

The Potential Challenges of Artificial Intelligence: A Political-Economics Perspective

by Roger Myerson, July 2026

<https://home.uchicago.edu/~rmyerson/humagnif.pdf>

Artificial intelligence (AI) has enormous potential to transform our world, for better or worse, in ways that defy our understanding. But we must try to understand the possible impacts of AI, lest its potential for harm be discovered only after it is too late to avert. In our efforts to anticipate the potential challenges from such extraordinary technological advances, we should seek a broad perspective, enlightened by general conceptual frameworks from the social sciences, and informed by some recognition of how human society has been transformed by technological developments in the past.

In this note, I try to offer such a perspective, to show how AI's promise of increased economic productivity may come with a threat to concentrate global power in a way that could undermine vital political institutions, unless we take some due care to reinforce their local foundations. Although I approach these questions as a social scientist, guided by basic models from economic and political theory, I agree with Pope Leo XIV on the potential value of insights from Biblical references to the Tower of Babel and the walls of Jerusalem under Nehemiah. My conclusions here will also concur with his *Magnifica Humanitas* encyclical in emphasizing the importance of subsidiarity as a fundamental principle for a society that can be strong enough to meet the challenges of AI today. But finally, as all our current predictions about AI's potentially disruptive social impact can be taken only as informed speculation, I must close by observing that the need for deliberate caution in implementing this new technology will be greatest in the leading nations that implement it first.

Rethinking the economics perspective on technological change

Since the beginnings of modern economics, economists have argued that the benefits to society of increased productivity from technological advances can be greater than the losses to people whose skills and former livelihoods have been made obsolete by the new technology. Such arguments are less persuasive, however, when it is unclear whether those who gain from the new technology will actually pay any commensurate compensation to those who lose from the obsolescence of their personal investments in jobs that disappear.

Economists can also observe that losses of jobs in one sector should, in the long run, be replaced by new jobs in other sectors, as wage rates adjust to equalize supply and demand for labor. But economic theory admits the possibility that technological advances could lead to lower wages in the new equilibrium, as may happen if the new technology enables owners of firms and capital to get increased profits and rents that claim a larger share of the national product. Thus, to evaluate the social impact of technological changes, we need to ask what kinds of jobs may be lost, and how the distribution of wealth and power in society may be affected by these changes.

The first basic point that we should recognize is that artificial intelligence will generally to be a substitute for human intelligence. If the availability of artificial intelligence decreases the premium that employers are willing to pay for human intelligence, then humans may have less incentive to make costly investments in their own education. Furthermore, when firms compare investments in training humans or AI machines, the possibility of new skills enabling human employees to demand higher competitive wages could add a motivation for firms to invest more in artificial intelligence instead of human intelligence. Thus, if we do not want to live in a society with less human intelligence, then we might appropriately call for a tax on artificial intelligence to generate revenues for subsidizing human education.

Such subsidies may be particularly needed in communities where so many jobs are lost that the next generation could grow up without adult role models for future livelihoods. In such communities, the cost of current job losses may be amplified by negative externalities of dislocation that could persist over generations. We humans evolved a propensity to find dignity in our work as adults when we see its similarity to work that we saw respected elders doing when we were young. So economic analysis should not ignore the adverse impact on long-term human welfare when technological change disrupts entire communities.

The impact of AI on jobs in the networks of economic power

There is reason to believe, however, that the impact of artificial intelligence could be substantially different from that of other recent technological advances. To see why, we may begin by considering a distinction among jobs that became evident during the recent covid-19 pandemic.

In the depths of the pandemic, many people found that they could continue to do their jobs from home by communicating with emails and video conferences over the global internet.

But many other people still had to go out to work, risking infection in the pandemic, because they had responsibilities to help take care of local matters that required their physical presence outside their home. We may refer to these two broad groups as the networkers and the caretakers. Many well-paid jobs with responsibilities for guiding decisions and policies in large organizations were included among those of the networkers who could work from home.

Although networkers might have been seen as a privileged group during the pandemic, jobs like theirs could be particularly at risk from the future applications of artificial intelligence. A person working at home through the internet would be regularly receiving digitally transmitted texts and data and images, applying personal knowledge and experience to process this information, and then generating new texts or data or images to help guide others in the network. But recent advances in AI technology have introduced powerful new methods for training computers to respond to digital information in ways that conform to patterns exhibited by human experts, when enough data about human expert responses is available. Thus, in principle, generative AI models could be trained to replace any kind of jobs that networkers could do remotely through the internet.

Technological changes in previous generations have tended to be very different qualitatively. Great technological advances from the 19th century included motors and electricity. Motors could substitute for human muscles in physical caretaking jobs, and electric circuits enabled the introduction of rapid long-distance information transmission, which could complement human brains' information-processing abilities. As a result, developments in 20th century permitted a remarkable increase in fraction of humanity whose jobs involved processing information in global networks that supported decision-making and resource allocation in large organizations. This expansion of a highly literate middle class may have contributed to the unprecedented spread of democracy in nations throughout the world.

Now, however, artificial intelligence threatens to reverse this trend, displacing people from jobs in the networks of economic and political power that manage the allocation of resources across the world, back towards jobs that focus on taking care of local resources. By reducing wage costs in the administration of large corporations and government departments, AI can lead to a world that is dominated by larger organizations which are directed by a smaller number of people. We should be concerned now about the possibility of a growing fraction of

humanity feeling disconnected and alienated from the global networks that exercise economic and political power over people in every community.

Reducing costs of moral hazard and concentrating benefits of power

In this regard, a relevant concept from economic theory is that of a *moral-hazard rent*. Large organizations regularly delegate substantial powers over resources and policies to agents whom they employ. In many cases, a job may include delegated powers that the agent could be tempted to abuse for personal profit, in ways that could be harmful to the employer's interests in the long run but would be difficult for the employer to observe in the short run. For such jobs, the employer may need to motivate good behavior by promising that, in addition to a competitive wage, the agent can expect long-term rewards that will be greater than the agent's potential benefits from abusing the delegated powers, if the agent maintains a record of good service. These supplemental rewards are moral-hazard rents for the responsible agent.

Economists understand that, in many cases, firms could avoid paying moral-hazard rents for local resource management jobs by letting the agent have a stake of ownership over the local resource. But delegated power over administrative networks cannot be so easily privatized. Thus, many of the most desirable jobs in our economy today get substantial moral-hazard rents from positions of power in global economic and political networks.

Now consider what could happen if an AI computer system could be trained to respond to complex network inputs like the best human experts in these positions, but without any human interest in abusing the delegated power of these positions for personal profit. Then employers that substitute computers for people in these positions could save the cost of their moral-hazard rents, which may be much more than the cost of a competitive living wage. From an economic perspective, eliminating the cost of these moral-hazard rents would be a clear productivity gain, allowing the employers to produce more goods and services for less resources.

From a political perspective, however, we should have serious concerns about a sweeping elimination of moral-hazard rents from large political and economic organizations. Throughout history, moral-hazard rents have been a substantial factor for compelling governing elites to share the benefits of power with broader groups of people in widespread administrative networks. Even before the development of constitutional democracy, national leaders understood that their ability to govern depended on a reputation for reliably delivering such promised rewards to political supporters and executive agents of the state; and fear of losing this

reputation could be a basic force for compelling national leaders to respect constitutional constraints that were broadly recognized in their society.

Thus, by enabling organizations to replace self-interested human administrators by artificial intelligence systems that are programmed to serve the interests of their corporate masters, AI technology may offer substantial gains of economic efficiency even as it undermines traditional mechanisms for ensuring that benefits of power are distributed more widely in society. These are fundamental reasons for us to have serious concerns about AI's potential to support historically unprecedented concentrations of power in the hands of narrow elites.

Lessons from Nehemiah about limiting centralized power by local empowerment

In countries where principles of constitutional democracy have been well established, we may hope that democratic accountability could also compel leaders to share the benefits of power more broadly. But national political coalitions are becoming increasingly dependent on AI-administered algorithms that can determine who gets wide attention on the internet. Thus, any hope for people to organize politically against a concentration of wealth and power in the hands of those who control the great AI-managed networks must begin with local political groups that can recruit members without the internet, using town meetings and face-to-face interactions among neighbors.

From this perspective, we can appreciate the wisdom of Pope Leo's suggestion that our thinking about the challenges of AI may be usefully guided by a comparison of two Biblical stories about ancient construction projects: the Tower of Babel, and Nehemiah's restoration of the walls of Jerusalem. Here we take the story of Nehemiah as a timeless reminder that even the greatest empires may still need autonomous local institutions for protecting people's welfare.

Among lessons from antiquity that can offer a perspective on the AI revolution today, particularly close parallels can be found in the Biblical story of the Tower of Babel, where people in a unified global society undertook the construction of a tower to reach the divine heavens, until their sacrilegious project was halted by a breakdown of their unified state. In the *Magnifica Humanitas* encyclical, Pope Leo suggests that this story should be contrasted with the Biblical story of Nehemiah's project to rebuild of the walls of Jerusalem, after the Persian emperors of Persia permitted a Jewish return from the Babylonian exile.

Both construction projects were undertaken within vast empires that seemed global to their inhabitants. Where the story of Babel warns us to respect some limits on what can be

achieved with centralized power in a great empire, the story of Nehemiah tells us about a devolution of power to autonomous local authorities within the great Persian empire. Nehemiah, a Jewish servant in the Persian emperor's court, asked the emperor for permission to rebuild the walls around Nehemiah's ancestral hometown of Jerusalem. With the emperor's approval, rebuilding these walls would enable the local leaders of this minor provincial town to provide some basic protection for their people.

So the walls of Jerusalem were a local initiative that could get imperial permission because the emperor listened to Nehemiah as a representative of his local provincial interests. In return for the emperor's support, his empire could be strengthened by the enhanced prosperity of the local community and loyalty of its trusted local leaders. Thus, the Persian Empire could be more durable than Babel because the Persian emperors encouraged local leaders to make decentralized local investments for promoting peace and prosperity in each part of their vast domain.

The general principle here is called *subsidiarity*, limiting the centralization of power by devolving greater responsibilities and capabilities to local communities and associations. But as AI enables power over global networks to be concentrated more narrowly, it becomes even more important that people everywhere should have locally accountable local leaders who can offer their communities some basic protection, as well as some voice to advocate local interests to national leaders.

Thus, concerns about the centralizing effects of AI may lead us also to ask, like Nehemiah, what can be done to strengthen local self-government for people everywhere. The modern analogues of walls for ancient Jerusalem could start with provisions that give a solid fiscal basis for autonomously elected local governments. It might also be worth considering a greater potential role for local governments to help certify their citizens' identities on the internet, and possible ways to boost local news sources in social networks on the internet. We could also suggest that, when AI engineers formulate basic constitutional principles to guide the behavior of their AI agents, respect for community leaders in autonomous local institutions should be listed as an essential principle.

The advantages of going slow on mass implementation of AI

The difficulty of identifying appropriate responses to the new challenges of AI has been increased by common arguments that policy-makers should do nothing that could slow down the

development of artificial intelligence in their country, lest other countries gain some crucial advantages from leading in the AI-development race. To conclude here, we should offer some general response to these arguments.

Let us first set aside the national-security component of these arguments. The tragically compelling logic of an arms race might be considered relevant for a few specific applications of artificial intelligence technology, such as in cybersecurity or in tactical direction of new weapons. But the great military powers that participate in arms races generally do not expect that their investments in new weapons systems would ever be scaled up for sale to people in society at large. The development of new AI capabilities that are considered necessary for national security could be financed from military budgets for the narrow purposes of military procurement, without building a wide national network of data centers for transforming the whole economy. So it is not clear why cutting-edge development of AI systems for national-security purposes would depend on broader national leadership in the development of general-purpose large language models and generative AI.

We have seen many good reasons to fear that the mass implementation of artificial intelligence technology could have a significant potential to cause social disruption, by eliminating responsible jobs, by reducing incentives to invest in human capital, and by enabling a greater concentration of economic and political power. The costs of such social disruption could be particularly acute in nations that become early leaders in implementing AI throughout their economies. By waiting to see the impact of AI in the nations that introduce it first, other countries that introduce AI later might find ways to regulate it so as to reduce its most harmful impacts. Indeed, we can imagine a possible future time when the world would be dominated economically and politically by people from the countries that introduced AI later and thus were better able to preserve their strengths of human capital and social capital. Certainly, the nations that introduce AI first should want to invest most carefully in trying to anticipate its potential adverse effects. Our inability to confidently predict all the potentially adverse effects of AI on our society should be a fundamental reason for caution.