

Economics and Business Review

Volume 12 (2) 2026

CONTENTS

Editorial introduction

Monika Banaszewska, Konrad Sobański

INVITED PERSPECTIVE

AI's biggest secret: We can shape it

Daron Acemoğlu

ARTICLES

Unlocking digital trade in Southeast Asia: The roles of infrastructure and regulation

Zuzana Brinčiková, Andrej Přivara

Relative institutional maturity of export markets and export performance: The moderating roles of managerial capabilities and export intensity

Piotr Trąpczyński, Tilo Halaszovich

Unlocking innovation: Governance and firm-level innovation in selected Asian countries

Aloysius Gunadi Brata, Pristanto Silalahi, Ronny Prabowo, Billie Ariyanta Hartono

Micro and small enterprise survival in contexts of poverty: Evidence from Peru

Ricardo Rafael Díaz Calderón, Carmen Graciela Arbulú Pérez Vargas, Juan Fernando Yalta Vallejos, Dante Hartman Cieza León

Beyond normality: Capital market Value-at-Risk modelling using symmetric and asymmetric Laplace distributions

Jan Kaczmarzyk

Economics and Business Review

Volume 12 (2) 2026

Editorial Board

Monika Banaszewska (Editor-in-Chief), *Ivo Bischoff*, *Horst Brezinski*,
Gary L. Evans, *Niels Hermes*, *Witold Jurek*, *Tadeusz Kowalski*, *Joanna Lizińska*,
Ida Musiałkowska, *Paweł Niszczoła*, *Michał Pilc* (Deputy Editor-in-Chief),
Katarzyna Schmidt-Jessa, *Konrad Sobański*

Editorial Advisory Board

Edward I. Altman – NYU Stern School of Business
John Cantwell – Rutgers Business School, Rutgers University
Conrad Ciccotello – University of Denver, Denver
Wojciech Florkowski – University of Georgia, Griffin
Oded Galor – Brown University, Providence
Binam Ghimire – Northumbria University, Newcastle upon Tyne
Linda Gonçalves Veiga – University of Minho, Braga
Christopher J. Green – Loughborough University
Eduard Hochreiter – The Vienna Institute for International Economic Studies
Mark J. Holmes – University of Waikato, Hamilton
Andreas Irmen – University of Luxembourg
Bruce E. Kaufman – Georgia State University, Atlanta
Robert Lensink – University of Groningen
Steve Letza – The European Centre for Corporate Governance
Robert McMaster – University of Glasgow
Tomasz Mickiewicz – Aston University
Victor Murinde – SOAS University of London
Dani Rodrik – Harvard Kennedy School
Hugh Scullion – National University of Ireland, Galway
Yochanan Shachmurove – The City College, City University of New York
Thomas Taylor – School of Business and Accountancy, Wake Forest University, Winston-Salem
Ari Van Assche – HEC Montréal
Thomas D. Willett – Claremont Graduate University and Claremont McKenna College
Habte G. Woldu – School of Management, The University of Texas at Dallas

Thematic Editors

• **Economics:** *Monika Banaszewska*, *Ivo Bischoff*, *Horst Brezinski*, *Niels Hermes*, *Witold Jurek*,
Tadeusz Kowalski, *Ida Musiałkowska*, *Michał Pilc*, *Konrad Sobański* • **Finance:** *Monika Banaszewska*,
Gary Evans, *Witold Jurek*, *Joanna Lizińska*, *Paweł Niszczoła*, *Katarzyna Schmidt-Jessa*, *Konrad Sobański*
• **Statistics:** *Marcin Anholcer*, *Maciej Beręsewicz*, *Elżbieta Gołata*

Language Editor: *Robert Pagett*

Paper based publication

© Copyright by Authors, Poznań 2026

© Copyright for this edition by Poznań University of Economics and Business, Poznań 2026



This work is licensed under a Creative Commons Attribution 4.0 International License
<https://creativecommons.org/licenses/by/4.0>

<https://doi.org/10.18559/ebr.2026.2>

ISSN 2392-1641
e-ISSN 2450-0097

POZNAŃ UNIVERSITY OF ECONOMICS AND BUSINESS PRESS
ul. Powstańców Wielkopolskich 16, 61-895 Poznań, Poland
phone +48 61 854 31 54, +48 61 854 31 55
<https://wydawnictwo.ue.poznan.pl>, e-mail: wydawnictwo@ue.poznan.pl
postal address: Al. Niepodległości 10, 61-875 Poznań, Poland

Printed and bound in Poland by:
Perfekt – Gaul i wspólnicy sp. k.

Circulation: 80 copies

AI's biggest secret: We can shape it¹

 Daron Acemoğlu²

Introduction

Tech leaders and industry insiders are giddy with excitement about the advances under way in artificial intelligence, whether due to the scaling up of existing models and functionalities, new reasoning models, promises of capable AI assistants, or entirely new, emergent capabilities.

Many hopes—and some fears—are centred on “powerful AI” or artificial general intelligence (AGI), which will be reached when AI becomes as smart as humans in almost every way. Some view AGI as an inexorable step towards unprecedented abundance, while others are concerned that such powerful AI could turn against humanity or, at the very least, create dangers for us. Almost everyone in the industry and many in tech journalism see this future as inevitable.

There is one word missing from much of this discussion: choice.

What AI will do to society, to labour and jobs, to misinformation and democracy and to geopolitical tensions is intertwined with the choices we make about how to develop this promising but still nascent technology. There are at least four choices, all of them interrelated. The first choice concerns AGI

Suggested citation: Acemoğlu, D. (2026). AI's biggest secret: We can shape it. *Economics and Business Review*, 12(2), 9–15. <https://doi.org/10.18559/ebr.2026.2.4267>

Licence

This article was originally published in *Prospect Magazine* in July 2025 and is reprinted by permission. All rights reserved. <https://www.prospectmagazine.co.uk/ideas/technology/70087/ai-artificial-intelligence-biggest-secret>

¹ Invited perspective. The text discusses topics related to those presented in Daron Acemoğlu's lecture delivered during the ceremony at which he was awarded the title of Doctor Honoris Causa by Poznań University of Economics and Business. Headings were added to the text by the journal's editors.

² Massachusetts Institute of Technology, 77 Massachusetts Avenue Cambridge, MA 02139, USA, <https://orcid.org/0000-0003-0908-7491>.

compared to AI designed as a tool for humans. The second centres on the automation of human tasks versus the expansion of human capabilities. The third is about large, all-purpose models—including “foundation models” such as OpenAI’s GPT, which powers ChatGPT—versus domain-specific approaches, which are designed to do one thing well. The fourth turns on whether AI will be developed in international conflict or through a potentially semi-cooperative process.

Let me take each one of these in turn and explain why these choices are currently being made in line with the interests of a small group of companies and executives, intent on making money and boosting their power over the rest of humanity, and why we should object to them.

1. Artificial General Intelligence versus AI designed as a tool for humans

While many industry figures make repeated claims that AGI is just around the corner, there are legitimate doubts about how easily existing models and approaches can achieve truly human-like high-level capabilities in every domain. But the bigger question is whether this is even desirable.

To explore this issue, we can turn to a very different vision of AI, powerfully articulated 65 years ago by JCR Licklider, a leading computer scientist and psychologist. Licklider deserves the accolade “grandfather of the Internet” more than anybody else due to his foundational ideas for the packet-switching architecture, which breaks information into smaller bits and transmits them separately over the network. This architecture was put into practice in an early computer network called Arpanet (Advanced Research Projects Agency Network). In 1960, Licklider wrote: “The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.”

Licklider’s aspiration was one in which computers would become information systems capable of expanding the potential of the human brain, by providing effective and useful information to human decision-makers. At the time, this was just an aspiration and no more; it continued to be so for the next six decades. The internet would come to contain much of the information of humanity, but nobody could use that information in a reliable, timely and useful manner, because it wasn’t easy for a person to find and process the bits of information that would be useful to them in that moment. Generative

AI, capable of creating new content or information on the basis of prompts, can change that.

With AGI, humans are sidelined and it is machines (or more often, the people who control them) that rule. Such dreams about AGI were present from the beginning of computer science and AI. Leading figures, such as Alan Turing (significant not just for his important mathematical work but for formulating how we can conceptualise computers reaching human-level capabilities) were articulating the possibility of such an AGI as early as 1949. Others, such as Marvin Minsky (one of the early defining figures in the field of AI) were trying to put it into practice in the 1960s.

Minsky's passion overrode repeated setbacks. In 1970, Minsky was still telling *Time*: "In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point the machine will begin to educate itself with fantastic speed. In a few months it will be at genius level and a few months after that its powers will be incalculable."

Generative AI has revived these dreams and given them a greater realism, although whether AGI can be achieved with the current architecture and approaches remains up for debate.

A sharp choice is therefore whether to focus predominantly on AGI or try to make Licklider's vision of machines amplifying human capabilities a reality. You might think that, technically, the cards are stacked against the latter. That would be wrong. In fact, even if overshadowed in the popular press and funding rounds, a drive like Licklider's has been at the root of many of the most important innovations in digital technologies over the last several decades. Ideas centred on human-machine complementarity, not inchoate theories about general intelligence, led to the technologies we routinely depend on today: computer mice, menu-driven computers, hyperlinks, hypertext and, of course, the internet.

It is also important to note that, even if AGI is not feasible in the short term, the relentless pursuit of it could cause significant harm. An important aspect of this harm would be the neglect of alternative directions that could powerfully enhance human capabilities. In this way, even if AGI is nowhere to be found, the mad dash towards it could become self-fulfilling, conjuring its harms into reality.

AGI and its associated aspirations are also bad for democracy, which requires an informed public. The conceit that machines are better decision-makers than humans, along with the centralisation of information that current approaches entail, is inherently anti-democratic. Once you accept that machines—or their designers—are much better than regular people and can gainfully sideline them, it becomes easy to understand why social media was likely to evolve in a way that cultivated misinformation and manipulation.

2. The automation of human tasks versus the expansion of human capabilities

The second choice is related but distinct. It concerns whether, when AI-powered commercial products are developed, they will predominantly strive to automate human tasks or to create new tasks and capabilities for human workers.

The costs of the former can be seen from previous rounds of digital technologies that have been extensively used for automating office work and factory jobs. This automation drive has increased profits and productivity to some degree, but has also boosted inequality. It underpins the declining real wages of a large fraction of the labour force in the United States. In contrast, my work shows that, during periods in which new technologies are used to create new tasks and capabilities for workers, there is an even bigger boost to productivity and a greater likelihood of accompanying wage growth. The bottom line is that an excessive focus on automation would not be consistent with shared prosperity, in which workers of different skill levels also enjoy some of the benefits of improved productivity.

To make matters worse, automation can be pursued relentlessly, even when its effects on productivity are minimal. Just like AGI, large-scale automation can become self-fulfilling even when AI isn't up to the job. Indeed, it is not difficult to find examples of digital technologies being adopted wholesale without a clear idea of how they can increase productivity.

With all the hype surrounding AI, it isn't hard to imagine how many businesses will feel greater pressure to jump on the bandwagon before they know how AI can help them; before they can contemplate how their organisation needs to be restructured for AI and human employees to work together; and before AI is ready for automation—or in fact for anything else.

3. Large, all-purpose AI models versus domain-specific approaches

The third choice is about architecture. Currently, much of the focus is on very large models such as GPT, which are all-purpose and aim to mimic human intelligence. Such foundation models can simultaneously write Shakespearean sonnets and produce marketing material for skin cream. The alternative is the domain-specific models where the industry has practised a degree of bait and switch. Some of the most celebrated successes, such as AlphaFold, are

domain-specific models: AlphaFold has been designed for predicting the 3D structure of proteins; it cannot market your skin cream and you should not ask it for dating advice (then again, you shouldn't ask that of ChatGPT either).

It is not easy to get deep expertise from AI when it is a jack of all trades. It also seems plausible that AI will be more likely to create hallucinations—giving users false, distorted or misleading answers—when asked to process and combine vastly different bits of information from a range of sources. Hence, it is worth investing in domain-specific models, especially given the difficulties the industry faces in monetising the foundation models (in part because they have not produced any applications, beyond coding, that appear particularly useful for the business sector).

4. Conflict versus semi-cooperative process in developing AI

Finally, the direction of AI could be pursued in a zero-sum manner between different countries, especially between China and the US, or there could be more information sharing and common regulatory approaches. Currently, the industry is opting for the first.

While these four choices about AI are distinct, they are also closely linked. AGI is a bedfellow with the foundation model (you wouldn't be able to generate human-like intelligence across the domains with domain-specific models). Both foundation models and AGI encourage automation—after all, if AI is on its way to achieving high-level human capabilities, it should take over tasks from humans. Last but not least, if AGI is the ultimate and near-term destination, a zero-sum approach makes a lot of sense, as it would be reasonable to presume that whichever country reaches AGI (and then artificial superintelligence) first will have a significant strategic advantage. This leaves little room for cooperation.

Conclusions

I disagree with all of these choices. AI can enhance human capabilities, resulting in significantly better social consequences. This objective is more likely to be achieved with domain-specific models that use high-quality data embedding expertise from the best-trained and most experienced human workers. Dreams of AGI can get in the way, and if the aim is to improve productivity

and find solutions to shared issues from pandemics to cancer, there is a lot of room for cooperation between the US and China—and any other country.

So, why are we locked into the AGI-automation, huge-model-conflictual approach to AI? I believe there are two main reasons. The first is economic, and the second is ideological.

Any technology requires corporate champions, and Silicon Valley companies are the natural investors in new digital technologies, possessing the capabilities and experience to scale them up effectively. Silicon Valley giants know how to monetise automation technologies and tools that can be used for digital advertising. It is less clear how to make big bucks from human-complementary technologies in the absence of either the business know-how or a clear market for them.

This mainstream approach also receives a significant boost from tech giants collecting vast amounts of data from people without paying for it, thereby receiving a substantial implicit subsidy for their key input. The antagonistic relationship between capital and labour in the US provides additional support for automation, which reduces the managerial dependence on workers. The fact that automation technologies cut wages and cause joblessness does not seem to concern them.

The ideological factor may be even more important. The belief that machines can and should ultimately surpass humans is an ideology. It might be based on mathematical and philosophical foundations, as Alan Turing's thoughts were. It could be derived from excessive exposure to science fiction books and movies. But more often, it is a different, veiled agenda: the actual subtext may not be so much that machines should surpass and dominate all humans. Rather, that machines designed and controlled by a few people should have dominion over the rest. As such, hidden in the AI decisions being made right now is an elitist, essentially authoritarian ideology, turning technology into a vehicle for the grandiose dreams of a small cadre.

This analysis suggests that the tech sector is making a series of connected choices on the type of AI to develop and its intended purpose. These choices carry a heavy social cost, but are still bolstered by economic and ideological factors. It is the tech sector making these choices, and yet we, as a society, are allowing powerful companies to step up to the plate.

There is an alternative: it is technically feasible and socially desirable to have a different trajectory for AI. This can be achieved if, instead of AGI, we aim for AI tools which are at the service of humans—just as Licklider hoped. Instead of automation, we should strive for pro-worker AI that can expand human capabilities and empower workers. Rather than an endless quest for bigger and bigger models, we should strive to build domain-specific models with greater expertise and higher-quality data. If we eschew the facile zero-sum thinking prevailing in policy circles and Silicon Valley today and attempt to find a more cooperative approach to the development and regulation of AI

globally, we may achieve more effective outcomes. Just like the choices currently being made in Silicon Valley, there are alternative options that could complement each other.

Yet, there is no automatic self-correction to our current path. Democracy, civil society, the labour movement, regulators and independent analysts must work together to improve the narrative (it's not a question of US AGI versus Chinese AGI, but AI that is good versus AI that is bad for humanity) and leverage this to build a research agenda and a regulatory framework to support the more socially beneficial direction of AI.

There is no silver bullet regulation that can take us from here to there, but essential steps can be taken. First, the path of least resistance for the tech industry—to collect as much data as possible and then monetise it via manipulative digital ads—needs to be made less attractive. This requires a digital ad tax imposed on tech platforms for their revenues from digital advertisements (but crucially not from other sources of income, such as subscription or software sales).

Second, domain-specific, pro-worker models require high-quality data. This underscores the need for a well-functioning data market, where companies pay for the day-to-day use of data and individuals are incentivised to generate the high-quality data on which pro-worker AI will have to depend.

Third, existing distortions resulting partly from asymmetries in the tax code—when lower taxes are imposed on capital income than labour income—and from the intense conflict between management and labour need to be resolved with fiscal tools. Both this conflict and existing tax distortions encourage excessive automation, to the detriment of workers.

Most importantly, AI researchers themselves need to recognise that every choice they make has social and ethical consequences. As in many other things, there may not be an absolute right and an absolute wrong, but choices have consequences, some good, some bad, some very bad indeed. Those who make decisions about AI also have responsibilities, whether they are engineers—or entrepreneurs in hoodies.

Aims and Scope

The aim of **Economics and Business Review** is to provide a platform for academicians from all over the world to share, discuss and integrate their research in the fields of economics and finance, including both behavioural economics and finance, with a key interest in topics that are relevant for emerging market economies. The journal welcomes submissions of articles dealing with micro, mezzo and macro issues that are well founded in modern theories or based on empirical studies and which are valuable for an international readership.

Your paper your way policy

The authors are initially expected to adjust their manuscripts to meet the basic requirements presented in the submission checklist below. Once the text has been accepted for publication, authors must adhere to all guidelines available on our website: <https://journals.ue.poznan.pl/ebr>

Basic requirements

- The submission has not been previously published nor is it under consideration for publication elsewhere (or an explanation has been provided in Comments to the Editor).
- The submitted manuscript must be anonymous. A separate title page must also be submitted, specifying each author's affiliation, email address, and ORCID iD. Acknowledgements and references to research grants should be included on the title page.
- The manuscript should be prepared in OpenOffice, Microsoft Word, or RTF document file format.
- The length of the manuscript should not exceed 8,000 words (including figures and tables, references, and footnotes).
- The manuscript includes an abstract of 100 to 150 words and is divided into clearly distinctive sections, including Introduction and Conclusions. The Introduction should state the aim of the study, research methods, main results, and particularly the study's contribution to international literature. The final paragraph should outline the article's content.
- All tables and figures should be numbered and presented consecutively according to their order in the text. Tables and figures should be as self-explanatory as possible, so readers do not need to refer to the main text to understand the information presented. The sources of all data used in tables and figures must be specified.
- The authors should use a consistent referencing style throughout the text.
- Please add an AI declaration (as a separate section in the manuscript) specifying whether and how you used generative AI during the preparation of the manuscript.

The submission must be made via the **submission system**: <https://journals.ue.poznan.pl/ebr/submission>

More information and advice on the suitability and formats of manuscripts can be obtained from:

Economics and Business Review

Al. Niepodległości 10

61-875 Poznań

Poland

e-mail: secretary@ebr.edu.pl

<https://journals.ue.poznan.pl/ebr>

Subscription

Economics and Business Review (EBR) is published quarterly and is the successor to the Poznań University of Economics Review. The EBR is published by the Poznań University of Economics and Business Press.

Economics and Business Review is indexed and distributed in Scopus, Clarivate Analytics, DOAJ, ERIH plus, ProQuest, EBSCO, CEJSH, BazEcon and Paradigm.

Subscription rates for the print version of the EBR: institutions: 1 year – €50.00; individuals: 1 year – €25.00. Single copies: institutions – €15.00; individuals – €10.00. The EBR on-line edition is free of charge.