



THE RACE WORTH WINNING

Middle Powers in the Age
of Machine Intelligence

Anton Leicht & Dean Ball

February 2026

Table of Contents

	Executive Summary	1
	Introduction	3
1	Between Worlds: The Middle-Power Challenge	6
2	The Substrate of a Revolution: Frontier AI Access	12
3	Ghosts in Your Government: Institutional Transformation	19
4	Bottlenecks for Abundant Intelligence: Economic Transformation...	25
	<i>Securing a Share of AI Growth</i>	26
	<i>Securing a Durable Position</i>	31
5	Navigating the Transition	33
	Conclusion	37



Executive Summary

Artificial intelligence presents the greatest transformation in human affairs since the industrial revolution, and it may be greater still. It promises material and informational abundance: machines to carry out the dangerous and arduous, leaving humanity wealthier, safer, and freer as a result.

Yet challenges come with that: AI looks likely to upend the economic and industrial strategies of all countries. It may erode or reinforce the informational advantage that has underwritten the state's claim to govern. And it may throw off or bolster the favorable balance between defense and offense that secures a state's ability to protect its citizens from harm.

The nations with the most to lose and gain from this transformation will be the middle powers: nations with significant institutional and industrial capacity but without frontier AI development capability. Unlike great powers, they do not control whether and how AI gets built, and whether it rips through society or not. Yet unlike the least developed nations, middle powers possess institutional substrate—functioning economies and bureaucracies. They have something to transform. Unburdened by the need to win a race they were never going to win, they are free to—but must—ask different questions on how to deploy the assets they still have to secure the future they want.

The race that remains open, the race where the middle powers can compete and even lead, is the race to favorably adapt frontier artificial intelligence into their governments and economies. This fact cannot stand as an excuse to remain inactive, but instead has

to be understood as a generational effort. It will require finding answers to encompassing and complicated questions that have no comfortable answers:

- **ACCESS ARCHITECTURE:** What contractual, technical, and institutional arrangements can durably guarantee middle powers' access to imported frontier AI?
- **INSTITUTIONAL TRANSFORMATION IN PRACTICE:** How can the public sector be reshaped to be amenable to AI-driven transformation? What separates the risks of falling behind and the opportunities for leapfrogging?
- **ECONOMIC BOTTLENECKS:** Where along the flow of AI through the global economy can a given middle power occupy a bottleneck valuable enough to safeguard its fiscal and economic prospects?
- **REDISTRIBUTION MECHANICS:** How can middle powers support their citizens when AI-driven revenue shifts simultaneously from labour income to capital returns and from domestic economies to the countries where AI is built?

Middle powers that fail will do so in familiar ways with unfamiliar dramatic consequences: allowing the legitimacy of the state to erode, taking half measures and easy roads, shutting their eyes and closing their borders to a global phenomenon, and waiting for a sovereign solution that never arrives. Those that move with clarity and resolve will shape the transformation.

Introduction

Artificial intelligence presents the greatest transformation in human affairs since the industrial revolution, and it may be much greater still. While many today pay lip service to this notion, few have internalized it. The revolution is not some distant hypothetical whose contours can be shaped from within the halls of government, academia, or think tanks alone. Instead this revolution is unfolding, and will unfold ever more rapidly, now—with the *status-quo* political economy, institutions, technical capabilities, industrial power, and diplomatic relations particular to each nation-state on Earth. Policymakers will not grapple with this revolution using some imagined future toolkit but rather the toolkit in their hands today. The era of machine intelligence will not “begin soon.” It has begun.

Zoomed out far enough, the stakes of this era can seem existential: we face generational risks, but should we navigate them, what awaits on the other side is something like material abundance—machines to shoulder the dangerous and arduous work that has always fallen to human hands, leaving us wealthier, safer, and freer for it. Once again, and at a scale without precedent, a technological revolution that would make the world genuinely better appears to be afoot. But even in futures where AI proves to be positively transformative, it places unprecedented pressures on the statecraft of all nations. Nowhere is that pressure stronger than in middle powers—those countries that hold just enough cards to shape their response to this transformation for incredible gain, but that have just as much to lose.

Almost no government or policymaking apparatus in any country anywhere on Earth has really internalized this reality. It is a difficult reality to internalize. In the abstract, technological revolutions promise greater wealth, freedom, and abundance to the world. But in the concrete, this future will not be evenly distributed. AI may well upend the economic and industrial strategies of many countries worldwide, particularly those which staked their futures on growth in the knowledge and service sectors. Suddenly,

the path forward for many industries within those sectors looks more precarious than it once did. At the same time, as we have learned repeatedly since the COVID pandemic snarled global supply chains, controlling manufacturing capacity, especially but not exclusively at supply-chain chokepoints, is a source of tremendous international leverage.

Even beyond the much-discussed economic effects of AI, there are precious few domains of human activity that the revolution in machine intelligence will not touch. The adaptations of human beings and our institutions to these transformations may collectively give rise to a new order of human affairs—one which is capitalistic, but not quite “capitalism” as we know it today; one in which business is vigorously conducted, but by organizations that do not quite resemble what we now think of as “a business.” The role of the nation-state in this transformation is uncertain, and even the fate of the nation-state is a live object on the field of play.

The outcomes in every respect are deeply unclear, often in more fundamental ways than the leaders of present-day institutions would care to admit. What this means is that many institutions—and especially the institutions of government, insulated as they are against the forces of the market—will probably stumble badly as they encounter the new terrain. Some may fall down altogether. But where peril abounds, there is also great opportunity. The leader who navigates his institution wisely through the transformation to come will face rich rewards indeed. The leader who holds steady at the helm may deliver a new era of prosperity, abundance, and joy to his people.

Those who emerge successfully out of the revolution to come will not be the ones who dreamt up the most exquisite regulation to impose on a foreign purveyor of technology. They will be the ones who navigate the transition to the era of machine intelligence with grace, savvy, cunning, and aplomb. Doing this will require the creative use of AI and other technologies, arguably more than it will require regulating them: regulation has a way of constraining options and creating path dependency. There is a risk that regulating today is less akin to playing chess than it is to taking chess pieces off one’s own side of the board before the game has even begun. The asymmetry of such a move is particularly stark since governments can strip themselves and their law-abiding citizens

of AI, but cannot guarantee to bind their competitors in the global marketplace or the adversaries that threaten them. As much is becoming more and more clear as great powers are less and less willing to have their firms and institutions abide by the rules set by foreign markets.

In the face of that realisation, wise policymakers in the rest of the world will shift their regulatory attention—to methods to grapple with the harms and risks of AI deployment that both of this paper’s authors have discussed in the past, perhaps to long-term prospects of multilateral governance, but primarily to how the deployment of AI to governance and government can uplift states’ capacity to deal with the ramifications of the transformation. Clever governments will remember that their bureaus and agencies themselves may yet be invigorated by a technology that makes careful attention to and drafting of complex, text-based documents functionally free. The compounding benefits to legislation and enforcement could be the tool that governments need to keep up with the rapidly changing risk profiles that have them worried even today. Indeed, such a technology could change the nature of “regulation” itself, and this is just one of the many profound new capabilities machine intelligence delivers.

The states that first grasp the transformative effects of AI as a technological instrument of governance, who mechanize the cognitive machinery of public institutions just as once in the past they were electrified, who understand not only “AI in government” but *AI as governance*, will also be the first to discover profound strategic insights for navigating the world ahead. They will have an advantage over all other nations, including the hegemons, if they were to stumble and miss a step. If AI is a race among governments, it is a race to discover the contours of the institutional revolution that is brewing. Many of the obvious approaches are wrong, many of the easy paths are treacherous. This paper hopes to elucidate some of the hard questions middle powers ought to ask to begin competing in that race.

Between Worlds: The Middle-Power Challenge

This report deals with one group of nations in particular. Those with the most to gain—and perhaps the most to lose—from navigating this transformation wisely may be neither the hegemons nor the smallest states, but rather the middle powers: nations with significant industrial and institutional capacity but usually without frontier AI development capability.

The middle powers occupy a distinctive position in this transformation. Unlike the least developed nations, they possess what we might call *institutional substrate*—functioning bureaucracies, judicial systems, regulatory bodies, universities, and civil services. They have much to transform. But unlike the United States and China, they cannot realistically build frontier AI ecosystems from whole cloth. The capital requirements, the talent concentrations, the existing industrial bases, and the sheer momentum of the leaders make this path essentially foreclosed. It is mandatory for them to seek agreements and even dependencies, to chart their course through a transformation governed by others.

This recuses middle powers from some of the most pressing questions that the American—and in parts, the Chinese—policy environment face in dealing with the thorny challenges of regulating the development of frontier AI. Questions of safety in development, of security against the risks of internal deployment, of susceptibility to misuse and loss of control, are not immediately in scope for middle powers: not because they should be disinterested in them or because we suggest they should take the accelerationists' response to them, but because the middle powers currently simply lack the leverage to exert their will over the modalities of deployment in either of the great powers.




Figure 1. G20 countries (excepting U.S. and China) and a selection of further middle powers.

These middle powers differ enormously from one another—in institutional capacity, in economic structure, in alliance commitments, and in their existing exposure to AI-driven growth. Germany’s challenge is not India’s; South Korea’s options are not Brazil’s. A country with a position in the semiconductor supply chain faces a fundamentally different starting point from one whose economy is built on services or commodities. This paper does not pretend otherwise, and the research agenda outlined at the end is deliberately designed to be pursued with regional and sectoral specificity. But we believe these nations share a structural condition that justifies treating them as a category: none of them will build frontier AI, all of them must import it, and each of them possesses enough institutional or industrial substance that the choices they make will genuinely matter far beyond their borders. The framework that follows is intended to clarify those shared choices—not to substitute for the country-specific strategies that must follow.

Middle Powers and the World

This reality confronts middle powers with a changing geopolitical order, even at a time where they already face a shift of alliances and dependencies. There is more to be gained from *using* AI to transform the mechanics of the status quo in imaginative ways than there is to be gained from fearing it, even when the AI in question originates in another country, and even when that country is the United States. The alternative is to take on the immense and heretofore unsurmounted challenge of creating a frontier AI development ecosystem, on top of all the other challenges this transformation is likely to pose for policymakers and political leaders—which seems impossible.

Yet if we ask middle powers to take on a role as AI importer, they must be embedded in a geopolitical environment that allows them to trust that setting, long-term. On one hand, a cornerstone of the Trump administration's technology policy agenda, for which one of this paper's authors served as a key architect, is to let American companies do productive business with countries around the world, not to keep AI shackled within American borders. Yet on the other hand, recent efforts of the administration to unilaterally redefine the structure of its alliances and the balance of power abroad have made allies understandably hesitant to entrust themselves to a hegemonial order. They rightly feel that they will need leverage to back up agreements with the U.S., and refuse to rely on good faith alone.

Against that backdrop, no middle power will chart a successful course entirely on its own, nor will it prevail by exposing all its assets to any single bilateral relationship, lest it grow parasitic and extractive. The solution to this challenge lies in coordination among the middle powers, in very much the way that insightful Western leaders have recently remarked. The most nimble version of this coordination does not take the detour through multilateral institutions, but arises from narrow agreements between like-minded middle powers in pursuit of similar goals.. Coordination, done right, reassures partners that they can focus on their comparative advantage instead of pursuing full autarky; and it enables powers to see eye to eye with AI exporters instead of being picked off by great powers offering rewards for defection one by one.

The U.S. government, too, should ultimately welcome this increase in coordination, and not understand it as a rally of anti-American sentiment. Allies coordinated among each other serve the American interest as well—they become more capable partners because their coordination allows them to concentrate on their comparative advantage, and they become more reliable importers because coordination quells the sovereignty concerns that would otherwise make them vulnerable: to failed domestic aspirations or to trade with China. America should therefore meet the prospect of middle power coordination with equanimity, and stand ready to export its AI to any willing group of middle powers.

The Middle-Power Opportunity

For the middle powers, then, the alternative path—using AI to transform the mechanics of governance and society—is not merely preferable. It is mandatory. This is not a limitation. It may be an advantage. The hegemons are deeply invested in the current paradigm of AI development: the race to foundation models, the competition for chips and data centers, the scramble for talent. This investment can become a form of blindness. The middle powers, unburdened by the need to win a race they were never going to win, are free to ask different questions—not “how do we build the next frontier model?” but “how do we use the models that exist to reinvigorate our courts, our tax administration, our schools, our agriculture, our urban planning?”

The middle powers must also reckon with the new salience of the physical economy. Chips, pharmaceuticals, critical minerals, energy infrastructure: the nations that make and move physical things occupy a position that no amount of cognitive automation can displace. For the middle powers, this creates both peril and opportunity. The service-sector strategies that once promised a shortcut to prosperity may be foreclosing. But many middle powers possess manufacturing bases, natural resources, strategic geography, or some combination thereof.

These two imperatives—institutional transformation and industrial development—are not in tension, but are instead complementary. The middle powers that transform

their institutions most effectively will also be best positioned to attract the manufacturing investment, manage the complex logistics, quickly adjudicate regulatory matters, and administer the industrial policy that physical-economy strategies require. The middle powers that grasp this—that pursue both tracks simultaneously, using AI to reinvigorate governance while positioning themselves in physical supply chains—will find themselves with options that others lack.

Yet the middle powers also face distinct temptations. These are sophisticated states, often proud of their regulatory traditions, with capable technocracies eager to demonstrate mastery over a new domain. The temptation to regulate—to impose frameworks, to set standards, to define what AI may and may not do within their borders—is strong precisely because they have the institutional capacity to do so. But as we have argued, this temptation is dangerous. Regulatory frameworks constructed now, before the contours of the transformation are remotely clear, risk constraining options that leaders do not yet know they will need: to leverage data from the physical or digital world, to allow machines to work in close conjunction with humans, or to provide value precisely by not requiring humans in the loop. The middle powers have more chess pieces to remove from their own side of the board, and a greater temptation to remove them.

The stakes for the middle powers are substantial and concrete. Consider what institutional transformation could mean: court systems where cases now pending for years or decades could be resolved in months, and perhaps in the future days or even hours; economic data collected and processed live; agricultural knowledge transmitted to farmers with a precision and timeliness previously impossible; civil services tailored for each recipient and administered with genuine conscientiousness; public health systems that identify emerging threats in real time and rapidly mitigate them; educational institutions that adapt to each student rather than forcing each student to adapt to them; city streets that are safer and more free. All of this is possible with the technology of the present day. What does not yet exist is the institutional imagination to seize them—or the wisdom to avoid foreclosing them through premature regulatory constraint.

The middle powers also face a question of talent and ambition. The brightest minds in every nation look for opportunity. If the institutions of the middle powers remain static—or worse, if they become mired in regulatory complexity while the frontier nations race ahead—talented young people will continue to leave for the places where the future is being built. But if the middle powers become laboratories for institutional transformation, if they demonstrate that the most interesting questions are not about building models but about remaking society with them, they may find that the current of talent reverses. The nation that first discovers how to truly govern with AI, rather than merely govern AI, will become a destination for those who wish to participate in that discovery, and a beacon to inspire and attract nations the world over.

The Substrate of a Revolution: Frontier AI Access

If middle powers are to seize the institutional and industrial opportunities outlined above, they must first secure the most basic input required to do so. Before questions of adoption, reform, or economic positioning can even be meaningfully asked, there is a more fundamental constraint to confront: access to frontier artificial intelligence, built in China and America. Many middle powers still reject this uncomfortable reality for its implications for geopolitical dependency and capital expenditure, and instead consider a bet on second-best technology. This would be a fatal mistake, betraying a marked ignorance of historical reality: whenever an infrastructural input became the driving force of a historical era, those who failed to secure an abundance of it have fallen behind. In the electric age, there have been no low-electricity, high-income countries; during industrialisation, no power without abundant steam power has prevailed; and when the energy costs of geopolitical competition still came down to feeding standing armies, the world's bread baskets were masters of their times. Only middle powers with access to abundant frontier AI will stand a chance.

Energy use per capita




Figure 2. Energy consumption per capita vs. GDP per capita, 2023.

Data source: IEA Energy Statistics Data Browser, IEA, via World Bank (2026); Eurostat, OECD, IMF, and World Bank (2026)

Note: GDP data is expressed in international-\$** at 2021 prices. [OurWorldinData.org/energy](https://ourworldindata.org/energy) | CC BY

By this, we mean access to the best AI systems. Not just models that perform reasonably, not even just models that are good—the best models there are. As it stands, chances are simply too high that second-best is not enough, that the winners of the AI era will be those with the best AI. As much might be true in the absolute and relative sense. In the absolute sense, countries should seek out the best simply because the most powerful AI models will be the most useful to any national ambition. In the relative sense, it is worth understanding that middle powers find themselves in a threefold competition that means that whatever deployment tactics we describe below must necessarily relate to the deployment not of any AI model, but of the very best.

The first dimension of this competition is between its public institutions and those that seek to use and sometimes abuse it—that would flood the administrative state in requests and complaints and overwhelm a human bureaucracy by moving at the pace of machine intelligence, leaving the state unable to act and react, eroding trust and capacity as it goes. On that, this paper comments in much greater detail in the next section.

The second dimension is between its security apparatus and those that would seek to harm their citizens. Attackers foreign and domestic will be equipped with near-leading AI systems, whether it is through stolen models, adversarial jailbreaks, or near-frontier open-source capabilities. They will use these capabilities to devise and execute new kinds and scales of crime and terror: cyberattacks, chemical and biological weapons, strategic and organisational support for criminal operations, and as a force multiplier for hostile militaries and non-state actors. Nation-states' abilities to defend against these novel threats will depend on their ability to wield AI defensively, to scale up their detection of attacks and engagement of adversaries. The race between attackers and defenders might often be on a knife's edge: researchers are as of yet unsure which way the offense-defense balance in AI points; but they find that more advanced models are consistently more effective at both [offensive and defensive](#) operations. In the interest of the safety of their citizens, countries cannot afford to use weaker AI systems than their adversaries.

The third dimension is between its own private sector and international competitors. AI will serve as an accelerant to today's already-dynamic market forces, exacerbating

the pace and international nature of competition and further shrinking markets. As customisation becomes cheaper and information more accessible, institutional moats will shrink, and incumbents without a substantive edge will be disrupted by their global competitors. The quality of inputs is one of the major differentiators left in that world—and the quality of artificial intelligence will be one very important input. We can see glimpses of this reality in the areas that are most like tomorrow's economy today: In San Francisco, leading AI developers are already turning off their competitors' access to their own models, because they think the marginal acceleration from *slightly* better software engineering systems is a relevant factor in competition. And even before AI, particularly fast-paced competitive environments have become a question of marginal technology access: in high-frequency financial trading, slightly faster decisions matter so much that proximity to crosscontinental cables has become a prized asset for trading firms.

In a digital economy stripped of friction and pretense by AI, we expect these marginal differences to matter more and more. Much of this is not a question of government use, but of private-sector use—but government access plays a meaningful role nevertheless: failed government policies could force even the private sector to employ subpar technological solutions, such as has been suggested in Europe time and time again. And quite frequently, the responsiveness of the government itself is a bottleneck for the competitiveness of a national economy: whether it is for permitting and licensing, for enforcing market protections, or for enabling imports and exports, economic policy can either be AI-empowered in an AI-accelerated world or grow to be a substantial bottleneck.

Foregoing access to the best AI systems puts governments at risk of failing to fulfil their promise of efficiency to their citizens, of protection to their charges, and of support to their economies.

Ensuring Access

However, the best AI models might not be abundantly available, and their provision might be constrained by geopolitical factors. Middle powers are growing more and more skeptical of importing critical economic inputs, having been burnt by their lack of

leverage in unilateral trade relationships with great powers and by recent high costs to their energy dependency.

They are not entirely wrong to be worried for their enduring access to AI models developed in the U.S. and China. If middle powers' access to AI capabilities is only guaranteed through individual subscriptions and API contracts, it remains uninsured against disruptions: sudden price hikes by developers that think they have locked down a market, service disruptions because limited computational capacity is first routed to other priority markets, or great powers that take policy action to restrict the outflow of artificial intelligence—either for strategic and security reasons or to exert leverage to extract concessions in other areas.

That means middle powers will need a plan to secure durable access to frontier AI capabilities. Any further strategy for economic competitiveness and strategic sovereignty requires this necessary precondition: without frontier capabilities, any middle power's position is at great peril.

Yet middle powers will not ensure this access through building their own AI models. Sovereign development of frontier AI systems faces a highly restrictive market and an unfavorable political economy. The key inputs for frontier AI—computational power and exceptional research talent—are limited, and they follow agglomeration effects and the ability to mobilise capital. Only [China and the U.S.](#) have been able to sow the seeds of this agglomeration in their respective hotspot regions, and only state-directed funds in China and the unrivaled venture capital ecosystem in America have been able to provide for scaling it to frontier performance. For everyone else, sovereign development of frontier AI systems is far out of reach. Past attempts of others to catch up have fallen far short—landmark infrastructure projects such as the EU's [gigafactories](#) lag about three to five years behind the American expansion, for instance. We are doubtful that this will change in time: 2028's infrastructure is determined by today's buildout decisions, where middle powers lag far behind; and any further acceleration would compete for often-scarce energy, chips, researchers, and so on.



Figure 3. Map of frontier AI datacenter buildouts.
Source: Epoch AI Frontier Data Centers Hub

One permutation of this hope holds that one could sidestep this race and bootstrap to a new paradigm—as is reportedly the basis of a recently announced European frontier AI initiative. This, too, is unlikely to succeed; both because paradigm innovation also benefits from the scale effects of major developers, and because the last years of AI development drive home the veracity of the “[bitter lesson](#),” which holds that technical approaches that can leverage computational power at scale ultimately always outperform any domain-specific sophistication. There seems to be no viable middle-power path to sovereign frontier AI.

The good news, however, is that American AI companies *want* to export their advanced capabilities, and the U.S. government is very interested in supporting and enabling such exports. Multiple large developers have created and advertised for-countries initiatives promising the export of AI capabilities to willing governments at scale. The U.S. government’s commitment to the same has been affirmed through an executive order providing for the creation of an export promotion program for AI capabilities; and through the [Pax Silica initiative](#) that seeks to exchange semiconductor supply chain integration for access to American AI capabilities. America’s AI industry is open for global business.

But middle powers are understandably hesitant to take the deal that U.S. industry offers. They are concerned about the limitations of *ad hoc* access: will they be subject to price hikes or strategic access limitations? Are they guaranteed to be provided the next generation of frontier models, and the one after that? The path forward is the negotiation of longer-term contracts backed up by real leverage, conferred by leveraging bottlenecks that interest the great powers that can provide frontier AI access. We provide the economic roadmap to occupy such bottlenecks later on—but whatever they are,

they have to be exchanged for long-term access conditions, enforceable by middle powers through a combination of local compute, market access conditions, and prohibitive government-level contracts that can compel the continued long-term provision of AI capabilities. In short: middle powers must be strong enough to be able to enforce the contracts that secure their AI access. They might find that strength in their idiosyncratic sources of leverage along the semiconductor supply chain, in their effective deployment of AI systems, or in the numbers of a broader coalition to negotiate the same—but find it they must.

Open Source and Chinese Imports

Some middle powers will object to this logic, arguing they might access models from elsewhere, in particular through competing exports from China that might or might not be open-source models. However, the availability of powerful open-source models will be constrained by economic reality and strategic necessity: deploying them is not always worthwhile for developers seeking to recoup their development investment, and seeing them released is not always in the interest of great powers that either seek to leverage AI access or that are justly concerned about misuse pathways. As a result, one should expect open-source models to remain a generation behind the frontier, as they are [today](#). In that scenario, middle powers might choose to deploy open source for select applications that put a premium on customisability over quality, but it is unlikely that they will be sufficiently effective to substitute for genuine frontier access.

One permutation of this question concerns the role of Chinese AI capabilities—both open source and proprietary—which looms particularly large for middle powers that have sought to play both sides of great power competition, extracting concessions from having America and China compete for their markets. China has aggressively released capable open-source models, in part to undercut American market share and establish dependency in regions where U.S. export frameworks have been slow to deliver. For middle powers seeking to diversify their sources of AI access, Chinese models present a superficially attractive option: capable, freely available, and unburdened

by the conditions they fear American firms might attach to their exports. But building a national AI strategy around Chinese open source carries its own dependencies and shortcomings.

There is no guarantee that China will continue to release its most capable models openly once middle powers have structured their digital economies around them; the current generosity in open-source releases appears driven by a competitive strategy to establish market share, not by an enduring commitment to openness. Well-documented security concerns attend the deployment of Chinese-developed models in sensitive government and infrastructure applications. Moreover, and perhaps importantly, the use of foreign-built open-source models comes with counterintuitively greater dependencies than the use of proprietary models: if a middle power uses open-source models for free, their provider has nothing to lose from cutting them off, while a revenue-driving contract comes with mutual leverage over a developer that does not want to lose a customer.

The perils of choosing China are compounded by the additional complication that deepening reliance on Chinese AI may jeopardise a middle power's standing within American-led technology partnerships, including access to the semiconductors and compute infrastructure on which all AI deployment ultimately depends. China, now and for many years in advance, cannot provide the same infrastructure access, and until that changes, cannot be a viable alternative. The strategic logic, then, points toward American AI as the primary import source for allied middle powers—not out of naïvety about American intentions, but because the broader architecture of alliances, supply chains, and security relationships makes the American stack the more natural foundation for nations wishing to chart their own course.

Once they have so ensured their access to an abundance of artificial intelligence, middle powers must deploy it well: to supercharge and revolutionise their institutions to remain and become capable in the new world, and to channel it into durable economic pathways that guarantee long-term participation in AI-driven growth and flourishing.

Ghosts in Your Government: Institutional Transformation

Middle powers contemplating the adoption of AI face a tremendous opportunity—to lead the world in institutional innovation and transformation. But they also face a genuine threat: that AI will erode the foundations of the nation-state itself.

The modern state emerged in large part because it could do something that no other institution could do as well: collect, aggregate, and act upon information [at scale](#). The census, the tax roll, the statistical office, the regulatory agency—these were instruments of knowledge, and knowledge was the precondition of effective administration. Private actors could know their own affairs, and perhaps the affairs of their immediate trading partners, but only the state could know the society as a whole. This informational advantage underwrote much of what governments distinctively did: allocating resources, managing externalities, coordinating the provision of public goods, enforcing contracts and resolving disputes, responding to crises. The state could see what others could not see, and so the state could act where others could not act. In doing so, it was able to retain a position of trust and authority—even where others could have taken its place in providing a piece of information or rendering a service, societies trusted their governments to act as able and responsible stewards instead.

This advantage has been eroding for some time. The rise of large corporations, and later of global platforms, began to shift the balance. Credit bureaus came to know the financial lives of citizens in ways that tax authorities could not match. Retailers learned to predict purchasing behavior with a precision that no economic planning office ever achieved. Technology platforms accumulated data on the movements, preferences, relationships, and sentiments of billions of people—data that dwarfed anything in the administrative records of even the most sophisticated governments. The private sector learned to see, and what it saw was often more granular, more current, and more actionable than what the state could see.

Artificial intelligence does not initiate this trend, but it accelerates it in ways that may prove decisive. The informational advantage of the state rested not only on the collection of data but on the capacity to process complexity—to make sense of vast, heterogeneous, and often contradictory information, and to translate that understanding into policy and action. This was cognitive work, and it was expensive. It required large bureaucracies staffed by trained professionals. It required time. The scarcity of this cognitive capacity was itself a barrier to entry that protected the state's position. AI erodes that barrier. When the processing of complexity becomes cheap—when any sufficiently resourced private actor can ingest, analyze, and act upon information at a scale and speed that was once the exclusive province of government—the state's distinctive role comes into question. Private actors may soon understand populations better than the governments that nominally serve them: their needs, their behaviors, their vulnerabilities. They may understand the mechanics of regulation and the pathways through bureaucracy better than the officials who administer them. The state's claim to govern has always rested in part on its capacity to know. What becomes of that claim when others know more?

Something deeper, more structural, is at play in the dawning of the AI revolution. Like earlier information revolutions—including writing and the printing press—the scope, scale, and structure of institutions will evolve. Grappling with these changes is not solely, or even primarily, a regulatory challenge. The instinct to regulate is understandable—it is what governments know how to do—but it mistakes the nature of the challenge. The challenge is not primarily to constrain AI but to ensure that public institutions do not fall irrecoverably behind the private alternatives that AI enables. This is a capability race, and governments are at risk of losing it.

The race is already underway across nearly every domain of government activity. Consider healthcare: a public health system that cannot match the diagnostic accuracy, the personalized treatment recommendations, and the administrative efficiency of AI-augmented private alternatives will find its citizens seeking care elsewhere. Those who can afford private options will take them. Those who cannot will receive worse care than the technology of the age makes possible, and they will know it. Or consider

education: when a private AI tutor can adapt to each student's pace, identify gaps in understanding, and provide individualized instruction around the clock, what becomes of the public school that offers one teacher for 30 students and a static curriculum? Those able will exit. The public system will become a residual institution for those with no other choice.

Finally, consider the courts: when private dispute resolution—AI-assisted arbitration, contracts that execute automatically, platforms that resolve conflicts in hours rather than years—offers speed and predictability that public courts cannot approach, commerce and even personal disputes will migrate out of the legal system. The courts will not be abolished; they will be bypassed. Consider the bureaucracy itself: when citizens can navigate permit applications, tax filings, benefit claims, and regulatory requirements through AI intermediaries that understand the system better than the officials who run it, the government becomes a legacy backend—still technically sovereign, but substantively hollowed out.

These are not distant hypotheticals. They are emerging now, unevenly but unmistakably, in every country with meaningful AI deployment. There is reason to fear a threshold effect: public institutions that miss the AI acceleration curve may **never catch up**. The gap between an AI-augmented institution and a legacy one is not stable; it widens over time as the augmented institution learns, iterates, and improves. The laggard does not merely remain behind; it falls further behind with each passing year. For governments, this means that the decision not to adopt—or to adopt slowly, cautiously, experimentally—is not a neutral choice. It is a choice to accept relative decline, perhaps permanently. There remains the hope of “leap-frogging” instead—of adopting AI at just the right time to bridge that widening gap in an efficient manner. But the time to jump is *today*—the rapidly accelerating iteration cycles of AI technology spell compounding effects which imply that leapfrogging timing is measured in weeks and months, instead of years and decades.

The full suite of political, economic, public policy, and institutional solutions to this conundrum lies beyond the scope of this report. But one thing is clear: if governments do not adopt AI with alacrity, they will not “lag behind.” They will grow obsolete.

The middle powers, then, must adopt and adopt aggressively. But they face an obvious difficulty: they cannot build frontier AI themselves. The models, the compute, the talent, and the capital are concentrated elsewhere—overwhelmingly in the United States, and to a lesser extent in China. Middle powers must import what they cannot build. The question is how to do so without surrendering the very sovereignty they seek to preserve.

Wise adoption should follow several principles. First, architectural flexibility: no middle power should adopt systems that lock it to a single AI provider. The frontier is volatile; the best model today may not be the best model in six months. Governments should build or procure platforms and interfaces that allow rapid switching between underlying models as the competitive landscape shifts. Dependence on one provider is dependence on one set of strategic interests, one set of potential failure modes, one set of terms that can be altered unilaterally—and, increasingly, dependence on a model that may no longer represent the state of the art. Second, sovereign compute for critical applications: the most sensitive functions of government—national security, judicial systems, core administrative processes—should run on infrastructure that the government controls or that resides with trusted allies. This does not mean autarky; it means prudent hedging. Third, interoperability and portability as design requirements: every system adopted should be structured to permit switching between providers with minimal friction. Lock-in transfers leverage permanently to vendors; portability preserves options and ensures that governments can always avail themselves of the best available capabilities. Fourth, data governance as an explicit element of negotiation: access to government data, and the terms on which private AI systems can operate on citizen data, are concessions of real value. They should be traded for real value in return.

None of this is to suggest that adoption should be slow or tentative. The opposite is true. Speed matters more than perfection, and the costs of inaction exceed the costs of

failed experiments. The volatility of the frontier is itself an argument for moving quickly: a government that delays adoption until a clear winner emerges will wait forever, because clear winners do not emerge for long in a landscape this dynamic. Governments should identify high-impact domains where AI can deliver rapid, visible improvements: document processing and administrative workflow, citizen service interfaces, data collection and integration, permit processing and regulatory compliance. These are not the most glamorous applications, but they are the ones where gains can be realized quickly, where institutional capacity can be built through practice, and where the risks of failure are manageable. The point is to learn by doing. Strategic awareness of what AI can accomplish does not come from white papers and consultations; it comes from deployment, from iteration, from the hard-won knowledge of what works and what does not in a particular institutional context. A government that waits until it fully understands AI before adopting it will be overtaken by those less fastidious.

None of this is to suggest that governments should adopt AI in a defensive posture. Indeed, adoption of AI should be undergirded by a fundamentally optimistic, if also realistic, vision that moves beyond the defensive logic of staving off erosion. Every middle power has governance failures that AI is distinctively suited to address. [Backlogs of court cases](#) stretching years or decades. Agricultural extension services that cannot reach farmers with timely information. Public health surveillance systems that detect outbreaks too late. Educational systems that leave millions of children behind. Tax administration that is both burdensome and porous. These are not new problems, but AI offers new solutions—solutions that were not possible five years ago and that will be more powerful still five years from now. The nations that move first will not merely preserve their own institutions; they will discover solutions that others will want to emulate. They will become models, exporters of institutional innovation, destinations for those who want to learn how governance works in the age of machine intelligence.

This is the race worth winning. Not the race to build frontier models—that race is already decided, or nearly so, and the middle powers were never serious competitors. The race that remains open, the race where the middle powers can compete and even lead,

is the race to discover what governance looks like when machine intelligence is woven into its fabric. The nation that wins this race will not merely have better public services. It will have demonstrated something important to the world: that the era of machine intelligence need not mean the erosion of the state, the hollowing out of public life, the retreat of self-government before the advance of private power. It can mean, instead, the reinvigoration of public institutions—a new era in which governments are more capable, more responsive, and more worthy of the trust their citizens place in them.

Bottlenecks for Abundant Intelligence: Economic Transformation

The fundamental challenge most middle-power economies will face is to secure a share in the rampant economic growth we expect as a result of the widespread diffusion of truly advanced AI systems. This participation in extraordinary growth is not a foregone conclusion—rivaling economic effects pull toward and against it. On one hand, just as past technological progress has, advanced AI systems are likely to increase the quality and depress the prices of goods and services that make up the baskets we use to measure the wealth and welfare of individuals and nations. In very material ways, this would make lives better beyond the macroeconomic headlines: fewer would go hungry, sick, or compelled by economic circumstances; more would have access to information, education, and the tools to actualize their plans and desires. On the other hand, though, the revenue from this influx, and therefore the access to this abundance, could accrue asymmetrically. Experts still disagree on where exactly: perhaps on the layer of infrastructure and chips, perhaps on the layer of model development, perhaps on the layer of applications and deployment software, perhaps on the layer of industry that turns intelligence into goods. But any such concentration could render some middle powers—those that do not have a stake in that respective layer—comparatively behind the curve.




Figure 4. Service sector as share of GDP across middle powers in 2010 and 2022. Many middle powers' workforce is increasingly concentrated in service sectors, with no obvious "next sector."
Source: [World Bank](#).

At the same time, they could also render economic sectors superfluous that are load-bearing for middle powers today: labour could be replaced by AI agents, research and development work could be conducted by “geniuses in a datacenter” instead of innovators in hard tech hotspots, customer service could be provided by customised chatbots instead of service centers. One scenario holds that cheap agentic AI systems run a hyperefficient economy—perhaps resolving any need for large digital that today make up a growing part of many middle powers’ economies, eliminating large swathes of jobs and thereby also dramatically reducing taxable income in these powers. The consequence could be a world of material abundance that many middle powers—their firms, citizens, even governments—would not be able to afford. Notably, this gap would ultimately imply that these middle powers might not even continue to be able to afford the very AI models that would make the diffusion and adoption outlined above possible.

There is precedent—not just for sectoral displacement driven by technological transformation, but also for sector-driven widening of economic gaps. In the past 15 years, “software has eaten the world,” meaning some of the highest growth and highest margins of the 2010s have accrued to software firms. This has opened a widening gap between previously comparable economies, a contrast particularly visible in [the US-EU economic divide](#) that is principally driven by absolute American dominance in the field of technological innovation: the world has gotten richer, but at the same time, Europeans have relatively gotten poorer, have lost relative leverage, and have borne the cost of economic stagnation as a result. The same effect, exacerbated by orders of magnitude, might open a gap between middle powers and great powers as AI, too, “eats the world.”

Securing a Share of AI Growth

To ward off that prospect, governments need to allow their economies to develop exposure to AI-driven growth—that is, they should make sure that a substantial share of AI revenue accrues in their respective economies.

SEMICONDUCTOR POWERS

This challenge and opportunity is best explained by drawing on the example of semiconductor powers. A select few middle powers—some savvy, some lucky—are cut in on AI-driven growth by default, by virtue of their position in the complex supply chain around frontier AI systems. The most immediately obvious examples are countries with a position in [semiconductor manufacturing](#): Taiwan’s TSMC, the Netherlands’ ASML, Germany’s Zeiss, and some others provide absolutely essential machinery to produce the frontier chips that make AI systems possible. Contracts with these firms are coveted resources, and as much is reflected in both their revenue and valuations as well as the keen strategic interest that great powers have taken in securing access to this equipment. One degree further removed, though no less strategically important, is the provision of critical raw materials. This cluster of countries has already received increased policy attention from the U.S., with the Department of State launching the Pax Silica initiative aiming to secure American access to these important inputs.

We believe that, rapid and expansive as current buildouts may be, the world might still be undersupplied by advanced chips by the end of this decade. If AI becomes a central economic resource, but remains fundamentally constrained by the availability of advanced chips, positions in the semiconductor supply chains are sure to greatly appreciate in value.

THE REST OF THE WORLD

Any other middle power must ask how it might build what semiconductor powers already have. The rest of the world has the disadvantage of uncertainty: no matter how exactly the value accretion dynamics unfold, semiconductor powers are likely to capture value because semiconductors are the fundamental basis for both the development and deployment for AI. Other nations will have to make riskier bets, and hope they play out. This, if anything, is the place for active industrial policy: those middle powers that are unfortunate enough to lack sectors exposed to AI-driven growth might need to make sure to create them. Three sectors of economic activity in particular qualify

for ensuring this participation, and most middle powers are in a position to strengthen their position in at least one of them. We want to point toward these sectors as an attempt to spur a conversation—the minutiae of strengthening them through economic policy, in the meantime, differ deeply from region to region.

The first nexus of economic value in the era of advanced AI forms around *inputs into the broader AI supply chain*. The semiconductor supply chain is far and away the most obvious example here, but other inputs matter, too: energy provision at scale matters to deploy chips, high-quality data is needed to feed the progress of artificial intelligence—and even the sheer capability to get datacenters up and running without uncertainty or delay could be unevenly distributed, putting countries that are good at infrastructure buildouts far ahead of the curve.

The second nexus is made of *sectors that retain their relative value* even in the face of world-changing AI deployment: sectors where human presence is priced highly, or that cannot be meaningfully expanded by productivity gains. As a shorthand, these sectors have sometimes been described as being particularly afflicted by Baumol’s cost disease—tourism, healthcare, artisanal craft, and so on.

AI AND INDUSTRY

The third and most important nexus is around *downstream bottlenecks* for generating real-world value from abundant intelligence. This is perhaps the most underrated part of the AI puzzle, and one that is particularly responsive to policy action: no matter how capable these systems get, many of their real-world effects will only come from arduous manufacturing and real-world iteration around physical outputs: medicine and compounds, materials and machines, drones and robots. The barrier to do so effectively is high: not only do you need industrial capacity as such, but it needs to be technologically advanced enough to integrate deeply with AI, providing data and responding to suggested changes quickly; modular enough to upend entire workflows at the arrival of another

innovation; and staffed by workforces capable enough to work in effective conjunction with highly intelligent machines, able to tell the good recommendations from the bad.

Many middle powers know—or at least have forgotten—how to do just this; whatever is left of their industrial bases can be rebuilt at pace. This is an idea that faces a more favorable political economy than any other attempt at restructuring your economy: revitalising industrial regions and reemploying skilled blue-collar workers has long been a dream of many powers, and it is made more economically attractive again by the advent of advanced AI. And while there are no policy blueprints for an AI revolution, the powers that have prevailed until today once knew a lot about carrying out an industrial revolution.

This manufacturing nexus gains particular importance from a geopolitical asymmetry: while [China commands industrial capacity at scale](#), the U.S. industrial base is atrophied and only beginning to rebuild. While it does, America will price very highly any contribution to transforming San Francisco-built artificial intelligence into the hardware required to provide material abundance and, if need be, to fight a war.




Figure 5. Middle powers can capture economic and strategic value at upstream and downstream bottlenecks for advanced AI.

Sources: SIA/BCG 2024 · OECD 2025 · IEA Critical Minerals Outlook 2025 · Georgetown CSET ETO · IFR World Robotics 2024 · World Bank · Stanford HAI 2025

Securing a Durable Position

Once a middle power has secured a niche that ensures its economic participation in AI-driven growth, it needs to leverage it into an enduring position and protect it against intrusions and attacks. Much of this will depend on iteration as we go, and on responding to whatever the specific features of tomorrow's technical paradigms will be. But there are at least three principles worth recognizing ahead of time.

First, no niche is durable unless it is fortified by aggressive adoption of AI systems. Whether it is upstream or downstream of capable AI, close hardware-software-integration is likely to remain essential. Just as in past technological stacks from heavy industry to consumer software, the most competitive products are reliably produced where customised layers interlock: data from manufacturing feeds the quality of AI systems, analysis from customised AI systems improves manufacturing processes, and so on. The same interplay exists between chips and AI models, between computing components and finished chips—in short, across the entire stack. Middle powers must seek close integration with products along the stack, and they must avoid attempting to retain full modularity that inevitably costs them the efficiency gains to be had from close integration.

Second, great powers will seek to “reshore” most bottlenecks that middle powers will occupy. The U.S. will somewhat appreciate allies that contribute to its industrial stack, that offer their equipment to build American chips—but it would always prefer to instead have these capabilities on American soil. So where great powers enter into deals for access to bottlenecks, they will attempt to include an element of technology transfer back to their shores, and where they cannot achieve these transfers, their large firms are likely to leverage capital markets to outright acquire the bottleneck businesses instead. Middle powers must resist these attempts. They will be tempted to “cash out” at times where China and America have already understood the enormous value of bottlenecks, but middle powers have not yet. To middle powers and their private sectors, an acquisition offer for a semiconductor firm or an industrial robotics firm might look like an unexpected windfall—but once one accounts for these very firms' role as the ticket guaranteeing their home country's stake in the AI economy, the calculus changes

entirely. Middle powers must realise the counterintuitively high importance of a select few kinds of businesses—many of which will be marginal in today’s economies—in navigating investment screening, tech transfer dealing, and acquisition licensing, or they will lose their niches before they can leverage them.

Third, no niche will last forever. Some are more robust than others: semiconductors will retain their value as long as AI systems remain similar to today’s—but no one knows what the next decades will bring. At the pace of technological transformation we predict, very little economic structure seems certain to remain static. No middle power can stop innovating, or it risks going the way of Finland past Nokia, of Nauru past the phosphate, of Detroit past automobiles. For any middle power that successfully occupies a niche, there will be lessons to be learnt from the rise of the Gulf’s oil states, and how they have attempted—sometimes managed, sometimes failed—to diversify their economy away from the resource that has gotten them so far. Much success has had to do with the ability to identify the next “big thing” in time.

Navigating the Transition

Insofar as one primary objective of a middle power government is to allow this dramatic restructuring of its economic portfolio to happen—and sometimes even to support it—it needs a good response to the grievances that result from rapid restructuring with winners and losers. This is both to help the whole of the country into the new paradigm, and to insure whatever ambitious choices it might have made against the inevitable political backlash that comes with ambitious change.

To address the likely disruption, middle powers first need to understand the kind of challenge their economies might face. That challenge is not absolute impoverishment, or economic loss in general terms—we believe it is highly likely that AI ultimately presents as beneficial to all nations; and as long as they find a niche according to the above categories, no nation is likely to fall through the cracks. But the internal relative balance of powers is nevertheless prone to disruption. Large sectors of economies might not remain viable, leaving regions without economic drivers just as we have seen in past technological transformations—just that this time, it might be the coastal cities and not the rural agrarian and industrial towns that face the shock. Cohorts of workers and entire demographics might be asymmetrically affected, and sometimes without clear prospects: just as not all former autoworkers were able to become coders, not all junior software engineers and copywriters might find their way into a bottleneck sector. Millions will face uncertainty and unemployment, cities will be empty, other regions overrun. Such is the reality of a changing world.

In the past, nations sometimes had the luxury to introduce artificial frictions to slow transformations like these. In the breakneck competition between national economies that might arise from the deployment of advanced AI systems, they might not be afforded the same luxury. Absent meaningful international coordination, adding artificial frictions domestically just further exposes workers to external disruption: if one power stalls the costly augmentation of its workforce through AI because this will mean

substantial job losses, its service industry might be cannibalized entirely by another power that moves more quickly to make the same cuts and accelerate AI deployment in its own service economy; or simply be outperformed by workforces purely assembled from artificial agents altogether. The same threat of competition hangs over any other attempt to unilaterally slow down the effects of AI in a field not certainly monopolized.

Middle powers cannot artificially protect what cannot earn its keep. Instead, their response requires making progress on two related questions, fast.

First, they will need to find some other policy to satisfy the losers of this transition. They will absolutely need to provide assurances ahead of time, and they need to calibrate them right. Too specific, and they will be unrealistic and unbelievable: guarantees to keep any one job or any one day-to-day life cannot succeed. But too vague, and they will do little to reassure the vulnerable. Almost necessarily, these reassurances will have to include a message of redistribution, and one that cannot be an afterthought. The design space for such reassurances is broader than the policy debate typically acknowledges, but every corner of it is constrained. Direct redistribution—transfers, wage insurance, expanded safety nets—requires fiscal capacity that middle powers may struggle to muster precisely when they need it most.




Figure 6. Share of tax revenue from labor and corporate income tax.
Source: OECD Revenue Statistics 2025.

As AI-driven revenue shifts from labour income to corporate profits and from domestic economies toward the countries where AI is built, [the tax bases that would fund ambitious social policy risk eroding rather than expanding](#). Middle powers do not enjoy the fiscal position of the United States, where AI-driven growth at least accrues

domestically and where the taxable base may shift but does not vanish; nor will most middle powers easily levy meaningful taxes on AI systems or their developers, who sit beyond their jurisdiction. The most credible path to funding transition support is therefore the same path outlined in the industrial strategy above: securing a genuine stake in AI-driven growth through bottleneck positions in the physical supply chain. There are some easier pathways for ensuring a merely fiscal stake in AI growth, too—simply exposing one’s wealth funds and investments to AI-driven growth by taking stakes in likely winners of the AI revolution might be almost as good and just as required as securing a physical stake. As much could be achieved through folding this kind of fiscal opportunity into import deals or market access conditions. Countries that succeed in capturing a share of the revenue will be the ones with the fiscal room to cushion their populations through the transition. At the end of the day, a successful social strategy and a successful economic strategy must converge. Any alternative approach that puts them in tension is exposed to the forces of accelerating competition.

Their concrete implementation of the economic value gained and the fiscal leeway established into the space of redistribution, then, will be one for *ad hoc* iteration—safety nets, retraining, chances for workers to become entrepreneurs in their own right, and direly needed progress toward actual prospects for individuals whose skillset might simply not lend itself to any kind of labor past the advent of advanced AI. Two pieces of good news are that, first, the provision of these services, too, can be greatly accelerated by the deployment of AI; and that, second, any AI advanced enough to cause such disruptions would necessarily also provide material abundance and exert deflationary pressures. These factors will make it much easier than ever to secure a standard of living and a level of wealth that we would call rich by today’s standards. Yet the bad news is that the great uncertainty around the precise economic effects of advanced AI makes it essentially impossible to predict the exact scope of policy required in advance. We have strategic notions of how to secure the preconditions of gainful redistribution, but the space of how the gains from AI could be captured and distributed to compensate the losers of the upcoming transition is direly underexplored. Exploring it is in the obvious

interest of anyone in middle powers—especially those that seek to accelerate the AI revolution, or they will be confronted with powerful backlash at every turn.

But second, leaders of middle powers must also muster the political courage and forge the political consensus to endure temporary backlash against the restructuring that happens. Fears will always play well, the reality of positive returns will only set in in time, and staying the course laid out above will be threatened by political backlash on all sides. This will be a moment for capable and courageous leadership—at the helm, with leaders who understand the challenge and are willing to see the response through; but also in the bow, in the broader political environments surrounding these leaders that need to refrain from scoring cheap political wins by derailing costly but successful AI strategies. For this to hold, both in the face of opportunistic political oppositions and of real questions of democratic legitimacy, it requires convincing electorates and populaces with a substantive message to rally around, and for leaders to be avatars of this message. Many middle powers will fail this test dramatically. But now is the time of the savvy statesman who realizes this future.

Conclusion

This analysis presents a picture of the future that many fear, often rightfully so. The future we expect is surely bright, but without doubt it is dark too. AI will enable a world of higher contrasts, a larger dynamic range. Much of the resulting variance will be between different middle powers, shaped by their choices. The temptation among governments, civil society organizations, and academic institutions worldwide will be to ignore the reality of what is unfolding, to talk themselves into complacency rather than action. This, more than AI itself, is the biggest threat of them all. The events unfolding in the field of AI are stark, titanic, and well underway. It is not “too late” to begin taking serious strategic action, but neither is it early. Further delays only complicate an already narrow and fraught path.

Countries that thrive in this era will all do so for idiosyncratic reasons and in idiosyncratic ways. It is for that reason that we hope this report constitutes a contribution in two ways: by clarifying the lay of the land and the depths of the pitfalls to leaders in middle powers and those inclined to help them; and by laying out the right questions to ask for researchers and organisations with greater sectoral and regional expertise than us. These questions we have raised throughout this report, but in particular, the following strike us as essential:

- **ACCESS ARCHITECTURE:** What contractual, technical, and institutional arrangements can durably guarantee middle powers’ access to frontier AI—and what minimum threshold of local compute, market leverage, or coalition size is required to make such guarantees enforceable against the unilateral interests of a providing great power?
- **INSTITUTIONAL TRANSFORMATION IN PRACTICE:** Which public-sector functions are most amenable to AI-driven transformation, what procurement and governance models allow governments to adopt at speed without locking into

a single vendor, and where runs the line between insurmountable gaps and opportunities for leapfrogging?

- **ECONOMIC BOTTLENECKS:** Where along the AI stack—infrastructure, model development, applications, or downstream manufacturing—will economic value durably concentrate, and how quickly can today’s bottleneck positions be eroded by great-power reshoring, paradigm shifts, or substitute technologies?
- **REDISTRIBUTION MECHANICS:** How can middle powers fund transition support when AI-driven revenue shifts simultaneously from labour income to capital returns and from domestic economies to the countries where AI is built; what novel fiscal instruments or inter-middle-power arrangements can bridge the need for social policy and the erosion of the tax bases that would traditionally finance it?
- **THE POLITICAL ECONOMY OF ACCELERATION:** What political coalitions, communications strategies, and institutional safeguards allow democratic middle powers to sustain ambitious AI-driven restructuring in the face of displaced workers, opportunistic opposition, and legitimate questions of democratic mandate, without resorting to artificial frictions that would undermine their competitive position?

Middle powers that fail to find answers to these questions will often do so in similar ways: allowing the legitimacy of the state to erode in the face of technological competition, taking half measures rather than decisive action, attempting to close their eyes and borders to the transformative effects of AI, and resisting rapid adoption in the name of some hoped-for future “sovereign” solution. In truth there is no greater risk to sovereignty than waiting for that which is perfect instead of accepting that which is tolerable.

The era of machine intelligence will not wait for the middle powers to be ready for it. It will arrive—is arriving—on the terms set by the technology itself, by the industrial realities of the present day, and by the strategic decisions of the great powers. But within these constraints lies a space of genuine agency, and it is larger than many suppose. The middle powers that move with clarity and resolve—securing frontier AI

access, transforming their institutions from within, and positioning themselves at the bottlenecks where intelligence meets the physical world—will not merely survive the transformation. They will shape it. They will discover, through the hard work of deployment and iteration, institutional forms and governance capabilities that the hegemons, consumed by the race to build, may be slow to develop themselves. In doing so they will demonstrate something that matters greatly for the future of self-government everywhere: that a nation need not build the revolution's tools in order to lead the revolution.

This will require statesmanship of an uncommon kind—leaders who can hold in mind simultaneously the enormity of what is changing and the practical, often unglamorous work of making government function better with new instruments. It will require the courage to adopt before the path is certain and the discipline to resist the false comfort of regulation as a substitute for capability. It will require, above all, a refusal to mistake caution for wisdom. The middle powers have been here before, at prior hinges of history, and the ones that prevailed were never those that waited for certainty. They were the ones that acted with the tools at hand, in the world as it was, and built something worthy of the world to come.

About the Authors

Anton Leicht

Anton Leicht is a visiting scholar at the Carnegie Endowment for International Peace and author of the Substack *Threading the Needle*.

Dean Ball

Dean Woodley Ball is a senior fellow at FAI. He most recently served as senior policy advisor for artificial intelligence and emerging technology at the White House Office of Science and Technology Policy and strategic advisor for AI at the National Science Foundation. Previously he was a research fellow in the Artificial Intelligence & Progress Project at George Mason University's Mercatus Center and a policy fellow at Fathom.

The Foundation for American Innovation champions the technology, talent, and ideas essential to American prosperity, security, and flourishing.

thefai.org

