F&D

FINANCE & DEVELOPMENT INTERNATIONAL MONETARY FUND

DECEMBER 2023

ARTIFICIAL INTELLIGENCE

What AI means for economics

Contents

Finance & Development
A Quarterly Publication of the International Monetary Fund
December 2023 | Volume 60 | Number 4

Artificial Intelligence

Also in this issue

20

JUN CEN

Departments

On the Cover
AI can develop in many possible ways—it is up to society to determine which path prevails. New York-based illustrator Jun Cen's December 2023 cover evokes AI's deeply unpredictable future as well as our role in deciding its direction.

Subscribe at www.imfbookstore.org/f&d
Read at www.imf.org/fandd
Connect at facebook.com/FinanceandDevelopment

PORTER GIFFORD; ZUMA PRESS, INC./ALAMY STOCK PHOTO

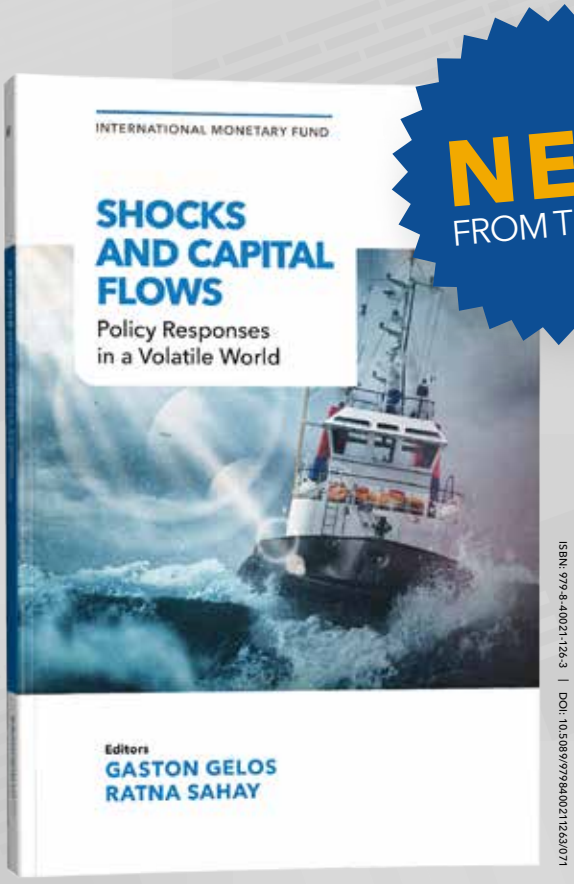

A must-read for policy makers, academics, and financial market participants interested in understanding how policy makers respond to external shocks

- Highlights work to support the development of an integrated policy framework.
- Presents cross-country analyses assessing which policies are most effective in dealing with recurring boom-bust capital flow cycles.
- Brings together recent empirical studies on what countries did in practice and what worked.

"...a valuable and timely contribution on the topic of global capital flows, presenting rich cross-country evidence..."

NOR SHAMSIAH MOHD YUNUS
Governor, Bank Negara Malaysia

PUBLICATIONS

Global economics at your fingertips
IMF.org/pubs

F&D

FINANCE & DEVELOPMENT
A Quarterly Publication of the International Monetary Fund

Subscriber services, changes of address, and advertising inquiries:

IMF Publication Services
Finance & Development
PO Box 92780
Washington, DC 20090, USA
Telephone: (202) 623-7430
Fax: (202) 623-7201
E-mail: publications@imf.org

Postmaster: send changes of address to *Finance & Development*, International Monetary Fund, PO Box 92780, Washington, DC 20090, USA.

The English edition is printed at Dartmouth Printing Company, Hanover, NH.

Finance & Development is published quarterly by the International Monetary Fund, 700 19th Street NW, Washington, DC 20431, in English, Arabic, Chinese, French, Russian, and Spanish.
English edition ISSN 0145-1707

FSC FPO

Editor's Letter

The AI Awakening

FULL DISCLOSURE: THIS ISSUE OF F&D WAS PRO-duced entirely with human intelligence. But someday soon at least parts of this magazine may be assisted by artificial intelligence—a topic that has dominated global discourse since ChatGPT's introduction one year ago.

Generative AI has introduced tantalizing new possibilities in both the public and private spheres. Think how these "machines of the mind" can improve health care diagnoses, close education gaps, tackle food insecurity with more efficient farming, drive planetary exploration—not to mention eliminate the drudgery of work.

Yet the initial excitement surrounding AI has given way to genuine and growing concerns—including about the spread of misinformation that disrupts democracy and destabilizes economies, threats to jobs across the skills spectrum, a widening of the gulf separating the haves and have-nots, and the proliferation of biases, both human and computational.

This issue is an early attempt to understand AI's implications for growth, jobs, inequality, and finance. We bring together leading thinkers to explore how to prepare for an AI world.

In our lead article, Stanford's Erik Brynjolfsson and Gabriel Unger sketch possible "forks in the road" that lead to very different outcomes (beneficial or detrimental) for AI and the economy. The future that emerges will be a consequence of many things, including technological and policy decisions made today, they note.

For MIT's Daron Acemoglu and Simon Johnson, AI's ultimate impact depends on how it affects workers. Innovation always leads to higher productivity, but not always to shared prosperity, depending on whether machines complement or replace humans. The economists outline policies, such as giving labor a voice, that can redirect efforts away from pure automation toward a more "human-complementary" path that creates new and higher-quality tasks.

AI progresses by leaps and bounds. Given its inherent unpredictability, Anton Korinek, of the University of Virginia, recommends scenario planning. He lays out how different technological paths, depending on whether—and how soon—AI exceeds human intelligence, would lead to vastly different outcomes for the economy and workers. Policymakers should prepare reforms for these multiple scenarios and revise as the future unfolds, he notes.

"Ultimately, it's about what AI can do to help people."

This leads us to AI governance. Ian Bremmer, president of Eurasia Group, and Mustafa Suleyman, CEO of Inflection AI, point to regulatory challenges amid a race for AI supremacy among governments. They warn that governing AI will be among the international community's most difficult challenges in coming decades and outline principles for AI policymaking.

The IMF's Gita Gopinath urges balancing innovation and regulation in developing a unique set of policies for AI. Because AI operates across borders, we urgently need global cooperation to maximize the enormous opportunities of this technology while minimizing the obvious harms to society, she writes.

In other thought-provoking articles, Daniel Björkegren and Joshua Blumenstock show how Kenya, Sierra Leone, and Togo adapted AI to benefit the poor. Nandan Nilekani describes how India is on a cusp of an AI revolution to address pressing economic and social challenges. And we profile Harvard labor economist Lawrence F. Katz, whose defining work on inequality illuminates the discussion on AI.

AI can develop in very different directions, underscoring the role of society in actively and collectively determining its future. What is clear is that the technology must be guided as tools that can enhance, rather than undermine, human potential and ingenuity. Ultimately, it's about what AI can do to help people. F&D

Gita Bhatt, editor-in-chief

PHOTO: IMF

Kaleidoscope

A global view, in brief

THE BIG PICTURE: Still emerging from the pandemic, countries in sub-Saharan Africa have been hit by a sluggish global economy, worldwide inflation, high borrowing costs, and a cost-of-living crisis, says the IMF's latest *Regional Economic Outlook for Sub-Saharan Africa*. Moreover, political instability is an ongoing concern. To ensure that the rebound is more than just a transitory glimpse of sunlight, the region's governments must focus on reforms to recover lost ground from recent crises and create space to address pressing development needs, the report says. *Above, Gladys Lampey at Jamestown Harbor, Accra, Ghana. IMF Photo/Andrew Caballero-Reynolds.*

Reset in Fiscal Policy Thinking

THE TEMPTATION TO finance all spending through debt must be resisted, IMF Deputy Managing Director Gita Gopinath warns in a *Financial Times* op-ed. She calls for renewed focus on fiscal policy, and with it, a reset in fiscal policy thinking.

We need to rethink what governments can do, Gopinath says. They cannot be the insurer of first resort for all shocks. They should rebuild depleted fiscal buffers and target future shock responses—which should be temporary by design—to the most vulnerable. Revenues must keep up with spending, and carbon pricing policies should be pursued. Moreover, fiscal frameworks need strengthening; countries must be able to respond to shocks but with clear exit mechanisms.

"These are demanding times for policymakers," Gopinath writes. "Putting fiscal houses in order is essential to ensure governments can deliver for their people."

> "For several advanced economies with aging populations, entitlement reforms are inescapable."
>
> —Gita Gopinath, *The Financial Times*, October 27, 2023

Overheard

"Many economists, particularly over the last 15–20 years, take central bank independence for granted. They have these sterile, technocratic models and don't realize that it is something that central bankers have to fight for every day. It's a very political environment in which central bankers have to protect their right to do monetary policy independently."

—Kenneth Rogoff, Maurits C. Boas Chair of International Economics, Harvard University, in a session on monetary policy challenges at the IMF Annual Research Conference

"What India needs is a strategic plan to chase down the most important opportunities for AI to help. The trick is not to look too hard at the technology but to look at the problems people face that existing technology has been unable to solve."

—Nandan Nilekani, chairman and cofounder of Infosys (see "Unlocking India's Potential with AI," in this issue of F&D)

IN THE NEWS: "Ahead of COP28, we must raise ambition to decisively reduce emissions to prevent grave risks to economic well-being and macro-financial stability," IMF Managing Director Kristalina Georgieva said at the G20 Leaders' Summit in September, citing the United Nations climate change conference that takes place this month in the United Arab Emirates. *Workers at an electrical substation at the Aurora wind farm in Chile. IMFPhoto/Tamara Merino.*

By the numbers

So-called de-risking strategies by China and the United States and other OECD countries that aim to reshore production domestically or friend-shore away from one another can result in a significant drag on growth. (GDP, percent deviation from baseline)

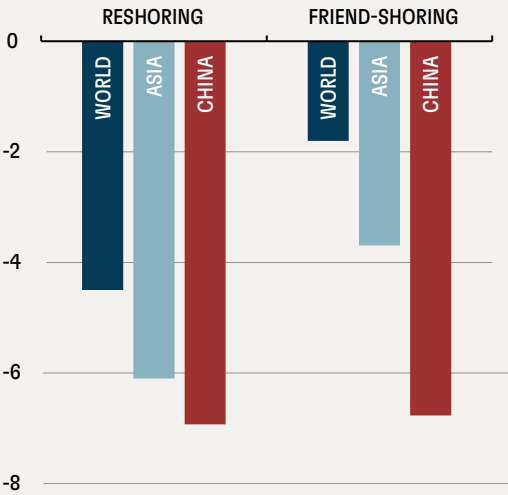

4.5%

of global output could be lost in the long term under a reshoring scenario, in which OECD countries and China increase non-tariff barriers on all countries to reduce dependence on foreign inputs

SOURCE: IMF Regional Economic Outlook for Asia and Pacific, October 2023.

Back to Basics

Artificial Intelligence's Promise and Peril

Generative AI is poised to unleash a wave of creativity and productivity but poses important questions for humanity

Hervé Tourpe

PICTURE A WORLD WHERE machines are artists, storytellers, or even economists producing content that imitates human intelligence. Alan Turing, the pioneering computer scientist, first envisioned the possibility of machines reaching such levels of mastery in a 1950 paper. With ChatGPT and other so-called generative artificial intelligence tools, his prediction of an "imitation game" is now reality. It feels as if we've been catapulted into a universe once reserved for science fiction. But what exactly is generative AI?

GenAI represents the most impressive advance in machine-learning technologies yet. It marks a significant leap in AI's ability to understand and interact with complex data patterns and is poised to unleash a new wave of creativity and productivity. But it also raises important questions for humanity. Key innovation milestones marked the path to its current sophistication.

In the 1960s, a program called ELIZA impressed scientists with its ability to generate human-like responses. It was basic and operated by set rules, but it was the precursor of what we now know as "chatbots." Two decades later, artificial neural networks appeared. These networks, inspired by human brains, gave machines new skills, such as understanding the nuances of language and recognizing images. But a limited pool of data for training and inadequate computing power held back real progress. Remarkably, these twin resources kept doubling each year, setting the stage for the third wave of AI in the 2000s: deep learning.

Deep learning

With innovations such as Google Translate, digital assistants like Alexa and Siri, and the emergence of self-driving cars, machines started to understand and interact with the world. Yet for all this progress, a piece of the puzzle was still missing. Machines could assist and predict, but they couldn't truly understand the intricacies of human conversation, and they were poor at generating human-like content.

Then, in 2014, generative adversarial networks (GANs) leveraged the ability of two competing neural networks to sharpen each other's skills continuously. The "generator" created imitation data, text, or images, while the "discriminator" tried to differentiate between real and simulated content. This dual-network competition revolutionized the way AI understood and replicated complex patterns.

The last piece of the puzzle arrived in 2017 with a groundbreaking paper, "Attention Is All You Need." By teaching the AI to pay attention to relevant parts of the input, it suddenly seemed that the machine started to get it—to grasp the essence of the input. This generative AI produced eerily human-like content, at least in labs.

Together, GANs and attention mechanisms, supported by ever-growing information and computing power, set the stage for ChatGPT—the most astonishing chatbot ever. It was launched by OpenAI in November 2022, and other big-tech firms soon followed with GenAI chatbots of their own.

Economics and finance

AI is not, of course, a new concept in economics and finance. Traditional AI (advanced analytics, machine learning, predictive deep learning) has been

DALE CROSBY-CLOSE

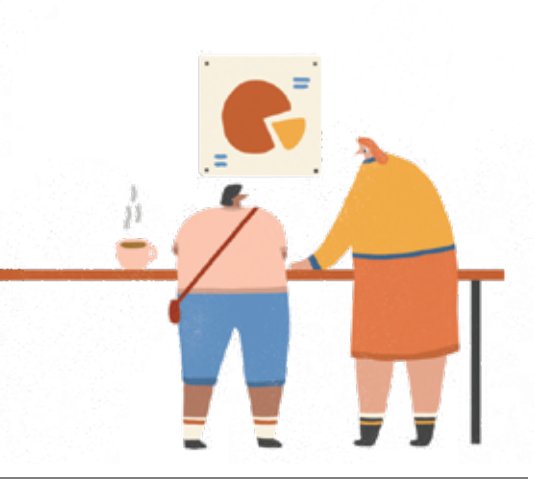

"GenAI creations can be so convincing that they create a false sense of reality. This has the potential to spread misinformation, incite panic, and even destabilize economic or financial systems."

crunching numbers, gauging market trends, and customizing financial products for a long time. What sets GenAI apart is its ability to delve deeper and interpret complex data in a more creative manner. By dissecting intricate relationships between economic indicators or financial variables, it spits out not just forecasts but alternate scenarios, insightful charts, and even snippets of code that could significantly change how the sector operates.

The evolution from traditional to generative AI has introduced a new era of possibilities into both the public and private spheres. Governments are beginning to employ these smarter tools to improve citizen services and overcome workforce shortages. Central banks are taking note, seeing in GenAI an enhanced capacity for sifting through vast amounts of banking data to refine economic forecasts and better monitor risks, including fraud.

Investment firms are turning to GenAI to detect subtle shifts in stock prices and market sentiment, drawing from a larger body of knowledge to propose more creative options, paving the way for potentially more lucrative investment strategies. Meanwhile, insurance companies are exploring how generative models can create personalized policies that align more closely with individual needs and preferences.

GenAI is evolving at a breakneck pace, pushing the boundaries of AI capabilities in economics and finance and introducing novel solutions to old challenges. Some people are skeptical. They say that, like a stochastic parrot, AI can create nonsensical and untrue facts, a phenomenon called "hallucination," and it doesn't really know the meaning behind the words. ChatGPT's knowledge, they point out, is limited to its latest training date. Possibly. But given the mind-boggling pace of innovation, how long will these arguments remain relevant?

Still, the initial excitement surrounding GenAI has given way to growing and genuine concerns. Traditional challenges associated with AI, such as the amplification of existing biases in training data, or the lack of decision transparency, have taken on renewed urgency. New concerns have also arisen.

AI weaponized

One particularly alarming risk is GenAI's remarkable ability to tell stories that resonate with individuals' preexisting beliefs and viewpoints, potentially reinforcing echo chambers and ideological silos. Malicious actors can leverage this ability not only through the written word: in March 2022, an AI-generated video purported to show Ukrainian President Volodymyr Zelenskyy surrendering to Russian forces. Such incidents demonstrate how GenAI can be weaponized to manipulate politics, markets, and public opinion.

Whether it's a fabricated story, doctored image, or synthetic video, GenAI creations can be so convincing that they create a false sense of reality. This has the potential to spread misinformation, incite panic, and even destabilize economic or financial systems with unprecedented efficiency and intensity. It may not always be deliberate: machines may spread misinformation unintentionally as a result of hallucinations.

The threat of AI is not limited to manipulation. Job displacement is another concern as GenAI continues to advance, potentially automating tasks that were previously performed by humans, leading to many job losses and requiring strategies for employment and retraining.

Earlier this year, leading AI experts, including ChatGPT's creator, cosigned a letter warning that "mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." They were echoing concerns expressed decades earlier by Turing, who warned that "there is a danger that machines will eventually take control of our lives."

We stand at a crossroads of technology and ethics. GenAI, with its vast promise and profound, existential questions, cannot be uninvented. As we leverage its transformative power, it's imperative to remember Turing's enduring counsel. GenAI is a monumental shift that demands vigilant oversight, new regulatory frameworks, and an unwavering commitment to ethical, transparent, controllable innovations that harmonize with human values. F&D

HERVÉ TOURPE *is head of the IMF's Digital Advisory Unit.*

Point of View

Building Blocks for AI Governance

Ian Bremmer and Mustafa Suleyman

Policymakers must adhere to five guiding principles to govern AI effectively

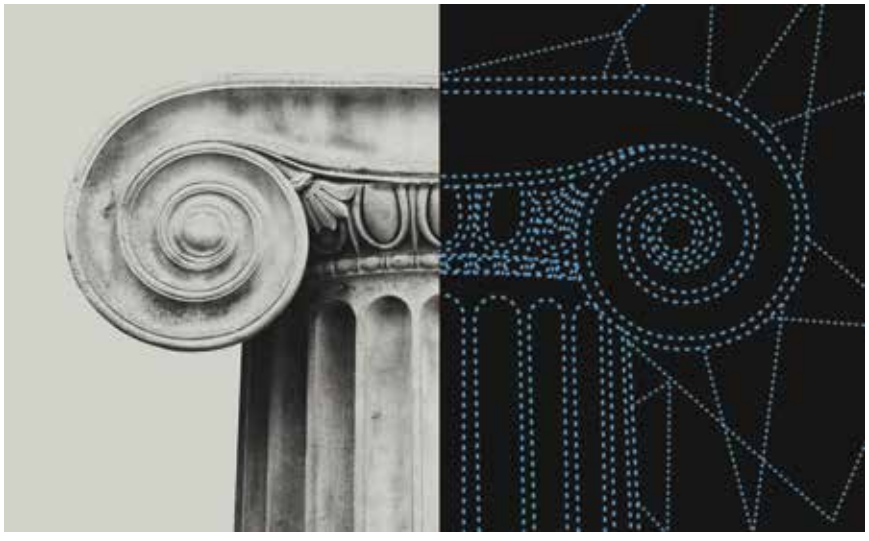

Artificial intelligence will open people's lives and societies to groundbreaking scientific advances, unprecedented access to technology, toxic misinformation that disrupts democracies, and economic upheaval. In the process, AI will trigger a fundamental shift in the structure and balance of global power.

This creates an unparalleled challenge for political institutions around the world. They will have to establish new norms for a dynamic novel technology, mitigate its potential risks, and balance disparate geopolitical actors' interests. Increasingly, these actors will come from the private sector. And it will require a high level of coordination from governments, including strategic competitors and adversaries.

In 2023, governments around the world woke up to this challenge. From Brussels to Beijing to Bangkok, lawmakers are busy crafting regulatory frameworks to govern AI, even as the technology itself advances exponentially. In Japan, Group of Seven leaders launched the "Hiroshima Process" to tackle some of the trickiest questions raised by generative AI, while the UN launched a new AI high-level advisory body. At the Group of Twenty summit in New Delhi, Indian Prime Minister Narendra Modi called for a new framework for responsible human-centric AI governance, and European Commission President Ursula von der Leyen advocated for a new AI risk monitoring body modeled on the Intergovernmental Panel on Climate Change.

In November, the UK government hosted the world's first leader-level summit dedicated to addressing AI safety risks. Even in the US, home of the biggest AI companies and traditionally hesitant to regulate new technology, AI regulation is a question of when, not if, and a rare instance of bipartisan consensus.

This flurry of activity is encouraging. In a remarkably short amount of time, world leaders have prioritized the need for AI governance. But agreeing on the need for regulation is table stakes. Determining what *kind* of regulation is just as important. AI doesn't resemble any previous challenge, and its unique characteristics, coupled with the geopolitical and economic incentives of the principal actors, call for creativity in governance regimes.

AI governance is not just one problem. When it comes to climate change, there may be many routes to achieving the ultimate objective of lowering greenhouse gas emissions, but there is a single overriding objective. AI is different, as an

AI may become the first technology with the means to improve on itself.

RICHARD JOPSON (LEFT) AND COURTESY OF MUSTAFA SULEYMAN (RIGHT)

Illustration by Joan Wong

AI policy agenda must simultaneously stimulate innovation to solve intractable challenges and avoid dangerous proliferation, and it must help attain geopolitical advantage without sleepwalking the world into a new arms race.

The AI power paradox

The nature of the technology itself is a further complication. AI can't be governed like any previous technology because it's unlike any previous technology. It doesn't just pose policy challenges; its unique characteristics make solving those challenges progressively harder. That is the AI power paradox.

For starters, all technologies evolve, but AI is hyper-evolutionary. AI's rate of improvement will far surpass the already powerful Moore's Law, which has successfully predicted the doubling of computing power every two years. Instead of doubling every two years, the amount of computation used to train the most powerful AI models has increased by a factor of 10 every year for the past 10 years. Processing that once took weeks now happens in seconds. The foundation technologies that enable AI are only going to get smaller, cheaper, and more accessible.

But AI's uniqueness is not just about expanded computing capacity. Few predicted AI's evolution, from its ability to train large language models to being able to solve complex problems or even compose music. These systems may soon be capable of quasi-autonomy. This would on its own be revolutionary but would come with an even more dramatic implication: AI may become the first technology with the means to improve on itself.

AI proliferates easily. As with any software, AI algorithms are far easier and cheaper to copy and share (or steal) than physical assets. And as AI algorithms get more powerful—and computing gets cheaper—such models will soon run on smartphones. No technology this powerful has ever been so accessible so widely so quickly. And because its marginal cost—not to mention marginal cost of delivery—is zero, once released, AI models can and will be everywhere. Most will be safe; many have been trained responsibly. But, as with a virus, all it takes is one malign or "breakout" model to wreak havoc.

Incentives point toward ungoverned AI

The nature of AI suggests different incentives as well. Dual-use technologies are nothing new (there's a reason civilian nuclear proliferation is closely monitored), and AI is not the first technology whose civil and military uses are blurred. But whereas technologies such as nuclear enrichment are highly complex and capital-intensive, AI's lost cost means it can be deployed endlessly, whether for civil or military use. This makes AI more than just software development as usual; it is an entirely new and dangerous means of projecting power.

Constraining AI is hard enough on a technological basis. But its potential for enriching and empowering powerful actors means that governments and the private companies developing AI are incentivized to do the opposite. Simply put, AI supremacy is a strategic objective of every government and company with the resources to compete. If the Cold War was punctuated by the nuclear arms race, today's geopolitical contest will likewise reflect a global competition over AI. Both the US and China see AI supremacy as a strategic objective that must be achieved—and denied to the other. This zero-sum dynamic means that Beijing and Washington focus on accelerating AI development, rather than slowing it down.

But as hard as nuclear monitoring and verification were 30 years ago, doing the same for AI will be even more challenging. Even if the world's powers were inclined to contain AI, there's no guarantee they'd be able to, because, as in most of the digital world, every aspect of AI is currently controlled by the private sector. And while the handful of large tech firms that currently control AI may retain their advantage for the foreseeable future, it is just as likely that the gradual proliferation of AI will bring more and more small players into the space, making governance more complicated. Either way, the private businesses and individual technologists who will control AI have little incentive to self-regulate.

Any one of these features would strain traditional governance models; all of them together render these models inadequate and make the challenge of governing AI unlike anything governments have faced before.

Governance principles

If global AI governance is to succeed, it must reflect AI's unique features. And first among those is the reality that as a hyper-evolutionary technology, AI's progress is inherently unpredictable. Policymakers must consider that given such unpredictability, any rules they pass today may not be effective or even relevant in a few months, let alone a few years. To box in regulators with inflexible regimes now would be a mistake.

Instead, good governance would be best served by establishing a set of first principles on which AI policymaking can be based:

- **Precautionary:** The risk-reward profile of AI is asymmetric; although there are vast benefits to AI's potential, policymakers must guard against its potentially catastrophic downsides. The already widely used precautionary principle needs to be adapted to AI and enshrined in any governance regime.
- **Agile:** Policymaking structures tend to be static, prizing stability and predictability over dynamism and flexibility. That won't work with a technology as unique as AI. AI governance must be as agile, adaptive, and self-correcting as AI is fast-moving, hyper-evolutionary, and self-improving.
- **Inclusive:** The best industry regulation, especially when it comes to technology, has always worked collaboratively with the commercial sector, and this is especially true for AI. Given the exclusive nature (at least for now) of AI development—and the complexity of the technology—the only way for regulators to properly oversee AI is to collaborate with private technology companies. To reflect the borderless nature of AI,

governments should make companies parties to international agreements. Including private companies in high diplomacy may veer toward unprecedented, but excluding those who have so much control would doom any governance structure that excludes them before it even starts.

- **Impermeable:** For AI governance to work, it must be impermeable; given AI's ability to easily proliferate, just one defection from the regime could allow a dangerous model to escape. Therefore, any compliance mechanisms should be watertight, with easy entry to compel participation and costly exit to deter noncompliance.
- **Targeted:** Given AI's general-purpose nature and the complexities involved in governing it, a single governance regime is insufficient to address the various sources of AI risk. In practice, determining which tools are appropriate to target which risks will require developing a live, working taxonomy of discrete potential AI impacts. AI governance must therefore be targeted, risk-based, and modular rather than one-size-fits-all.

Governing AI will be among the international community's most difficult challenges in the coming decades. As important as the imperative to regulate AI is the imperative to regulate it correctly. Current debates on AI policy too often tend toward a false debate between progress and doom (or geopolitical and economic advantages versus risk mitigation). And rather than think creatively, solutions too often resemble paradigms for yesterday's problems. This will not work in the age of AI.

Good policymaking will be vital, but getting there rests on good institutions. To build these institutions, the international community will need to agree on a conceptual framework for how to think about AI. We offer these principles as a start. F&D

IAN BREMMER *is president and founder of Eurasia Group and GZERO Media.* **MUSTAFA SULEYMAN** *is CEO and cofounder of Inflection AI.*

Fostering More Inclusive Democracy with AI

Hélène Landemore

AI can enhance democratic institutions by ensuring citizens' voices are truly heard

People worry that artificial intelligence is, or will soon be, undermining democracy. They fear AI will take away jobs, destabilize the economy, and widen the divide between the rich and the poor. This could further concentrate power in the hands of a few tech companies and weaken government structures designed to regulate them. Some also fear that tech giants and government may increasingly delegate human decision-making to machines, eventually replacing democracy with "algocracy," rule not by the people but by algorithm.

This dystopian vision misses our current capacity to shape AI development. We, as human societies, have the political ability (at least for now) and the responsibility to address the harm AI could inflict on us. We also have the technological opportunity to harness AI to enhance our democracy in a way that strengthens our collective ability to govern—rather than simply regulate—AI.

Like other ethical and political challenges, such as gene editing, AI governance requires not just more expert intervention and regulation but more citizen voice and input—for example, on how to navigate the distributive impact of AI on the economy. Like other global concerns, such as climate change, AI governance requires this democratic voice to be heard at the level of international institutions. Luckily, AI has the potential to usher in a more inclusive, participatory, and deliberative form of democracy, including at the global scale.

Participatory experiments

For 40 years many governments have engaged in experiments aiming to include ordinary citizens in policymaking and lawmaking in richer ways than through voting alone. These experiments have mostly been local and small-scale, much like the citizens' assemblies and juries that have proliferated on climate and other issues. A 2020

PHOTO COURTESY OF HELENE LANDEMORE

"AI regulation is likely to be better enforced and more effective in AI-empowered democracies."

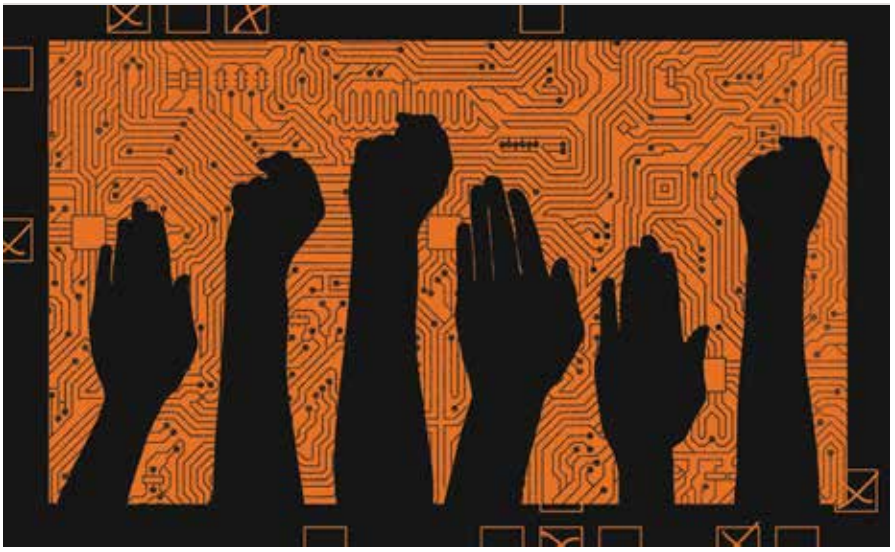

Organisation for Economic Co-operation and Development report found close to 600 such cases in which a random sample of citizens engages deeply with an issue and formulates informed policy recommendations (and in one case even proposals).

But some of these political experiments have also aimed for mass participation, as in the participatory constitutional processes organized in Brazil, Kenya, Nicaragua, South Africa, and Uganda in the 1980s and 1990s, and more recently in Chile, Egypt, and Iceland, which have used mass consultations and crowdsourcing to reach out to ordinary people. Not every attempt has been successful, of course, but all are part of a significant trend.

Some governments have also rolled out broad multi-format consultation campaigns. The 2019 Great National Debate launched by French President Emmanuel Macron in response to the yellow vest movement, with some 1.5 million participants, is one example. Another is the EU-wide Conference on the Future of Europe, which invited citizens from EU member countries to weigh in on reforms to EU policies and institutions, prompting 5 million people to visit the website and 700,000 to engage in debate.

Despite some online elements, these have been mostly low-tech, analog processes, involving no AI whatsoever. Politicians, overwhelmed by the raw and multifaceted data or unsure of its meaning, have as a result easily ignored the citizens' input. People were allowed to speak but were not always heard. And the level of deliberation, even for those involved, was often superficial.

Enhanced deliberation

We now have the chance to scale and improve such deliberative processes exponentially so that citizens' voices, in all their richness and diversity, can make a difference. Taiwan Province of China exemplifies this transition.

Following the 2014 Sunflower Revolution there, which brought tech-savvy politicians to power, an online open-source platform called pol.is was introduced. This platform allows people to express elaborate opinions about any topic, from Uber regulation to COVID policies, and vote on the opinions submitted by others. It also uses these votes to map the opinion landscape, helping contributors understand which proposals would garner consensus while clearly identifying minority and dissenting opinions and even groups of lobbyists with an obvious party line. This helps people understand each other better and reduces polarization. Politicians then use the resulting information to shape public policy responses that take into account all viewpoints.

Over the past few months pol.is has evolved to integrate machine learning with some of its functions to render the experience of the platform more deliberative. Contributors to the platform can now engage with a large language model, or LLM (a type of AI), that speaks on behalf of different opinion clusters and helps individuals figure out the position of their allies, opponents, and everyone in between. This makes the experience on the platform more truly deliberative and further helps depolarization. Today, this tool is frequently used to consult with residents, engaging 12 million people, or nearly half the population.

Corporations, which face their own governance challenges, also see the potential of large-scale AI-augmented consultations. After launching its more classically technocratic Oversight Board, staffed with lawyers and experts to make decisions on content, Meta (formerly Facebook) began experimenting in 2022 with Meta Community Forums—where randomly selected groups of users from several countries could deliberate on climate content regulation. An even more ambitious effort, in December 2022, involved 6,000 users from 32 countries in 19 languages to discuss cyberbullying in the metaverse over several days. Deliberations in the Meta experiment were facilitated on a proprietary Stanford

Illustration by Joan Wong

University platform by (still basic) AI, which assigned speaking times, helped the group decide on topics, and advised on when to put them aside.

For now there is no evidence that AI facilitators do a better job than humans, but that may soon change. And when it does, the AI facilitators will have the distinct advantage of being much cheaper, which matters if we are ever to scale deep deliberative processes among humans (rather than between humans and LLM impersonators, as in the Taiwanese experience) from 6,000 to millions of people.

Translation, summarization, analysis

The applications of AI in deliberative democracy are still in the exploratory phase. Instantaneous translation among multilinguistic groups is the next frontier, as is summarization of collective deliberations. According to recent research, AI is 50 percent more accurate than human beings when it comes to summarization (as evaluated by trained undergraduates comparing AI summaries and human coders' summaries of deliberation transcripts). Some amount of human judgment will, however, likely be necessary for many of these tasks. In such cases AI can still serve as a useful aid to human analysts, facilitators, and translators.

More ways that AI can enhance democracy are on the horizon. OpenAI, the company that launched ChatGPT, recently introduced a grant program called Democratic inputs to AI. The grants subsidized the 10 most promising teams in the world working on algorithms that serve human deliberation (full disclosure: I am on the board of academic advisors who helped formulate the grant call and select the winners). These tools can hopefully soon be deployed to serve, among other goals, global deliberation on AI governance, in line with the vision of OpenAI CEO Sam Altman.

Addressing risks

Deploying AI in democracy has its risks—like data bias, privacy concerns, potential for surveillance, and legal challenges—in almost every field. It also raises the problem of the digital divide and the potential exclusion of illiterate and techno-skeptical groups. Many of these problems will need to be addressed politically, economically, legally, and socially first and foremost, rather than through technology alone. But technology can help here too.

For example, privacy and surveillance concerns may be remediated by something such as zero-knowledge protocols (also called zero-knowledge proofs, or ZKP), which aim to verify or "prove" identity without collecting data on participants (for example, through text messaging authentication or through blockchain). ZKP can be used both for online voting and in deliberative contexts—for example, to share sensitive information or play the role of whistleblower. Meanwhile, generative AI can make previously scarce knowledge and tutoring resources available to everyone who needs them. As a custom-tailored interlocutor for citizens, it can explain technical policy issues in people's particular cognitive style (including through images) and convert their oral input into written input as needed.

Despite its limitations and risks, AI has the potential to bring about a better, more inclusive version of democracy, one that would in turn equip governments with the legitimacy and knowledge to oversee AI development. AI regulation is likely to be better enforced and more effective in AI-empowered democracies.

Still, there is a risk that democracy itself could be a casualty of the AI revolution. Urgent investment is needed in AI tools that safely augment the participatory and deliberative potential of our governments. F&D

HÉLÈNE LANDEMORE *is a professor of political science at Yale University. She is also a fellow at the Ethics in AI Institute at the University of Oxford and an advisor to the Democratic inputs to AI program at OpenAI.*

Unlocking India's Potential with AI

Nandan Nilekani and Tanuj Bhojwani

India is on the brink of a transformation that could change its economic and social future

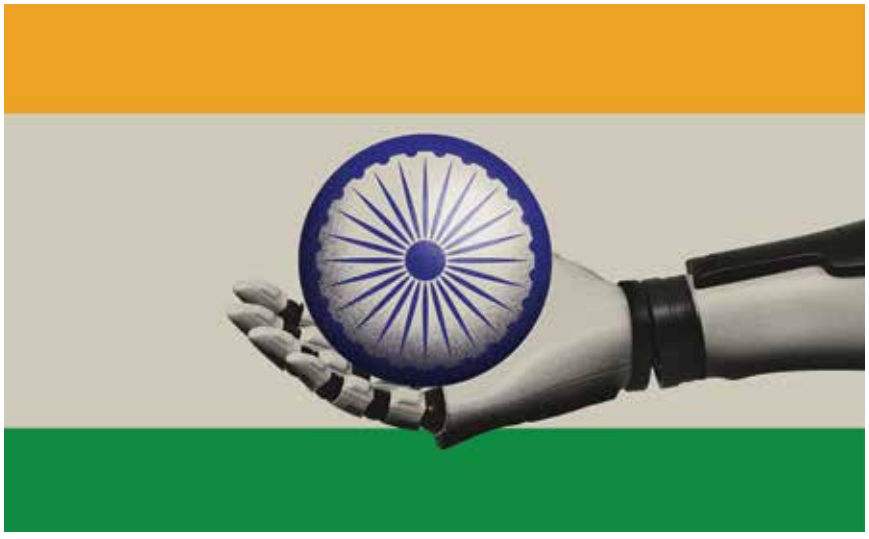

Before the end of this decade, more Indians will use AI every day than in any other country in the world. What's more, people in advanced economies will be surprised by the ways the country will use AI. India is on the cusp of a technological revolution that could alter the trajectory of its social and economic future, and in this revolution there are lessons for the rest of the world.

Our prediction hinges on three facts: India needs it, India is ready for it, and India will do it.

India needs it

The concept of "China plus one" has been gaining traction, with its admonition that global companies should not depend inordinately on China for their manufacturing and software needs. India, with its growing infrastructure investments, favorable policies, and young working population, is the most likely beneficiary of this shift. It is perhaps the only country poised to match the scale of China.

With 1.4 billion people, India is closer to a continent than a country. Its population is almost twice that of Europe. But the average age in India is 28, compared with Europe's 44, which means a higher share of the population is of working age. This is the starting point: India is a very large country of very young people.

The rest of the world has been eyeing AI with curiosity, waiting for real-use cases. In India, we see potential today.

This demographic dividend, favorable global trends, and the unlocking of decades of suppressed potential are starting to show returns. Even as the macroeconomic projections for most of the world seem modest or bleak, India remains a bright spot. These young Indians are aspirational and motivated to use every opportunity to better their lives.

What really sets India apart from the West are its unique challenges and needs. India's diverse population and complex socioeconomic concerns mean that AI there is not just about developing cutting-edge technology. It's about finding innovative solutions to address pressing problems in health care, education, agriculture, and sustainability.

Though our population is just double the size of Europe's, we are much more diverse. Indians, like Europeans, are often bi- or multilingual. India recognizes 19,500 dialects spoken by at least 10,000 people. Based on data from the Indian census, two Indians selected at random have only a 36 percent chance of speaking a common language.

This language barrier is complicated by the fact that the official literacy rate in the country hovers near 77 percent, varying vastly between states. This means that roughly 1 in 4 people can't read or write. Even though the government tries to provide welfare assistance for its most vulnerable, it's hard to spread awareness about the service and reach the last mile. Filling out a simple form to access welfare can be daunting for someone who is illiterate. Determining eligibility for assistance means depending on someone who can read, write, and navigate the bureaucracy.

Illustration by Joan Wong

Actually receiving services means assistance seekers must have an agent helping them who is not misinformed—or worse, corrupt. These barriers disproportionately affect those who need government assistance the most.

We have the ability to solve a lot of problems for our population, but the hard part has always been in the *distribution,* not the solution. In India, we believe that AI can help bridge this access gap. AI enables people to access services directly with their voice using natural language, empowering them to help themselves. As Canadian writer William Gibson aptly said, "The future is already here—it's just not evenly distributed." Nowhere is this more glaringly evident than in India.

The rest of the world has been eyeing AI with curiosity, waiting for real-use cases. In India, we see potential today. While this may be true of many other developing economies, the other important factor is that. . .

India is ready for it

India's population isn't just young, it is *connected*. According to the country's telecommunications sector regulator, India has more than 790 million mobile broadband users. Internet penetration continues to increase, and with the availability of affordable data plans, more and more people are online. This has created a massive user base for AI applications and services.

But where India has surpassed all others is in its digital public infrastructure. Today, nearly every Indian has a digital identity under the Aadhaar system. The Aadhaar is a 12-digit unique identity number with an option for users to authenticate themselves digitally—that is, to prove they are who they claim to be.

Further, India set up a low-cost, real-time, interoperable payment system. This means that any user of any bank can pay any other person or merchant using any other bank instantly and at no cost. This system—the Unified Payments Interface—handles more than 10 billion transactions a month. It is the largest real-time payment system in the world and handles about 60 percent of real-time payment transactions worldwide.

With the success of these models, India is embracing innovation in open networks as digital public infrastructure. Take the example of Namma Yatri, a ride-hailing network built in collaboration with the union of auto-rickshaw drivers in Bangalore and launched in November 2022. These drivers have their own app, with a flat fee to use it, no percentage commission, and no middleman. The app has facilitated close to 90,000 rides a day, almost as many as ride-hailing companies in the city.

Unlike Western countries, which have legacy systems to overhaul, India's tabula rasa means that AI-first systems can be built from the ground up. The quick adoption of digital public infrastructure is the bedrock for these technologies. Such infrastructure generates enormous amounts of data, and thanks to India's Account Aggregator framework, the data remain under the citizens' control, further encouraging public trust and utilization. With this solid footing, India is well positioned to lead the charge in AI adoption.

India will do it

In September 2023, the Indian government, in collaboration with the EkStep foundation, launched the PM-Kisan chatbot. This AI chatbot works with PM-Kisan, India's direct benefit transfer program for farmers, initiated in 2019 to extend financial help to farmers who own their own land. Access to the program, getting relevant information, and resolving grievances was always a problem for the farmers. The new chatbot gives farmers the ability to know their eligibility and the status of their application and payments using just their voice. On launch day more than 500,000 users chatted with the bot, and features are being released slowly to ensure a safe and risk-managed rollout.

These steps are part of an encouraging trend of early adoption of new technology by the Indian government. But the trend extends beyond the government. India's vibrant tech ecosystem has taken off as well, a direct offshoot of its booming IT exports—currently at nearly $250 billion a year. Next to those from the US, the largest number of developers on GitHub, a cloud-based service for software development, are from India. This sector not only innovates but also widely adopts digital public infrastructure. The effect is cyclical: start-ups feed the growing tech culture and, in turn, leverage the data to build more precise and beneficial AI tools. India's dynamic start-up ecosystem, moreover, is actively working on AI solutions to address various challenges.

AI can be a game changer in education as well, helping close the literacy gap. AI technologies are uniquely positioned to help students learn in their native languages, as well as learn English. AI's applications are useful not only for students; they extend to teachers, who are often overwhelmed by administrative tasks that detract from teaching. As AI takes over routine tasks in government and start-ups, the roles of teachers and students evolve, and they form dynamic partnerships focused on deep learning and meaningful human interaction.

What India needs is a strategic plan to chase down the most important opportunities for AI to help. The trick is not to look too hard at the technology but to look at the problems people face that existing technology has been unable to solve. And organizations such as EkStep have stepped up with a mission called People+AI. Instead of putting AI first, they focus on the problems of people. This has led to surprising new uses unique to India.

India's emerging status as a technological powerhouse, combined with its unique socioeconomic landscape, puts it in a favorable position to be the world's most extensive user of AI by the end of this decade. From streamlining education to aiding in social protection programs, AI has the potential to deeply penetrate Indian society, effecting broad and meaningful change. F&D

NANDAN NILEKANI *is the chairman and cofounder of Infosys and founding chairman of UIDAI (Aadhaar).* **TANUJ BHOJWANI** *is head of People+AI.*

IMF WEEKEND READ

Cutting-edge analysis of global economics, finance, development and policy issues shaping the world delivered straight to your inbox.

SIGN UP TODAY!

GLOBAL FINANCIAL SAFETY NET

In a more shock-prone world, strengthening the financial safety net is more important than ever

IN TIMES OF economic crisis, countries can tap various financial resources, both internal and external. The global financial safety net is a set of institutions and mechanisms that provide insurance for economies against crises to lessen their impact.

This safety net consists of four main layers: countries' own international reserves, bilateral swap lines whereby central banks exchange currencies to provide liquidity to financial markets, regional financing arrangements by which countries pool resources to leverage financing in a crisis, and the IMF.

International reserves are the first line of defense in a crisis; however, because of their high cost, they are unevenly distributed, with most held by advanced economies and larger emerging market economies.

A more efficient way of insuring against crises is through pooled resources, such as the IMF, swap lines, and regional financing arrangements. Although the latter two have grown considerably over the past two decades, they are still available only to a limited group of countries.

This is why the IMF is so important to this system. It is the ultimate global crisis lender and insurer of the uninsured. Yet the IMF's lending capacity as a share of global external liabilities has gradually diminished over time. And the share of borrowed resources has increased.

To continue to play this critical role at the center of the global financial safety net, permanent quota IMF resources need to be boosted. This will bolster the capacity to protect against future crises and, in particular, support members with smaller financial buffers, who need them most. F&D

ANDREW STANLEY *is on the staff of* Finance & Development.

First line of defense

Countries prepare for and respond to shocks by building up foreign exchange reserves, but these reserves are very costly.
(international reserves, trillions of US dollars)

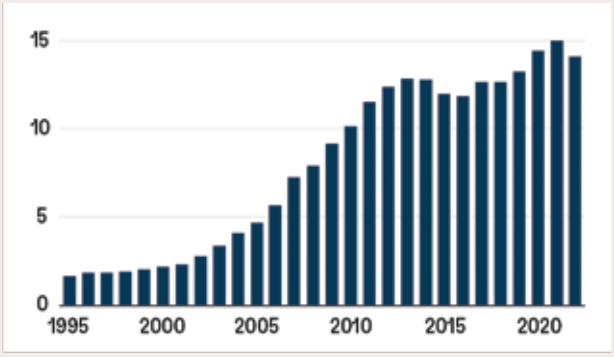

SOURCE: IMF.

97%

of international reserves are held by approximately half the world's economies, with a group of about 90 vulnerable emerging market and low-income countries accounting for the remaining 3%.

International insurance mechanisms

Pooling reserve holdings between countries and drawing on them when needed is more efficient; this is where the other components of the global financial safety net come in.
(trillions of US dollars)

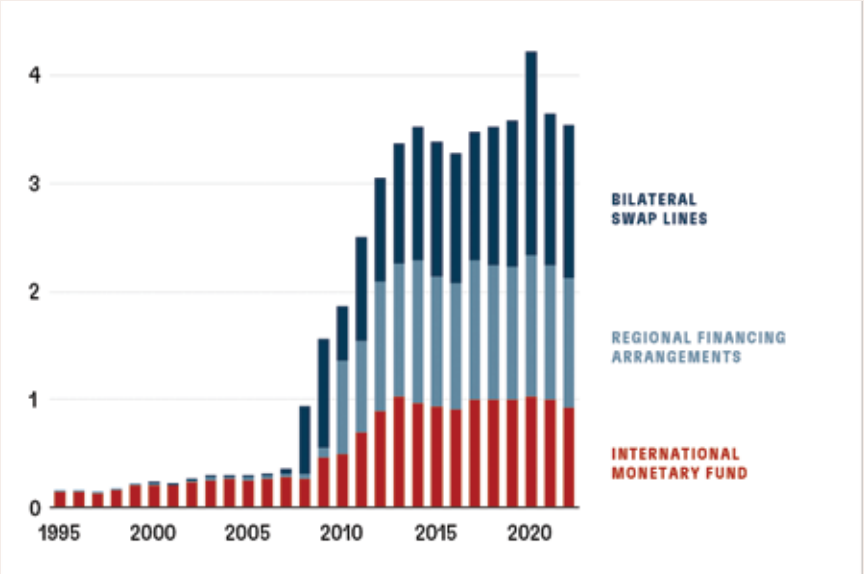

SOURCES: Central bank websites; regional financing arrangements' annual reports; and IMF staff estimates.
NOTE: Since the safety net is composed of various currencies, its US dollar value fluctuates with exchange rates.

Shrinking safeguard

Although international reserves have risen rapidly, that increase in self-insurance has been highly uneven. Poorer countries remain underinsured, leaving them vulnerable to shocks. Meanwhile, the IMF, at the center of the safety net, has shrunk in relation to the total size of global external liabilities and is now far more reliant on resources temporarily borrowed from a few member states. The IMF's traditional quota resources, its permanent capital contributed by all member states, have decreased in relative terms.

Distribution of international reserves

(US dollars, latest available data)

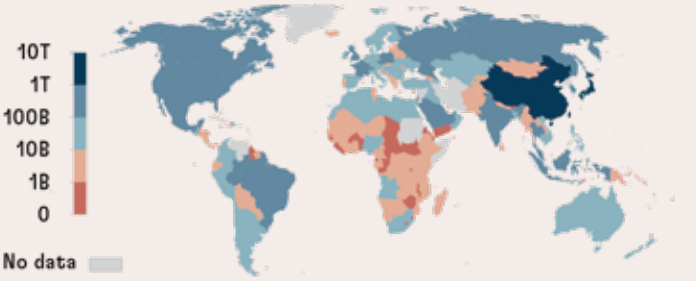

International reserves

(trillions of US dollars)

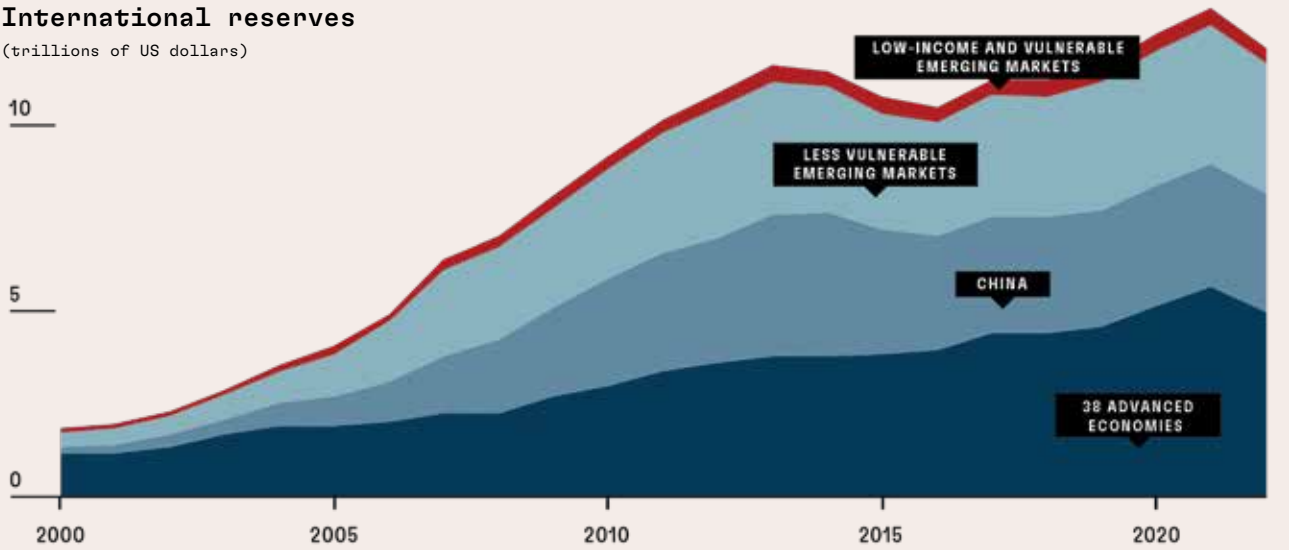

Composition of the global financial safety net, share of total

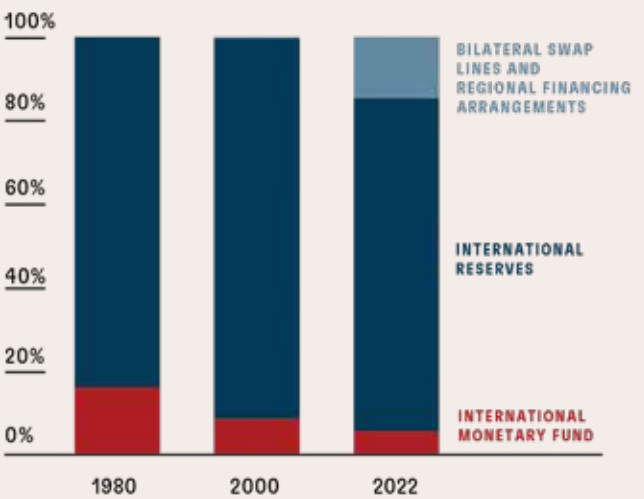

IMF resources relative to global external liabilities

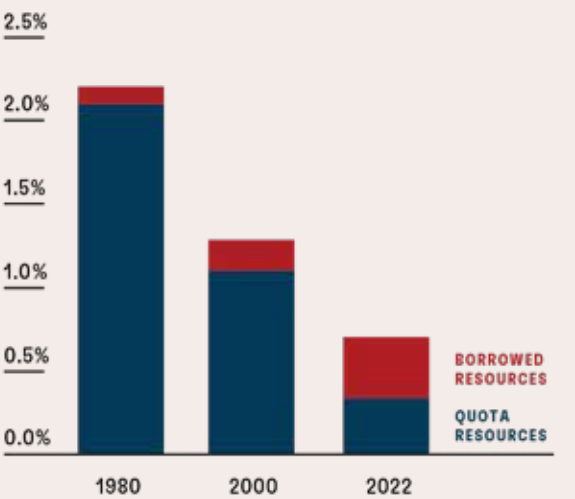

SOURCES: Central bank websites; regional financing arrangements' annual reports; and IMF staff calculations.

NOTE: Reserves chart and map above exclude gold. Data are not available for all economies. The boundaries, colors, denominations, and any other information shown on the maps do not imply, on the part of the IMF, any judgment on the legal status of any territory or any endorsement or acceptance of such boundaries.

THE MACROECONOMICS OF

Illustration by Jun Cen

ARTIFICIAL INTELLIGENCE

The collective decisions we make today will determine how AI affects productivity growth, income inequality, and industrial concentration

Erik Brynjolfsson and Gabriel Unger

Economists have a poor track record of predicting the future. And Silicon Valley repeatedly cycles through hope and disappointment over the next big technology. So a healthy skepticism toward any pronouncements about how artificial intelligence will change the economy is justified. Nonetheless, there are good reasons to take seriously the growing potential of AI—systems that exhibit intelligent behavior, such as learning, reasoning, and problem-solving—to transform the economy, especially given the astonishing technical advances of the past year.

AI may affect society in a number of areas besides the economy—including national security, politics, and culture. But in this article, we focus on the implications of AI on three broad areas of macroeconomic interest: productivity growth, the labor market, and industrial concentration. AI does not have a predetermined future. It can develop in very different directions. The particular future that emerges will be a consequence of many things, including technological and policy decisions made today. For each area, we present a fork in the road: two paths that lead to very different futures for AI and the economy. In each case, the bad future is the path of least resistance. Getting to the better future will require good policy—including

- Creative policy experiments
- A set of positive goals for what society wants from AI, not just negative outcomes to be avoided
- Understanding that the technological possibilities of AI are deeply uncertain and rapidly evolving and that society must be flexible in evolving with them

First fork: Productivity growth

The first road concerns the future of economic growth—which is largely the future of productivity growth. The US economy has been stuck with disturbingly low productivity growth for most of the past 50 years, except for a brief resurgence in the late 1990s and early 2000s (Brynjolfsson, Syverson, and Chad 2019). Most advanced economies now have the same problem of low productivity growth. More than any other factor, productivity—output per unit of input—determines the wealth of nations and the living standards of their people. With higher productivity, such problems as budget deficits, poverty reduction, health care, and the environment become far more manageable. Boosting productivity growth may be the globe's most fundamental economic challenge.

Low-productivity future

On one path of the productivity fork, AI's impact is limited. Despite the rapidly improving technical capabilities of AI, its adoption by businesses may continue to be slow and confined to large firms (Zolas and others 2021). The economics of AI may turn out to be of a very narrow labor-saving variety (what Daron Acemoglu and Simon Johnson call a "so-so technology," such as an automated grocery checkout stand), instead of one that enables workers to do something novel or powerful (see "Rebalancing AI" in this issue of F&D). Displaced workers might disproportionately end up in even less productive and less dynamic jobs, further muting any aggregate benefit to the long-term productivity growth rate of the economy.

Like so many of Silicon Valley's recent technological enthusiasms (3D printers, self-driving cars, virtual reality), AI may also end up being less promising or less ready to bring to market than initially hoped. Any real economic gains, even modest ones, may show up in the data many decades after the first moments of technological promise, as has often been the pattern. The famous paradox identified by economist Robert Solow in 1987—"You can see the computer age everywhere but the productivity statistics"—may become more extreme, as everyone seems to have an AI chatbot that amazes their friends, but businesses do not seem more productive for their increased use of AI. Firms may further blunt any economic benefits from AI by failing to figure out the organizational and managerial changes they need to best leverage it.

And, as in the case of self-driving cars, the technological challenges of going from an exciting proof of concept to a highly reliable product may be further compounded by a legal regime that was not designed to accommodate this new technology and may seriously hinder its development. In the case of AI, there is tremendous uncertainty over what current laws concerning intellectual property imply when models are trained on millions of data points that may include the

"The path that leads to a worse future is the one of least resistance and results in low productivity growth, higher income inequality, and higher industrial concentration."

protected intellectual property of others. Intellectual property law may eventually respond by creating something analogous to a "patent thicket" that effectively prevents models from being trained on data to which the developers do not have clear rights. At the same time, the wrong choices could undermine the incentives of creative professionals to produce more of the novel content that powers machine learning systems.

In addition, national regulators, driven by any number of concerns, may impose strict regulations that slow the speed of AI development and dissemination. They may even be urged on by the early developers of AI who are eager to protect their lead. Moreover, some countries, businesses, and other organizations may totally ban AI.

High-productivity future

But there is an alternate scenario in which AI leads to a higher-productivity-growth future. AI might be applied to a substantial share of the tasks done by most workers (Eloundou and others 2023) and massively boost productivity in those tasks. In this future, AI lives up to its promise of being the most radical technological breakthrough in many decades. Moreover, it ends up complementing workers—freeing them to spend more time on nonroutine, creative, and inventive tasks rather than just replacing them. AI captures and embodies the tacit knowledge (acquired through experience but hard to articulate) of individuals and organizations by drawing on vast amounts of newly digitized data. As a result, more workers can spend more time working on novel problems, and a growing share of the labor force increasingly comes to resemble a society of research scientists and innovators. The result is an economy not simply at a higher level of productivity, but at a permanently higher growth rate.

In this future, the successful integration of AI with robots also means that much more of the economy is amenable to AI-related progress. And AI enables society not just to do better the things it already does but to do things and envision things previously unimaginable. AI-backed research in medicine enables radical advances in knowledge of human biology and drug design. AI becomes capable of helping the engine of creativity and scientific discovery itself—math, science, further AI development—a kind of recursive self-improvement that was once just a science fiction thought experiment.

Second fork: Income inequality

The increase in income inequality between individual workers over the past 40 years is a major concern. A large body of empirical research in labor economics suggests that computers and other forms of information technology may have contributed to income inequality by automating away routine middle-income jobs, which has polarized the labor force into high-income and low-income workers. Although the CEO and the janitor remain, computers have replaced some of the middle tier of office workers (Autor, Levy, and Murnane 2003). We consider two scenarios for AI's effect on inequality.

Higher-inequality future

In the first scenario, AI leads to higher income inequality. Technologists and managers design and implement AI to substitute directly for many kinds of human labor, driving down the wages of many workers. To make matters worse, generative AI starts to produce words, images, and sounds, tasks formerly thought of as nonroutine and even creative—enabling machines to interact with customers and create the content for a marketing campaign. The number of jobs under threat from AI competition eventually grows much larger. Entire industries are upended and increasingly replaced (a threat to labor perhaps foreshadowed by the recent strikes of screenwriters and actors in the United States, who demanded that studios restrict their use of AI).

This is not a future of mass unemployment. But in this higher-inequality future, as AI substitutes for high- or decently paying jobs, more workers are relegated to low-paying service jobs—such as hospital orderlies, nannies, and doormen—where some human presence is intrinsically valued and the pay is so low that businesses cannot justify the cost of a big technological investment to replace them. The final bastion of purely human labor may be these types of jobs with a physical dimension. Income inequality increases in this scenario as the labor market is further polarized into a small, high-skilled elite and a large underclass of poorly paid service workers.

Lower-inequality future

In the second scenario, however, AI leads to lower income inequality because its main impact on the workforce is to help the least experienced or least knowledgeable workers be better at their jobs. Software coders, for instance, now benefit from the assistance of AI models, such as Copilot, which effectively

draw on coding best practices from many other workers. An inexperienced or subpar coder using Copilot becomes more comparable to a very good coder, even when both have access to the same AI. A study of 5,000 workers who do complex customer assistance jobs at a call center found that among workers who were given the support of an AI assistant, the least skilled or newest workers showed the greatest productivity gains (Brynjolfsson, Li, and Raymond 2023). If employers shared these gains with workers, distribution of income would become more equal.

In addition to creating a future of lower income inequality, AI may help labor in another more subtle, but profound, sense. If AI is a substitute for the most routine and formulaic kinds of tasks, then by taking tedious routine work off human hands, AI may complement genuinely creative and interesting tasks, improving the basic psychological experience of work, as well as the quality of output. Indeed, the call center study found not only productivity gains, but reduced worker turnover and increased customer satisfaction for those using the AI assistant.

Third fork: Industrial concentration

Since the early 1980s, industrial concentration—which measures the collective market share of the largest firms in a sector—has risen dramatically in the United States and many other advanced economies. These large superstar firms are often much more capital-intensive and technologically sophisticated than their smaller counterparts.

There are again two divergent scenarios for the impact of AI.

Higher-concentration future

In the first scenario, industrial concentration increases, and only the largest firms intensively use AI in their core business. AI enables these firms to become more productive, profitable, and larger than their competitors. AI models become ever more expensive to develop, in terms of raw computational power—a massive up-front cost that only the largest firms can afford—in addition to requiring training on massive datasets, which very large firms already have from their many customers and small firms do not. Moreover, after an AI model is trained and created, it can be expensive to operate. For example, the GPT-4 model cost more than $100 million to train during its initial development and requires about $700,000 a day to run. The typical cost of developing a large AI model may soon be in the billions of dollars. Executives at the leading AI firms predict that the scaling laws that show a strong relationship between increases in training costs and improved performance will hold for the foreseeable future, giving an advantage to the companies with access to the biggest budgets and the biggest datasets.

It may be, then, that only the largest firms and their business partners develop proprietary AI—as firms such as Alphabet, Microsoft, and OpenAI have already done and smaller firms have not. The large firms then get larger.

More subtly, but perhaps more important, even in a world in which proprietary AI does not require a large fixed cost that only the largest firms can afford, AI might still disproportionately benefit the largest firms, by helping them better internally coordinate their complex business operations—of a kind that smaller and simpler firms do not have. The "visible hand" of top executives managing resources inside the largest firms, now backed by AI, allows the firm to become even more efficient, challenging the Hayekian advantages of small firms' local knowledge in a decentralized market.

Lower-concentration future

In the lower-industrial-concentration future, however, open-source AI models (such as Meta's LLaMA or Berkeley's Koala) become widely available. A combination of for-profit companies, nonprofits, academics, and individual coders creates a vibrant open-source AI ecosystem that enables broad access to developed AI models. This gives small businesses access to industry-leading production technologies they could never have had before.

Much of this was foreshadowed in an internal memo leaked from Google in May 2023, in which a researcher said that "open-source models are faster, more customizable, more private, and pound-for-pound more capable" than proprietary models. The researcher said that processes in small open-source models can be quickly repeated by many people and end up better than large private models that are slowly iterated by a single team and that open-source models can be trained more cheaply. In the Google researcher's view, open-source AI may end up dominating the expensive proprietary models.

It may also be that AI encourages the kind of broad, decentralized innovation that better flourishes across many small firms than within one large firm. The boundaries of the firm are the outcome of a series of tradeoffs; a world in which more AI-backed innovators need

the residual control rights to their work might be one in which more innovators decide they would rather be owners of small firms than be employees of large ones.

The result is that the long rise in industrial concentration starts to run aground, because some nimble smaller businesses close or even reverse the technology gap with their larger counterparts and win back more market share.

Toward a policy agenda

For each of the forks in the road, the path that leads to a worse future is the one of least resistance and results in low productivity growth, higher income inequality, and higher industrial concentration. Getting to the good path of the fork will require hard work—smart policy interventions that help shape the future of technology and the economy.

It is also important to appreciate a broader point about policy. Much of the discourse around AI regulation now takes place along a kind of hydraulic model: should we have more AI or less AI—or even ban AI. This discussion happens when AI is perceived as somewhat of a fixed thing, with a predetermined future. AI can come fast or slow. There can be more or less of it, but basically it is what it is.

However, if policymakers understand that AI can develop in different directions, the discourse will be framed differently. How can policies encourage the types of AI that complement human labor instead of imitating and replacing it? What choices will encourage the development of AI that firms of all sizes can access, instead of just the largest ones? What kind of open-source ecosystem might that require, and how do policymakers support it? How should AI labs approach model development, and how should firms approach AI implementation? How does society get an AI that unleashes radical innovation, instead of marginal tweaks to existing goods, services, and systems?

Many different actors have power to affect the direction of the AI future. Major corporations will have to make important decisions about how they choose to integrate AI into their workforce. The largest of these companies will also develop in-house AI. AI/computer science labs at universities will also develop AI models, some of which they will make open-source. Federal legislators and regulators will have a large impact, as might more local ones. Voters have a voice. Labor unions must figure out what kind of relationship they want with AI and what their demands will be.

Although we have sketched a number of possible futures for AI, we want to emphasize not only how deeply unpredictable the future of this technology is but also the agency society has in actively and collectively determining which AI future emerges.

We have raised more questions than we have answered, which reflects, in part, the nascent stage of AI adoption and impact. But it also reflects a deeper imbalance between research efforts advancing the frontier of the technology and the more limited research aimed at understanding its economic and social consequences.

This imbalance was of less significance when the technology had limited macroeconomic consequences. But today, when the effects of AI on society are likely to be measured in trillions of dollars, far greater investment should be made in research on the economics of AI. Society needs innovations in economic and policy understanding that match the scale and scope of the breakthroughs in AI itself. Reorienting research priorities and developing a smart policy agenda can help society move toward a future of both sustained and inclusive economic growth. F&D

"Society needs innovations in economic and policy understanding that match the scale and scope of the breakthroughs in AI itself."

ERIK BRYNJOLFSSON *is the Jerry Yang and Akiko Yamazaki Professor at the Stanford Institute for Human-Centered AI, where he directs the Stanford Digital Economy Lab.* **GABRIEL UNGER** *is a postdoctoral fellow at the Stanford Digital Economy Lab.*

REFERENCES

Autor, David, Frank Levy, and Richard Murnane. 2003. "The Skill Content of Recent Technological Change." *Quarterly Journal of Economics* 118 (4): 1279–333.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2019. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press.

Brynjolfsson, Erik, Danielle Li, and Lindsay Raymond. 2023. "Generative AI at Work." NBER Working Paper 31161, National Bureau of Economic Research, Cambridge, MA.

Eloundou, Tyna, Sam Manning, Panels Mishkin, and Daniel Rock. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." *arXiv preprint arXiv:2303.10130*.

Zolas, Nicholas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2021. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." NBER Working Paper 28290, National Bureau of Economic Research, Cambridge, MA. https://www.nber.org/papers/w28290.

REBALANCING AI

Daron Acemoglu and Simon Johnson

The drive toward automation is perilous—to support shared prosperity, AI must complement workers, not replace them

A farmer checks tomatoes in a smart greenhouse in Yantai, East China's Shandong Province, January 2022.

Optimistic forecasts regarding the growth implications of AI abound. AI adoption could boost productivity growth by 1.5 percentage points per year over a 10-year period and raise global GDP by 7 percent ($7 trillion in additional output), according to Goldman Sachs. Industry insiders offer even more excited estimates, including a supposed 10 percent chance of an "explosive growth" scenario, with global output rising more than 30 percent a year.

All this techno-optimism draws on the "productivity bandwagon": a deep-rooted belief that technological change—including automation—drives higher productivity, which raises net wages and generates shared prosperity.

Such optimism is at odds with the historical record and seems particularly inappropriate for the current path of "just

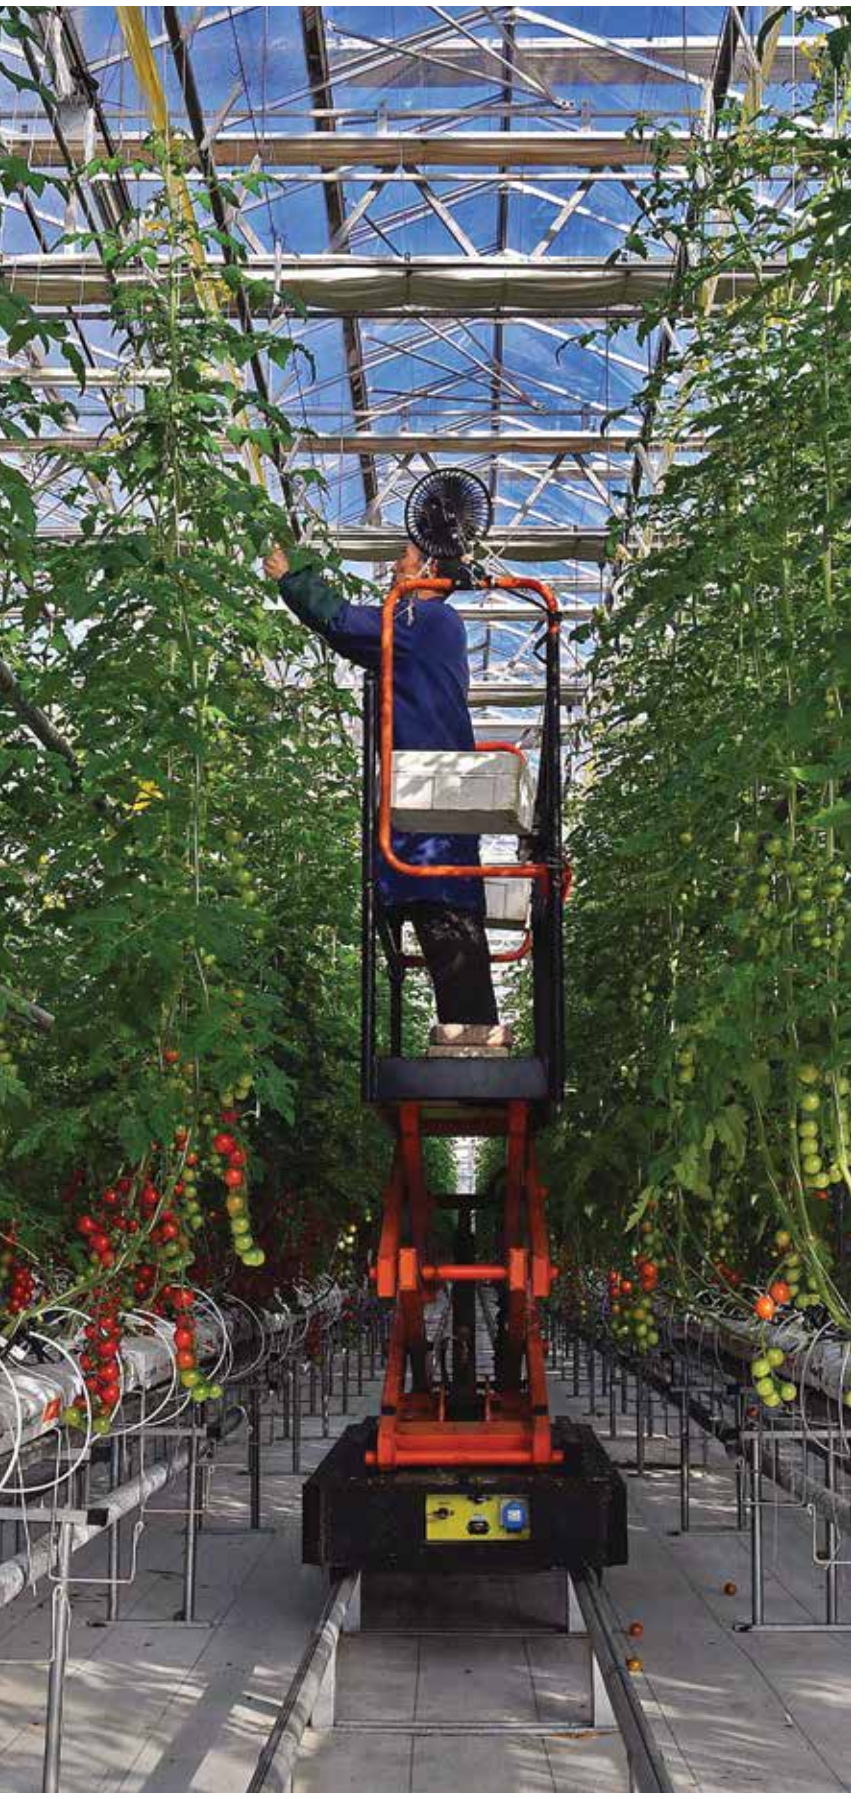

SUN WENTAN/COSTFOTO/FUTURE PUBLISHING VIA GETTY IMAGES

let AI happen," which focuses primarily on automation (replacing people). We must recognize that there is no singular, inevitable path of development for new technology. And, assuming that the goal is to sustainably improve economic outcomes for more people, what policies would put AI development on the right path, with greater focus on enhancing what *all* workers can do?

The machinery question

Contrary to popular belief, productivity growth need not translate into higher demand for workers. The standard definition of productivity is "average output per worker"—total output divided by total employment. The hope is that as output per worker grows, so will the willingness of businesses to hire people.

But employers are not motivated to increase hiring based on average output per worker. Rather, what matters to companies is *marginal productivity*—the additional contribution that one more worker brings by increasing production or by serving more customers. The notion of marginal productivity is distinct from output or revenue per worker; output per worker may increase while marginal productivity remains constant or even declines.

Many new technologies, such as industrial robots, expand the set of tasks performed by machines and algorithms, displacing workers. Automation raises average productivity but does not increase, and in fact may reduce, worker marginal productivity. Over the past four decades, automation has raised productivity and multiplied corporate profits, but it has not led to shared prosperity in industrial countries.

Replacing workers with machines is not the only way to improve economic efficiency—and history has proved this, as we describe in our recent book, *Power and Progress*. Rather than automating work, some innovations boost how much individuals contribute to production. For example, new software tools that aid car mechanics and enable greater precision can increase worker marginal productivity. This is completely different from installing industrial robots with the goal of replacing people.

New functions

The creation of new tasks is even more important for raising worker marginal productivity. When new machines open up new uses for human labor, this expands workers' contributions to production and increases their marginal productivity. There was plenty of automation in car manufacturing during the momentous industry reorganization led by Henry Ford starting in the 1910s. But mass-pro-

duction methods and assembly lines simultaneously introduced a range of new design, technical, machine-operation, and clerical tasks, boosting the industry's demand for workers.

New tasks have been vital in the growth of employment and wages over the past two centuries. And many of the fastest-expanding occupations in the past few decades—those of MRI radiologists, network engineers, computer-assisted machine operators, software programmers, IT security personnel, and data analysts—did not exist 80 years ago. Even people in occupations that have been around longer, such as bank tellers, professors, and accountants, now work on many relatively new tasks using technology. In almost all these cases, new tasks were introduced because of technological advances and have been a major driver of employment growth. These new tasks have also been integral to productivity growth—they have helped launch new products and enabled more efficient production processes.

Productive automation

Automation in an industry can also drive up employment—in that sector or in the economy broadly—if it substantially increases productivity. In this case, new jobs may come either from nonautomated tasks in the same industry or from the expansion of activities in related industries. In the first half of the 20th century, the rapid increase in car manufacturing stimulated massive expansion of the oil, steel, and chemical industries. Vehicle output on a mass scale also revolutionized the possibilities for transportation, enabling the rise of new retail, entertainment, and service activities.

The productivity bandwagon is not activated, however, when the productivity gains from automation are small—what we call "so-so automation." For example, self-checkout kiosks in grocery stores bring limited productivity benefits because they merely shift the work of scanning items from employees to customers. When stores introduce self-checkout kiosks, fewer cashiers are employed, but there is no major productivity boost to stimulate the creation of new jobs elsewhere. Groceries do not become much cheaper, there is no expansion in food production, and shoppers do not live differently.

Even nontrivial productivity gains from automation can be offset when they are not accompanied by new tasks. For example, in the American Midwest, the rapid adoption of robots has contributed to mass layoffs and ultimately prolonged regional decline.

The situation is similarly troubling for workers when new technologies focus on surveillance. Increased monitoring of workers may lead to some small improvements in productivity, but its main function is to extract more effort from workers.

All this underscores perhaps the most important aspect of technology: *choice*. There are often myriad ways of using our collective knowledge to improve production and even more ways to direct innovation. Will we invent and implement digital tools for surveillance, automation, or to empower workers by creating new productive tasks?

When the productivity bandwagon is weak and there are no self-correcting mechanisms to ensure shared benefits, these choices become more consequential—and a few tech decision-makers become economically and politically more powerful.

Complementing humans

New technology may complement workers by enabling them to work more efficiently, perform higher-quality work, or accomplish new tasks. For example, even as mechanization gradually pushed more than half of the US labor force out of agriculture, a range of new blue-collar and clerical tasks in factories and newly emerging service industries generated significant demand for skilled labor between about 1870 and 1970. This work was not only better paying but also less dangerous and less physically exhausting.

"AI offers an opportunity to complement worker skill and expertise if we direct its development accordingly."

This virtuous combination—automation of traditional work alongside creation of new tasks—proceeded in relative balance for much of the 20th century. But sometime after approximately 1970, this balance was lost. While automation has maintained its pace or even accelerated over the ensuing five decades, the offsetting force of new task creation has slowed, particularly for workers without four-year college degrees. As a result, these workers are increasingly found in low-paying (though socially valuable) services such as in cleaning, food service, and recreation.

The critical question of the new era of AI is whether this technology will primarily accelerate the existing trend of automation without the offsetting force of good job creation—particularly for non-college-educated workers—or whether it will instead enable the introduction of new labor-complementary tasks for workers with diverse skill sets and a wide range of educational backgrounds.

It is inevitable that AI systems will be used for some automation. A major barrier to automation of many service and production tasks has been that they require flexibility, judgment, and common sense—which are notably absent from pre-AI forms of automation. Artificial intelligence, especially generative AI, can potentially master such tasks.

It is unclear how much this type of automation will contribute to aggregate productivity growth while these technologies are immature, but they could contribute to sizable productivity gains as costs fall and reliability improves.

The dominant intellectual paradigm in today's digital tech sector also favors the automation path. A major focus of AI research is to attain human parity in a vast range of cognitive tasks and, more generally, to achieve artificial general intelligence that mimics and surpasses human capabilities. This intellectual focus encourages automation rather than the development of human-complementary technologies.

However, AI offers an opportunity to complement worker skill and expertise if we direct its development accordingly.

Human productivity is often hampered by lack of specific knowledge or expertise, which could be supplemented by next-generation technology. For example, AI holds great potential for training and retraining expert workers, such as educators, medical personnel, and those in modern crafts (such as electricians and plumbers). AI could also create new demands for human expertise and judgment in overseeing these processes, communicating with customers, and enabling more sophisticated services.

Five principles

Redirecting technological change is not easy, but it is possible. Governments everywhere—especially in the US and other countries where technology is under active development—should take the following five steps to help put AI development onto a human-complementary, rather than human-displacing, path:

- ***Reform business models:*** The dominant developers of AI easily expropriate consumer data without compensation, and their reliance on digital advertising incentivizes grabbing consumers' attention through any means possible. Governments need to establish clear ownership rights for all consumers over their data and should tax digital ads. Enabling a more diverse range of business models—or even requiring more competition—is essential if AI is to be helpful to all humans.
- ***Tax system:*** The tax code in the US and many other countries places a heavier burden on firms that hire labor than on those that invest in algorithms to automate work. To shift incentives toward human-complementary technological choices, policymakers should aim to create a more symmetric tax structure, equalizing marginal tax rates for hiring (and training) labor and for investing in equipment and software.
- ***Labor voice:*** Given that workers will be profoundly affected by AI, they should have a voice in its development. Government policy should restrict deployment of untested (or insufficiently tested) AI for applications that could put workers at risk, for example in high-stakes personnel decision-making tasks (including hiring and termination) or in workplace monitoring and surveillance.
- ***Funding for more human-complementary research:*** Research and development in human-complementary AI technologies require greater support. Governments should foster competition and investment in technology that pairs AI tools with human expertise to improve work in vital social sectors. Once there is sufficient progress, governments can encourage further investment with advice on whether purported human-complementary technology is appropriate for adoption in publicly funded education and health care programs.
- ***AI expertise within government:*** AI will touch every area of government investment, regulation, and oversight. Developing a consultative AI division within government can help agencies and regulators support more timely, effective decision-making.

Potential macroeconomic impact

AI could increase global GDP over the next five years, although not as substantially as enthusiasts claim. It might even modestly raise GDP growth in the medium term. However, on our current trajectory, the first-order impact is likely to be increased inequality within industrial countries.

Middle-income countries and many lower-income countries also have much to fear from the existing path. New capital-intensive technology will soon be applied everywhere. There is no guarantee that, on its current path, AI will generate more jobs than it destroys.

If we can redirect AI onto a more human-complementary path, while using it to address pressing social problems, all parts of the planet can benefit. But if the just-automate approach prevails, shared prosperity will be even harder to achieve. F&D

DARON ACEMOGLU *is Institute Professor at the Massachusetts Institute of Technology (MIT).*
SIMON JOHNSON *is the Ronald A. Kurtz Professor of Entrepreneurship at MIT Sloan and a former IMF chief economist.*

This article is adapted from the authors' book, Power and Progress: Our 1000 Year Struggle over Technology and Prosperity, *and also draws on joint work with David Autor.*

SCENARIO PLANNING FOR AN A(G)I FUTURE

Anton Korinek

AI may be on a trajectory to surpass human intelligence; we should be prepared

ARUN SANKAR/AFP VIA GETTY IMAGES

Artificial intelligence is rapidly advancing, and the pace of progress has accelerated in recent years. ChatGPT, released in November 2022, surprised users by generating human-quality text and code, seamlessly translating languages, writing creative content, and answering questions in an informative way, all at a level previously unseen.

Yet in the background, the foundation models that underlie generative AI have been advancing rapidly for more than a decade. The amount of computational resources (or, in short, "compute") used to train the most cutting-edge AI systems has doubled every six months over the past decade. What today's leading generative AI models can do was unthinkable just a few years ago: they can deliver significant productivity gains for the world's premier consultants, for programmers, and even for economists (Korinek 2023).

Conjecture about AI acceleration

Recent advances in artificial intelligence have prompted leading researchers to project that the pace of current progress may not only be sustained but may even accelerate in coming years. In May 2023, Geoffrey Hinton, a computer scientist who laid the theoretical foundations of deep learning, described a significant shift in his perspective: "I have suddenly switched my views on whether these things are going to be more intelligent than us." He conjectured that artificial general intelligence (AGI)—AI that possesses the ability to understand, learn, and perform any intellectual task a human being can perform—may be realized within a span of 5 to 20 years.

Some AI researchers are skeptical. These divergent perspectives reflect tremendous uncertainty about the speed of future progress, whether progress is accelerating or may eventually plateau. In addition, we face significant uncertainty about the broader economic implications of advances in AI and the prospective ratio of benefit to harm from increasingly sophisticated AI applications.

At a fundamental level, the uncertainty also relates to profound questions about the nature of intelligence and the capabilities of the human brain. Chart 1 shows two competing perspectives on the complexity distribution of work tasks the human brain can perform.

Panel 1 illustrates one perspective, that the capabilities of the human brain in solving ever more complex tasks are unbounded. This aligns with our economic experience since the Industrial Revolution: as the frontier of automation advances, humans have automated simple tasks (both mechanical and cognitive) and reallocated workers to perform more of the remaining more complex tasks—that is, they have moved into the right tail of the complexity distribution illustrated in the chart. Straightforward extrapolation would suggest that this process will continue as AI advances and automates a growing number of cognitive tasks.

Another perspective, illustrated in panel 2 of Chart 1, holds that there is an upper bound to the complexity of tasks the human brain can perform. Information theory suggests that the human brain is a computational entity, constantly processing a plethora of data. The brain's inputs include sensory perceptions—sights, sounds, and tactile sensations, among others—and its outputs manifest as physical actions, thoughts, and emotional responses. Even complex facets that make us human, such as emotions, creativity, and intuition, can be viewed as computational outputs, emerging from intricate interactions of neural circuits and biochemical reactions. Although these processes are highly elaborate and involve complexities we do not fully understand, this perspective suggests that there is a definitive upper limit to the intricacy of tasks the human brain can perform.

The two perspectives have dramatically different implications for the potential scope of future automation. As of 2023, the human brain is the most advanced computing device when it comes to the ability to perform a broad range of intellectual tasks in a robust manner. However, if the second perspective turns out to be correct, modern AI systems are catching up fast. In fact, many measures of the computational complexity of cutting-edge foundation models are already close to those of the human brain. The computational complexity of human brains is bounded by biology, and the brain's ability to transmit information to other intelligent entities (humans or AI) is limited by the slow speed of information transmission of our senses and our language. Nevertheless, AI systems continue to advance rapidly and can exchange information at speeds that are significantly faster.

Preparing for multiple scenarios

Economists have long observed that the optimal way of dealing with uncertainty is to use a portfolio approach. Given the starkly differing perspectives on future progress in AI by world-renowned experts, it would be unwise to put all eggs in one

Robot waiters carry food for customers at a robot-themed restaurant in Chennai, India.

basket and formulate economic plans for a single scenario. Instead, the uncertainty about what the future will look like should motivate us to hedge our bets and engage in careful analysis of a range of different scenarios that may materialize, from business as usual to the possibility of AGI. Aside from doing justice to the prevailing level of uncertainty, scenario planning makes the potential opportunities and risks tangible and helps us to develop contingency plans and be prepared for multiple possible outcomes.

Following are three technological scenarios spanning a wide range of possible outcomes that economic policymakers should pay attention to:

Scenario I (traditional, business as usual): Advances in AI boost productivity and automate a range of cognitive work tasks, but they also create new opportunities for affected workers to move into new jobs that are, on average, more productive than those from which they were displaced. This view is encapsulated by panel 1 of Chart 1.

Scenario II (baseline, AGI in 20 years): Over the next 20 years, AI gradually advances to the point of AGI, resulting in its ability to perform all human work tasks by the end of the period, devaluing labor (Susskind, forthcoming). This would correspond to the perspective of finite brainpower captured by panel 2 of Chart 1, together with the assumption that it would take 20 years for the most complex cognitive tasks to be accessible to AI.

Scenario III (aggressive, AGI in five years): This scenario replicates Scenario II but on a more aggressive timeline, such that AGI with all the associated consequences for labor would be reached within five years.

Although I am highly uncertain, at the time of writing, I estimate that each of these scenarios has a greater than 10 percent probability of materializing. To account for the uncertainty and adequately prepare for the future, I believe that policymakers should take each of these scenarios seriously, stress-test how our economic and financial policy frameworks would perform in each scenario, and where necessary reform them to ensure that they would be adequate.

The three scenarios have the potential to lead to markedly different economic outcomes across a wide range of indicators, including economic growth, wages and returns to capital, fiscal sustainability, inequality, and political stability. Moreover, they call for reforms to our social safety nets and systems of taxation and affect the conduct of monetary policy, financial regulation, and industrial and development strategies.

Korinek and Suh (2023) analyze the implications of the scenarios described for output and wages in a

CHART 1

AI pushing the limits

Automation has already taken over many tasks that previously only humans could perform, as indicated by the gray bars. As AI advances, human beings may be able to continue to move into more complex tasks indefinitely (yellow bars), or AI may eventually overtake all human capabilities (blue bars).

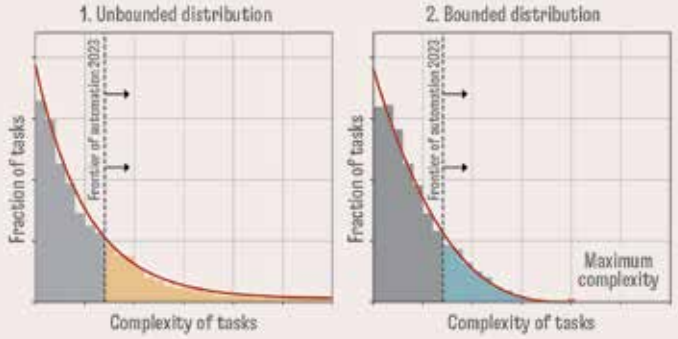

SOURCE: Anton Korinek.

CHART 2

Scenarios for output and wages

The advent of AGI leads to a takeoff in growth at the expense of plummeting competitive market wages.

(initial year = 100)

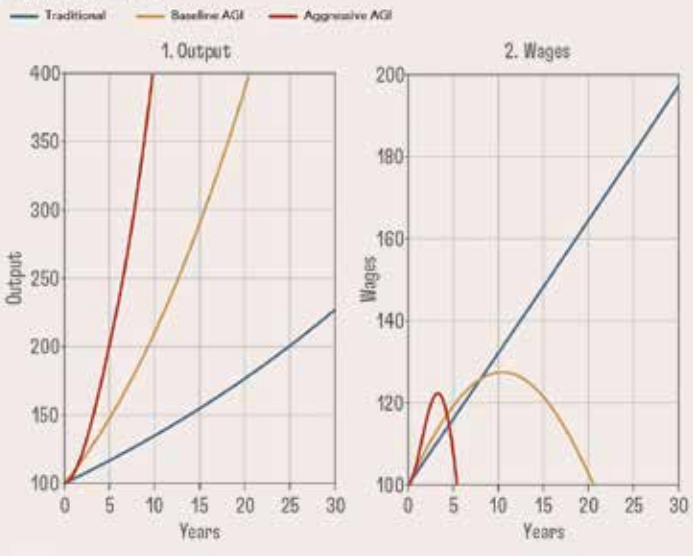

SOURCE: Anton Korinek.
NOTE: AGI = artificial general intelligence.

mainstream macroeconomic model of automation. The results for all three scenarios are illustrated in Chart 2, in which the path of output for each scenario is displayed in panel 1 and the path of competitive market wages in panel 2.

Three main insights stand out:

First, whereas growth continues along the trajectory we are used to from past decades in the conservative business-as-usual scenario, output growth in the two AGI scenarios is much faster, as the scarcity of labor is no longer a constraint on output.

Second, wages initially rise in all three scenarios—but only as long as labor is scarce. They plummet as the economy is close to reaching AGI.

Third, the takeoff in output and the collapse in wages in the two AGI scenarios are both driven by the same force: the substitution of scarce labor by comparatively more abundant machines. This suggests that it should be possible to design institutions that compensate workers for their income losses and ensure that the gains from AGI lead to shared prosperity.

Chart 2 illustrates the broad contours of how unprecedented technological changes may affect the macroeconomy, but it is best understood as an illustration of possibilities rather than as a precise prediction. A long list of caveats applies. First, the model underlying the chart is cast in an efficient economy in which labor earns competitive returns. A range of factors may slow the rollout of AGI compared with what is technologically possible, from organizational frictions, regulations, and constraints on capital accumulation—such as chip supply chain bottlenecks—to societal choices on the implementation of AGI. Even when it is technologically possible to replace workers, society may choose to keep humans in certain functions—for example, as priests, judges, or lawmakers. The resulting "nostalgic" jobs could sustain demand for human labor in perpetuity (Korinek and Juelfs, forthcoming).

"Given the starkly differing perspectives on future progress in AI, it would be unwise to put all eggs in one basket and formulate economic plans for a single scenario."

To determine which AI scenario the future most resembles as events unfold, policymakers should monitor leading indicators across multiple domains, keeping in mind that all efforts to predict the pace of progress face tremendous uncertainty. Useful indicators span technological benchmarks, levels of investment flowing into AI development, adoption of AI technologies throughout the economy, and resulting macroeconomic and labor market trends. Technological benchmarks offer the most direct measure of how well AI systems perform a wide range of labor tasks. Levels of investment, such as investment in research and development, talent, and computer chips, capture how much of our resources are flowing into AI development. Indicators of growing AI adoption through all sectors of the economy would capture whether the resulting systems are usefully deployed in practice. Finally, the macroeconomic implications would eventually become visible in productivity statistics and labor market trends.

Tracking these complementary signals allows policymakers to tailor policy responses to the realities of AI as they manifest. But we must remain humble—the future is likely to surprise us.

The starkly different economic trajectories implied by the three scenarios described earlier underscore the importance of developing adaptive policy frameworks that can respond nimbly as the future unfolds. Policymakers should stress-test existing institutions against each scenario and reform them where necessary to ensure they are resilient. This may involve gradual steps, such as reforming systems of taxation and expanding social safety nets, or new programs, such as introducing small basic incomes that can be scaled up when necessary.

Policymakers should charge teams of experts with iterative scenario planning to help them regularly update their views on how the probabilities of the various scenarios evolve. Embracing the uncertainty through an adaptable, scenario-based approach will allow us to maximize the benefits and mitigate the risks in the economic sphere from AI's continuing evolution. F&D

ANTON KORINEK *is a professor in the Department of Economics and Darden School of Business at the University of Virginia. He is economics of AI lead at the Centre for the Governance of AI, a research associate at the National Bureau of Economic Research, and a research fellow at the Centre for Economic Policy Research.*

REFERENCES

Korinek, Anton.Korinek, Anton. 2023. Generative AI for Economic Research: Use Cases and Implications for Economists. *Journal of Economic Literature* 2023, 61(4).

Korinek, Anton, and Megan Juelfs. Forthcoming. "Preparing for the (Non-existent?) Future of Work." *Oxford Handbook of AI Governance*. Oxford, UK: Oxford University Press.

Korinek, Anton, and Donghyun Suh. 2023. "Scenarios for the Transition to AGI." University of Virginia working paper, Charlottesville, VA.

Susskind, Daniel. Forthcoming. "Technological Unemployment." *Oxford Handbook of AI Governance*. Oxford, UK: Oxford University Press.

TECHNOLOGY'S BIFURCATED BITE

Andrew Berg, Chris Papageorgiou, and Maryam Vaziri

Some workers will win, others will lose as the use of artificial intelligence grows

Technological developments—such as factory robots, smart home devices, and self-driving cars—transform the way we live and work. Such developments are exciting in many ways, because they promise higher productivity and standards of living. But they can also be frightening: when the machines take over, how will many people make a living?

This is an old question, of course. Fears about technology destroying jobs, displacing workers, and damaging lifestyles arose during the Industrial Revolution—best exemplified, perhaps, by the Luddites in England, who fought life-altering changes in the textile industry. These fears persist today. As then-US Senator John F. Kennedy said in 1960, at the dawn of the computer revolution, "Today we stand on the threshold of a new industrial revolution—the revolution of automation. This is a revolution bright with the hope of a new prosperity for labor and a new abundance for America, but it is also a revolution which carries the dark menace of industrial dislocation, increasing unemployment, and deepening poverty."

In retrospect, Kennedy's concern about lost jobs seems misplaced. In the years after his speech, the US economy created millions of net new jobs, and mass technological unemployment did not emerge—as demonstrated by today's unemployment rate of about 3.5 percent and the multi-decade-high ratio of employment to population.

These labor market developments would seem to assuage the concerns of a modern-day Luddite: with the benefits of technology and the power of the market, people will find new jobs, and rising productivity will raise living standards—which ultimately happened during the Industrial Revolution of the 18th and 19th centuries. Indeed, the standard of living has increased enormously since 1900. Technologies such as electricity, internal combustion engines, telephones, and modern medicine have improved the quality of life and increased life expectancy.

That is not to say, however, that Kennedy's concerns were unfounded. Only a few years after his speech, wage inequality began to worsen sharply (see Chart 1), and the share of income going to workers fell.

HENRY NICHOLLS/REUTERS

A person holds a sign during a rally in support of striking warehouse workers in Coventry, England, January 2023.

Economists have developed frameworks for thinking about the implications of artificial intelligence (AI)—which simulates human intelligence in machines—and, more generally, the impact of technological change, automation, and robots on inequality. In this respect we will highlight four key channels that affect inequality:

- Technological change that improves the productivity of skilled more than unskilled workers
- Reductions in the cost of capital that complement chiefly skilled labor
- Increased ability of machines to replace workers entirely for particular tasks
- Increased concentration of market power in a few firms as a result of technology

Regarding the first channel, Katz and Murphy (1992) explained the evolution of relative wages in the United States as the outcome of a race between increases in the demand and supply of skilled workers. They focused on aggregate productivity and factor-augmenting technological change. Increases in the supply of skilled workers reduced the skill premium, whereas persistent increases in demand for such workers had the opposite effect. These forces explain both the dip in the skill premium in the early 1970s—when the supply of educated workers rose sharply because more people went to college—and the rise in skill premiums after the 1980s.

In the second channel, capital, especially machinery and equipment, tends to *complement* skilled workers and *substitute* for unskilled workers—for example, machine tools require more programmers but replace other workers in factories. Berg, Buffie, and Zanna (2018) extend this approach to look at AI and robots as a new type of capital—additional to traditional machinery and structures—that substitutes for some groups of workers and complements others. Over the past 30 years the substitutability between information and communications technologies (ICT)—a proxy for new technologies, including computers and early AI—and unskilled workers seems to have increased (see Chart 2). In other words, ICT capital apparently is now better able to perform the tasks of unskilled workers.

CHART 1

Losing ground

Workers' share of total income has eroded.

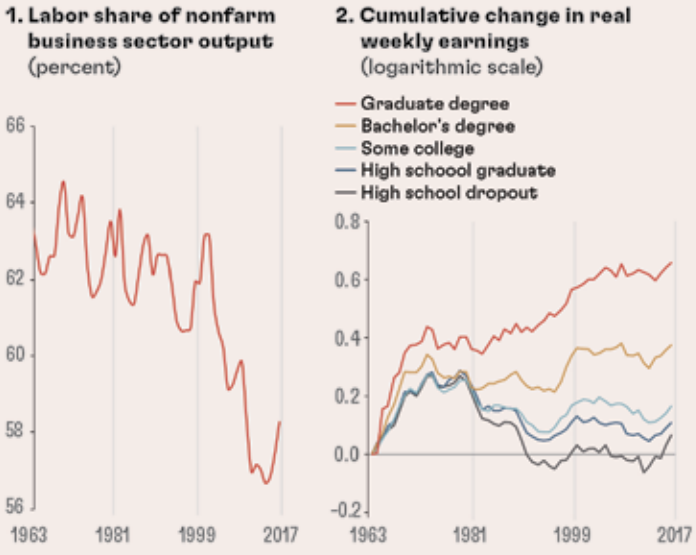

SOURCES: Panel 1: US Bureau of Labor Statistics; panel 2: Autor (2019).
NOTE: Panel 2 depicts earnings of working-age adults (ages 18–64). The values show log differences and are approximately equal to the percent change in average real weekly earnings compared with the base year (1963).

CHART 2

Squeezed out

The ability of information and communications technologies to replace unskilled workers has steadily increased over the past three decades.

(elasticity of substitution of technology to replace unskilled workers)

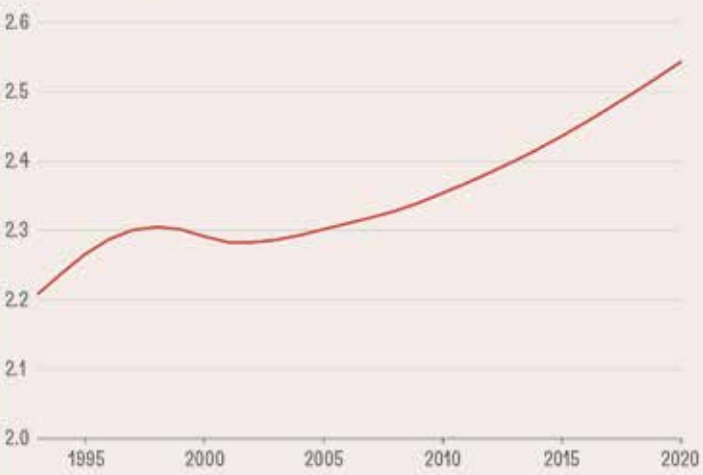

SOURCE: Berg and others (forthcoming).
NOTE: The higher the elasticity of substitution, the more technology can be substituted for workers.

The higher substitutability of workers with machines and AI increases wage inequality and the share of total income that goes to the owners of capital—raising the question of how the benefits of AI technologies should be distributed or, put differently, who owns AI. In the long run, society may well be better off with the higher overall productivity that ensues, but there would be many losers, concentrated among those already less well off. And during a possibly decades-long transition, many could see real wage declines.

Acemoglu and Restrepo (2020) point out that technology has increasingly replaced workers in routine tasks, even as it has enhanced the creativity of other workers' roles. The race between these new creative tasks and automation of routine tasks affects the demand for different types of workers and ultimately determines wages and overall productivity. Acemoglu and Restrepo (2020) show that exposure of different labor groups to automation explains most changes in relative wages—without much of a role for skill-based technological change or for foreign trade and outsourcing-related replacement of workers.

A fourth dimension of technological change extends beyond the labor market to firms' market power. Corporations such as Alphabet and Microsoft clearly dominate leading AI technologies. Developing these technologies is costly and depends heavily on big data, to which only a few firms have access. Yet it also means that as owners of the AI capital those few firms will take a larger slice of the pie. As they rent their technologies to firms in other industries, labor share will continue to fall, while the income from AI technologies will increase.

But the implications of corporate market power are not limited to owning AI. So far, we have discussed technological change as a process that happens naturally. In reality, however, companies innovate, and their innovations shape both the speed of growth and the kinds of new technologies that emerge. Once firms are big enough, they can purchase and bury possible competitors—potentially stifling competition, limiting innovation, and worsening inequality.

Moreover, large corporations with access to leading AI technologies may be able to influence the regulatory framework to align with their interests and to direct innovation toward corporate goals rather than social welfare. For example, Acemoglu and Restrepo (2022) note that the automation observed in recent decades may have been the sort that displaces workers without producing much in the way of overall productivity growth. They show that machines can displace workers without being all that much better at the relevant tasks. Moreover, higher inequality and lower labor share in income may be permanent features, and any transition could be very difficult. The

short run could be a lifetime for some workers (Berg, Buffie, and Zanna 2018).

The First Industrial Revolution reflected both the optimistic long-term and worrisome short-term perspectives. Few would want to give up the benefits from earlier industrial revolutions—from indoor toilets to cell phones—but the transition was both economically and politically wrenching. Carl Benedikt Frey argues in *The Technology Trap* that for certain "vulnerable" groups three entire generations were worse off as a result. Joseph Stiglitz argues in the December 6, 2011, issue of *Vanity Fair* that the technology-driven transition from agriculture to manufacturing in the 1920s set the stage for the Great Depression. More recently, the distributional implications of technological change are arguably an important factor in the rise of populism and anti-globalization sentiment.

AI is rapidly evolving in unpredicted directions—perhaps making it impossible to draw any historical lessons. The early 2023 emergence of ChatGPT-4—an AI model that seeks to generate human-like language—marks a significant acceleration in the pace of change, highlighting AI's ability to extend far beyond routine tasks. Experts in AI surveyed by McKinsey in 2019 expected computers to be able to write at the level of the top 25 percent of humans by 2050 and perform human-level creative tasks by 2055. However, they have revised their estimates to 2024 and 2028, respectively.

It is easy to see why the projections have changed so sharply. Generative pretrained transformers (GPT) seem to have the potential for widespread labor market impact—one estimate suggests that once GPT is introduced into the work environment, about 20 percent of workers could see at least half of their tasks affected. GPT seems to increase productivity in more creative tasks, such as writing, legal analysis, and programming. These studies compare the productivity of groups using GPT with a control group in the given task and find big jumps in productivity with GPT. Just as remarkable, though, is the observation that the least-skilled participants benefit most and that at least in some cases the GPT-augmented input is more creative; moreover, there are signs that GPT-4 alone may exceed human-level output. These findings contrast with earlier emphasis on the automation of *routine* tasks and the substitution of AI and robots for *unskilled* labor. Such shifts in the impact of new technologies on skilled and low-skilled workers seem to be a key difference between GPT and previous waves of technology, such as digitalization.

All this suggests major implications for both growth and inequality, but it also suggests that the past may not be prologue. Will some wage inequality be reversed as lower-skilled workers benefit more? Or will major corporations—with the best access to data, computers, and top talent—gain more economic and political power? The so-far-hypothetical prospect of artificial general intelligence (AGI) adds another dose of uncertainty. AGI would presumably be capable of any human intellectual effort. How all this would play out will clearly depend on both the evolution of the technology and the policy and the broader societal response. There are optimistic and pessimistic AI scenarios, but under any of them, economic, social, and political upheaval seems a safe prediction, and policymakers must do their best to understand the distributional implications of the rapid changes that are underway.

As we navigate the transition to widespread use of AI, it is crucial to acknowledge the global implications of AI technologies—which so far have not been largely studied. Previous research suggested that the substitution of AI for unskilled labor could widen global income disparities, putting lower-income countries at a disadvantage (Alonso and others 2022). But the advent of generative AI suggests that the impact of these technologies on different countries is uncertain. Developing economies may benefit from AI as a tireless universal tutor and expert programming assistant that strengthens their workforces. Conversely, limited access to data and expertise and technological gaps could widen divergence. F&D

DATA

25%

Computers may be able to write at the level of the top 25 percent of humans by 2024, according to AI experts surveyed by McKinsey.

ANDREW BERG *is deputy director of the IMF's Institute for Capacity Development, where* **MARYAM VAZIRI** *is an economist.* **CHRIS PAPAGEORGIOU** *is a division chief in the IMF's Research Department.*

REFERENCES

Acemoglu, D., and P. Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128 (6): 2188–244.

Acemoglu, D., and P. Restrepo. 2022. "Demographics and Automation." *Review of Economic Studies* 89 (1): 1–44.

Autor, D. H. 2019. "Work of the Past, Work of the Future." In *AEA Papers and Proceedings,* vol. 109, 1–32. Nashville, TN: American Economic Association.

Berg, Andrew, Edward F. Buffie, Mariarosaria Comunale, Chris Papageorgiou, and Luis-Felipe Zanna. Forthcoming. "Searching for Wage Growth: Policy Responses to the AI Revolution." IMF Working Paper, International Monetary Fund, Washington, DC.

Berg, A., E. F. Buffie, and L. F. Zanna. 2018. "Should We Fear the Robot Revolution? (The Correct Answer Is Yes)." *Journal of Monetary Economics* 97: 117–48.

Katz, L. F., and K. M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1) 35–78.

AI'S REVERBERATIONS ACROSS FINANCE

Jeff Kearns

Financial institutions are forecast to double their spending on AI by 2027

Artificial intelligence tools and the people to use them are the new must-haves for the world's financial institutions and central banks.

In June 2023, JPMorgan Chase & Co. had 3,600 AI help-wanted postings, according to Evident Insights Ltd., a London-based start-up tracking AI capabilities across financial services companies.

"There's a war for talent," said Alexandra Mousavizadeh, the founder of Evident Insights. "Making sure you are ahead of it now is really life and death."

Like other technological breakthroughs, AI offers fresh potential—accompanied by novel risks. The financial services industry may be among the biggest beneficiaries of the technology, which may enable them to better protect

MATT CHINWORTH

assets and predict markets. Or the sector may have the most to lose if AI spurs theft, fraud, cybercrime, or even a financial crisis that investors can't conceive of today.

The debut of OpenAI's ChatGPT in November 2022 is rippling through finance and other industries. It quickly topped 100 million users to become the fastest-growing application in internet history.

In finance, the demand for people who know how to tap into AI is global. Three of the top 10 cities in Evident's talent index are in India, said Mousavizadeh, an economist, mathematician, and former cohead of country risk at Morgan Stanley.

Financial embrace

The money flowing into AI from financial and other enterprises underscores the new priorities. Sales of software, hardware, and services for AI systems will climb 29 percent this year to $166 billion and top $400 billion in 2027, according to International Data Corp. Financial sector spending will more than double to $97 billion in 2027, with a 29 percent compound annual growth rate—the fastest of five major industries—according to the market researcher.

Hedge funds, long the pioneers of cutting-edge tech, are embracing generative AI. Nearly half of them use ChatGPT professionally, and more than two-thirds of those use it to write marketing text or summarize reports or documents, according to a BNP Paribas survey of funds with $250 billion in combined assets.

Investment businesses are using and investigating AI's potential across various business lines. Europe's largest investment company, Amundi SA, is building out its own AI infrastructure for research on macroeconomics and markets. It's also using the technology for applications such as robo-advising tools for individual customers.

Paris-based Amundi, with €2 trillion ($2.1 trillion) under management, uses AI-based tools to customize portfolios for some of its more than 100 million clients by asking their preferences about risk. Responses help shape portfolios and provide a real-time sentiment gauge.

Aggregate view

"This kind of algorithm allows us to see the behavior of the clients," said Monica Defend, chief strategist at the Amundi Investment Institute, the company's research and strategy unit. "There's a benefit to the customer, but you also have an aggregate view of how attitudes are changing across this user base."

> "AI tools may exacerbate a crisis, whatever the cause, because they are trained on past data that may not reflect reality in an unprecedented situation."

For other uses, such as making institutional decisions on investment and trading, AI can be limited by data that proves unreliable or by unprecedented high-impact situations, she said. It's also a priority to avoid abuses and ensure that AI is used within a secure, ethical, and compliant framework.

"Artificial intelligence cannot replace the brain," Defend said. A wholly AI-driven process could be dangerous, she said. "It's equally important the interpretation, the understanding, and the check of what the algorithms are providing."

JPMorgan, the largest US lender, spends more than $15 billion a year on technology, where it deploys almost a fifth of its approximately 300,000 employees. An AI research group employs 200, and AI enables hundreds of uses ranging from prospecting and marketing to risk management and fraud prevention. AI also runs across payment processing and money movement systems worldwide.

"It is an absolute necessity," Chief Executive Officer Jamie Dimon told shareholders in April.

Monetary world

Much more is at stake for policymakers safeguarding economies. Central banks, which by design are slower-moving and more risk-averse, are learning to use AI in a much different context—and weighing potential risks.

AI has shown promise in a range of central bank applications, such as supervision. Brazil's central bank built a prototype robot to download consumer complaints about financial institutions and categorize them through machine learning. The Reserve Bank of India this year hired consulting firms McKinsey and Accenture to help deploy AI and related analytics in its supervision work.

The Basel Committee on Banking Supervision found that AI can make lending more efficient in credit decisions and in thwarting money laundering. The committee of central bankers and bank supervisors acts as one of the world's top standard-setters for regulation. It also cited risks such as understanding outcomes from opaque models and the potential for bias and greater cyber risks.

"Supervisory processes for judging what is safe and sound, and being able to distinguish between responsible and irresponsible innovation, will no doubt improve," Neil Esho, the panel's secretary general, said last year. "For now, we still have some way to go."

The Bank for International Settlements (BIS), the group of global central banks that hosts the committee secretariat in Basel, Switzerland, has tested a variety of potential uses. The BIS Innova-

tion Hub's Project Aurora, for example, showed that neural networks, a type of machine learning, can help detect money laundering by sniffing out patterns and anomalies in transactions that traditional methods can't identify.

Signal in noise

The Bank of Canada built a machine learning tool to detect anomalies in regulatory submissions. Data Science Director Maryam Haghighi said its automated daily runs catch things people wouldn't—while freeing up staff to follow up on the analysis.

"This is an example of where AI can really shine for central banks," Haghighi said. "It's something rather tedious, and it's something that you can train AI to do well and do better and faster than humans."

The European Central Bank (ECB) is using AI for applications such as automating classification of data from 10 million business and government entities, scraping websites to track product prices in real time. It is also using the technology to help bank supervisors find and parse news stories, supervisory reports, and corporate filings.

With the data universe growing exponentially, cleaning it up to be intelligible is a key issue, especially for unstructured data, said Myriam Moufakkir, the ECB's chief services officer. AI can help humans make important distinctions. The ECB is also exploring large language AI models to help write code, test software, and even help make public communications easier for people to understand.

Financial stability

London School of Economics researcher Jon Danielsson, who studies how AI affects the financial system, sees the technology's capabilities on a continuum from basic to advanced, he said. On the basic side, there's chess, with pieces on a board and rules known to all. AI easily beats humans there, but its advantage diminishes with complexity. People in unexpected situations can draw on a range of knowledge to make better-informed decisions, from economics and history to ethics and philosophy. And this, he said, is where humans beat AI—for now.

AI is already making important financial decisions, such as handling credit card applications, and it's making rapid inroads in the public and private sectors. The technology can help ensure that banks don't misbehave by, for example, taking advantage of clients or allowing fraud or money laundering, he said. At the same time, such expanded uses may introduce danger, he said.

"The technology creeps up on us when we start trusting it but using it more and more," Danielsson said.

AI could spark a financial crisis, according to US Securities and Exchange Commission Chair Gary Gensler. He is charged with protecting a $46 trillion stock market that makes up two-fifths of the world total. Financial stability risks from AI demand "new thinking on systemwide or macroprudential policy interventions," he told reporters in July. "AI may heighten financial fragility, as it could promote herding—with individual actors making similar decisions because they are getting the same signal from a base model or data aggregator."

The warning reflected Gensler's work as global economics and management professor at the Massachusetts Institute of Technology, where he published a 2020 paper with Lily Bailey on deep learning. That subset of AI offers "previously unseen predictive powers enabling significant opportunities for efficiency, financial inclusion, and risk mitigation," they wrote. But they cautioned that financial regulations rooted in earlier eras "are likely to fall short in addressing the systemic risks posed by broad adoption of deep learning in finance."

'Polycrisis' factor

Another danger is that AI tools may exacerbate a crisis, whatever the cause, because they are trained on past data that may not reflect reality in an unprecedented situation, according to Anselm Küsters, head of the digitalization and new technologies department at the Centre for European Policy in Berlin. Küsters has cited the term polycrisis, popularized by fellow economic historian Adam Tooze, referring to interaction of different shocks that together add up to be worse than the sum of their parts.

Increased use of opaque AI applications "creates new systemic risks," as they can quickly amplify negative feedback loops, Küsters wrote, urging the European Parliament to "focus on the additional risks of algorithmic prediction arising in crises."

Such questions posed by the rapidly evolving technology will confront central bankers and other policymakers in coming years as benefits and threats become clearer.

"We are not yet at the point where we know what makes sense for central bankers," the ECB's Moufakkir said. "We are at the beginning." F&D

JEFF KEARNS *is on the staff of* Finance & Development.

TECHNOLOGY FOR DEVELOPMENT

Daniel Björkegren and Joshua Blumenstock

AI must be carefully adapted to benefit the poor, shows research in Kenya, Sierra Leone, and Togo

As artificial intelligence (AI) reshapes developing economies, it raises familiar risks of disruption, misinformation, and surveillance—but also promises many potential benefits. Recent examples illustrate how AI-based technologies can target aid and credit better and improve access to tailored teaching and medical advice. But balancing these risks and opportunities means more than just plug and play of existing technology—it calls for local innovation and adaptation.

Most recent advances in artificial intelligence originated in wealthy nations—developed in those countries for local users, using local data. Over the past several

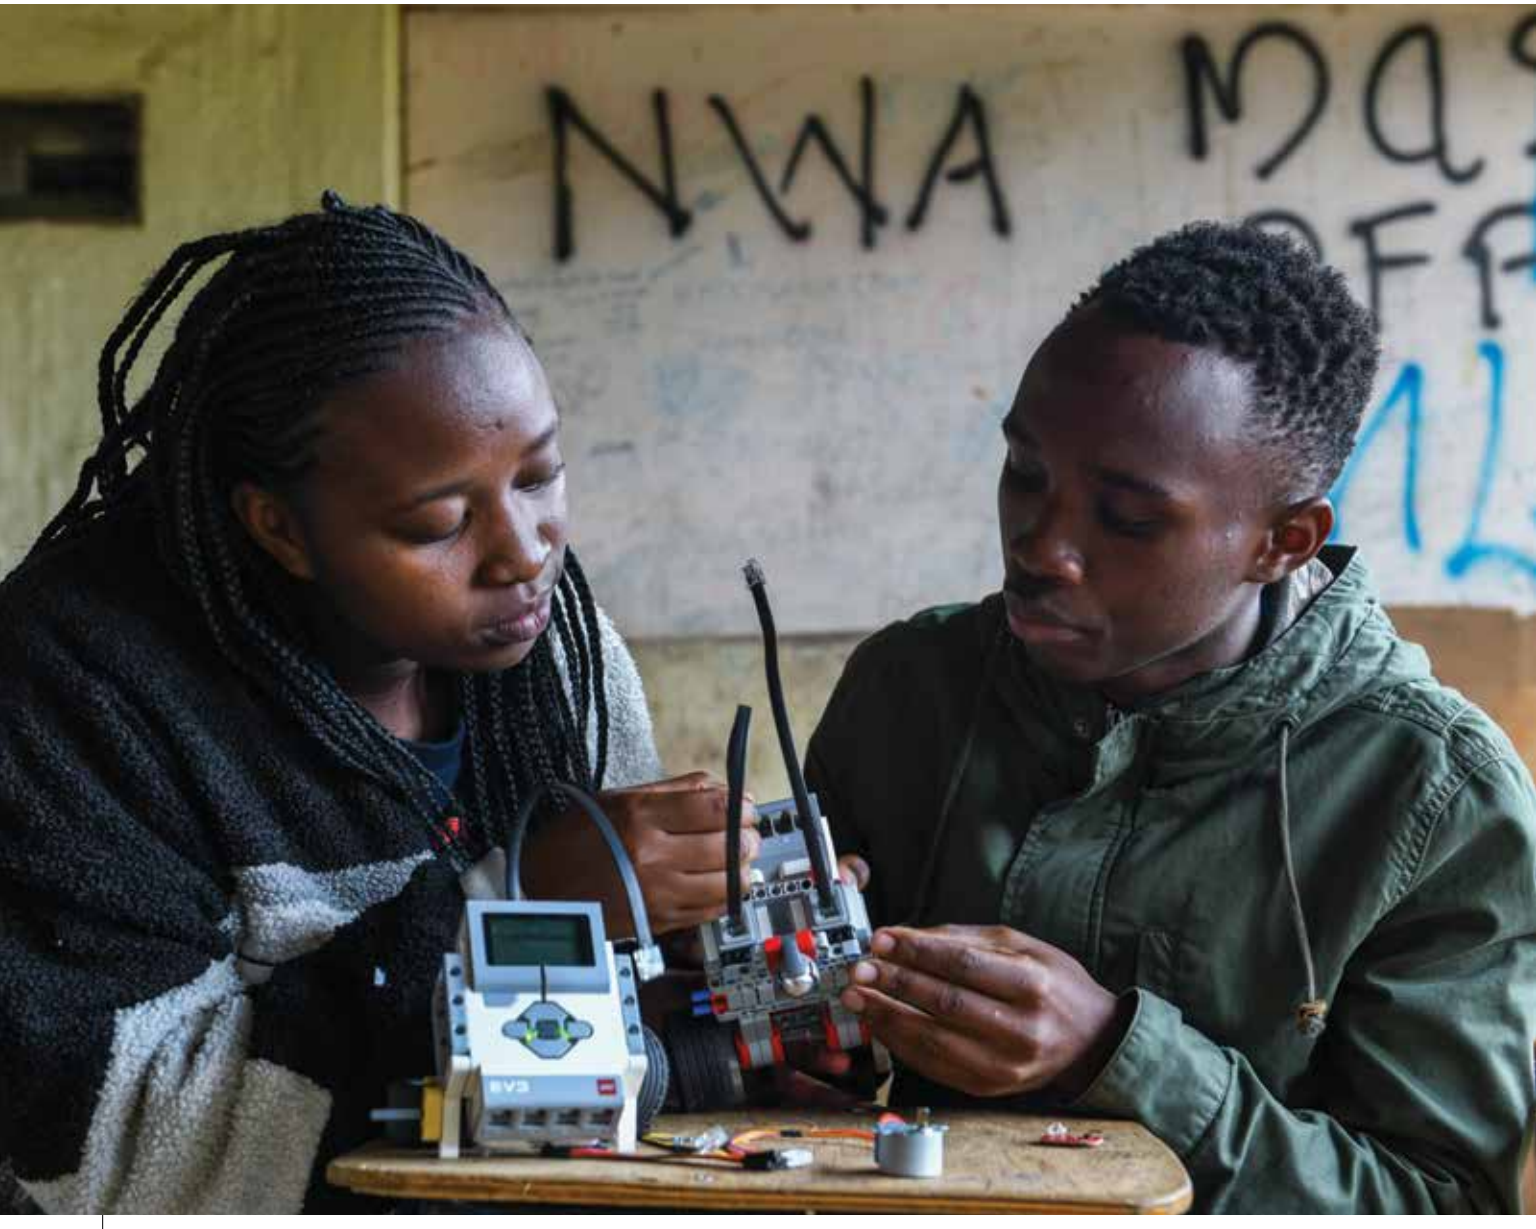

GERALD ANDERSON/ANADOLU AGENCY VIA GETTY IMAGES

Kenyan youth attend a robotics and coding class to learn future technologies.

years, we have conducted research with partners in low-income nations, working on AI applications for those countries, users, and data. In such settings, AI-based solutions will work only if they fit the local social and institutional context.

In Togo, where the government used machine learning technology to target cash aid during the COVID-19 pandemic, we found that adapting AI to local conditions was the key to successful outcomes. The government repurposed technology originally designed to target online advertising to the task of identifying the country's poorest residents. Using AI, the system processed data from satellites and mobile phone companies to identify signatures of poverty—such as villages that appeared underdeveloped in aerial imagery and mobile subscribers with low balances on their phones. Targeting based on these signatures helped ensure that cash transfers reached people with the greatest need (Aiken and others 2022).

This application worked in Togo only because the government, in collaboration with researchers and nonprofit organizations, customized the technology to meet local needs. They built a system for distributing mobile money payments that worked for all mobile subscribers, adapted existing machine learning software to target cash transfers, and interviewed tens of thousands of beneficiaries to ensure that the system reflected the local definition of poverty. And even then, the AI-based solu-

tion was not designed to be permanent; it was to be phased out after the pandemic ended.

The AI-based program also raised another concern: algorithms that perform well in a laboratory may not be reliable when deployed for consequential decisions on the ground. For instance, in an aid-targeting system like the one in Togo, people might adapt their behavior to qualify for benefits, thereby undermining the system's ability to direct cash to the poor.

Elsewhere, machine learning is used to determine eligibility for microloans, based on mobile phone behavior (Björkegren and Grissen 2020). For example, in Kenya over a quarter of adults have taken out loans using their mobile phones. But if those with more Facebook friends are likelier to be approved for a loan, some applicants may consider adding friends quickly. Ultimately, this can make it hard for systems to target the intended people.

In a study with the Busara Center in Kenya, we found that people were able to learn and adjust their smartphone behavior in response to such algorithmic rules (Björkegren, Blumenstock, and Knight, forthcoming). We showed how a proof-of-concept adjustment to the algorithm, which anticipates these responses, performed better. However, technology alone cannot overcome problems that arise during implementation; much of the challenge of building such systems is ensuring that they are reliable in real-world conditions.

On the other hand, some systems require adaptation before they will be useful. For instance, in many lower-income countries, teachers must handle large classes with limited resources. In Sierra Leone, a local partner piloted an AI chatbot system for teachers, called TheTeacher.AI, which is similar to ChatGPT but tailored to local curriculum and instruction and accessible even when internet connections are poor. In the pilot phase, many teachers couldn't phrase questions in a way that yielded useful answers, but a small group began to use the system regularly to help with teaching concepts, planning lessons, and creating classroom materials (Choi and others 2023). It took training and experimentation for teachers to use it in practice. Uses of AI may not be immediately obvious to those who stand to benefit; discovering the many uses will depend on trial and error and sharing applications that help.

Communication barriers

Grasping the potential of AI is likely to be harder for people in lower-income countries, where literacy and numeracy are lower and residents are less familiar with digital data and the algorithms that process this information. For instance, in our field experiment in Nairobi, Kenya, we found it difficult to explain simple algorithms with negative numbers and fractions to low-income residents. But our team found simpler ways to communicate these concepts. It was clear when people responded to the algorithm that they grasped the concept. Still, complex AI systems are difficult to understand, even for AI researchers.

"AI systems will require investment in knowledge infrastructure, especially in developing economies, where data gaps persist and the poor are digitally underrepresented."

Some applications don't require that users know how algorithms work. For instance, Netflix movie recommendations can benefit users even if they do not understand how the algorithm selects content it thinks they will like. Likewise, in a humanitarian crisis, policymakers may deem it acceptable to use an inscrutable "black box" algorithm, as Togo's government did in response to the COVID-19 crisis.

Transparency is sometimes critical. When targeting social protections in nonemergency settings, explaining eligibility criteria to potential beneficiaries is essential. This is easier said than done: scores of interviews and focus groups showed us how norms and values around data and privacy are fundamentally different in a setting such as rural Togo than in wealthy nations, where AI-based systems are more common. For instance, few people we spoke to were worried about the government or companies accessing their data (a dominant concern in Europe and the United States), but many

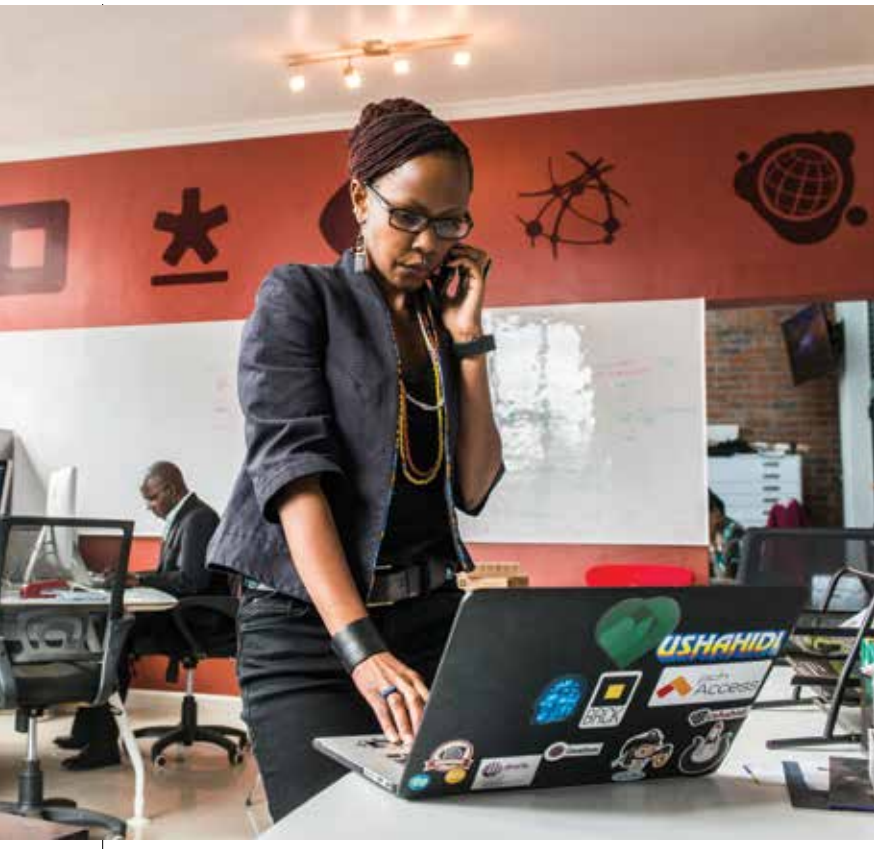

Juliana Rotich works on a laptop at the i-Hub technology innovation center in Nairobi, Kenya.

WALDO SWIEGERS/BLOOMBERG VIA GETTY IMAGES

wondered if and how such information would be shared with their neighbors.

As AI is more commonly deployed, populations must understand its broader societal effects. For instance, AI can generate provocative photographs that are entirely false and robocalls that mimic voices. These rapid changes will affect how much people should trust information they see online. Even remote populations must be informed about these possibilities so that they are not misled—and to ensure that their concerns are represented in the development of regulations.

Building connections

AI solutions rest on existing physical digital infrastructure: from massive databases on servers, to fiber-optic cables and cell towers, to mobile phones in people's hands. Over the past two decades, developing economies have invested heavily in connecting remote areas with cellular and internet connections, laying the groundwork for these new applications.

Even though AI applications benefit from digital infrastructure, some could make better use of existing resources. For example, many teachers in Sierra Leone struggle with poor internet access. For some tasks, it may be easier to get ideas from a chatbot and then validate the response than to collate information from several online resources.

Some AI systems will, however, require investment in knowledge infrastructure, especially in developing economies, where data gaps persist and the poor are digitally underrepresented. AI models there have incomplete information about the needs and desires of lower-income residents, the state of their health, the appearance of the people and villages, and the structure of lesser-used languages.

Gathering these data may require integrating clinics, schools, and businesses into digital record-keeping systems; creating incentives for their use; and establishing legal rights over the resulting data.

Further, AI systems should be tailored to local values and conditions. For example, Western AI systems may suggest that teachers use expensive resources such as digital whiteboards or digital slide presentations. These systems must be adjusted to be relevant for teachers lacking these resources. Investing in the capacity and training of local AI developers and designers can help ensure that the next generation of technical innovation better reflects local values and priorities.

Artificial intelligence promises many useful applications for the poor across developing economies. The challenge is not in dreaming big—it's easy to imagine how these systems can benefit the poor—but in ensuring that these systems meet people's needs, work in local conditions, and do not cause harm. **F&D**

DANIEL BJÖRKEGREN *is an assistant professor at the Columbia University School of International and Public Affairs.* **JOSHUA BLUMENSTOCK** *is a chancellor's associate professor at the University of California, Berkeley, School of Information and the Goldman School of Public Policy.*

REFERENCES

Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E. Blumenstock. 2022. "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid." *Nature* 603 (7903): 864–70. https://doi.org/10.1038/s41586-022-04484-9.

Björkegren, Daniel, and Darrell Grissen. 2020. "Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment." *World Bank Economic Review* 34 (3): 618–34. https://doi.org/10.1093/wber/lhz006.

Björkegren, Daniel, Joshua E. Blumenstock, and Samsun Knight. Forthcoming. "Manipulation-Proof Machine Learning."

Choi, Jun Ho, Oliver Garrod, Paul Atherton, Andrew Joyce-Gibbons, Miriam Mason-Sesay, and Daniel Björkegren. 2023. "Are LLMs Useful in the Poorest Schools? theTeacherAI in Sierra Leone." Paper presented at the Neural Information Processing Systems (NeurIPS) Workshop on Generative AI for Education (GAIED).

AI IN PRACTICE

Technology is reshaping the way we cultivate food, care for our health, and preserve national security

On the ground and in the air, big changes are brewing, powered by the latest technological disruption, artificial intelligence. AI's potential to transform society is vast; here we examine how it is impacting three key sectors—agriculture, medicine, and defense.

AI is revolutionizing these (and many other) fields, from optimizing crop yields and improving precision healthcare to enhancing military capabilities and national security. The following real-life examples from around the world highlight some of the technology's tangible benefits but also raise questions about the ethical considerations, policy lags, and training gaps associated with the integration of AI into critical industries. They attempt to shed light on the boundless possibilities and challenges that lie at the intersection of technology and human progress.

COSTFOTO/NURPHOTO VIA GETTY IMAGES

Tea plantation in Hangzhou, China: AI-driven smart agriculture provides tremendous potential for boosting food security and to reduce or even end hunger in many regions.

Growing More with Less

In the battle against hunger, AI can help fewer farmers generate more food

Robert Horn

ARTIFICIAL INTELLIGENCE (AI) IS MAKING its mark on food and agricultural production chains. The groundbreaking technology is already in use to engineer new varieties of climate-resistant rice; provide data on soil; guide drones that precision-spray fertilizers and pesticides; and sort, inspect, and grade produce. "AI-driven smart agriculture provides tremendous potential for boosting food security and to reduce or even end hunger in many regions of the world," said Channing Arndt, of the Consultative Group on International Agricultural Research, a global research partnership.

Policymakers in Thailand agree. In 2014, they unveiled Thailand 4.0, a 20-year national strategy for advanced development. Among its priority sectors are food, agriculture, and digital industries. Those are woven together in government-run Smart Farmer and Young Smart Farmer programs that encourage growers to adopt precision agriculture by connecting with new technologies. These include AI-controlled drones and software for intelligent and targeted spraying to increase yields while protecting the environment and ecosystem.

Food and agriculture have long been sources of strength for Thailand. The Southeast Asian kingdom of 70 million people is the world's 15th largest food exporter and the only net food exporter in Asia. With a projected $44.3 billion in shipments this year, Thailand plays a crucial role in regional and global food security and the campaign to end hunger.

But that campaign, despite impressive gains during the past decade, has recently suffered setbacks. The pandemic, war in Ukraine, and resulting disruptions left an estimated 735 million people (9.2 percent of the global population) undernourished in 2022, according to *The State of Food Security*

"For all the promise, some researchers warn about risks. If the data are bad, AI's results will be bad."

and Nutrition in the World, published by the United Nations. Even food-rich Thailand saw hunger rise for the first time in a decade. In response, policymakers in several regions are exploring how digital technologies can make agriculture more productive and food chains more efficient to turn the tide on malnutrition and food scarcity.

USING AI TO FIGHT HUNGER

Thailand is just one example of how countries are using AI to combat rising hunger, food insecurity, and poverty, which take a toll on economies. Undernourished people need more public assistance and are less productive workers, which can affect per capita income, growth, and sometimes political stability. At the same time, younger farmers are migrating to better-paying jobs in cities, leaving fewer hands to produce an increasing volume of food for a growing global population. Those combined trends could spell crisis, but analysts and policymakers see hope in new technologies, including AI, to help fewer farmers generate more food.

Nurturing the digital ecosystem is foundational in that effort, said Krithpaka Boonfueng, executive director of the National Innovation Agency (NIA). In October, Thailand launched THEOS-2, the first Earth observation satellite jointly designed by Thai and British engineers, which will gather data for smart agriculture. NIA has incubator and accelerator programs that help source private sector investment for agricultural technology start-ups to deliver the data to the field. The Digital Economy Promotion Agency (DEPA), another technology arm of the government, manages the One Community, One Drone program, which has farmers in 500 communities sharing drone services to manage their fields. "Even farmers want technology, but the technology has to be simple enough for them to use," said Preesan Rakwatin, executive vice president of DEPA, which helps match tech businesses with markets and also funds start-ups.

One of those start-ups, Ricult, is already helping farmers in Pakistan, Thailand, and Vietnam. Founded in 2015, Ricult is a dual fintech and agritech firm. Its AI-driven app, with more than 800,000 downloads in Thailand, provides information and tools that help smallholder farmers choose the right crop varieties and precision methods to increase productivity and profitability. Its portal also assists farmers with a perennial pressing problem: access to finance. Meanwhile, Mitr Phol Group, the largest sugar producer in Asia, has partnered with IBM for AI-driven data solutions for farmers, and Chia Tai, one of Thailand's biggest agri-food companies, is using autonomous drones made by XAG of China.

SMARTER POLICIES

But smart farming is still relatively rare. Ricult cofounder Aukrit Unahalekhaka said that government agencies trying to micromanage while working in insulated teams hinder uptake. It's a problem, he said, throughout the region. "Governments should be creators of policies and facilitators of funding for start-ups, innovators, and farmers. It is much more efficient to let the market work," Unahalekhaka said.

That doesn't always happen. Several governments in Africa, another continent struggling with hunger and food security, have passed restrictive drone regulations, and acquiring a license for one is difficult, according to "Empowering Africa's Food Systems for the Future," a report by the International Food Policy Research Institute (IFPRI). However, countries such as Kenya, Rwanda, Tanzania, and others are devoting resources to building a digital ecosystem and literacy for farmers so that they can access online extension services, weather forecasts, market information, and financing. Obstacles such as connectivity and digital literacy, however, remain. "While the digital revolution holds immense promise for African food systems, addressing these challenges is pivotal for its success," according to IFPRI.

For all the promise, some researchers warn about risks. If the data are bad, AI's results will be bad. And AI can be programmed to increase yields while ignoring negative impacts on the environment. "AI can be fine-tuned to match your goals. It is not perfect," Unahalekhaka said, adding that he hasn't seen farmers misuse AI so far. He is one of many who believe that the benefits outweigh the risks and that results are likely to be positive. That's because of the motivation he believes he shares with others in agricultural technology: "We want to make the world a better place."

ROBERT HORN *is a Bangkok-based freelance writer who previously worked at* Fortune *magazine,* Time *magazine, and the Associated Press.*

TONY KARUMBA/AFP VIA GETTY IMAGES

A technician scans the eye of a woman in Kianjokoma, Kenya, with a smartphone application. The "eyephone app" provides Kenyans the chance to get a quick and effective diagnosis, even in remote rural areas.

AI's Healing Powers

Artificial intelligence shows promise in medicine, but there are recognized drawbacks and risks

Kerry Dooley Young

IT'S EASY TO IMAGINE THE POTENTIAL FOR artificial intelligence (AI) to help people around the world live healthier lives.

Some already use AI to spot early signs of disease quickly, as recently reported in a study done in Rangpur, Bangladesh. In this study, the nonprofit Orbis International, which seeks to address preventable causes of blindness, and local physicians used the LumineticsCore system, from Coralville, Iowa–based Digital Diagnostics. The system uses a special kind of camera designed to capture images of the eyes and evaluates them with AI.

This product already has an impressive track record. In 2018, it was the first AI-driven device to gain US Food and Drug Administration (FDA) clearance to check for diabetic retinopathy. In 2020, the giant US Medicare health program agreed to pay for use of the device in primary care offices.

In the Bangladesh study, researchers tracked the productivity of a retina clinic whose patients with diabetes were randomly assigned to the AI or the control group.

The result? When the AI tool was used, an estimated 1.59 patients an hour received what was deemed a high-quality visit, versus 1.14 in the control group, wrote Digital Diagnostics founder Michael Abramoff and his coauthors in an article published in October in Nature Portfolio's *npj Digital Medicine* journal.

The test showed that LumineticsCore could help more people get screened for eye damage from diabetes, even in developing economies, said Abramoff, who is also a professor of ophthalmology and engineering at the University of Iowa.

He distinguishes between what he calls "impact AI" in medicine and "glamor AI," meaning products that garner tantalizing headlines but do not yet show hard evidence of benefit to patients.

"We like what we now call impact AI," which is shown to help improve people's health, Abramoff said. "I'm an engineer, so I love technology, but we shouldn't be paying too much for it if it doesn't improve outcomes."

And there's a need for vigilance with AI because its application in medicine has already shown potential to harm as well as help people.

For example, a 2019 *Science* paper reported that an algorithm widely used by large health systems and insurers underestimated the severity of Black patients' illness, thus setting the stage to deny them care. Researchers and policy experts have raised concerns about developing AI tools based on data drawn with a tilt toward relatively wealthy people, who are often White and have good access to health care.

Greater diversity is essential among the patients whose data help train AI tools, as well as among the people who build these products, said Jerome Singh, one of the advisors on a 2021 World Health Organization (WHO) guidance report on ethics and governance of AI for health.

"You're going to need to have multiracial, multicultural coders," Singh said. "The interpretation is quite important. AI is only as good as its coding."

That need for more diversity is one of the chief challenges ahead for attempts to use AI in medicine globally, especially in the Global South, Singh said.

The need for AI may be greater in less developed economies, where the ratio of medical personnel to patients tends to be much higher than in affluent areas. In the United States, there are about 36 doctors for every 10,000 people and in the United Kingdom, about 32, but in India, there are about 7 per 10,000 people, according to WHO data.

Yet these less affluent countries also face challenges when it comes to the infrastructure and insti-

tutional knowledge needed to successfully deploy AI, Singh said. These include lack of electricity and computer servers as well as a shortage of workers able to translate AI-assisted diagnoses into effective treatment.

"In some settings, it's going to be more of a sprint" to successfully integrate AI into health care, Singh says. "In other settings, it's going to be a marathon."

AI adoption in medical practice is inevitable at this point, said Partha Majumder, who served as the cochair of an expert group that provided guidance on the 2021 WHO report.

"We have to accept that this is reality," he said. "Checks and balances need to be hammered in such a way that inappropriate predictions and diagnoses are not made. That's all we can do. We actually can't hold back the rolling out of the AI methods."

Regulators and policymakers around the world are wrestling with ways to make sure AI is applied safely and effectively in health care. Much of this work centers on trying to address bias in how algorithms are developed and trained.

In October, the WHO issued a new report outlining the challenges of regulating AI in medicine. It cited particular concerns about rapid deployment of tools derived from large language models, a class that includes chatbots, without a full understanding of whether these programs will help or harm patients. A European Parliament report issued last year noted that there are also concerns about lack of transparency and privacy and security issues. The FDA is working to refine its approach to regulating AI in medical products through formal guidance. These show companies what kind of evidence they will need to produce to win FDA clearance of products.

AI can eliminate many of the frustrating setbacks that have long been a hallmark of pharmaceutical research, said Tala Fakhouri, associate director for policy analysis at the FDA's Center for Drug Evaluation and Research Office of Medical Policy. It's becoming easier to understand in the initial stages how compounds will work in the body, reducing the chance of side effects that often crop up in later testing. Researchers can now quickly analyze information about experimental drugs with AI that would have taken years to synthesize in the past, she said.

"The efficiencies that have been built now on the discovery side are exponential," Fakhouri said. We're going to see a lot coming to the market soon."

KERRY DOOLEY YOUNG *is a freelance journalist who specializes in health care.*

PAULA BRONSTEIN/GETTY IMAGES

A homemade prototype drone is tested with a fake RPG -7 grenade in a field outside Kyiv, Ukraine.

New Model Army

AI has accelerated shifts in battlefield dynamics, and policymakers are playing catch-up

Jeremy Wagstaff

RUSSIA'S WAR IN UKRAINE HAS BECOME A testing ground for new technology, in particular demonstrating how artificial intelligence (AI) can be used to great effect. But it has also highlighted weaknesses in how governments and the defense industry adopt, deploy, and control AI-based technology.

AI has been used in several ways in the Ukraine war, from broad strategic decision-making—through how to act on real-time or recent intelligence at the local level—to handling more mundane tasks, such as predicting logistical challenges. A fourth use involves information warfare. This is a way of leveraging AI to, in the words of Matthew Ford of the Swedish Defence University in Stockholm and coauthor of a book on the battlefield digital explosion, *Radical War: Data, Attention and Control in the 21st Century*, "shape how narrative construction works."

But even though the war has shown that AI can help armies monitor enemy movements and deliver payloads remotely and autonomously, it has also accelerated shifts in battlefield dynamics. Forces soon alter tactics, techniques, and procedures either to leverage the new technology or to mitigate its impact.

A failure to adapt quickly can be seized on by an agile foe. When Russian soldiers and pilots communicated without encrypting their conversations, Ukraine developed AI-based voice recognition and translation software to monitor these communications and extract actionable intelligence. And even when countermeasures are adopted, each side must be ready to rethink and enhance its technology as rapidly as the other. When Russia introduced electronic jamming to thwart Ukraine's combat drones, for example, Kyiv's cadre of programmers developed an AI tool to help its drones evade Russian jamming and stay locked on target.

UNPILOTED DRONES

This technological arms race is strikingly different from how many military thinkers saw the deployment of AI. For one thing, the principles behind unpiloted aerial vehicles, or UAVs, have not changed significantly since the 1990s. But in Ukraine the range of drones, and their capabilities, has evolved rapidly, largely by coupling them with continuous advances in AI. While the military-grade Turkish-made Bayraktar TB2 drone played a key role in Ukraine's defense in the early months of the 2022 invasion, it became less useful as Russia upgraded its air defense and electronic warfare capabilities.

With more permanent battle lines drawn later in the year, Ukraine pushed its drone makers to adapt. The result has been a succession of improved and diverse devices. In September, for example, Kyiv approved the deployment of homegrown Saker Scout drones, which can detect enemy targets often missed by the human eye, even when hidden under camouflage.

This emphasis on rapid evolution has helped change thinking among military strategists, says Lauren Kahn, senior research analyst at Georgetown University's Center for Security and Emerging Technology (CSET). Despite excitement about AI in military circles since 2021, if not earlier, practical examples were either hypothetical or project based. "That changed after Ukraine," she says.

Planners began to see that AI was not just a box to tick but raised a series of searching questions about what would make it useful: data, knowledge about your own side and the other, testing and evaluation procedures. The creative way Ukraine has developed drone technology is something "no one could have imagined," she said.

DATA

The Ukraine war has highlighted the importance of data—the fuel that powers AI—but has also raised troubling questions for policymakers and planners. Ukraine understood early that what constituted data in a war had shifted. It quickly reconfigured a government app for filing taxes to also allow citizens to upload photos, videos, and other details about Russian troops and positions to a database run by the military.

It combined commercially available satellite images with classified data from its allies, as well as from hacking into Russian surveillance cameras and from its own fleet of drones. But all this data needed to be turned into actionable intelligence, for which Kiev turned to private tech companies—the most visible being Palantir, a US company specializing in big-data analytics. Palantir's involvement extended the role a private company might play in processing sensitive data, especially during a war. Its chief executive, Alex Karp, is on record as saying the company is responsible for most of the targeting in Ukraine. According to CSET's Kahn, "It's almost like a full service they provide, which I think has proved invaluable."

What hasn't been fully considered, at least publicly, are the implications. Private companies, says the Swedish Defence University's Ford, are going to be crucial, because they are the only organizations that can develop the kind of AI armed forces can use. But, he asks, "Once it's out there, where does it go next? How's it going to be controlled, shaped, or directed?"

DIGITAL BATTLEFIELD

The war also introduces another aspect of AI and data. "The Ukraine-Russia war is the most documented war in history," says Andrew Hoskins, professor of global security at the University of Glasgow and Ford's *Radical War* coauthor. Telegram, the social media platform now used by three-quarters of Ukrainians and well over a third of Russians to share videos and photos as the war plays out in front of them, "is the digital battleground of this war," he says.

That information is not being uploaded only to army and intelligence servers, but also to NGOs and investigators mining it to catalog human rights abuses for future war crimes trials. AI, too, is improving what can be seen and extracted, says Hoskins. When you apply AI to these archives, "you start to find things you never anticipated." F&D

JEREMY WAGSTAFF *is a technology and media consultant and former journalist at the BBC, Reuters, and the* Wall Street Journal.

AI LEXICON

Artificial Intelligence

A field of computer science that focuses on building systems to imitate human behavior and demonstrate machine intelligence.

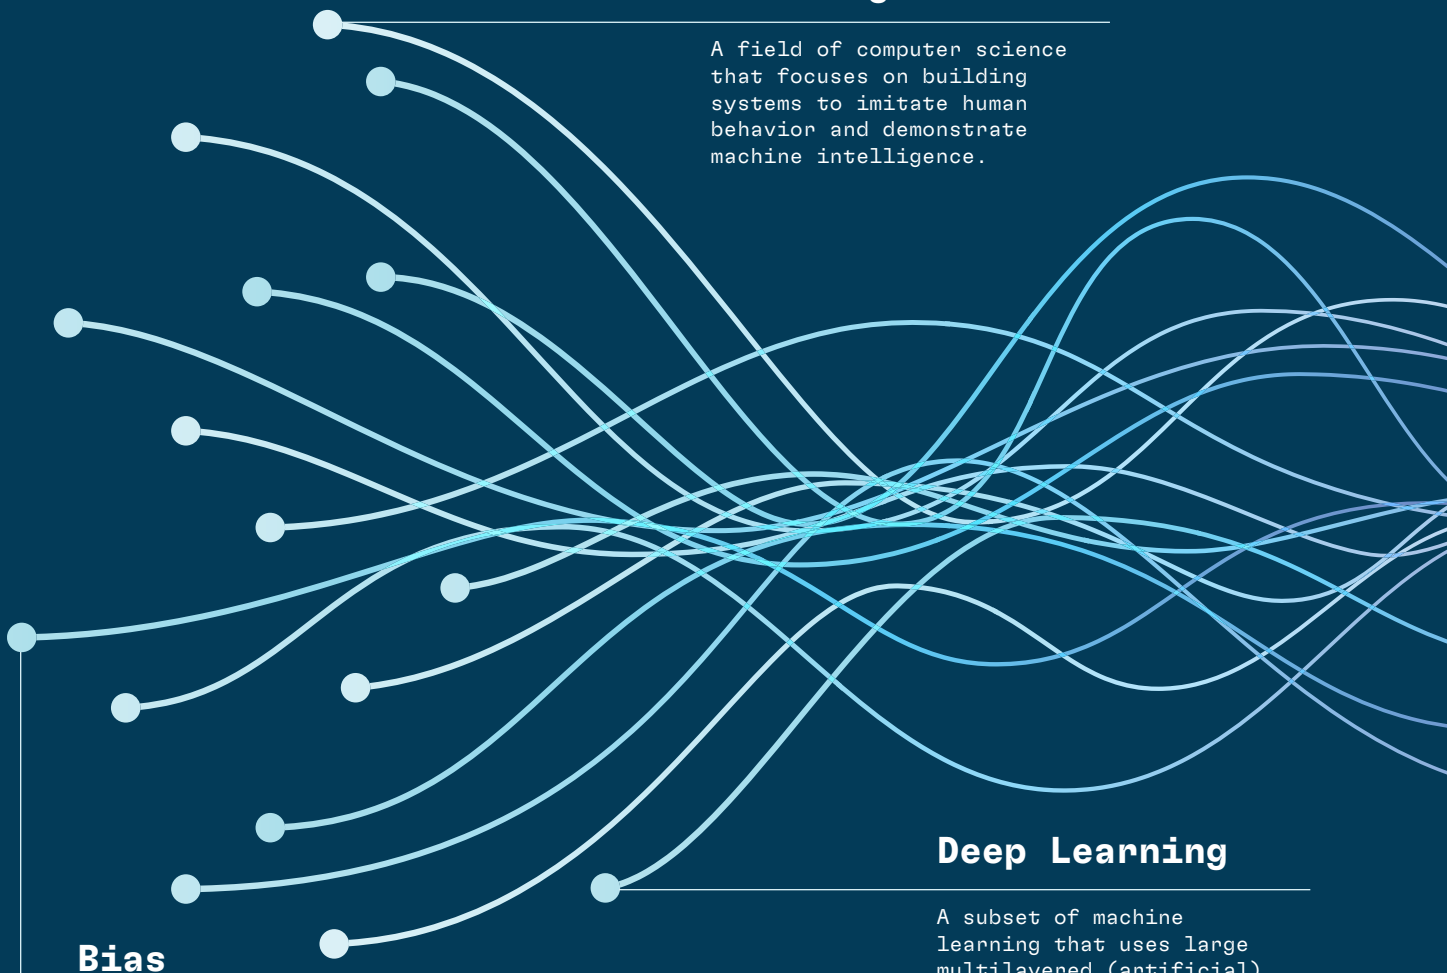

Deep Learning

A subset of machine learning that uses large multilayered (artificial) deep neural networks that compute with continuous (real-number) representations, a little like the hierarchically organized neurons in the human brain. It is especially effective at learning from unstructured data such as images, text, and audio.

Bias

A phenomenon that occurs when an AI system produces results that are systematically unfair or inaccurate due to erroneous assumptions or influences in the machine learning process. Bias in AI can have negative impacts on individuals and society, such as discrimination, misinformation, or loss of trust. There are different types and sources of bias in AI, such as data bias, algorithm bias, human bias, and societal bias.

Fine-tuning

The process of adapting a pretrained foundation model to perform a specific task better. This entails a relatively short period of training on a labeled dataset that is much smaller than the dataset on which the model was initially trained. This additional training allows the model to learn and adapt to nuances, terminology, and specific patterns.

Generative AI

A form of machine learning whereby AI platforms can generate new output in response to prompts based on the data on which it has been trained.

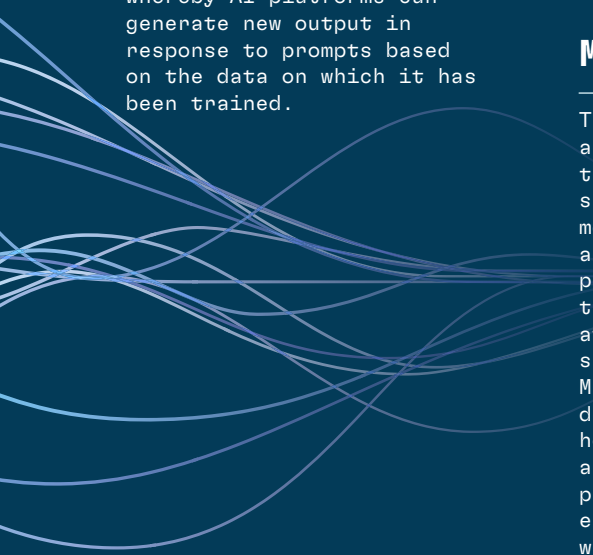

Hallucination

A phenomenon in which an AI system produces outputs that are not based on reality or the given context. For example, an AI chatbot might make up facts or stories, or an AI image recognition system might see objects or patterns that are not there.

Large Language Models

A neural net trained on large amounts of text to imitate human language. This class of foundation models can process massive amounts of unstructured text and learn the relationships between words or portions of words, known as tokens. This enables them to generate natural-language text to perform tasks such as summarization or knowledge extraction. GPT-4 (which underlies ChatGPT) and LaMDA (the model behind Bard) are examples of LLMs.

Machine Learning

The study of how AI acquires knowledge from training data. It is a subset of AI in which a model gains capabilities and improves its perception, knowledge, thinking, or actions after it is trained on or shown many data points. Machine learning algorithms detect patterns and learn how to make predictions and recommendations by processing data and experiences. In this way, the system learns to provide accurate content over time.

Neural Network

A computational model inspired by the structure and function of biological neurons.

Prompt Engineering

A technique used in artificial intelligence to optimize and fine-tune language models for particular tasks and desired outputs. Also known as prompt design, it refers to the process of carefully constructing prompts or inputs for AI models to enhance their performance on specific tasks.

Prompts

Instructions given to an AI system using natural language rather than computer language. For example, generative AI can be prompted to create content that appears novel or interesting.

Supervised Learning

A type of machine learning that uses labeled datasets to train algorithms to classify data or predict outcomes. Labeled datasets are collections of data that have been assigned a label or a category by humans.

Unsupervised Learning

A type of machine learning in which algorithms learn patterns from unlabeled data, without any human guidance or feedback.

SOURCES: Coursera. 2023. "Artificial Intelligence (AI) Terms: A to Z Glossary." Mountain View, CA; Manning, Christopher. 2020. "Artificial Intelligence Definitions." Stanford University Human-Centered Artificial Intelligence glossary, Stanford, CA; and McKinsey Digital. 2023. "The Economic Potential of Generative AI: The Next Productivity Frontier."

China's Bumpy Path

Eswar Prasad

GROWTH SLOWS, RISKS ABOUND, BUT ECONOMIC AND FINANCIAL COLLAPSE CAN BE AVOIDED

China's economic performance has been stellar over the past three decades, with remarkable and persistent high growth that lifted the economy from low-income to upper-middle-income status. Measured at market exchange rates, China's GDP was $18.3 trillion in 2022, 73 percent of the GDP of the United States and 10 times more than the 7 percent of US GDP it registered in 1990. China's per capita income is now roughly $13,000, approximately 17 percent of US per capita income—compared with less than 2 percent in 1990. Over the past decade and a half, China has been the main driver of the world's economic growth, accounting for 35 percent of global nominal GDP growth, while the United States accounted for 27 percent.

China accomplished this without many attributes that economists have identified as being crucial for growth—such as a well-functioning financial system, a strong institutional framework, a market-oriented economy, and a democratic and open system of government. Until the COVID-19 pandemic rocked it back on its heels, the Chinese economy powered through periods of domestic and global turmoil seemingly unscathed.

But detractors have long argued that China's economic collapse was imminent, pointing to numerous fragilities. The country's growth has been powered by investment in physical capital, especially real estate, that has been financed by an inefficient banking system. With domestic debt levels high and rising, the property market unraveling, and the labor force shrinking, some analysts say the day of reckoning has finally arrived.

They are likely wrong. Unbalanced reforms that have kept the institutional structure weak, a schizophrenic approach to the role of the market versus that of the state, and strains in

Illustration by Álvaro Bernis

financial and property markets could result in significant volatility in coming years. But none of this means a financial or economic collapse is inevitable.

Sources of growth

China's economic performance has relied largely on investment growth financed by an inefficient banking system. This pattern intensified after the global financial crisis that began in 2008. Increased investment accounted for about two-thirds of GDP growth during 2009–10. Because China is a labor-rich economy and has a capital-to-labor ratio much lower than that of advanced economies, more rather than less investment is probably desirable. However, much of the investment has been driven by the public (state) sector rather than the nongovernmental sector. This is not inherently a problem. Investment in private sector firms, especially smaller ones, can be much riskier than in large, state-owned enterprises. But in China, state-owned enterprises, which collectively receive a disproportionate share of bank credit, typically have not generated strong returns on those investments.

Recognizing that its growth model has been inefficient and financially risky, the Chinese government set itself the objective of rebalancing the economy. This means

- Reducing reliance on investment-heavy growth and getting household consumption to be the key contributor to GDP growth
- Generating more growth from the services sector than from low-skill, low-wage manufacturing
- Shifting away from physical-capital-intensive growth in a manner that improves employment growth

In recent years, household consumption has in fact become the main contributor to growth. The services sector now accounts for more than half of annual GDP and close to half of aggregate employment.

Thus, while the trajectory has been uneven, there has been significant progress toward the objective of growth rebalancing, with household consumption becoming the key driver of growth and the services sector displacing investment as more prominent than manufacturing.

Growth prospects

Prognostications about China's growth prospects are a fraught exercise, and at best forecasters can use the growth of various factors that go into the creation of output as indicators of what the future might hold.

"Over time debt has risen relative to the size of the economy—although gross debt levels are not out of line with those of other major economies, such as the United States and Japan."

China's labor force, the population in the 15–64 age range, is shrinking. By 2030, it is expected to decline by about 1 percent a year. Higher investment growth could pick up some of the slack, but that carries many risks. The recent decline in nongovernmental investment growth—state investment accounted for much of the growth in overall fixed asset investment outside the property sector in 2022—is a sign that private businesses are wary of increasing investment when they see the economic and political environment as unfavorable.

That leaves productivity, or the amount of output per unit of input, as a growth engine. For all the inefficiencies that pervade its economy, over the past few decades China has averaged a decent 3 percent growth in total factor productivity—which is growth that cannot be attributed to increased inputs, such as labor and capital, and is a general indicator of efficiency. But productivity growth has slowed to about 1 percent a year over the past decade. China's growth will run aground without an improvement in productivity growth.

Recognizing the need to improve productivity and shift away from low-skill manufacturing, the government recently articulated a "dual circulation" growth policy, which augments continued engagement with global trade and finance with greater reliance on domestic demand, technological self-sufficiency, and homegrown innovation. But the approach has run into difficulties. China still needs foreign technology to upgrade its industry, and rising economic and geopolitical rifts with the United States and the West could limit China's access to foreign tech-

nology and hi-tech products, as well as to markets for its exports. Moreover, the government's recent crackdown on private firms in sectors such as technology, education, and health has had a chilling effect on entrepreneurship.

Potential pitfalls

There are concerns that China's economy is headed for a crash similar to those experienced by other high-flying Asian economies—such as Malaysia and Thailand. China's overall debt has been a significant concern for many years. Over time debt has risen relative to the size of the economy—although gross debt levels are not out of line with those of other major economies, such as the United States and Japan. Moreover, public borrowing as a percentage of nominal GDP is lower in China than in other major economies. China has a high level of corporate debt—about 131 percent of GDP. But most of it is denominated in China's own currency and owned by domestic banks and investors, which presents less of a threat than were the debt owed to foreign investors and denominated in foreign currencies, such as the US dollar.

There are, however, specific sectors in which the concentration of debt could be a problem—especially the real estate sector. Real estate investment has become a bulwark of the economy, helping to keep growth on an even keel when other sectors floundered. Local government officials are eager to sell land to developers, boosting public revenues and enabling a range of government expenditures. So a fall in real estate prices—or the emergence of other factors that restrain real estate activity—could have knock-on effects across other sectors, local government finances, and even household wealth.

Household exposure to the real estate sector has created additional vulnerabilities that could affect economic and social stability. Easier access to residential mortgages, which the government encouraged, boosted housing demand and contributed to a surge in household debt, from about 30 percent of GDP a decade ago to more than 60 percent. Property has also become a mainstay of Chinese household wealth. Households are exposed in multiple ways to house price fluctuations. Still, total household debt is less than total household deposits in the banking system.

Because debt accumulation in China has been financed mostly by domestic savings, overall financial risk is limited. The state owns many of the key creditors and debtors, which means a financial shock is unlikely to set off a financial crisis or a collapse in growth. The more pertinent issues are major inefficiencies and waste because of a broken system of allocating capital.

How debt and assets are distributed throughout the economy matters. Tumbling house prices have caused several major property developers, such as Country Garden and Evergrande Group, to run into financial trouble recently, and many others are similarly exposed—with high debt and vulnerable balance sheets. So are some of the financial institutions that lent to them. But a systemic meltdown is not in the cards. Most major Chinese banks are under state control and can provide infusions of cash to troubled corporations, even if that only pushes problems off into the future. Stumbles are inevitable as China tries to give market forces freer rein, but the government has enough control and resources to prevent broader financial crashes.

External risks

Many emerging market economies have run into distress from high levels of external debt, particularly foreign currency debt, which can cause balance sheet problems when a country's economy and exchange rate deteriorate simultaneously. But China's external debt is estimated to be a modest 16 percent of GDP, and less than half of it is denominated in foreign currencies.

Still, economic and political uncertainty have created concerns about capital flight, which could bring down the financial system and cause the currency's value to crash. But this is an unlikely scenario, because much of the banking system is state owned and the government would probably back all deposits in the event of financial panic. Moreover, because the government directly controls much of the banking system, it can choke off the conduits for large capital outflows.

Although there have been reforms in recent years, many of them were related to the financial sector and capital markets, with far fewer in other areas, such as state enterprises and the institutional framework. This lack of balance creates risks.

The government seems to have grasped the need for financial sector reforms and liberalization to promote better resource allocation. Fixing the financial system is not just about managing risks and avoiding disaster but also about allocating capital to the more productive, dynamic, and employment-generating parts of the economy. China's financial system is still dominated by banks, whose loan portfolios are concentrated in the state enterprise sector. Fixing the banking system requires recognizing and removing bad loans from banks' balance sheets, as well as reform of the state enterprises themselves, including weaning

them off dependence on bank credit.

In recent years, as it dealt with episodes of housing market and stock market volatility, the government often found itself caught in a schizophrenic effort to balance maintaining confidence in the market with allowing the market to discipline itself—which had the perverse effect of heightening market turbulence. This on-off approach to intervention has sometimes injected a strong dose of uncertainty on top of already fragile investor sentiment and added to market volatility.

Moreover, market-oriented reforms can backfire, adding to volatility and generating more risks if they are not accompanied by broader reforms. China needs more transparency in its policymaking process, better corporate governance and accounting standards, and more operational independence for the central bank and regulatory authorities to supplement its financial and other market-oriented reforms.

The government has rightly encouraged the development of stock and corporate bond markets. But it has done little to improve corporate governance of Chinese companies or their accounting and auditing standards. The resulting opacity has contributed to large fluctuations in stock and bond markets, because investors have limited information about the companies they are investing in, leading them to follow and exacerbate market swings.

Reconciling the government's two contradictory impulses—more freedom for markets but with a heavy hand of government intervention to maintain "stability and order"—poses difficult challenges. Implementing even well-intentioned reforms in an economy with rampant inefficiencies involves transitional risks that might manifest in financial and economic volatility, especially if the government does not clearly communicate its policy intentions and leaves households and businesses guessing. So far, the government has had enough resources and policy space to cope with some of those transitional risks, but its actions and attempts to intervene directly in markets at difficult times might exacerbate problems, with long-lasting consequences.

"Reconciling the government's two contradictory impulses—more freedom for markets but with a heavy hand of government intervention to maintain 'stability and order'—poses difficult challenges."

What the future holds

The Chinese government has shown an uncanny ability to manage the severe economic and financial stresses that have built up from the highly inefficient and risky growth model it had embraced. At various points, the government has maneuvered the economy around the seemingly inevitable prospects of a banking crisis, massive currency devaluation, housing market meltdown, and economic collapse.

Yet each of these near misses has exacted a toll: a huge buildup in domestic debt, loss of $1 trillion in foreign exchange reserves during 2015–16, and highly volatile prices of stocks, property, and other assets.

The government now faces a number of policy dilemmas: how to continue reducing debt while maintaining growth, how to reduce energy-intensive production while the economy continues to rely on heavy industry, how to get markets to exert financial discipline even as the government tries to strengthen state control, how to restrain wealth inequality while relying on the private sector to generate more wealth, how to encourage private sector innovation while cutting successful private enterprises down to size.

The government's attempts to resolve these inherently contradictory impulses in the guise of market-oriented socialism will inevitably lead to further stumbles and accidents. Its policy approach, although driven by the right objectives, could generate more uncertainty and volatility in the short run, which in turn could reduce public support for much needed reforms to bolster long-term productivity and growth.

The underpinnings of China's growth seem fragile from historical and analytical perspectives. Even if no crises materialize, unfavorable demographics, high debt levels, and an inefficient financial system will constrain China's growth. Yet, if the government plays its cards right, one could equally well envision a more benign future for the Chinese economy—with moderate growth that is more sustainable from an economic, social, and environmental perspective. F&D

ESWAR PRASAD *is a professor in the Dyson School at Cornell University, senior fellow at the Brookings Institution, and author of* The Future of Money.

History's Inflation Lessons

People lining up for gas coupons in the Hanover Street Post Office in Liverpool, United Kingdom, November 29, 1973

WILLIAMS/MIRRORPIX/GETTY IMAGES

A STUDY OF 100 INFLATION SHOCKS SINCE THE 1970S PROVIDES VALUABLE POINTERS FOR POLICYMAKERS TODAY

Anil Ari and Lev Ratnovski

In the early 1970s, conflict in the Middle East set off a spike in oil prices that left central banks around the world scrambling to control inflation. After a year or so, oil prices stabilized and inflation started to retreat. Many countries believed they had restored price stability and loosened policy to revive their recession-hit economies only to see inflation return. Could history repeat?

World inflation reached historic highs in 2022 after Russia's invasion of Ukraine triggered a terms-of-trade shock akin to that of the 1970s. Disruptions to Russian oil and gas supplies added to COVID supply-chain problems to drive prices higher. In advanced economies, prices rose at the fastest pace since 1984. In emerging market and developing economies, the price increase was the largest since the 1990s.

Aided by the sharpest rise in interest rates in a generation, inflation has started to subside at last. Headline inflation in the United States and

across much of Europe has halved from about 10 percent last year to less than 5 percent today. The latest conflict in the Middle East has, for now at least, not had a large impact on oil prices. But it is still too soon for policymakers to celebrate victory over inflation.

Our recent study of over 100 inflation shocks since the 1970s offers two reasons for caution. First, history teaches us that inflation is persistent. It takes years to "resolve" inflation by reducing it to the rate that prevailed before the initial shock. Forty percent of countries in our study failed to resolve inflation shocks even after five years. It took the remaining 60 percent an average of three years to return inflation to pre-shock rates (Chart 1).

Second, countries have historically celebrated victory over inflation and loosened policy prematurely in response to an initial decline in price pressures. This was a mistake because inflation soon returned. Denmark, France, Greece, and the United States were among nearly 30 countries in our sample to loosen policy prematurely after the 1973 oil-price shock (Chart 2). In fact, almost all countries in our analysis (90 percent) that failed to resolve inflation saw price growth slow sharply in the first few years after an initial shock, only to accelerate again or become stuck at a faster pace.

Today's policymakers must not repeat their predecessors' mistakes. Central bankers are right to warn that the inflation fight is far from over, even as recent readings show a welcome moderation in price pressures.

Consistency and credibility

How should policymakers respond to persistent inflation? Again, history provides some lessons. The countries in our study that successfully resolved inflation tightened macroeconomic policies more in response to the inflation shock and, crucially, maintained a tight policy stance consistently over a period of several years. Examples here include Italy and Japan, which adopted tighter-for-longer policies after the 1979 oil-price shock. By contrast, countries that did not resolve inflation had looser policy stances and were more likely to change between tightening and loosening cycles.

Policy credibility matters, too. Countries where inflation expectations were more firmly anchored, or where central banks had more success maintaining low and stable inflation in the past, were more likely to defeat inflation.

Today's policymakers can take some solace from this finding. Central bankers in many countries may find it easier to defeat inflation this time because of the policy credibility they have built up over several decades of successful macroeconomic management.

CHART 1

Long inflation fight

History shows that many countries fail to defeat inflation.

Years until inflation falls to within 1 percentage point of its pre-shock rate
(number of episodes)

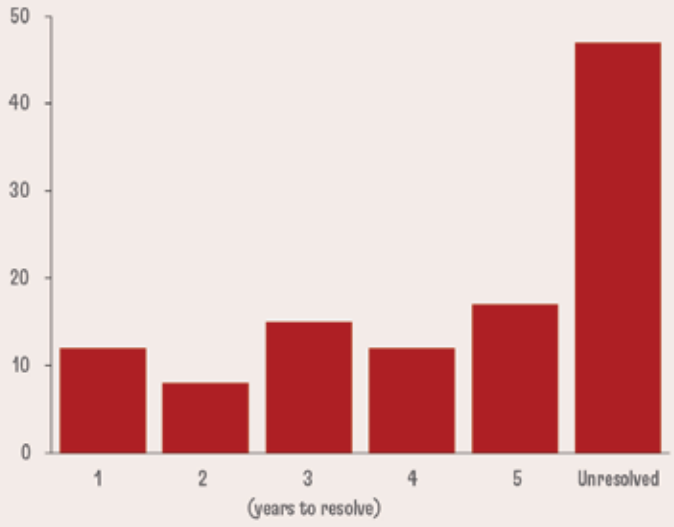

SOURCE: IMF staff calculations.

CHART 2

Premature celebrations

Inflation often slows after an initial shock only to accelerate again or become stuck at a faster pace.

(consumer price inflation, percent)

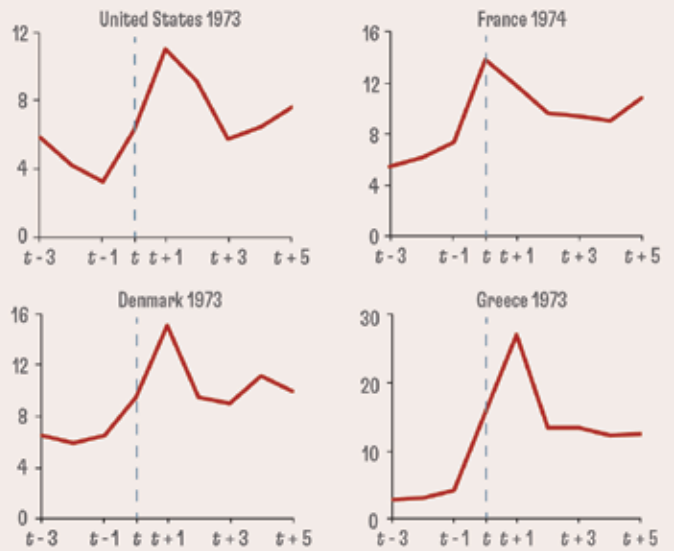

SOURCE: IMF, *World Economic Outlook*.

CHART 3

Lasting growth gains

Countries that defeat inflation grow faster after five years than those that allow inflation to linger.

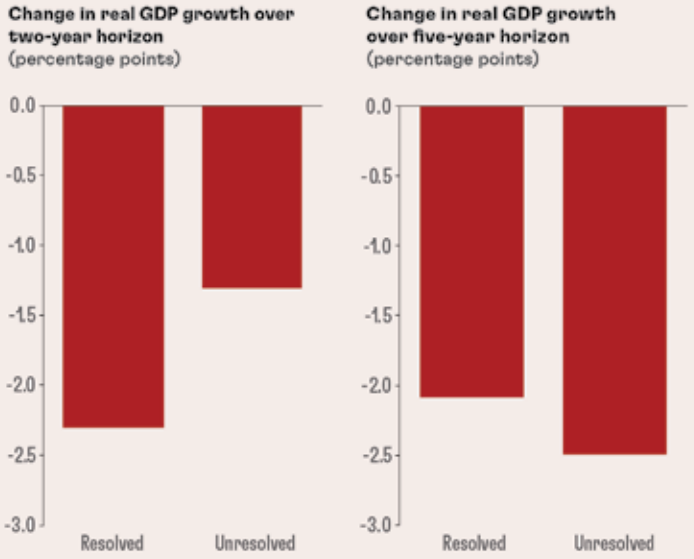

SOURCE: IMF staff calculations.

With the right policies in place, countries could resolve inflationary pressures sooner than in the past.

But it won't be easy. Conditions in the labor market in particular require close attention. In many countries, workers' wages have fallen in real inflation-adjusted terms and may need to rise again to catch up with higher prices. Yet wage growth could fuel inflation if it is too high and could lead to pernicious wage-price spirals.

Historically, countries that resolved inflation successfully tended to have lower nominal wage growth. Importantly, this did not translate into lower real wages and a loss of purchasing power, because lower nominal wage growth was accompanied by lower price growth. The implication for policymakers here is to remain focused on real wages, not nominal wages, when responding to developments in the labor market.

Countries that resolved inflation successfully were also better at maintaining external stability. Free-floating currencies were less likely to depreciate sharply, and currency pegs were more likely to survive. This is not a call for currency intervention. Instead, it appears that countries' success in fighting inflation—through tighter monetary policy and greater policy credibility—was instrumental in shoring up exchange rates.

"Countries that allow inflation to linger ultimately pay a higher price."

The ultimate prize

Fighting inflation is difficult. But it is important to recognize the benefits of price stability. Historically, countries that resolved inflation had lower economic growth in the short term than those that did not. But this relationship reversed over the medium and long term. Five years after the inflation shock, countries that resolved inflation had higher growth and lower unemployment than economies that allowed inflation to linger.

The economics behind this finding are intuitive. There is a trade-off between bringing inflation down on one hand and achieving higher growth and lower unemployment on the other. But this trade-off is temporary: growth recovers and jobs are created once inflation is brought under control.

By contrast, leaving inflation unresolved comes with its own costs of macroeconomic instability and inefficiency. These costs accumulate for as long as inflation remains high. Consequently, cumulative welfare losses from unresolved or permanently high inflation dominate over the medium to long term (Chart 3). Countries that allow inflation to linger ultimately pay a higher price.

Central bankers are on the front line of the fight against inflation and should pay the most attention to these lessons. But governments must not make the task of monetary authorities harder by adding to price pressures with loose fiscal policy. To make fiscal support during a cost-of-living crisis less inflationary, governments should target relief to the most vulnerable, where it will alleviate suffering most.

The past is never a perfect guide to the present, because no two crises are precisely alike. All the same, history offers clear lessons to policymakers today. Fighting inflation is a marathon, not a sprint. Policymakers must persevere, demonstrate policy credibility and consistency, and keep their eyes on the prize: macroeconomic stability and stronger growth brought about by returning inflation firmly to target. If history is a guide, inflation's recent decline could be transitory. Policymakers would be wise not to celebrate too soon. F&D

ANIL ARI *and* **LEV RATNOVSKI** *are economists in the IMF's Strategy, Policy, and Review Department and European Department, respectively.*

This article draws on IMF Working Paper 2023/190, "One Hundred Inflation Shocks: Seven Stylized Facts," by Anil Ari, Carlos Mulas-Granados, Victor Mylonas, Lev Ratnovski, and Wei Zhao.

A Critical Matter

Christopher Evans, Marika Santoro, and Martin Stuermer

FRAGMENTATION OF CRITICAL MINERAL MARKETS WOULD SLOW THE SHIFT TO CLEAN ENERGY

A scramble by competing powers to secure strategic minerals could add to price pressures and increase the costs of the climate transition. New trade restrictions in commodity markets more broadly have doubled since Russia's invasion of Ukraine as producers impose curbs on shipments. Critical minerals used to make everything from electric vehicles (EVs) to solar panels and wind turbines are highly vulnerable to more severe trade restrictions. A slide toward opposing trading blocs could substantially delay the energy transition.

Even without the added complication of geopolitically motivated export controls, countries will need unprecedented supplies of critical minerals to stave off the worst effects of climate change and reach net zero emissions. The International Energy Agency predicts that demand for copper will need to grow by a factor of 1.5, for nickel and cobalt to double, and for lithium to increase six times by 2030 (Chart 1). This will drive

In an aerial view, salt evaporation ponds are seen on Bristol Dry Lake in California.

DAVID MCNEW/GETTY IMAGES

up prices and could make these minerals as important as crude oil for the world economy over the next two decades (Boer, Pescatori, and Stuermer, forthcoming).

Why are critical mineral markets particularly vulnerable in the event of fragmentation? And what could be the impact on the energy transition?

Extreme vulnerability

Minerals such as copper, nickel, cobalt, and lithium are critical inputs for the energy transition. They are used in EVs, batteries and wiring, and renewable-energy technologies such as solar panels and wind turbines. A typical EV battery pack, for example, needs about 8 kilograms of lithium, 35 kilograms of nickel, and 14 kilograms of cobalt. Charging stations require substantial amounts of copper.

Critical minerals are extremely vulnerable in the event of trade disruptions because their global production is highly concentrated. Two-thirds of the world's cobalt is mined in the Democratic Republic of the Congo alone. The top three producers of nickel and lithium control more than 60 percent of supply. Crude oil production is, by comparison, much more diversified (Chart 2).

The combination of concentrated supply and widespread demand has led to extensive commodity trading. Many countries rely heavily on imports from only a handful of suppliers. To make matters worse, mining production can be difficult to relocate. Even where there are deposits, it takes time and expensive investment to extract them from the ground. Minerals are often hard to substitute. For example, lithium is essential for many EV batteries. As a result, demand for them responds only slowly when prices rise amid shortages.

This trifecta of high concentration of production and low reactivity of supply and demand makes critical minerals for the energy transition highly vulnerable in the event of trade restrictions.

Transition delay

How would more severe fragmentation of critical mineral markets affect the energy transition? For illustrative purposes, a team of IMF researchers divided the markets for four critical minerals into two hypothetical blocs that refuse to trade with each other, along the lines of a 2022 UN vote on Ukraine.

Results show that the inability of the hypothetical China-Russia+ bloc to import copper, nickel, lithium, and cobalt from mining countries such as Chile, the Democratic Republic of the Congo, and Indonesia would lead to an additional price increase of 300 percent, on average. Acquiring minerals would be more expensive, which would

CHART 1

Spike in mineral demand

Mineral demand could increase up to six times as countries scramble to secure supplies for the clean-energy transition.

Demand increases for critical minerals until 2030
(index, 2022 = 1, net zero emissions scenario)

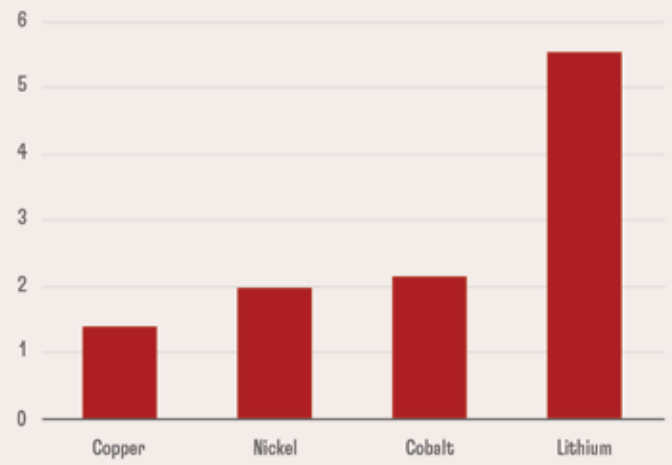

SOURCES: International Energy Agency, 2023; and IMF staff calculations.

CHART 2

Concentrated mineral market

The supply of critical minerals is more concentrated than that of crude oil and therefore more vulnerable to trade disruption.

Top three producers
(share of world production, percent)

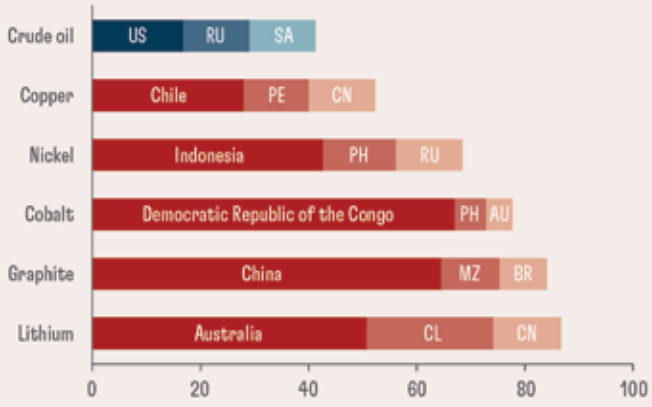

SOURCES: British Geological Survey; US Geological Survey; and IMF staff calculations.
NOTE: AU = Australia, BR = Brazil, CL = Chile, CN = China, MZ = Mozambique, PE = Peru, PH = Philippines, RU = Russia, SA = Saudi Arabia, US = United States.

CHART 3

Fragmentation penalty

The division of the world into two opposing trading blocs would slow the clean-energy transition.

Impact of critical mineral market fragmentation
(percent deviation from integrated markets scenario, 2030)

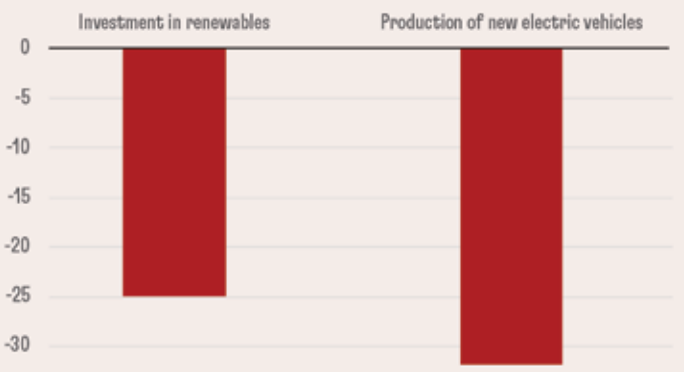

SOURCE: IMF staff calculations.
NOTE: The bars show the change in real investment in renewables and EVs due to fragmentation relative to the International Energy Agency's net zero emissions scenario with integrated copper, nickel, lithium, and cobalt markets. Country variables are aggregated using greenhouse gas weights.

lead to lower investment in solar panels and wind turbines and fewer EVs.

In the hypothetical US-Europe+ bloc, meanwhile, fragmentation would cause an oversupply of most of these mined minerals. However, the bloc's use of minerals would be constrained by the length of time it takes to scale up refining capacity. Fragmentation, therefore, generates only small gains in the US-Europe+ bloc by 2030: the bloc would produce slightly more EVs, but there would be no gains in renewable-energy capacity.

Decarbonizing the global economy would be more difficult if the market for minerals were fragmented. On balance, global net investment in renewable technology and production of EVs would be about 30 percent lower, if greenhouse gas emissions are used as weights to aggregate region-specific results (Chart 3). This measure accounts for the greater emissions intensity of activity in the China-Russia+ bloc and hence the greater effort needed to achieve global emissions mitigation goals.

International initiatives

Multilateral cooperation is essential to prevent vicious spirals where countries impose trade restrictions as a risk management tool. An agreement on enhanced World Trade Organization rules on export restrictions and tariffs as well as discriminatory subsidies would be the best solution.

If full cooperation is impossible, multilateral efforts should prioritize establishing a "green corridor," consisting at a minimum of agreement to maintain the free flow of critical minerals and not to discriminate between firms from different countries.

An international initiative to improve data sharing and standardization in mineral markets could also reduce market uncertainty. The international community should establish an institution or platform, similar to the International Energy Agency or the Food and Agriculture Organization, focused solely on critical minerals.

Individual countries can take proactive steps, too. Strategies could include diversifying sources of commodity supplies; greater investment in mining, exploration, and storage; and critical mineral recycling.

Industrial policies, meanwhile, must be designed carefully to ensure equal treatment of firms across competitive markets to prevent adverse cross-country spillovers, minimize distortions and inefficiencies, and mitigate fiscal risks and harmful political economy outcomes. "Friend-shoring" policies and local-content provisions can also distort markets and raise costs. Developing a framework for international consultation on friend-shoring could help identify negative cross-border spillovers and mitigate adverse consequences.

Fragmentation in critical mineral markets could make the clean energy transition more costly and potentially delay much-needed policies to mitigate climate change. Multilateral cooperation on trade policies and more data sharing would thwart additional obstacles to a cleaner global energy system. Critical minerals may someday be as important to the world economy as oil is today. We need a better understanding of their complex value chains. F&D

CHRISTOPHER EVANS *and* **MARTIN STUERMER** *are economists in the IMF's Western Hemisphere Department and Research Department, respectively.* **MARIKA SANTORO** *is a senior economist in the Strategy, Policy, and Review Department.*

This article draws on Chapter 3 ("Fragmentation and Commodity Markets: Vulnerabilities and Risks") of the IMF's October 2023 World Economic Outlook.

REFERENCE

Boer, Lukas, Andrea Pescatori, and Martin Stuermer. Forthcoming. "Energy Transition Metals: Bottleneck for Net-Zero Emissions?" *Journal of the European Economic Association.*

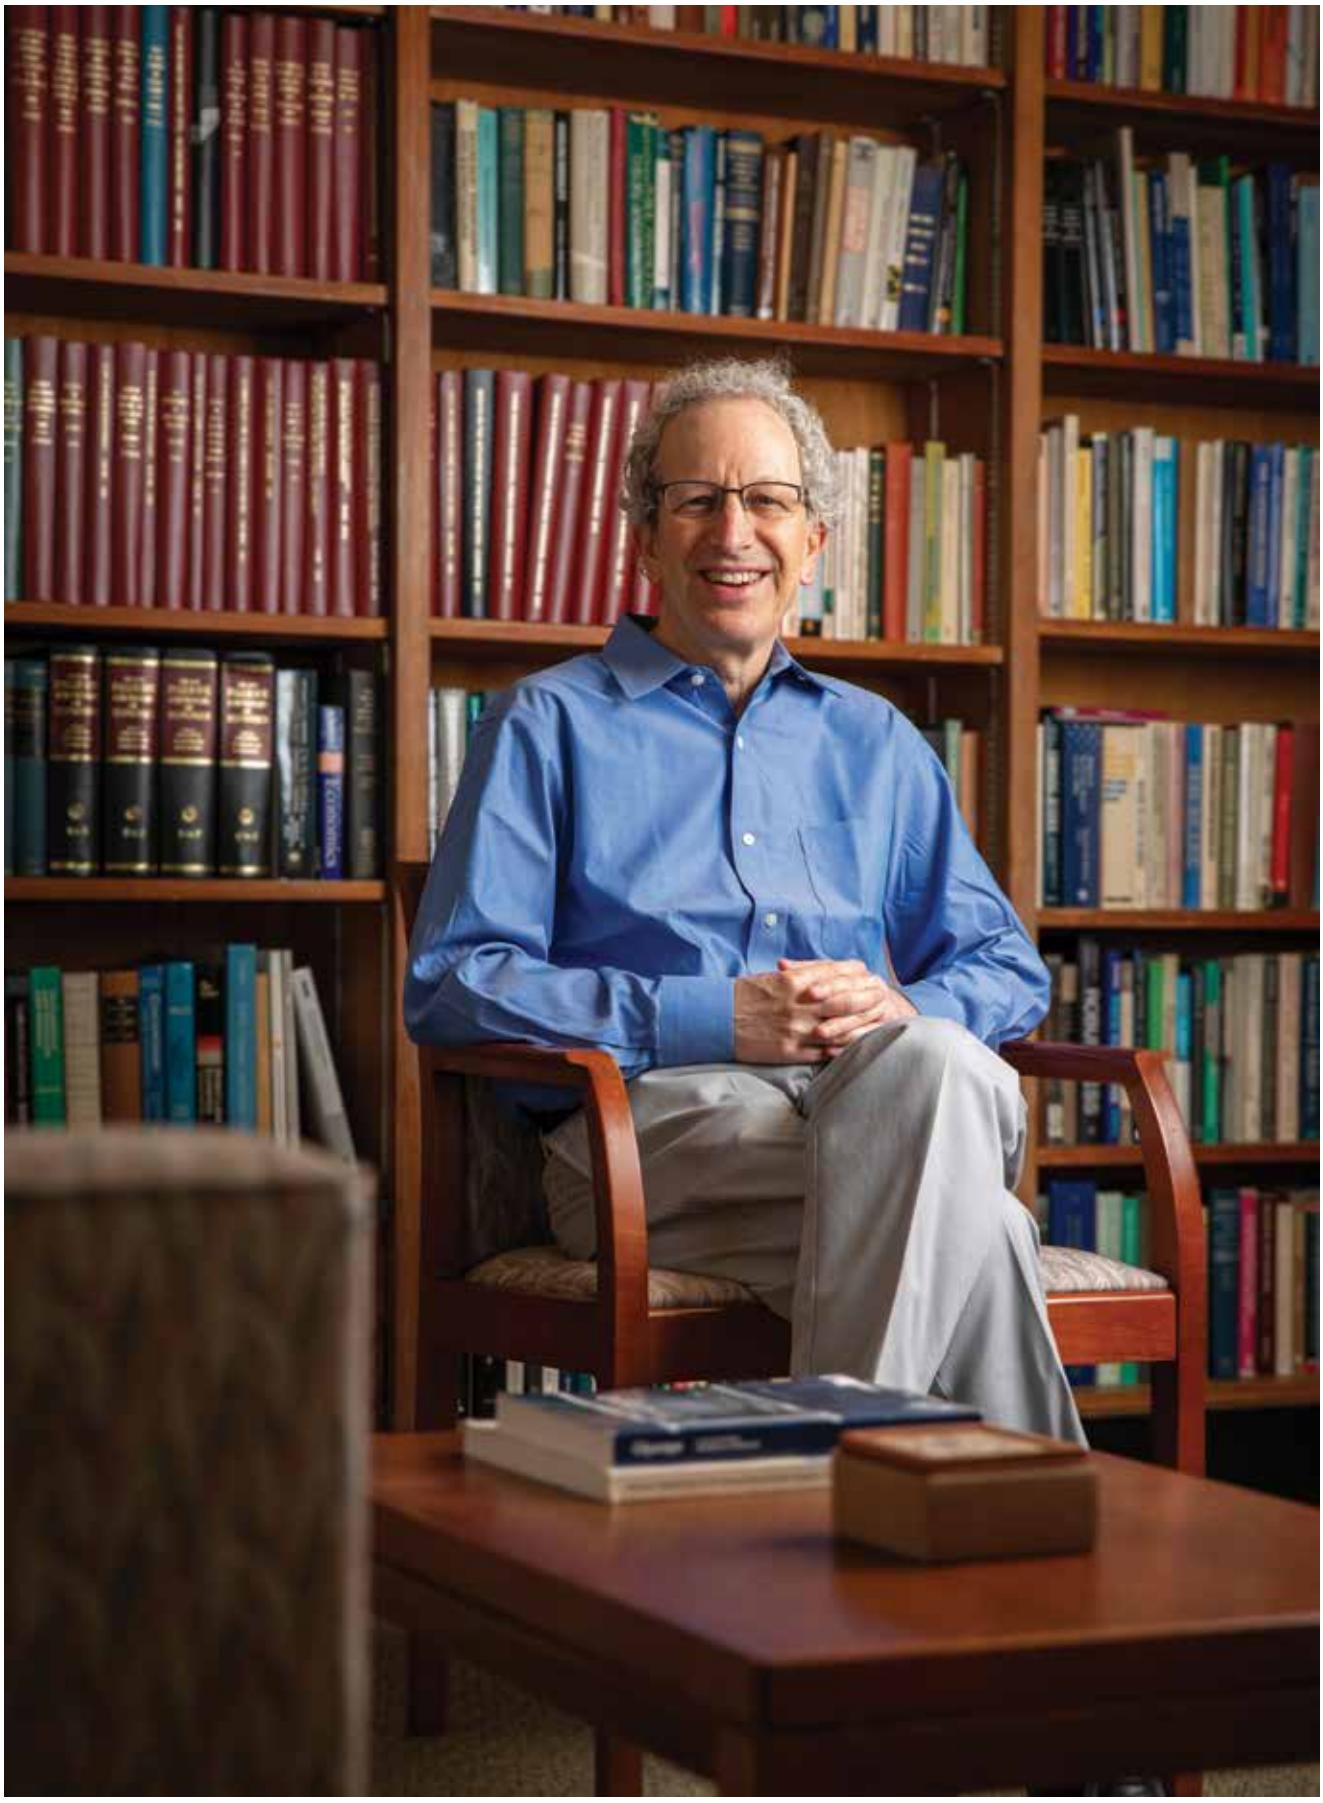

PORTER GIFFORD

People in Economics

The Inequality Economist

Bob Simison profiles Harvard's **Lawrence F. Katz**, whose research changed economists' understanding of economic disparity

LIKE THE REST OF US, HARVARD LABOR economist Lawrence F. Katz has been thinking about how artificial intelligence (AI) will change the future—especially what it will mean for inequality. Since the 1980s, he has made groundbreaking contributions to economists' understanding of the issue and what can be done about it.

Under one AI scenario, Katz says, the technology could help people who are already in advanced, high-paying professions, "thereby potentially exacerbating labor market inequality." Under another, it might help level the playing field for workers on the lower end of the scale.

"AI may increasingly substitute for elite expertise, making it less scarce and making the insights from elite expert knowledge more accessible to a broader range of workers," he says. "This scenario could help middle-skill workers versus elite professionals."

Whichever prevails, the 64-year-old Katz is likely to help lead the charge by academic economists—many of them his own protégés—to assess AI. Certainly, researchers will deploy his rigorous methods using big data and sophisticated analysis, reflecting his pervasive influence on economics over almost 40 years.

"He really casts a long shadow in economics," says the Massachusetts Institute of Technology's (MIT's) David Autor. "Larry was sounding the bell on the danger of rising inequality long before anyone else did." Autor is one of Katz's more than 200 former doctoral students. They include two winners of the John Bates Clark Medal, the top prize for economists younger than 40; three MacArthur Foundation grantees; and dozens of tenured scholars at top universities.

Katz's work sparked two intellectual revolutions in economics, according to a 2023 biographical sketch by Autor and Harvard's David Deming, another former Katz student. One was to apply economic theories of supply and demand to explain fluctuations in wage inequality over time. The other was to lead large-scale field experiments involving real people to answer big questions in social science, most prominently on the effects over multiple generations of moving to a higher-opportunity neighborhood.

"When the education system doesn't keep up, you have widening inequality," Katz says.

In addition, as the editor since 1991 of the august *Quarterly Journal of Economics*, "Katz has shaped the agenda of the economics profession over three decades," Autor and Deming write. They cite data suggesting that the QJE, as the journal is known, has had much greater influence on economics per paper published than any of the four other leading economic research publications, based on citations and other factors.

'Driven by social problems'

"Larry's really driven by social problems," says his wife and frequent research collaborator, 2023 Nobel laureate Claudia Goldin, another Harvard economist. "His passion is the underprivileged." (He does have another passion, she says: her champion scenting dog Pika, a 13-year-old golden retriever Katz walks several times a day.)

His passion for the underprivileged came from growing up as the son of a Los Angeles public school psychologist in the 1960s. His mother, born Vera Reichenfeld in 1938 in Belgrade, escaped the Holocaust with her family and grew up in Argentina and Uruguay. One of her teachers had studied at the University of Michigan. That led her to emigrate to Ann Arbor for college, where she met Katz's father.

As a Spanish speaker, she worked in some of the poorest neighborhoods in LA. Katz recalls that she took clothing and food to her schools for children from struggling families. He and his mother also discussed the hardships of attending schools without air-conditioning and whether that put kids in poorer schools at a disadvantage compared with students in richer, air-conditioned schools. Those encounters with poverty inspired Katz to focus on inequality, segregation, and race as a high school debater and as an undergraduate. (At the age of 85, his mother now works part-time as an actress in Spanish- and English-speaking roles.)

Katz likes to tell interviewers that he chose economics because the introductory class during his first quarter at Berkeley didn't meet until 10 a.m., and the beginning political science course was at 8 a.m. As an undergraduate, he began developing his data-driven approach to economic research when what's now known as the Fisher Center for Real Estate and Urban Economics hired him in 1979 as its first researcher.

He surveyed land-use officials from the 93 San Francisco Bay Area jurisdictions to collect reams of data showing how the recently passed property-tax-slashing Proposition 13 was leading to more restrictions on land use and driving up real estate prices. The findings turned into his senior thesis and his commencement address for the economics department.

In earning his PhD at MIT in 1985, Katz dug into the mechanics of unemployment. Rigorously analyzing US and UK data, he challenged an established theory that cyclical variations in joblessness grew out of shifts in labor demand requiring workers to move across sectors such as manufacturing and services. He showed that instead it had more to do with traditional business cycles from aggregate demand shocks. He further advanced understanding of the job-search behavior of workers who were temporarily idled and expected to be recalled, relying on longitudinal survey data.

This kind of large-scale data analysis opened a new frontier in economics long before advances in computing power enabled researchers to routinely crunch huge volumes of numbers. At that time, "public-use datasets came on nine-track tapes the size of deep-dish pizzas, and computer time was rented by the processor minute," according to Autor and Deming.

Katz put his approach on full display in 1992, when he published two influential papers. In one, he collaborated with the French macroeconomist Olivier Blanchard, who later served as chief economist of the International Monetary Fund. After joblessness nearly tripled in Massachusetts between 1987 and 1991 as a boom in tech and financial services went bust, they set out to understand what happens when there's a regional surge in unemployment.

Studying 40 years of state-by-state US data, Katz and Blanchard concluded that while it takes five to seven years for a state to rebound from an unemployment spike, the decline in the jobless rate reflects largely workers leaving the state rather than employers creating new jobs. It takes more than a decade for wages to return to normal.

"We found very solid patterns in the data, which gave a clear picture of labor mobility and regional evolutions," Blanchard says. The findings changed the way economists think about regional policies in other places, such as Europe, he says.

The other landmark 1992 paper directly addressed income inequality between people with and without college degrees. It overturned economists' thinking about earnings disparity. Katz and the University of Chicago's Kevin Murphy analyzed changes in US wages from 1963 to 1987, tapping into a vast Census Bureau dataset. They found that the income gap narrowed from 1970 to 1979 and widened dramatically after 1979. The conventional wisdom at the time attributed this to rising demand for workers with more education. But Katz and Murphy showed that it also reflected a sharp decline in growth of the supply of such workers relative to rising demand.

"When the education system doesn't keep up, you have widening inequality," Katz says. Those two research efforts set off "a work of passion" delving into inequality over the succeeding three decades, he says. One of the most important, longest-running projects was his collaboration with Goldin on their 2008 book, *The Race between Education and Technology*.

Education and inequality

The couple, who met in the late 1980s at the back entrance of the National Bureau of Economic Research in Cambridge, Massachusetts, started the research in the early 1990s. It grew out of some of Goldin's preliminary work on the history of education and its impact on wages, she says. "Larry was obsessed with changes in the wage structure," she says. "He was the first economist in the 1980s who saw the inequality gap expanding."

The researchers tapped into "tons of datasets" and manually tabulated Bureau of Labor Statistics data from the

early 20th century, Katz says. They dug up campaign materials created in the 1910s and 1920s by local school boards—in rural areas more than in cities—pushing the "high school movement" to prepare young people for better jobs.

This gave US workers a tremendous advantage as "America educated its youth to a far greater extent than did most, if not every, European country," they write. "By the 1930s, America was virtually alone in providing universally free and accessible secondary schools." Widening inequality in the US by the end of the 20th century reflected not so much the speed of technological change, they argue, as a shortfall in willingness to continue investing in education.

"We could have done the same thing with college and vocational education that we did with the high school movement," Katz says. "We have an incomplete postsecondary education revolution. We left it to families to provide that." Today he advocates investing more in state universities and in strong vocational education and sectoral employment training programs for high school graduates. (He and collaborators published a series of papers in the 2010s showing that employers put little value on degrees from costly for-profit colleges.)

Fifty to 60 percent of the rise in US wage inequality since 1980 grew out of the slowdown in educational advances relative to continuing growth in demand for college-educated workers, which widened the pay differential between those with and without college degrees, Katz says. Other factors include the decline of unions, the erosion of the federal minimum wage, the surge in executive and other top-end compensation, and the fissuring of supply chains with increased domestic outsourcing, greater use of the gig economy, and international offshoring, he suggests.

In 1993, Katz became chief economist of the Department of Labor during the administration of Bill Clinton. That put him in a position to help design what other economists call one of the most important social policy experiments in US history, the Moving to Opportunity housing mobility program.

After the LA riots over the 1991 police beating of Rodney King, "Congress felt a little responsible and passed a bill with some money for a demonstration project on neighborhoods," Katz says. The program began in 1994 in Boston, Baltimore, Chicago, New York City, and Los Angeles and included 4,604 families living in public housing in some of the country's poorest neighborhoods. The idea was to find out whether helping randomly assigned families move to a better neighborhood would benefit them economically.

It didn't, at first, Katz and other researchers found. But that was only part of the story. Participants did report improved physical and mental health, and as Katz and his colleagues continued following the group, something unexpected emerged. Children who were younger than 13 when their families moved to safer, lower-poverty neighborhoods had 30 percent higher earnings as young adults, were more likely to attend college, got into better colleges, and lived in lower-poverty neighborhoods as adults.

"Little did I know I would still be studying this more than 25 years later," Katz says.

The experiment has policy ramifications today as some local governments, such as Seattle's, apply the findings to recipients of housing vouchers. "Where you live affects how healthy you are and many other things," Katz says. "We could do a lot more using existing resources." The administration of Joe Biden sought to fund a broader program, but "it all got killed" in negotiations with Congress, Katz says.

The Katz effect

As editor of the QJE for the past 32 years, Katz has magnified his influence on economic research, other economists say. Under his leadership, the journal takes on big questions in social science and human welfare, extending the frontiers of economics, according to Autor and Deming. He pushes researchers to take risks and follow the data where it leads, others say.

"In the field, it's known as the Katz effect," says Harvard economist Raj Chetty, a John Bates Clark medalist and another Katz protégé. Chetty has played a leading role in studying the ramifications of the Moving to Opportunity project.

"He is highly respected by the authors, an exceptional feat," says France's Blanchard, who was coeditor of the journal with Katz for seven years. "Strong editors typically make many enemies. He has not." Katz reads and responds to every paper that's submitted, he says. The QJE receives about 2,000 submissions a year and publishes 48.

For the past 25 years, Katz has also played a role as mediator in labor negotiations and disputes between Harvard and various unions. He led what was unofficially known as the Katz Committee, which in 2001 issued a report on outsourcing that led to a wage and benefits parity policy between in-house and contracted-out workers. The policy aimed to allow Harvard to use outsourcing for efficiency gains but not to undercut unionized university employees.

Certainly a pillar of his legacy are the 239 PhD economists Katz has trained. He maintains an up-to-date nine-page list of them on his Harvard website, showing the year of each one's doctorate, initial posting, and current position. Many of them cite him as their professional inspiration.

"He is a prolific advisor who's had an enormous impact on public policy by nurturing so many leading economists," says the University of Michigan's Betsey Stevenson. "He was always available. He has an encyclopedic knowledge of research in the field and can instantly tell you where your project would fit in the literature."

As a graduate student, Stevenson was doing research on happiness and economics. She recalls telling Katz of her finding that winning the lottery usually makes people happier, at least initially.

"Winning the lottery probably wouldn't make me any happier," she says he told her. "It wouldn't help me write papers any faster." F&D

BOB SIMISON *is a freelance writer who previously worked at the* Wall Street Journal, *the* Detroit News, *and* Bloomberg News.

Straight Talk

Harnessing AI for Global Good

Gita Gopinath discusses how to maximize the benefits of artificial intelligence and manage its risks through innovative policies with global reach

Beginning in the 18th century, the Industrial Revolution ushered in a series of innovations that transformed society. We may be in the early stages of a new technological era—the age of generative artificial intelligence (AI)—that could unleash change on a similar scale.

History, of course, is filled with examples of technologies that left their mark, from the printing press and electricity to the internal combustion engine and the internet. Often, it took years—if not decades—to comprehend the impact of these advances. What makes generative AI unique is the speed with which it is spreading throughout society and the potential it has to upend economies—not to mention redefine what it means to be human. This is why the world needs to come together on a set of public policies to ensure AI is harnessed for the good of humanity.

The rapidly expanding body of research on AI suggests its effects could be dramatic. In a recent study, 453 college-educated professionals were given writing assignments. Half of them were given access to ChatGPT. The results? ChatGPT substantially raised productivity: the average time taken to complete the assignments decreased by 40 percent, and quality of output rose by 18 percent.

If such dynamics hold on a broad scale, the benefits could be vast. Indeed, firm-level studies show AI could raise annual labor productivity growth by 2–3 percentage points on average: some show nearly 7 percentage points. Although it is difficult to gauge aggregate effect from these types of studies, such findings raise hopes for reversing the decline in global productivity growth, which has been slowing for more than a decade. A boost to productivity could raise incomes, improving the lives of people around the world.

But it is far from certain the net impact of the technology will be positive. By its very nature, we can expect AI to shake up labor markets. In some situations, it could complement the work of humans, making them even more productive. In others, it could become a substitute for human work, rendering certain jobs obsolete. The question is how these two forces will balance out.

IMF PHOTO/KIM HAUGHTON

GITA GOPINATH *is first deputy managing director of the IMF.*

A new IMF working paper delved into this question. It found that effects could vary both across and within countries depending on the type of labor. Unlike previous technological disruptions that largely affected low-skill occupations, AI is expected to have a big impact on high-skill positions. That explains why advanced economies like the US and UK, with their high shares of professionals and managers, face higher exposure: at least 60 percent of their employment is in high-exposure occupations.

On the other hand, high-skill occupations can also expect to benefit most from the complementary benefits of AI—think of a radiologist using the technology to improve her ability to analyze medical images. For these reasons, the overall impact in advanced economies could be more polarized, with a large share of workers affected, but with only a fraction likely to reap the maximum productivity benefits.

Meanwhile, in emerging markets such as India, where agriculture plays a dominant role, less than 30 percent of employment is exposed to AI. Brazil and South Africa are closer to 40 percent. In these countries, the immediate risk from AI may be reduced, but there may also be fewer opportunities for AI-driven productivity boosts.

Over time, labor-saving AI could threaten developing economies that rely heavily on labor-intensive sectors, especially in services. Think of call centers in India: tasks that have been offshored to emerging markets could be re-shored to advanced economies and replaced by AI. This could put developing economies' traditional competitive advantage in the global market at risk and potentially make income convergence between them and advanced economies more difficult.

Redefining human

Then there are, of course, the myriad ethical questions that AI raises.

What's remarkable about the latest wave of generative AI technology is its ability to distill massive amounts of knowledge into a convincing set of messages. AI doesn't just think and learn fast—it now speaks like us, too.

"It's telling that even the pioneers of AI technology are wary of the existential risks it poses."

This has deeply disturbed scholars such as Yuval Harari. Through its mastery of language, Harari argues, AI could form close relationships with people, using "fake intimacy" to influence our opinions and worldviews. That has the potential to destabilize societies. It may even undermine our basic understanding of human civilization, given that our cultural norms, from religion to nationhood, are based on accepted social narratives.

It's telling that even the pioneers of AI technology are wary of the existential risks it poses. Earlier this year, more than 350 AI industry leaders signed a statement calling for global priority to be placed on mitigating the risk of "extinction" from AI. In doing so, they put the risk on par with pandemics and nuclear wars.

Already, AI is being used to complement judgments traditionally made by humans. For example, the financial services industry has been quick to adapt this technology to a wide range of applications, including introducing it to help conduct risk assessments and credit underwriting and recommend investments. But as another recent IMF paper shows, there are risks here. As we know, herd mentality in the financial sector can drive stability risks, and a financial system that relies on only a few AI models could put herd mentality on steroids. In addition, a lack of transparency behind this incredibly complex technology will make it difficult to analyze decisions when things go wrong.

Data privacy is another concern, as firms could unknowingly put confidential data into the public domain. And knowing the serious concerns about embedded bias with AI, relying on bots to determine who gets a loan could exacerbate inequality. Suffice it to say, without proper oversight, AI tools could actually increase risks to the financial system and undermine financial stability.

Public policy responses

Because AI operates across borders, we urgently need a coordinated global framework for developing it in a way that maximizes the enormous opportunities of this technology while minimizing the obvious harms to society. That will require sound, smart policies—balancing innovation and regulation—that help ensure AI is used for broad benefit.

Legislation proposed by the EU, which classifies AI by risk levels, is an encouraging step forward. But globally, we are not on the same page. The EU's approach to AI differs from that of the US, whose approach differs from that of the UK and China. If countries, or blocs of countries, pursue their own regulatory approach or technology standards for AI, it could slow the spread of the technology's benefits while stoking dangerous rivalries among countries. The last thing we want is for AI to deepen fragmentation in an already divided world.

Fortunately, we do see progress. Through the Group of Seven Hiroshima AI Process, the US executive order on AI, and the UK AI Safety Summit, countries have demonstrated a commitment to coordinated global action on AI, including developing and—where needed—adopting international standards.

Ultimately, we need to develop a set of global principles for the responsible use of AI that can help harmonize legislation and regulation at the local level.

"The advent of AI shows that multilateral cooperation is more important than ever."

In this sense, there is a parallel to cooperation on the shared global issue of climate change. The Paris Agreement, despite its limitations, established a shared framework for tackling climate change, something we could envision for AI too. Similarly, the Intergovernmental Panel on Climate Change—an expert group tracking and sharing knowledge about how to deal with climate change—could serve as a blueprint for such a group on AI, as others have suggested. I am also encouraged by the UN's call for a high-level advisory body on AI as part of its Global Digital Compact, as this would be another step in the right direction.

Given the threat of widespread job losses, it is also critical for governments to develop nimble social safety nets to help those whose jobs are displaced and to reinvigorate labor market policies to help workers remain in the labor market. Taxation policies should also be carefully assessed to ensure tax systems don't favor indiscriminate substitution of labor.

Making the right adjustments to the education system will be crucial. We need to prepare the next generation of workers to operate these new technologies and provide current employees with ongoing training opportunities. Demand for STEM [science, technology, engineering, and math] specialists will likely grow. However, the value of a liberal arts education—which teaches students to think about big questions facing humanity and do so by drawing on many disciplines—may also increase.

Beyond those adjustments, we need to place the education system at the frontier of AI development. Until 2014, most machine learning models came from academia, but industry has since taken over: in 2022, industry produced 32 significant machine learning models, compared with just three from academia. As building state-of-the-art AI systems increasingly requires large amounts of data, computer power, and money, it would be a mistake not to publicly fund AI research, which can highlight the costs of AI to societies.

As policymakers wrestle with these challenges, international financial institutions (IFIs), including the IMF, can help in three important areas.

First, to develop the right policies, we must be prepared to address the broader effects of AI on our economies and societies. IFIs can help us better understand those effects by gathering knowledge at a global scale. The IMF is particularly well positioned to help through our surveillance activities. We are already doing our part by pulling together experts from across our organization to explore the challenges and opportunities that AI presents to the IMF and our members.

Second, IFIs can use their convening power to provide a forum to share successful policy responses. Sharing information about best practices can help to build international consensus, an important step toward harmonizing regulations.

Third, IFIs can bolster global cooperation on AI through our policy advice. To ensure all countries reap the benefits of AI, IFIs can promote the free flow of crucial resources—such as processors and data—and support the development of necessary human and digital infrastructure. It will be important for policymakers to carefully calibrate the use of public instruments; they should support technologies at an early stage of development without inducing fragmentation and restrictions across countries. Public investment in AI and related resources will continue to be necessary, but we must avoid lapsing into protectionism.

An AI future

Because of AI's unique ability to mimic human thinking, we will need to develop a unique set of rules and policies to make sure it benefits society. And those rules will need to be global. The advent of AI shows that multilateral cooperation is more important than ever.

It's a challenge that will require us to break out of our own echo chambers and consider the broad interest of humanity. It may also be one of the most difficult challenges for public policy we have ever seen.

If we are indeed on the brink of a transformative technological era akin to the Industrial Revolution, then we need to learn from the lessons of the past. Scientific and technological progress may be inevitable, but it need not be unintentional. Progress for the sake of progress isn't enough: working together, we should ensure *responsible* progress toward a better life for more people. F&D

Book Reviews

Preparing for the Inevitable

Simon Sharpe

RAISING THE BASE OF A HOUSE AND mixing cement into its earth foundation make it less likely to collapse in a flood. Backup generators and pumping systems reduce the risk of power and water supplies failing, allowing hospitals to continue to operate. Robust health systems reduce the risk of disease following flooding. Education and communication systems can help people know when to stay home and when to seek higher ground. Functioning public services increase social trust, making it more likely that disaster response plans will be implemented successfully.

In *Risk and Resilience in the Era of Climate Change*, Vinod Thomas makes a strong case for a systemic approach to building climate resilience. For starters, risk assessment should be thorough and continuous. Investment is needed in early-warning systems, evacuation plans, and institutions that govern all elements of disaster preparedness and response. Resilience must be integrated from the outset into the design not only of buildings and infrastructure systems but also of national development strategies.

Above all, Thomas emphasizes the need to prepare, not just respond. Disaster proofing, he points out, adds less than a tenth to the cost of a new hospital. As climate change continues to progress, risks that once had low probability and high impact will become high-probability, high-impact events. Governments must anticipate this shift and continually prepare for the larger risks of the future, not merely for a repeat of events of the past.

Thomas presents emissions reduction and adaptation as intrinsically interdependent and goes so far as to describe decarbonization itself as an approach to building resilience. Time is at the heart of this relationship: just as in the COVID-19 pandemic, when slowing the spread of the disease was crucial to prevent hospitals' intensive care units from being overwhelmed, slowing the onset of climate change will give us more time to prepare for and cope with the extreme events it will bring.

Perhaps the most bracing part of the book, for some readers, is its chapter on economics. A distinguished economist himself, Thomas writes that "the mainstream economics profession has not been on board in the campaign for climate action." He criticizes leading economics journals for being largely silent on the issue, influential economic models for failing to adequately integrate scientific knowledge, and the profession in general for failing to appreciate the severity of the crisis and promoting a "persistently false dichotomy" between sustainability and growth. He advocates recognizing the value of natural capital, thinking about growth in terms of quality rather than quantity, and revising the curricula of economics and business schools to teach how growth can be made regenerative instead of destructive.

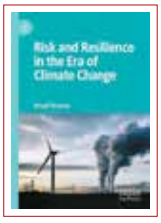

RISK AND RESILIENCE IN THE ERA OF CLIMATE CHANGE

Vinod Thomas

Palgrave Macmillan

London, UK, 2023, 201 pp., $27.99

If anything, this is one area the book could have pursued further. A shift in thinking on the economics of decarbonization has gathered pace in recent years. Whereas once there was consensus that carbon pricing was the most efficient solution, advances in the understanding of complex systems, studies of technology transitions of the past, and observations of what is happening now all suggest that approaches centered on investing in new solutions can be more cost-effective in driving innovation and structural change. It might have been interesting to explore the application of this new understanding to the transformational adaptation the author argues is now needed.

Based on three decades of the author's work at the World Bank and the Asian Development Bank, and underpinned by extensive academic research, the book is full of practical case studies as well as conceptual frameworks for understanding resilience. The importance of these issues will only increase as the climate crisis progresses and extreme weather events inevitably become more commonplace; this will serve as a valuable guide to those working in the field. F&D

SIMON SHARPE *is a senior fellow at the World Resources Institute and author of* Five Times Faster: Rethinking the Science, Economics, and Diplomacy of Climate Change.

Thinking about Inequality

Zia Qureshi

RISING ECONOMIC INEQUALITY in many countries, especially the rich ones, in recent decades has emerged as an important topic of political debate and a major public policy concern. Widening economic disparities and related anxieties are stoking social discontent and are a major driver of the increased skepticism about public institutions, political polarization, and populist nationalism that are so evident today. *Visions of Inequality,* a new book by Branko Milanovic, a leading scholar of inequality, places today's concerns and debate in context. It is an absorbing account of how thinking about inequality has evolved.

The book chronicles the way the economics discipline has viewed and analyzed inequality from the French Revolution to the end of the Cold War. It carefully distills the writings of six of the most influential economists of that time—François Quesnay, Adam Smith, David Ricardo, Karl Marx, Vilfredo Pareto, and Simon Kuznets—with a chapter devoted to each. It then reviews the work of neoclassical economists and inequality studies during the Cold War period. The book concludes by examining advances in the analysis of inequality in the more recent, post–Cold War period. In all, it is a sweeping, erudite treatise on the intellectual history of inequality.

The earlier economists' thinking about inequality was framed mainly around social classes and means of production—landowners, capitalists, workers. Their analysis focused more on the functional distribution of income—rents, profits, wages. Pareto framed inequality in terms of a social hierarchy of elites versus the rest of the population. With Kuznets and later neoclassical economists, the analysis shifted toward individuals and interpersonal distribution of income, a shift aided in part by greater availability of data on individual income. The new data and tools enabled the study of income distribution across individuals along various dimensions, such as educational attainment or urban versus rural location. Milanovic traces this evolution in economic thinking about inequality from classes to elites to people in rich detail, also showing how ideas about inequality were inextricably linked to historical context.

Milanovic calls the second half of the 20th century, spanning the Cold War, a "long eclipse of inequality studies." The relative lack of attention to distributional issues in part reflected the faith of neoclassical economists in the functioning of markets and their outcomes. In addition, inequality within Western economies initially moderated during this period, helped by rising demand for labor supported by stronger postwar economic growth, improvements in education, and the introduction of social welfare programs. According to Milanovic, these factors—which diminished attention to inequality in economists' work and public discourse during the Cold War—were reinforced by the politics of the era. Each side wanted to present itself as less class based and less unequal than the other: the competition between capitalism and communism pushed economics into the service of the ruling ideologies' political ends.

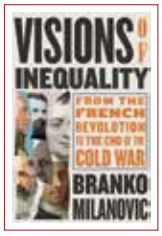

VISIONS OF INEQUALITY
From the French Revolution to the End of the Cold War

Branko Milanovic

Belknap Press

Cambridge, MA, 2023, 368 pp., $32.95

The picture has changed in recent decades. Inequality has resurged, driven by a combination of factors: the differential impacts of technological change and globalization across firms and workers and the current institutional and policy settings. These include the state's weakened redistributive role as tax progressivity declined and social programs were squeezed by tighter fiscal constraints. This has prompted economists to refocus their attention on inequality. Inequality has risen not only in Western economies, especially in the United States, but also in post-Soviet Russia and in major emerging market economies, such as China and India.

In the book's epilogue, Milanovic reviews how contemporary economists have expanded the frontiers of the study of inequality. The work of Thomas Piketty stands out in this context, especially in furthering the analysis of the role of wealth and nonlabor income in inequality. Inequality studies have a wider compass, reaching beyond a narrow neoclassical focus on markets to social and political power structures. These studies incorporate factors such as gender and race and examine inequality in broader dimensions than just monetary income. And the focus now extends beyond inequality among citizens within countries to include inequality among global citizens, an area of pioneering work by Milanovic.

There is a silver lining to recent inequality dynamics: inequality within countries has been rising, but global inequality (the sum of within-country and between-country inequality) has been falling. Developing economies are narrowing the income gap with rich countries. But, looking ahead, global economic convergence faces new challenges as the world's growth outlook weakens (especially for developing economies), geopolitical tensions and the risk of geoeconomic fragmentation threaten trade and investment between countries, and

technology alters the structure of international comparative advantage.

Visions of Inequality is an important scholarly work. But it is also a good read. Milanovic mixes his methodical examination of the evolution of economic thought about inequality with fascinating portraits of great economists and the society and polity of their times.

Renewed and deeper attention to inequality is timely. As Walter Scheidel documents in his book *The Great Leveler*, which reviews the history of the consequences of inequality, large and persistent increases in inequality can end up badly. F&D

ZIA QURESHI *is a senior fellow in the Global Economy and Development program at the Brookings Institution.*

The Flourishing Society

Vivek Arora

AMONG THE EARLY GRADUATE school memories of many macroeconomists are the "golden rule" of saving, the micro-foundations of wage and employment theory, and structural theories of unemployment. Edmund Phelps, the economist at the center of these fundamental insights and many others and the recipient of the 2006 Nobel Prize in economics, has written a new book that, as its captivating title suggests, describes his intellectual journeys from his earliest theories over six decades ago to his most recent ideas, how he came up with them, the people he met along the way, and the ideas they shared. He tells of his inspiration from the great minds he encountered, including many giants of modern economics, and from art, opera, and literature.

Creating any of the theories above could reasonably be considered a lifetime accomplishment, but Phelps is modest. He views these pillars of work as important but "[not requiring] ... a great deal of theoretical imagination" nor being "radical steps in economic theory."

Phelps' real passion, and what he sees as his crowning achievement, is his more recent theory of broad human "flourishing." Flourishing is about more than competing successfully in a free market and prospering in terms of money and material wealth. It is also about job satisfaction, rewarding work, and the wider notion of a good life that 19th and 20th century philosophers and economists envisioned.

A central concept in Phelps' theory of flourishing is "indigenous innovation." Such innovation, unlike innovation in the older tradition of Robert Solow or Joseph Schumpeter, is neither exogenous nor imported nor the sole preserve of famed inventors or entrepreneurs. Rather, it comes from the ingenuity of ordinary people going about their daily work. Phelps' key thesis is that when certain "modern" values—individualism, vitality, self-expression—are present, they tend to drive indigenous innovation, which in turn drives productivity, growth, and flourishing.

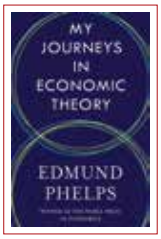

MY JOURNEYS IN ECONOMIC THEORY

Edmund Phelps

Columbia University Press

New York, NY, 2023, 248 pp., $27.95

The erosion of modern values in advanced economies across recent decades has, according to Phelps, contributed to the relative stagnation of productivity and real wages. A task for economic policy, therefore, is to help society regain these values in order to spark a new wave of dynamism and innovation. It will have to do so while dealing with overwhelming challenges, including climate change, digitalization, and the plight of low-income earners.

If the expanse of Phelps' vision seems overly broad for more conventional economists, it is worth knowing that Phelps is not alone. His more expansive view has been shared in different ways by notable economists. For example, in his recent book on inequality, Angus Deaton (another Nobel laureate) also notes the need for economics to take a broader view of human welfare than just performance in the marketplace.

There are caveats, of course. The goal of flourishing and job satisfaction, while inspiring, may strike some as more applicable to richer societies than to those where many people struggle simply to make a living. Phelps acknowledges that his theory of flourishing is not yet fully articulated in a formal model that can be tested. Moreover, the contention that a society's dynamism is driven by its culture and values, not simply by the incentives its people face, must deal with the observation that a society's economic performance can be transformed when its people's incentives change, for example through reforms. Experiences in Asia, including China in recent decades, and in Eastern Europe after the Cold War, are obvious examples. And the often high productivity of immigrants in their adopted countries testifies to the influence of the economic environment on people's fortunes.

Phelps' book is profound, far-reaching, and novel and combines analytical depth with a deep concern for economics to describe the lives not of "economic agents" but of actual human beings. The reader is guaranteed to emerge with a broader vision of economics and its possibilities. F&D

VIVEK ARORA *is deputy director of the IMF Independent Evaluation Office.*

On the Brink

Analisa R. Bala

Catalyzing support for the Philippine national bird could stem its rapid decline

The Philippines' new 1,000 peso banknote was released last year.

ZUMA PRESS, INC./ALAMY STOCK PHOTO

WITH A WINGSPAN OF MORE THAN SEVEN feet, steely gray-blue eyes, and a mane-like crest, the Philippine eagle is one of the world's largest and most striking raptors. It's also at high risk of extinction, according to the International Union for Conservation of Nature—a comprehensive information source of threatened animals, fungi, and plants.

Found on just four Philippine islands, the eagle is featured on the new 1,000 peso banknote to "highlight the importance of the preservation of this endangered species," says Sara Curtis, a director at the Philippine central bank. But because the birds are so difficult to track, there is still a "fundamental lack of information regarding their distribution and population size," says Dennis Salvador, executive director of the Philippine Eagle Foundation (PEF) and coauthor of a recent study in the journal *Animal Conservation*. Researchers estimate there are just 392 potential pairs left.

Eagle pairs need about 4,000–11,000 hectares of forestland to survive. Their range once covered 90 percent of the Philippines, or about 27.5 million hectares of forest cover. That span has dwindled to only 7 million hectares, according to data from the Department of Environment and Natural Resources.

PEF—an institution dedicated to the survival of these birds—is working with forest communities to reforest and protect cleared areas so that the birds can travel and hunt in safety. But if the eagles are to be saved, says Salvador, it will require a mass movement.

To raise awareness, two proclamations were issued by former Presidents Fidel Ramos and Joseph Estrada: one declared the eagle the national bird, the other established Philippine Eagle Week. In addition, a newly designed 1,000 peso polymer banknote was released in April last year, featuring the Philippine eagle and the national flower, sampaguita (a type of jasmine), on the front. The eagle exemplifies "Filipinos' strength and love for freedom," says Curtis.

The note's design and security features won it the International Banknote Society's 2022 "Banknote of the Year" award, and it is the first in the country to be printed on polymer.

Increased public awareness and a national wildlife conservation law have bought the eagle time, but the PEF is breeding the birds in captivity to keep them from going extinct. In an interview with *Living Bird* magazine, Salvador was hopeful: "I think we have a real chance of saving the eagles, even with the little we have left. It's just a matter of political will and attitude." F&D

ANALISA R. BALA *is on the staff of* Finance & Development.

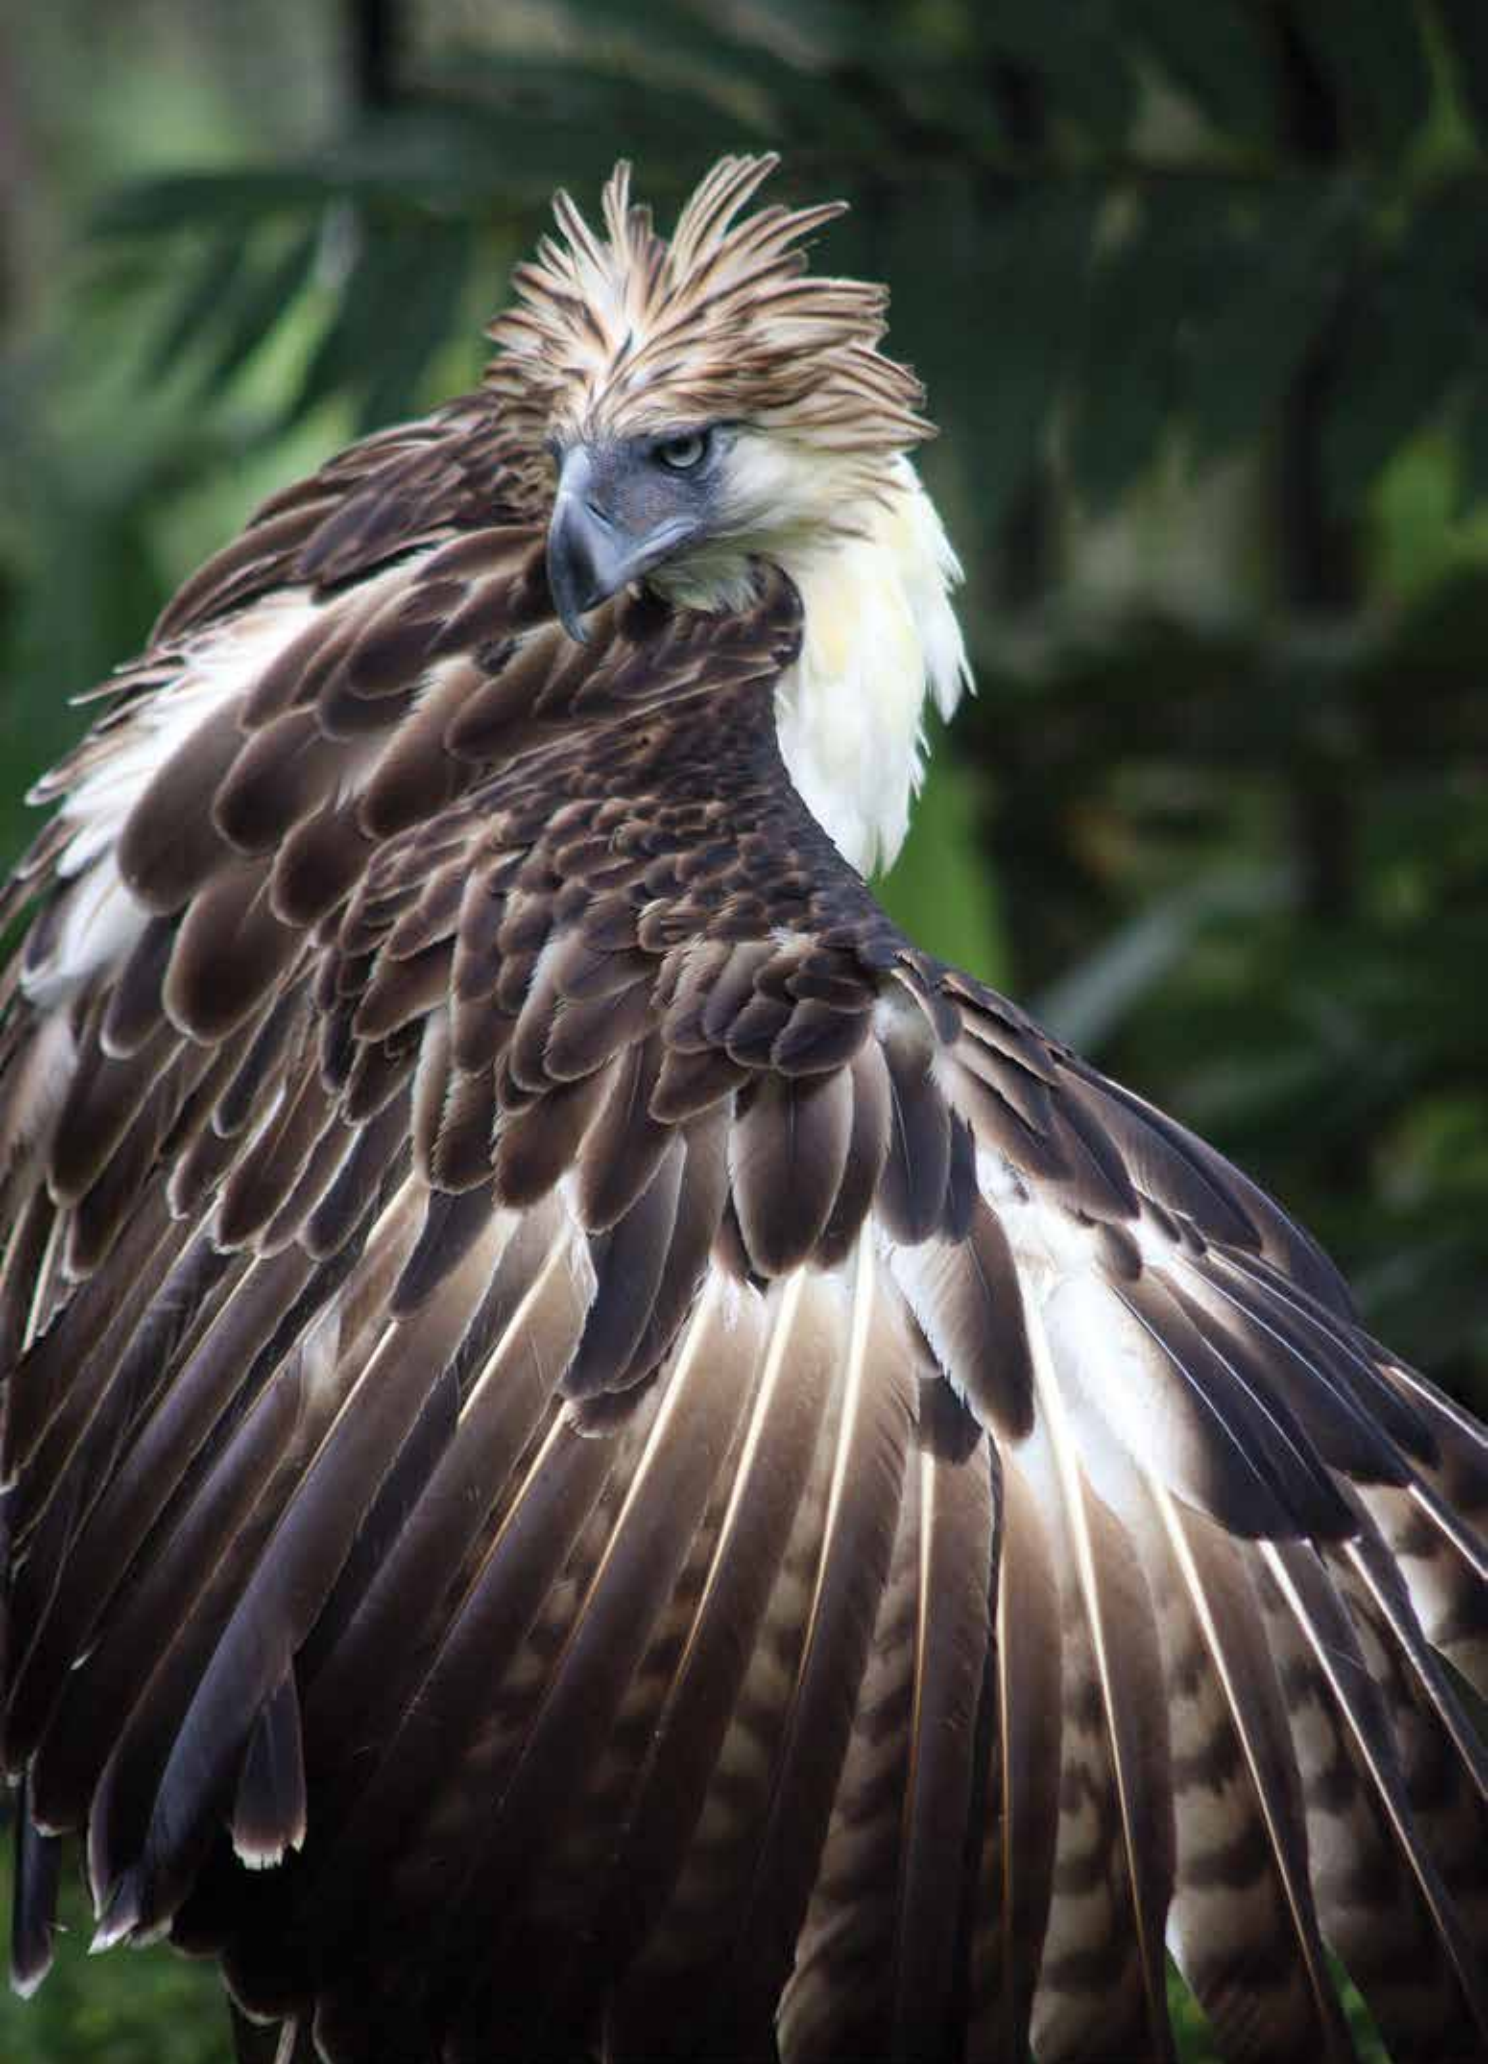

Timely. Topical. Free.

PUBLICATIONS

Global economics at your fingertips

IMF.org/pubs