


March 2020

VOLUME IX, ISSUE III

PEGASUS

A NEWSLETTER FOR THE CAUX ROUND TABLE FOR MORAL CAPITALISM
NETWORK LOOKING AT BUSINESS ABOVE THE CLUTTER AND CONFETTI



Pegasus

Introduction by Stephen B. Young	Page 2
The Age of Surveillance Capitalism – A Precis of its Disruptive Thesis by Stephen B. Young	Page 3
Ethical Principles for the Use of Artificial Intelligence (Draft)	Page 5
Other Voices: Facial Recognition Software Takes Personal Invasion into a New Dimension – The St. Louis Post-Dispatch	Page 9
AI at Google: Our Principles by Sundar Pichai by CEO of Google	Page 11
Will the Coming Combination of AI and Oligarchies Produce a New Feudalism? by Michael Wright	Page 15
Outro by Rich Broderick	Page 21

Introduction

Human history still marches on, coronavirus notwithstanding, and the challenge of doing what is right and proper continues in response to new developments.

Ethics, it seems, is a response to power – how we should use it to profit ourselves or to enhance well-being more generally.

When new technologies come about – fire, stone choppers, spears or knives with sharp killing edges to their blades, steam engines, energy from hydrocarbons, nuclear fission, Facebook and artificial intelligence (AI) – we are challenged by new power in our hands, making possible new realities and new opportunities to do good or evil.

In the last year or so, I have read more and more on what AI might mean for our future: will it condemn us to a dystopian world of state police surveillance 24/7, find cures for all diseases, give us lives of ease and enjoyment?

In any case, how we will use AI will make a difference. We will need laws and regulations. We can rely to some extent on the checks and balances of the marketplace, but, once again, we will also need ethics to constrain, person after person, the dark side of our natures.

This issue of Pegasus begins to consider how our approach of principled leadership in business should respond to AI. We have formulated some very preliminary draft ethical principles for the use of AI.

I am eager to learn from you what you think of this draft. What makes sense? Where does it fall short?


We will be devoting more time and attention in the coming months to consideration of AI and its contributions to the human experience.

Stephen B. Young
Global Executive Director
Caux Round Table for Moral Capitalism

The Age of Surveillance Capitalism A Precis of its Disruptive Thesis

Stephen B. Young

Shosana Zuboff, emerita Professor at the Harvard Business School, has written a passionate book about the use (abuse) of new social and cultural powers by Facebook, Google, Amazon and their ilk. The tone and scope of her book reminds me of Karl Marx in the Communist Manifesto and Das Capital where he subjected the new technologies and powers birthed in the industrial revolution to a moral critique of their impacts.



Zuboff considers the new internet facilitated data collection companies as holding a new form of power which she describes as “the instrumentation and instrumentalizing of behavior for purposes of modification, prediction, monetization and control.” This she names “surveillance capitalism.” It is something different from the market factors of land, labor and capital which Adam Smith and Karl Marx saw at work around them in their day.

This new form of power uses machines which we buy to “transform us into means to others’ market ends.”

The new market power has as its purpose “to fabricate predictions which become more valuable as they approach certainty.” This acquisition of certainty about us – individually and collectively – is a means of behavioral modification. We become objects, data points and so the moral milieu in which our lives unfold is objectification. “Under the regime of instrumentarian power,” we lose our mental agency and self-possession to a kind of automaticity. The digital order transforms our volition into conditioned responses. Whoever owns the means of such value and behavioral modification is our ruler, she proposes.

Instrumentarian power, she writes, replaces individual freedom with the intentions of others and society with certainty; it erodes democracy from within, eating away at the human capabilities and self-understanding required to sustain a democratic life. Instrumentarian power substitutes its machines for social relations. It abandons long standing organic reciprocities among people – products and services are “merely hosts for surveillance

capitalism's parasitic operations.” She concludes: “Surveillance capitalism must be reckoned as a profoundly anti-democratic social force.”

Instrumentarian power is morally agnostic. From its perspective, sin is being autonomous, having the audacity to reject the flows that herd us all toward predictability.

Instrumentarian power is transforming surveillance capitalists into society's self-appointed masters, a “privileged priesthood” which “tunes” people into one channel or another. These priests adopt an attitude of radical indifference to people and to values. “The best products don't win; the ones everyone uses win.” Connectivity is the de facto good.

Ethical Principles for the Use of Artificial Intelligence

(Draft)

The machines and mathematics which constitute Artificial Intelligence (AI) have brought forth a new form of social power, a capacity to replace human freedom with probabilistic certainty. Artificial Intelligence permits unprecedented control and manipulation of culture, economics, politics, law and regulation and individual psychology. Any such power – whether in the hands of individuals, companies, NGOs or governments – must be subordinated to ethical standards honoring human dignity.

Principle 1: Subordination of Artificial Intelligence to Human Purpose

Artificial Intelligence modalities and capabilities shall be disciplined by moral concerns and shall be subordinated to human reason and faith perspectives.

Principle 2: Discourse Ethics Should Guide Application of AI

Responsible AI shall rest its legitimacy in processes of communication and discourse among autonomous moral agents. Free and open discourse, embracing independent media, shall not be curtailed by AI.

Responsible AI acknowledges its duty to contribute value to society through the products and services it provides to consumers.

Responsible AI enhances society through effective and prudent use of resources, free and fair competition and innovation in technology and business practices.

Responsible AI respects the interests of and acts with honesty and fairness towards customers, employees, suppliers, competitors and the broader community using free and fair competition and innovation in technology and business practices.

Responsible AI, therefore, contributes to the economic, social and environmental development of the communities in which it operates in order to sustain its essential ‘operating’ capital – financial, social, environmental and all forms of goodwill.

Principle 3: Build Trust Beyond the Letter of the Law

Responsible AI recognizes that some business behaviors, although legal, can nevertheless have adverse consequences for stakeholders.

Responsible AI, therefore, adheres to the spirit and intent behind the law, as well as the letter of the law, which requires conduct that goes beyond minimum legal obligations.

Responsible AI always operates with candor, truthfulness and transparency and keeps its promises.

1. CUSTOMERS

A responsible AI business treats its customers with respect and dignity. Business, therefore, has a responsibility to:

- a. Provide customers with the highest quality products and services consistent with their requirements.
- b. Treat customers fairly in all aspects of business transactions, including providing a high level of service and remedies for product or service problems or dissatisfaction.
- c. Protect customers from harmful impacts of products and services.
- d. Respect the human rights, dignity and the culture of customers in the way products and services are offered, marketed and advertised.

2. COMPETITORS

A responsible AI business engages in fair competition which is a basic requirement for increasing the wealth of nations and ultimately for making possible the just distribution of goods and services. Business, therefore, has a responsibility to:

- a. Foster open markets for trade and investment.
- b. Promote competitive behavior that is socially and environmentally responsible and demonstrates mutual respect among competitors.
- c. Not participate in anti-competitive or collusive arrangements or tolerate questionable payments or favors to secure competitive advantage.

- a. Respect both tangible and intellectual property rights.
- b. Refuse to acquire commercial information through dishonest or unethical means, such as industrial espionage.

3. COMMUNITIES

As a global corporate citizen, a responsible AI business actively contributes to good public policy and to human rights in the communities in which it operates. Business, therefore, has a responsibility to:

- a. Respect human rights and democratic institutions and promote them wherever practicable.
- b. Recognize government's legitimate obligation to society at large and support public policies and practices that promote social capital.
- c. Promote harmonious relations between business and other segments of society.
- d. Collaborate with community initiatives seeking to raise standards of health, education, workplace safety and economic well-being.
- e. Promote sustainable development in order to preserve and enhance the physical environment, while conserving the earth's resources.
- f. Support peace, security and the rule of law.
- g. Respect social diversity, including local cultures and minority communities.
- h. Be a good corporate citizen through ongoing community investment and support for employee participation in community and civic affairs.

Principle 4: Improving the Well-being of Individual Citizens

Responsible private and public AI shall nurture and support all those social institutions most conducive to the free self-development and self-regard of the individual citizen. Responsible AI shall seek to avoid or to ameliorate conditions of life and work which exploit or deprive the individual citizen of dignity and self-regard or which permit powerful citizens to exploit the weak.

Principle 5: Transparency

To the extent that a responsible AI endeavor influences other parties in society to further its objectives, it has a fiduciary mission to be transparent regarding: its mission and objectives; its values and principles; its governance; its actions; and its means to achieve its objectives.

Scrutiny is only restricted to protect legitimate expectations of personal privacy or to sustain the confidentiality that is required in the organization's daily operations.

In addition, in dealing with governments, corporations and international organizations, responsible AI will always be clear and honest about the interests it represents and the extent to which it speaks on behalf of members, donors, beneficiaries or other stakeholders.

Principle 6: Stewardship

Responsible AI will recognize that its policies and activities are a legitimate subject of public comment and analysis. It is, therefore, willing to engage in reasoned discourse regarding its mission and objectives, values, principles, governance, actions and means used to achieve its objectives. When engaging in advocacy, a civil society institution (CSI) will always, in good faith, present accurate facts and truthful information. When planning its actions or executing its policies, responsible AI will demonstrate enlightened care and concern for those whose interests will be affected by its contemplated actions. In case a CSI inflicts damage upon a government, international organization, corporation or other party, it will be accountable for its actions.

Responsible AI takes care in the creation of quality social capital. From the rule of law to physical infrastructures, from the quality of a society's moral integrity and transparency of its decision-making to the depth and vitality of its culture, social capital demands investment of time, money, imagination and leadership.

Other Voices: Facial Recognition Software Takes Personal Invasion into a New Dimension *The St. Louis Post-Dispatch*

February 21, 2020

Law enforcers around the world are celebrating the rapidly expanding potential of facial recognition software to help them catch criminals faster and solve long-dormant cases. Such capabilities in the right hands could reap enormous public safety dividends. But in the wrong hands, an entirely new dimension in crime, extortion and mayhem could soon be unleashed.

Before it's too late, Americans need to think long and hard about the wisdom of venturing down this uncharted road. Few would argue against tools that help police capture dangerous criminals. It's the law-abiding among us who need to worry about what comes next when this software expands to general public use.



One of the most frightening technological advances is Clearview AI, a powerful facial recognition service currently available to U.S. and Canadian law enforcers. The company's software takes uploaded photos then scrapes the internet, including Facebook, for exact-match images. Its database now contains 3 billion photos and videos. But those aren't just photos of criminals. They're photos of children, grandmothers, families having fun at the beach. Anything and everything is fair game.

That's what makes it so useful to law enforcers. Precise comparison algorithms zip through the entire database at lightning speed to analyze frontal and profile facial photos along with any information connected to the targeted person. Whether it's your name, birthdate, hometown, children's names, high school, hobbies, favorite bar, political views — if it's on the internet, the algorithm uses it.

Most police officers would probably use such software only under authorized circumstances.

But we know from a few local cases that not all officers can be trusted. Some could use it for personal enrichment or to find out, say, who an ex-spouse is dating.

Now imagine such an app on a cellphone for general public use. You're walking down the street and a complete stranger greets you by name, identifies your spouse and kids, maybe mentions the name of your employer or how your family's Grand Canyon vacation went. Maybe the stranger mentions your address or your political leanings. Left unregulated, the threat and exploitation potential would be unlimited.

Clearview insists its software is closely monitored and secure and is designed to "identify child molesters, murderers, suspected terrorists and other dangerous people quickly, accurately and reliably to keep our families and communities safe."

But when a New York Times reporter looked into the company and contacted officers for a demonstration of the program, one officer received a call from Clearview and asked him why he uploaded a New York Times reporter's photo. A block was placed on searches of her. It was a clear demonstration of how the software is vulnerable to political manipulation. If ever there was a clarion call for Congress to impose tight restrictions on this technology, it's now — before the notion of privacy becomes a quaint memory of a bygone era.



AI at Google: Our Principles

Sundar Pichai

CEO

Published Jun 7, 2018

At its heart, AI is computer programming that learns and adapts. It can't solve every problem, but its potential to improve our lives is profound. At Google, we use AI to make products more useful — from email that's spam-free and easier to compose, to a digital assistant you can speak to naturally, to photos that pop the fun stuff out for you to enjoy.



Sundar Pichai.

Beyond our products, we're using AI to help people tackle urgent problems. A pair of high school students are building AI-powered sensors to predict the risk of wildfires. Farmers are using it to monitor the health of their herds. Doctors are starting to use AI to help diagnose cancer and prevent blindness. These clear benefits are why Google invests heavily in AI research and development, and makes AI technologies widely available to others via our tools and open-source code.

We recognize that such powerful technology raises equally powerful questions about its use. How AI is developed and used will have a significant impact on society for many years to come. As a leader in AI, we feel a deep responsibility to get this right. So today, we're announcing seven principles to guide our work going forward. These are not theoretical concepts; they are concrete standards that will actively govern our research and product development and will impact our business decisions.

We acknowledge that this area is dynamic and evolving, and we will approach our work with humility, a commitment to internal and external engagement, and a willingness to adapt our approach as we learn over time.

Objectives for AI applications

We will assess AI applications in view of the following objectives. We believe that AI should:

1. Be socially beneficial.

The expanded reach of new technologies increasingly touches society as a whole. Advances in AI will have transformative impacts in a wide range of fields, including healthcare, security, energy, transportation, manufacturing, and entertainment. As we consider potential development and uses of AI technologies, we will take into account a broad range of social and economic factors, and will proceed where we believe that the overall likely benefits substantially exceed the foreseeable risks and downsides.

AI also enhances our ability to understand the meaning of content at scale. We will strive to make high-quality and accurate information readily available using AI, while continuing to respect cultural, social, and legal norms in the countries where we operate. And we will continue to thoughtfully evaluate when to make our technologies available on a non-commercial basis.

2. Avoid creating or reinforcing unfair bias.

AI algorithms and datasets can reflect, reinforce, or reduce unfair biases. We recognize that distinguishing fair from unfair biases is not always simple, and differs across cultures and societies. We will seek to avoid unjust impacts on people, particularly those related to sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual orientation, ability, and political or religious belief.

3. Be built and tested for safety.

We will continue to develop and apply strong safety and security practices to avoid unintended results that create risks of harm. We will design our AI systems to be appropriately cautious, and seek to develop them in accordance with best practices in AI safety research. In appropriate cases, we will test AI technologies in constrained environments and monitor their operation after deployment.

4. Be accountable to people.

We will design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal. Our AI technologies will be subject to appropriate human direction and control.

5. Incorporate privacy design principles.

We will incorporate our privacy principles in the development and use of our AI technologies. We will give opportunity for notice and consent, encourage architectures with privacy safeguards, and provide appropriate transparency and control over the use of data.

6. Uphold high standards of scientific excellence.

Technological innovation is rooted in the scientific method and a commitment to open inquiry, intellectual rigor, integrity, and collaboration. AI tools have the potential to unlock new realms of scientific research and knowledge in critical domains like biology, chemistry, medicine, and environmental sciences. We aspire to high standards of scientific excellence as we work to progress AI development.

We will work with a range of stakeholders to promote thoughtful leadership in this area, drawing on scientifically rigorous and multidisciplinary approaches. And we will responsibly share AI knowledge by publishing educational materials, best practices, and research that enable more people to develop useful AI applications.

7. Be made available for uses that accord with these principles.

Many technologies have multiple uses. We will work to limit potentially harmful or abusive applications. As we develop and deploy AI technologies, we will evaluate likely uses in light of the following factors:

- *Primary purpose and use*: the primary purpose and likely use of a technology and application, including how closely the solution is related to or adaptable to a harmful use.
- *Nature and uniqueness*: whether we are making available technology that is unique or more generally available.
- *Scale*: whether the use of this technology will have significant impact.
- *Nature of Google's involvement*: whether we are providing general-purpose tools, integrating tools for customers, or developing custom solutions.

AI applications we will not pursue

In addition to the above objectives, we will not design or deploy AI in the following application areas:

1. Technologies that cause or are likely to cause overall harm. Where there is a material risk of harm, we will proceed only where we believe that the benefits substantially outweigh the risks, and will incorporate appropriate safety constraints.
2. Weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people.
3. Technologies that gather or use information for surveillance violating internationally accepted norms.
4. Technologies whose purpose contravenes widely accepted principles of international law and human rights.

We want to be clear that while we are not developing AI for use in weapons, we will continue our work with governments and the military in many other areas. These include cybersecurity, training, military recruitment, veterans' healthcare, and search and rescue. These collaborations are important and we'll actively look for more ways to augment the critical work of these organizations and keep service members and civilians safe.

AI for the long term

While this is how we're choosing to approach AI, we understand there is room for many voices in this conversation. As AI technologies progress, we'll work with a range of stakeholders to promote thoughtful leadership in this area, drawing on scientifically rigorous and multidisciplinary approaches. And we will continue to share what we've learned to improve AI technologies and practices.


We believe these principles are the right foundation for our company and the future development of AI. This approach is consistent with the values laid out in our original Founders' Letter back in 2004. There we made clear our intention to take a long-term perspective, even if it means making short-term tradeoffs. We said it then, and we believe it now.

Will the Coming Combination of AI and Oligarchies Produce a New Feudalism?

by Michael Wright


In this article we examine how the aggregation of Artificial Intelligence (AI) and Domain Specific Knowledge (DSK) is coalescing with the concentration of wealth.

Oligarchy is a form of power structure in which power rests with a small number of people. These people may be distinguished by membership in nobility, a wealth class, a family, an education level, a corporation, a religious affiliation, political group, or the military. Such oligarchic states are often controlled by families who typically pass their influence from one generation to the next, but inheritance is not a necessary condition for the application of this term.



The top 10% of American households, as defined by total wealth, now own 84% of all stocks in 2016, according to a recent paper by NYU economist Edward N. Wolff.(ignorance is accelerating)

Fewer people own and control fewer companies:



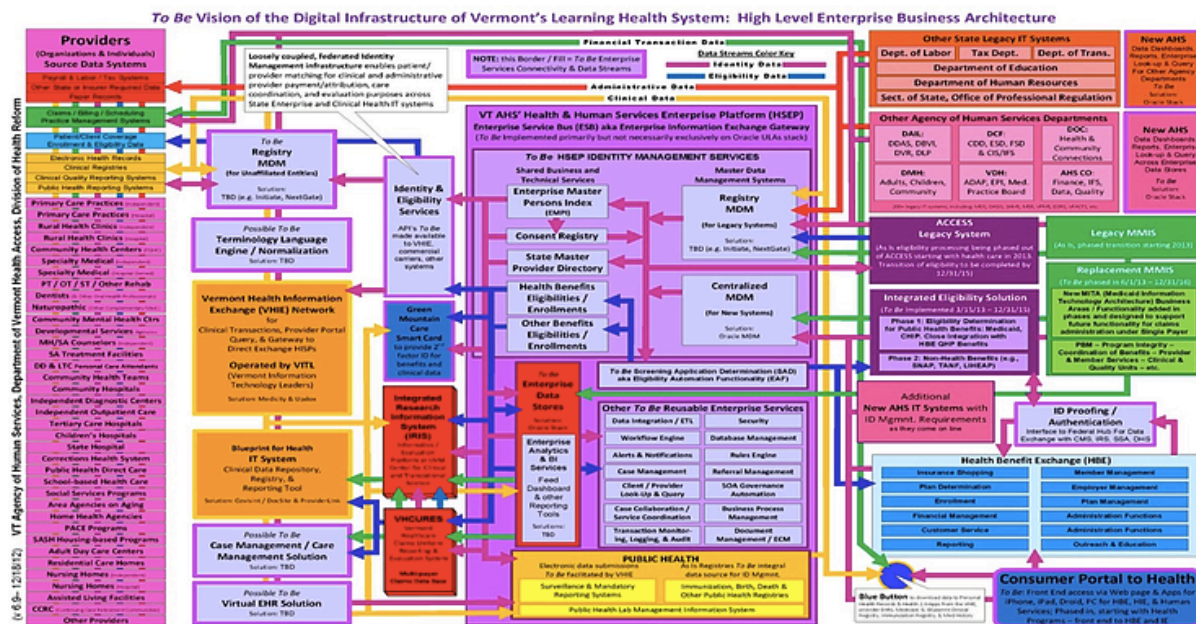
As we said in a previous article, “DSK at the expert level concentrates in the hands of fewer people and companies as it becomes more complex.” (ignorance is accelerating). In the book by Jagdish Sheth and Rajendra Sisodia (The Rule of Three: Surviving and Thriving in Competitive Markets 2002), the authors state that in a mature market, there will normally be three major competitors along with several minor competitors and these minor competitors will only succeed if they are able to operate in a niche market. Interestingly, they go on to say, “Ultimately, the Rule of Three is about the search for the highest level of operating efficiency in a competitive market. Industries with four or more major players, as well as those with two or fewer, tend to be less efficient than those with three major players.”

In looking at the ‘ownership’ (who will actually ‘own’ AI is another discussion) and development of AI, who are likely to be the three major companies (and their limited group of shareholders) and what will the niche players be focused on? It seems only logical that the companies who can generate the most operating income will prevail.

It is hard to imagine a scenario in which AI will not optimize operating efficiencies and follow the rule of three (and for now let’s ignore the possibility of a rule of one). That will mean the concentration of wealth, knowledge, and control of socio-economic activities on a global scale will be in the hands of very few people (by definition an Oligarchic structure). In at least one scenario a new version of feudalism for the 21st century will likely evolve at an unheralded pace beyond our capacity to imagine.

In earlier Insights, we touched on the growing gap being created by AI and the acceleration of ignorance for those without access to information. We said, “If half of human ‘work’ is being done by AI and if that AI is owned and controlled by less than 1% of the companies/people on the planet then the speed at which the 99% falls behind accelerates permanently beyond reach.” (ignorance is accelerating)

As a precursor just look at what is already happening as systems are developed to construct and manage ever larger portions of a complex set of specific activities. For example, the level of complexity inherent in the process of delivering a patient centric portal. This healthcare specific domain has layers of complexity that are well beyond one individual’s complete grasp. However, with entry level AI, this type of activity is easily managed and controlled, linked with other institutions, practices and databases.



Ownership of the flows of data and the subsequent formation of unique information by a few major companies creates tremendous value. And it is just one example of why the top ten players, who make up the early leaders in AI, are buying up Domain Specific Knowledge platforms and early stage AI companies.

What the data tells us is we are entrusting more of our future to fewer hands that are able to fund and control AI, while using technologies we are understanding less. (black box AI) Feng Xiang, a professor of law at Tsinghua University is one of China's most prominent legal scholars. He stated "the most momentous challenge facing socio-economic systems today is the arrival of artificial intelligence. If AI remains under the control of market forces, it will inexorably result in a super-rich oligopoly of data billionaires" (...next come the trillionaires?)

He looks to the one-party system in China where the state is being advertised as the ultimate oligarchic answer to eliminate market derived self-serving oligarchs. "Going forward, China's socialist market economy, which aims to harness the fruits of production for the whole population and not just a sliver of elites operating in their own self-centered interests, can lead the way toward this new stage of human development." He ups the optimism even further stating, "If properly regulated in this way, we should celebrate, not fear, the advent of AI. If (emphasis ours) it is brought under social control, it will finally free workers from peddling their time and sweat only to enrich those at the top." (end of capitalism WorldPost)

Similar to historical precedents prevalent throughout human civilization, wealth and knowledge are once again becoming ever more concentrated to a relative few. Only this time it is at an exponential rate of concentration that may generate a social construct that is potentially only responsive to a very few oligarchs, whether market driven or state run who ultimately may, or may not, have control of the decisions made by AI.

More importantly, history suggests the probability that an artificial superior intelligence will make mistakes beyond the control of its owners. As a result of the complex inter-connected decisions made for them by AI, it is highly unlikely that the new AI oligarchs will be able to anticipate the ramifications, and resulting unintended consequences on our economies and civilization.

**" LEAVE THE
THINKING
TO US. "**

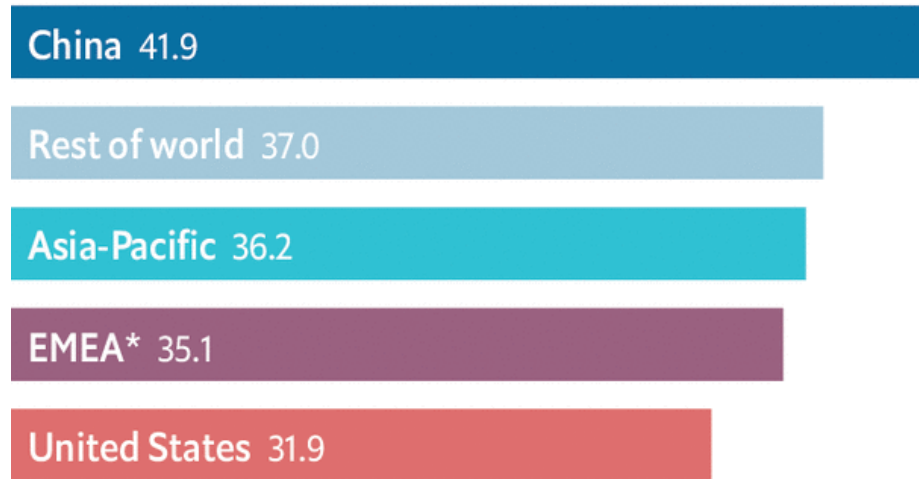
-THE OLIGARCHY

Deluged


Data generated, worldwide

Average annual increase

2010-18, %



Zettabytes[†]



Source: IDC, Seagate *Europe, Middle East and Africa †1ZB=1 trillion GB ‡Estimate §Forecast


The Economist

The bottom line

Average cost decrease and revenue increase from AI adoption, % of respondents* reporting


Costs decreasing

More than 20% 10-19% Less than 10%



Revenues increasing

Less than 5% 6-10% More than 10%



Source: McKinsey & Company

The Economist

*Surveyed November 2019

Outro

With the recent appearance of the coronavirus, all of humanity is once again brought to a place where the ultimate outcome of this problem will be handled in ways that will minimize pain and death or in a manner that ensures that this natural disaster is magnified by business, government and private citizens into a man-made catastrophe of unforeseeable scope.

At this moment, there is no certainty about the outcome, nor about the outcome of yet a new instrument of universal command and control, artificial intelligence (AI). Will AI deliver much needed benefits, especially in the field of medical research, or will it primarily serve as an inescapable method of control and surveillance by business and government? Or perhaps both?

Last year marked the 25th anniversary of our Principles for Business. Since then, we've issued similar guides to sustainable government and developed courses that disseminate these principles and that can be applied by individual organizations to measure their success on this score. All of our publications, round tables and consultation with governments at all levels, as well as individual companies, are grounded on the same underlying concept: in every case, the measure of sustainability is based upon the pursuit of stakeholder benefits, rather than the narrower measurements of shareholder value and partisan control.

What ultimately will be the outcome of today's epidemic? On a larger, even more critical scale, will we be able to adopt measures that minimize the risk of turning Earth into a planet unable to support human civilization – and perhaps even human survival?

We do not know the answer to these dire questions. What we can say, however, with some degree of certainty is that with principles of sustainability and the primacy of stakeholder values – stakeholders that include the physical environment on which we rely for our very survival – the chances that we will be able handle these crises increases exponentially.

Let us hope that as a species, we make the right choice now. And soon. Tomorrow, alas, may prove too late.

Rich Broderick
Director of External Relations
Caux Round Table for Moral Capitalism



www.cauxroundtable.org | Twitter: @cauxroundtable | Facebook: Caux.Round.Table

