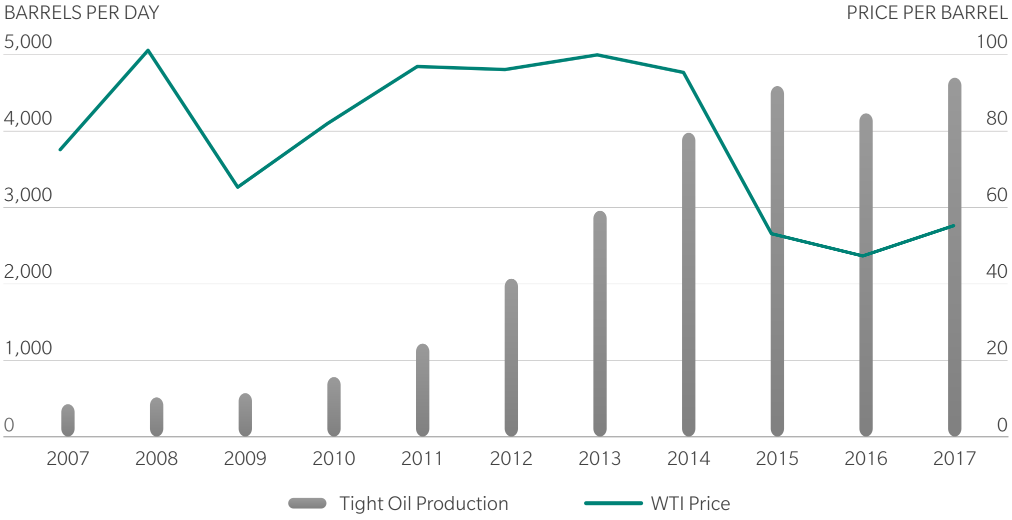

Source: US Energy Information Administration

While major oil companies plan to dramatically increase shale production in the Permian Basin in Texas and New Mexico, US shale production alone is unlikely to be enough to satisfy the world's growing oil needs – especially when oil reserves in shale may only get us another 10 years of oil and not necessarily 50. Oil companies will need to develop both new conventional and unconventional crude oil resources to keep up with current demand for roughly one million more barrels of oil every year, in addition to replacing the approximately four million barrels lost annually as reservoirs are naturally depleted. In total, we estimate that the oil and gas industry will have to replace about 40 percent of today's oil production over the next seven to nine years.

DIFFICULT DECISIONS

That means difficult decisions lie ahead for independent shale producers, national oil companies, and the major integrated companies. While they can start to tap into the global reserves of shale oil, which exist literally everywhere, developing the reserves in most places from China to Argentina will require a significant investment to build the shale ecosystem and supply chains needed, in addition to the infrastructure to gather, treat, transport, and store the crude oil. Or they can develop conventional reservoirs where it will require long-term investments in new technologies to bring the cycle times and costs more in line with those of nimble shale producers. Most major producers with large balance sheets will likely hedge their bets and attempt both.

The oil and gas industry will have to replace about 40 percent of today's oil production over the next seven to nine years

Generally cheaper oil is certain to have at least one short-term impact: It will compete with and potentially slow down the world's expected transition to renewable, clean energy. Transportation accounts for the majority of the world's oil demand, and as long as oil prices stay way below their 2008 peak crude oil price of $145 per barrel, there's less economic urgency to switch to electric vehicles and hybrids, even in China and Europe where there has been governmental support to move away from internal combustion cars. Electric vehicles will only account for 7 percent of the cars on the road over the next 12 years, Morgan Stanley estimated when oil prices were relatively low in May in a report titled, "*One Billion BEVS by 2050*."

Long term, however, oil demand to operate cars is likely to decline as fuel efficiency for all manner of transportation increases, car ownership continues to fall, and electric and autonomous vehicles become more popular. By 2050, more than half of the world's passenger cars are likely to be electric vehicles, according to the Morgan Stanley study. With the right combination of technological advances, cost reductions, and integration with renewable energy and storage, the tipping point for electric vehicle adoption could potentially be much earlier. These trends will require oil producers to shift their focus away from transportation and diversify towards innovative petrochemical products to capture market share in diverse end uses, such as clothing and construction materials.

EXHIBIT 2: THE WORLD'S SHALE RESOURCES BY THE BARRELS

The United States is the world's leading shale producer. But more than 1 trillion barrels of unconventional oil exists around the world. Here's how much tight oil could be recovered in 39 countries:

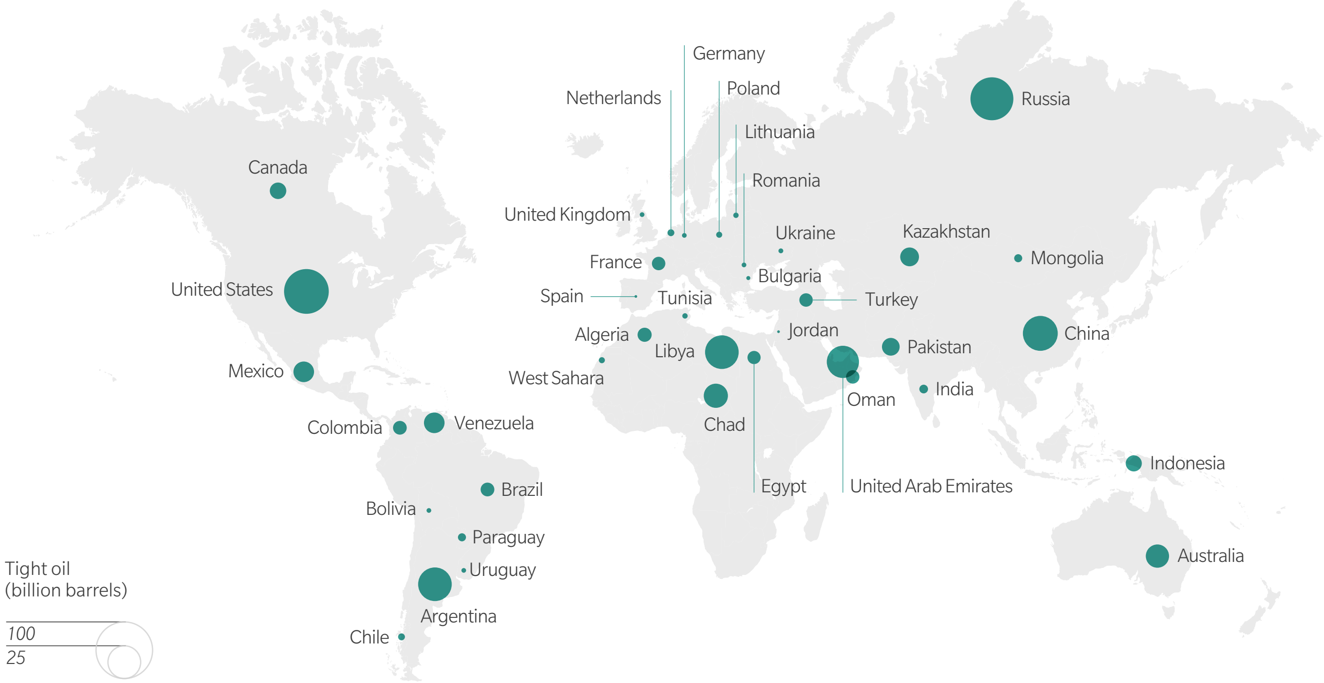

Source: US Energy Information Administration

To match the new environment of constant, low-grade volatility in both prices and supply, producers and consumers of oil may need to re-evaluate assumptions and continuously adjust their strategies. Here are several ways that some forward-looking producers and customers are already beginning to do this:

Diversifying oil suppliers and sources. Major oil and gas producers are preparing for greater uncertainty by shifting their reserve portfolios toward unconventional oil and gas in order to respond nimbly and competitively to market shifts. Companies such as Exxon, Chevron, and Shell have all said they expect to expand their production in shale assets in the US, Canada, and Argentina.

At the other end of the spectrum, refineries and other industrial customers are starting to broaden their sources of oil supply and seek more favorable terms. Last year, for example, India, which imports about 80 percent of its crude requirements, began importing oil from the US for the first time in its history. More recently, Poland signed its first ever contract for US crude oil to diversify its supplies from Russia. Some independent refineries in China, Japan, and Poland are trying to secure spot crude oil cargoes to supplement their supplies from traditional long-term supply contracts.

Major oil and gas producers are using cutting-edge technologies to reduce development cycle times and costs for offshore conventional oil projects.

Developing new digital efficiencies. Major oil and gas producers are now trying to apply lessons from the shale revolution's use of cutting-edge technologies to reduce development cycle times and costs for offshore conventional oil projects between 40 percent and 50 percent. Although the effort to digitize oil operations is still in its infancy, leading producers are working closely with oilfield services companies, engineering firms, and construction teams to incorporate artificial intelligence, robotics, and predictive maintenance into offshore operations. Drones are beginning to be used to check for pipeline leaks, self-driving trucks are moving tar sands, and Schlumberger is experimenting with a robotic drilling rig that will complete land wells in 30 percent less time than conventional rigs and require 30 percent fewer man-hours. All of this is with the aim to reduce the marginal cost of the barrel from the current $70 a barrel to around $40.

Major players are also optimizing their field development plans by tapping into new production data streams and developing three-dimensional digital models of their massive offshore platforms. By modularizing components, they hope that deepwater offshore developments can be prebuilt and assembled in three to four years instead of the current seven to nine years, at a fraction of the cost.

Investing in differentiating new services. At the same time, some national oil companies and oil majors are exploring new ways to differentiate themselves from shale producers by investing in refineries, pipelines, petrochemical production, and storage infrastructure close to their customers. Saudi Aramco, for example, is considering committing billions of dollars to expand its refining capacity in Malaysia and Indonesia, as well as a new refining and petrochemical plant in China in an effort to lock in customers.

Re-evaluating buffers to a more uncertain environment. As unpredictability becomes the industry's new normal, some oil and gas producers have started to rely more on hedging as a way to protect themselves from volatile crude oil prices by buying futures contracts that either lock in future prices or put limits on them all the way from their oil wellheads to their refined products. By placing upper and lower bounds on price volatility, producers can count on a more certain cash flow.

Some countries in Asia and the Middle East may come under social pressure to reinstate subsidies to shield their citizens from more frequent price swings at the fuel pump. Indonesia, the United Arab Emirates, and several others accelerated a process of retracting gasoline and diesel fuel subsidies and linking them to market price a few years ago to take advantage of the dramatic fall in oil prices while assuming that they would remain "lower for longer."

The first months of 2018 have shown that the oil industry has entered an era in which change will be the only constant for the foreseeable future. While oil prices will not spike to the peaks that they hit when OPEC and geopolitical events ruled the oil market, the oil markets will likely be unstable as the ranks of maverick shale producers swell, oil majors and national oil companies try out new digital techniques on their conventional fields, and new trade patterns emerge. As we have seen in other industries, to make the most of the new opportunities that lie ahead, oil companies will increasingly need to morph into agile organizations that can pivot to offset and even capitalize on disruptive new shifts. No one will be able to afford to stand still.

Saji Sam is a Dubai-based partner in Oliver Wyman's Energy practice.
Juan Trebino is a Houston-based partner and co-head of Oliver Wyman's Oil and Gas practice in North America.
Bob Orr is a Houston-based partner and co-head of Oliver Wyman's Oil and Gas practice in North America.

This article first appeared in Harvard Business Review on March 1, 2018.

WHAT BANKS CAN LEARN FROM DRIVERLESS CARS

Banks and insurers need to invest more in innovation to remain relevant

Ted Moynihan

IMAGINE IT'S 2006. Which seems more likely to happen by 2018: the emergence of affordable and real-time financial advice or a driverless car? And which would you consider more valuable?

Despite the fact consumers have been wary of self-driving cars, GPS, devices using live traffic updates, ride-hailing services, and autonomous vehicles continue to revolutionize travel. Yet easy, reliable financial advice remains unavailable to the mass consumer market.

The good news is that change is in the air. But there remains a question about whether it will be banks and insurance companies that deliver the change, or whether it will be the tech giants such as Google, Amazon, or Alibaba.

A PIVOTAL MOMENT

Banks and insurers, which have dominated financial services for generations, are at a pivotal moment in their histories. The pillars that supported them – particularly high, risk-free returns using deposits and insurance premiums to buy government debt – have eroded. And while banks and insurers are in decent health after a decade spent rebuilding their balance sheets, they are not spending enough on innovation. What they are spending is not well-harnessed, but scattered across the organizations, and returns are low.

Yet there is still time to channel investment to prevent tech giants and tech-driven newcomers from making banks and insurance companies yesterday's news. Crucially, according to a 2018 survey of 4,000 consumers by Oliver Wyman, customers continue to trust banks and insurers more than the tech giants. The report found that trust, while fragile, is still an asset that many financial services firms retain, and to remain relevant they must both protect it and find opportunities to build on it.

EXHIBIT 1: CONSUMERS TRUST FINANCIAL SERVICES COMPANIES MORE THAN OTHER COMPANIES

Percentage of respondents that strongly agree

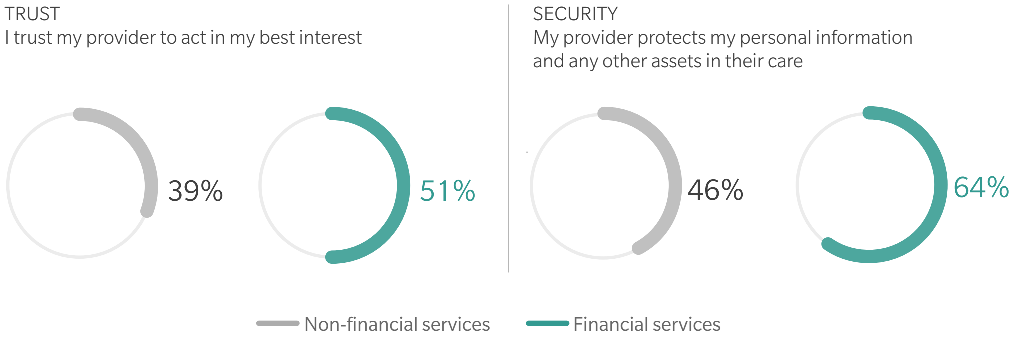

TRUST
I trust my provider to act in my best interest

SECURITY
My provider protects my personal information and any other assets in their care

Source: 2017 Oliver Wyman Global Consumer Survey

Key to taking advantage of this trust will be redefining the way banks and insurers frame customer value. When they have talked about customer value in the past, the focus has been on inward-looking metrics, essentially measuring the value that customers generate for the bank or insurer. Big tech firms on the other hand look at customer value differently, focusing on how to create value for customers, and consequently have emerged as leaders in many of the platforms we use on a daily basis. That is a fundamental mindset shift that offers companies the chance to convert engagement and loyalty into profit and shareholder value.

> *While banks and insurers are in decent health after a decade rebuilding their balance sheets, they are not spending enough on innovation*

FLYWHEEL MOMENTUM

To develop new products that consumers want, banks and insurers would do well to embrace the concept of "flywheel momentum" – which means to relentlessly develop and improve solutions as client engagement increases. The more they improve the customer experience, the more traffic they drive to their solutions. This, in turn, enriches the data they collect and the accuracy and relevance of their algorithms, ensuring better products and greater spending by customers, and making it harder for competitors to catch up.

EXHIBIT 2: THE FIVE MAIN DRIVERS OF VALUE TO CUSTOMERS

The different drivers of value that customers perceive from products and services across the industry

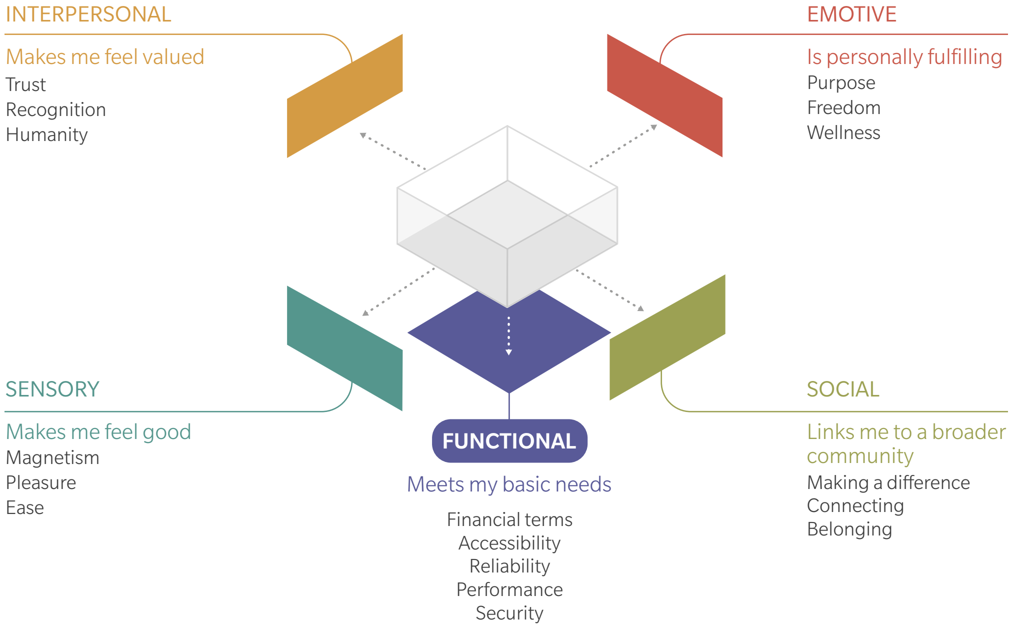

INTERPERSONAL
Makes me feel valued
Trust
Recognition
Humanity

EMOTIVE
Is personally fulfilling
Purpose
Freedom
Wellness

SENSORY
Makes me feel good
Magnetism
Pleasure
Ease

FUNCTIONAL
Meets my basic needs
Financial terms
Accessibility
Reliability
Performance
Security

SOCIAL
Links me to a broader community
Making a difference
Connecting
Belonging

Source: 2017 Oliver Wyman Global Consumer Survey

While incumbent banks and insurers have traditionally focused on the core financial needs of consumers, namely helping them borrow, safeguard, and grow their assets, the tech giants and fintech attackers have been solving other bigger problems faced by consumers. Research shows that consumers really value help to earn, spend, and transfer their wealth more cost-effectively. This is both an opportunity and a threat for traditional financial services firms. Banks and insurers that come to see their businesses as solving life's big problems have an opportunity to offer experiential solutions with real impact.

Banks must ask themselves how they can combine product features with an immersive, engaging, and interactive experience to help consumers earn more from their savings and spend less on necessities. For those that continue to offer merely traditional banking products with largely similar features to everyone else, the future is much more limited.

For consumers, these are good times, and getting better. A Google Maps equivalent for our financial lives is on its way, with a roadmap for its development already starting to emerge. The only question that remains is whether the banks and insurers can overcome their existential crisis to deliver audacious solutions to the big problems facing their clients today. If not, they will become footnotes to economic history, eclipsed by the tech giants that have leveraged data and algorithms to disrupt and seize dominant positions in so many industries already.

Ted Moynihan is a London-based partner and global head of Oliver Wyman's Financial Services practice.

This article first appeared in Financial News on August 21, 2018

RECENT PUBLICATIONS FROM OLIVER WYMAN

For these publications and other inquiries, please email riskjournal@oliverwyman.com. Or visit www.oliverwyman.com and our Oliver Wyman Ideas app: http://apple.co/1UBhSPE

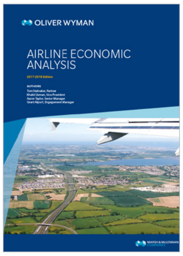

AIRLINE ECONOMIC ANALYSIS 2017-2018

Economic challenges and trends affecting the aviation industry's business models.

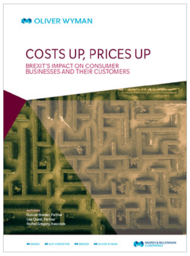

BREXIT: COSTS UP, PRICES UP

Brexit's impact on consumer businesses and their customers.

EXTENDING OUR HORIZONS

Assessing credit risk and opportunity in a changing climate: Outputs of a working group of 16 banks piloting the TCFD Recommendations.

MRO SURVEY 2018 TACKLING INDUSTRY DISRUPTION

Innovation, technology adoption, and challenges in the aviation maintenance, repair, and overhaul sector.

RETAIL'S REVOLUTION

A report on how to navigate through unprecedented changes in retail based on a survey of 3,200 consumers and retail executives.

TEN DIGITAL IDEAS FROM OLIVER WYMAN

In this collection of articles, we showcase ten digital ideas from across our firm for how business leaders can improve and grow their businesses.

THE OLIVER WYMAN AUTOMOTIVE MANAGER

Perspectives on the latest trends and issues in the automotive industry.

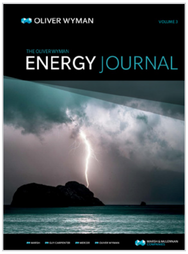

THE OLIVER WYMAN ENERGY JOURNAL, VOL. 3

The latest in thinking from across Oliver Wyman's Energy practice on how shifts underway will create new risks and opportunities not just for the energy sector, but also for every company and person that depends on it.

OLIVER WYMAN
THE OLIVER WYMAN
HEALTH INNOVATION JOURNAL
INSIGHTS FROM THE CUTTING EDGE OF AN INDUSTRY IN TRANSFORMATION

THE OLIVER WYMAN HEALTH INNOVATION JOURNAL

Insights into how the health industry is changing because of new technology and new attitudes.

OLIVER WYMAN
THE OLIVER WYMAN
RISK JOURNAL
PERSPECTIVES ON THE RISKS THAT WILL DETERMINE YOUR COMPANY'S FUTURE

THE OLIVER WYMAN RISK JOURNAL, VOL. 7

A collection of perspectives on the complex risks that are determining many companies' futures.

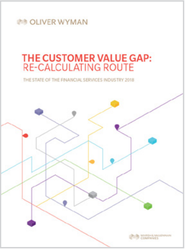

THE STATE OF THE FINANCIAL SERVICES INDUSTRY 2018

Transforming for future value.

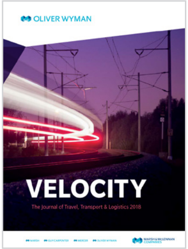

VELOCITY 2018

Perspectives on the issues facing the global travel, transport, and logistics industries.

WHOLESALE BANKS AND ASSET MANAGERS

Winning under pressure.

OLIVER WYMAN FOR SOCIETY ANNUAL SOCIAL IMPACT REPORT 2017

Our contributions to society in 2017 and the work of some nonprofit organizations we support.

OLIVER WYMAN