Economic policy for AGI

This essay evaluates eleven policies societies could implement to manage the potential economic disruption arising from increasingly advanced and pervasive AI.

Economist and Research Scientist at Google DeepMind

Growing anxiety about how AI will impact jobs and the economy is emerging as a major issue across broad swaths of society. It is already a top concern for political candidates and for younger generations – all before the effect of AI shows up in most economic data.

AI capabilities have increased tremendously over the last few years, and they show no signs of slowing down. A future economy infused with artificial general intelligence (AGI)—advanced AI systems capable of performing cognitive tasks at least at the level of the average person—could bring huge benefits. This includes improvements in health, leisure time, and higher standards of living. Yet concerns about AI's economic impacts make sense. In the past, major technological transformations have boosted *long-run* living standards, granting humanity prosperity and health that would have been unfathomable to our ancestors. Yet in the *short term,* significant transformations have often also

produced lower wages for many, hollowed out regional economies, and deepened inequality.

In other words, many people who lived through technological transitions were hurt by the periods of innovation that they lived through. The question is: Can we reduce or avoid such harm this time, while preserving the benefits of technological progress?

In the past, countries have struggled to adequately manage productivity-boosting economic transitions. During the Industrial Revolution and the early 20th century, policy often came after the fact, deploying redistributive interventions once the shock and disruption had occurred. In some instances, this may have been due to poor institutions and access to data. In others, slower responses were the result of concerns about acting too early in creating technology-adjustment policy.

Previous shocks have also differed from AGI in important ways. Unlike prior technologies that automated routine work and created new industries, by automating intelligence itself, AGI could allow for far more comprehensive automation of cognitive work. This would risk outpacing society's ability to adapt to the changes.

Yet precisely because AGI automates components of 'intelligence' itself, the economic impacts of widespread

diffusion are very uncertain. At one extreme, AGI's economic gains and disruption could fail to materialize as expected. Alternatively, it could more closely resemble familiar—though still very significant—prior technology shocks, which disrupted some jobs but eventually created others. Indeed, since the Industrial Revolution, labor has remained a relatively constant force in the economy–contrary to consistent predictions of mass labor displacement—and economists are currently not seeing definitive evidence of systemic employment or wage impacts. At the other extreme, however, the broad and dexterous nature of AGI capabilities could mark a break with prior technologies, reducing demand for human labor in most cognitive occupations.

Amid this uncertainty, there are at least a few reasons for hope that society can better manage a disruptive technological shock this time around. In the past, policymakers were limited by weaker and less knowledgeable institutions than we have today. Social science has made strides in understanding how technological shocks unfold and how to measure the impact of new innovations on labor markets. We also know more about the toolkit governments have to implement policies and how to scale them quickly.

Since we may not be able to preemptively determine which

AGI economic scenario will materialize —or its timescale—it is important to develop flexible economic policy responses tied to clear empirical triggers. This approach can help us avoid acting too late to cushion real disruption or overreacting prematurely with interventions that produce a cure worse than the disease. To develop these sorts of policies, we need to overcome three challenges.

The first is our lack of granular and timely labor-market data. Platform usage metrics cannot tell us how AI alters net employment, workplace field experiments rarely generalize to the macroeconomy, and official administrative datasets may lag real-world disruption. Better data on both current statistics and new measures—such as on consumer demand for AI and bottlenecks to adoption—may help us understand how governments ought to respond to the potential advent of AGI. The second challenge is determining *which* policies might actually distribute the benefits of AGI widely while ensuring that people thrive and have agency over their lives. And the third challenge is that there isn't currently a common rubric to assess the strength of one potential policy intervention versus another, across a set of standardized dimensions that society cares about. This includes clear data-based indications of when to deploy any particular policy response.

To help address this gap, we introduced a unified framework

for evaluating potential policies across four dimensions:

1. **Welfare and Resilience**: A policy's impact on material living standards, meaning and purpose, and overall economic stability.
2. **Agency and Voice:** A policy's impact on individual economic choice, direct ownership of AI-driven gains, and democratic participation.
3. **Feasibility and Efficiency**: A policy's political support, popular approval, economic cost, administrative simplicity, speed of rollout and institutional readiness. This includes whether policies create the conditions for AGI-enabled economic growth.
4. **Durability across potential AGI economic futures**: A policy's suitability across scenarios marked by mild disruption, broad displacement and complete economic transformation.

With this rubric in mind, we used a combination of manual literature reviews, surveys, and 51 AI agent raters trained on real economists to evaluate 11 commonly discussed policy ideas for managing the potential AGI economic transition. For the agent-deliberation analysis, each AI agent's personality was designed using survey data from 51 real economists, capturing a diversity of economic and political attitudes. This method, pioneered by John Horton, removes some degrees of researcher freedom, allowing us to test

how agents ‘rate’ particular interventions based on the available evidence and their own deliberative process.

The output of this methodology suggests that no single initiative is the answer to everything. Some policies enjoy popular support but lack practical feasibility. Others appear resilient to AGI disruption but could undermine human agency. Since each intervention carries distinct trade-offs between cost, administrative feasibility, and human agency, an effective strategy should be sequenced around observable thresholds where possible. This strategy may also require investing in infrastructure to track economic data and measure it against predetermined thresholds that can help society determine when to introduce already-designed policies.

Rather than endorsing a single static prescription, our methodology identifies three "least-regret" interventions matched to distinct scenarios:

1. **Preemptive Stabilizers for Mild Scenarios (Scenario 1)**: *Expanded Unemployment Insurance (UI), Earned Income Tax Credit (EITC) and employer-led retraining* - low-regret interventions that connect benefits to work-attachment, and automatically adjust in scale if disruption increases.
2. **For Moderate Displacement & Wage Compression (Scenario 2)**: *Transformation of EITC into Negative*

Income Tax (NIT)—more pronounced interventions that could be triggered if data reveals prolonged unemployment spells and falling median wages that outpace job reinstatement.

3. **For Structural Labor-Capital Decoupling (Scenario 3)**: *A Universal Basic Capital (UBC) backstop*—more significant interventions that could be designed now but deployed if macroeconomic indicators indicate a sustained decline in labor's share of GDP alongside increasing capital returns.

We assessed each of the 11 policies across measures of welfare, agency, feasibility and efficiency, and durability. The full analysis, methodology and evaluation of 14 potential funding mechanisms for each policy are described in our *Economic Policy for AGI* paper.

11 policies we evaluated

Taxonomy and operational scope of the evaluated household-facing interventions

Policy Intervention	Definition & Operational Scope
PANEL A: TARGETED LABOR-MARKET & WAGE INTERVENTIONS	
Active Labour Market Policies (ALMPs) & Retraining	State-sponsored programmes designed to reskill displaced workers and realign their human capital with AI-complementary occupations.
	A transitional subsidy that compensates displaced

Wage Insurance	workers for a portion of the income lost when they secure re-employment at a lower nominal wage.
Earned Income Tax Credit (EITC)	A means-tested, refundable tax credit functioning as a wage subsidy for low-to-moderate-income workers.
Federal Jobs Guarantee	The government becomes the employer of last resort, ensuring that anyone willing to work is provided with a publicly funded position.
Unemployment Insurance	A macroeconomic stabiliser providing temporary, partial wage replacement to displaced workers.
PANEL B: UNIVERSAL FLOORS, SERVICES, & STRUCTURAL ASSETS	
Negative Income Tax (NIT)	Establishes a guaranteed minimum income floor, automatically providing cash transfers to individuals whose income falls below a threshold.
Universal Basic Income (UBI)	An unconditional cash transfer delivered uniformly to all citizens regardless of their employment status or wealth.
Sovereign AI Dividend	A universal and unconditional cash payout funded by the taxation or public ownership of foundational artificial intelligence assets.
Universal Basic Capital (UBC)	The state-facilitated provision of an *ex-ante* asset or equity endowment (e.g., trust accounts or shares in a sovereign wealth fund).
Universal Basic Services (UBS)	The direct and unconditional public provision of essential services, such as healthcare, housing, and transport.
Industrial Policy	Strategic state intervention and targeted capital allocation intended to fund public goods and generate job growth in AI-complemented jobs.
Note: Evaluated interventions cover the primary redistributive, structural, and labor-market pathways for managing technological displacement.	

Welfare and agency

One evaluation criterion is how well a policy supports the welfare outcomes that society cares about the most. Does it

uphold strong standards of living for all people? Does it maintain conditions for economic autonomy and democratic participation? This requires looking not just at the impact of policies on household incomes, but also at whether they can maintain individual agency and the balance of power that is needed for economies and democracy to thrive.

Prior scholarship across the social sciences has shown that work has clear psychological and social value. As research—including our own forthcoming work—suggests, people often attach significant amounts of non-monetary value to the work that they do, and this is true across the income distribution. Research in sociology has argued, for instance, that work also plays an important function for *social cohesion*, with employment serving as a key factor in establishing relationships and deterring antisocial behavior. This research suggests that maintaining people's ability to contribute to the economy should be a salient dimension of a potential AGI policy.

Scenario 1: Preemptive Stabilizers for Mild Scenarios

We first considered a mild scenario, in which AGI boosts productivity and creates familiar automation and augmentation dynamics across labor markets, as we have seen in prior technological waves. Based on historical episodes, we could see wage polarization and shifting demand for skills throughout the economy, without a

collapse in aggregate labor demand.

Our mixed-methods, AI-agent-assisted research suggests that a range of economists would conclude that introducing automatically adjusting stabilizer policies would be helpful in handling potential mild disruption, particularly in a moment of uncertainty. For people who lose their job today and cannot find another right away, expanded unemployment insurance (UI) is one effective intervention that builds upon the welfare infrastructure most countries already have. To manage an AGI transition, however, our analysis suggests this infrastructure may need to be expanded or evolved. For example, in the US system, UI often provides replacement but at a lower wage, expires quickly, and typically excludes non-traditional workers, such as gig workers and freelancers. Existing evidence suggests this system may be inadequate for any scenario in which AGI results in significant labor market disruption.

By guaranteeing a majority of a person's prior wage for a set period after they lose their job, and while they look for new work, UI maintains standards of living during economic shocks while also creating capacity for genuine reskilling. In Denmark, for example, UI is a central component of the country's renowned flexicurity social security programming, in which people receive up to 90% of their prior salary, up to a maximum of $3,200 USD, for up to two years after losing

employment.

Worker retraining programs, particularly those that directly involve employers, such as apprenticeship models, could also be important. Public worker retraining programs have historically been disappointing in supporting labor adjustment to technological change. Instead, apprenticeships and employer-led models show the highest incidence of success by a considerable margin. If successfully deployed, retraining could be a useful intervention to kindle upskilling into AI-complementary jobs.

Alongside UI and retraining, expanding the Earned Income Tax Credit (EITC) could also play an important role as a preemptive stabilizer. The EITC is a refundable tax credit that supplements the wages of low- and moderate-income workers, rewarding employment while lifting households above the poverty line. Data has shown it to be one of the most successful poverty reduction programs in history. If AI compresses wages without eliminating jobs, expanding the EITC directly cushions living standards while maintaining incentives to work. Today's EITC largely excludes workers without dependents and cuts off at low wage levels. Expanding the credit to include a broader swath of low- and middle-income earners, and distributing it in regular payments rather than once a year at tax time, could directly buffer workers transitioning between jobs.

Evaluating AGI economic policies: Agency and democratic empowerment

Color-coded expert panel ratings (*N* = 51 simulated economist personas) across the three agency sub-criteria (0-100)

Policy Intervention	Economic Participation	Ownership of Gains	Democratic Voice	Composite Score
Universal Basic Capital (UBC)	66.7	94.9	67.2	76.3
Sovereign AI Fund / Dividend	52.7	54.9	58.0	55.2
Negative Income Tax (NIT)	64.2	39.8	57.0	53.7
Universal Basic Income (UBI)	61.9	51.3	46.2	53.1
Directed Industrial Policy	59.1	47.0	50.1	52.1
Earned Income Tax Credit (EITC)	64.7	20.6	56.4	47.2
Federal Jobs Guarantee (FJG)	51.1	33.5	54.0	46.2
Wage Insurance	59.8	21.3	53.0	44.7
Unemployment Insurance (UI)	63.9	19.2	48.1	43.7
Universal Basic Services (UBS)	61.4	12.4	52.0	41.9
Active Labour-Market Policies	52.1	9.4	49.4	37.0

Score Tiers: High (≥ 70) Moderate (55–69) Mixed (42–54) Low (< 42)

Source: Authors' Multi-Agent Expert Economist Panel Benchmark (*N* = 51 EDSL personas). Policies ordered by Composite Agency & Empowerment Score.

Economic Participation: preservation of individual worker autonomy, career self-determination, and economic agency. Ownership of Gains: broad-based societal distribution of AI-generated capital and productivity gains. Democratic Voice: individual influence over economic governance and workplace decisions.

Scenario 2: Moderate Displacement and Wage Compression

In scenarios where AGI capabilities advance and result in moderate-to-broad displacement, we would expect more substantial labor market churn, prolonged unemployment spells, and potentially permanent under/unemployment. In this world, jobs would be likely to change as efficiency gains drive new demand, particularly in sectors where human involvement is part of the value, and areas demanding human expertise or oversight. Yet labor reinstatement effects may take longer to unfold relative to the pace of automation.

If economic data indicates this scenario is emerging, our work suggests safety nets should shift from temporary transition support to durable income support. Universal basic income (UBI) programs dominate public discourse about economic policies for a world with AGI. Yet our analysis suggests it matters greatly *how* governments deploy direct transfers of money. In the case of UBI, the program risks kindling inflationary pressure, and, most importantly, is an expensive and blunt instrument that may fail to concentrate sufficient relief where it is needed most. In

a scenario where AGI radically transforms the economy, a universal income transfer like UBI may also fail to support economic agency because it would not alter the structural forces that shape economic participation, standards of living and opportunity. UBI may also leave people with no direct stake in the automated parts of the economy and no direct mechanism for economic participation.

Instead, an expanded Earned Income Tax Credit program that transitions to a negative income tax (NIT) if Scenario 2 thresholds are triggered may be a more economically efficient mechanism for redistributing technological gains, according to our analysis. If policymakers establish a tax code where individuals below a particular income threshold automatically receive cash transfers, which phase out as their earned wages rise, the state could guarantee a basic standard of living while preserving incentives for work.

An NIT provides support based strictly on *income level*, rather than the *cause* of income loss. This also makes it superior to targeted interventions like wage insurance, where administrators would face the difficult task of untangling whether a worker was displaced specifically by AGI, or by standard market churn. Additionally, NITs would take effect only when a person's income falls below a certain level, making the program much cheaper and more feasible than alternatives like UBI, which provide funds for everyone

regardless of their income level.

Scenario 3: Structural Labor Disruption

Finally, our research explores a hypothetical and extreme scenario in which AGI marks a fundamental discontinuity in the history of economics and human social order. In such a world, the marginal productivity of human labor could approach zero in many sectors. AGI that is versatile, adaptable, and recursively improving could theoretically erode the traditional mechanisms through which economies have created new jobs.

In such a scenario, policy should ensure that people continue to have a stake in the economy. In the absence of this, some scenarios of AI-enabled growth could calcify economic inequalities, where a narrow segment of the population controls the capital, captures the returns to innovation, and the broader public is locked out of a meaningful role in the productive parts of the economy. This would be particularly true if AGI were to result in a shift in economic growth from labor to capital, and if new AGI-complemented jobs took longer than expected to emerge.

Several ambitious policies have been discussed recently, including an AI sovereign wealth fund and a universal basic capital (UBC) program. An AI sovereign wealth fund is a state-owned investment pool based on a government

surplus. Within the context of AGI, this might involve the government capturing a share of AI-generated wealth—such as through increased corporate tax rates—investing the proceeds, and offering dividends to the public as the fund compounds over time. A good example of this intervention in practice is Alaska's sovereign wealth fund, where eligible residents receive an annual cash dividend (such as the $1,000 payout in 2025) funded by the realized investment earnings of the state's natural resource royalties.

Evaluating AGI economic policies: Welfare and resilience

Color-coded expert panel ratings (*N* = 51 simulated economist personas) across the three welfare sub-criteria (0-100)

Policy Intervention	Standards of Living	Meaning & Human Value	Macro Stabilisation	Composite Score
Negative Income Tax (NIT)	74.3	69.3	65.7	69.8
Earned Income Tax Credit (EITC)	67.1	71.8	67.6	68.8
Unemployment Insurance (UI)	63.2	61.4	75.6	66.7
Universal Basic Capital (UBC)	69.8	52.3	64.0	62.0
Universal Basic Services (UBS)	64.0	54.2	63.9	60.7
Sovereign AI Fund / Dividend	52.4	40.6	68.0	53.7

Universal Basic Income (UBI)	64.7	41.4	47.4	51.2
Direct Industrial Policy	58.3	64.3	24.6	49.1
Wage Insurance	55.5	58.9	27.8	47.4
Federal Jobs Guarantee (FJG)	63.0	41.2	28.9	44.4
Active Labour-Market Policies	25.6	47.5	37.7	36.9

Score Tiers: High (≥ 70) Moderate (55–69) Mixed (42–54) Low (< 42)

Source: Authors' Multi-Agent Expert Economist Panel Benchmark (*N* = 51 EDSL personas). Policies ordered by Composite Welfare & Resilience Score.

Standards of Living: protection against material deprivation and baseline consumption maintenance. Meaning & Human Value: individual psychosocial meaning, status, and purpose. Macroeconomic Stabilisation: counter-cyclical support for aggregate demand and household purchasing power during automation slumps.

Still, if AGI—let alone AI—does indeed boost economic productivity, these gains should appear throughout the economy, not just in the AI sector. This makes it fundamentally different from oil, which is a discrete, identifiable income stream.

Our research outputs suggested that a UBC program might be a more effective intervention for sharing capital growth. Unlike a UBI, which relies on transferring cash, or a SWF, which relies on direct government mediation, UBC works by granting people a direct ownership stake—such as shares in a publicly managed fund invested in equity markets and bonds—so their wealth compounds automatically alongside

the broader economy. This could take a form similar to recent proposals to invest Social Security and other trust funds in US equity markets. Such a diversified capital portfolio could grant people—regardless of work status—a stake in the economy's gains, creating a capital underpinning for the safety net.

In practice, our analysis suggests that it may be premature to deploy UBC as a reaction to advanced AI systems right now, before clearer structural impacts from the technology appear in the macroeconomic data. Instead, capital-sharing interventions like UBC could be treated as a backstop, triggered when macroeconomic data suggests signals of an economy-wide decoupling of economic growth from much of the labor market (Scenario 3). In the near term, the priority should be to design the institutional architecture–including governance, funding methods, and distributional mechanics–so that the policy could be ready to implement at the right moment.

Some interventions are more feasible than others

Of course, a policy intervention is worth little if society cannot effectively implement it. This is why we also assessed each intervention across measures of political support, implementation feasibility, popular support, and administrative capacity.

Political and popular support varies widely across the policies we evaluated. Worker retraining is a feature in most developed countries and is politically uncontroversial, though exact methods and resourcing for these programs diverge considerably. It is the most popular AGI economic policy intervention in the US, with 85% of Americans in support of public-funded worker retraining.

Meanwhile, versions of unemployment insurance are already deployed across many economies and are similarly popular, with 72% of Americans in support, according to our recent survey. By contrast, a universal basic capital (UBC) program is historically unprecedented (apart from sovereign wealth funds that are typically based around natural resources), and could face significant political headwinds, alongside possibly more tepid public support (54% of Americans surveyed were in favor of a UBC program).

Most importantly, UBC confronts entrenched political philosophies regarding the relationship between citizens and their state, may face legal challenges, and could require the development of new institutions or agencies. On the other hand, the underlying concept of individual, publicly managed investment accounts is not entirely new; it closely mirrors past bipartisan proposals for personal Social Security accounts.

Evaluating AGI economic policies:

Durability across scenarios

Color-coded expert panel ratings (*N* = 51 simulated economist personas) across the three AGI transition phases (0-100)

Policy Intervention	Mild Disruption	Broad Displacement	AGI Transformation	Composite Score
Universal Basic Services (UBS)	73.0	74.0	78.0	75.0
Universal Basic Capital (UBC)	63.0	66.0	93.5	74.2
Negative Income Tax (NIT)	74.0	69.8	65.0	69.6
Sovereign AI Fund / Dividend	60.4	62.0	68.0	63.5
Unemployment Insurance (UI)	70.0	65.9	50.1	62.0
Universal Basic Income (UBI)	51.0	57.0	65.0	57.7
Directed Industrial Policy	38.9	50.3	65.0	51.4
Earned Income Tax Credit (EITC)	68.9	52.3	31.9	51.0
Wage Insurance	35.7	38.9	33.0	35.9
Federal Jobs Guarantee (FJG)	18.0	37.0	50.0	35.0

Active Labour-Market Policies	42.5	20.6	5.9	23.0

Score Tiers: High (≥ 70) Moderate (55–69) Mixed (40–54) Low (< 40)

Source: Authors' Multi-Agent Expert Economist Panel Benchmark (*N* = 51 EDSL personas). Policies ordered by Composite Durability Score.

Mild Disruption: localized task automation with gradual reallocation. Broad Displacement: widespread displacement outrunning new job creation. AGI Transformation: comprehensive automation fundamentally decoupling human labor from baseline subsistence.

Building positive AGI economic futures

Since we cannot reliably predict when AGI might emerge, nor the precise nature of its economic impacts, we also measured the *durability* of all 11 potential policy interventions across a range of potential AGI economic scenarios. One contribution of our roadmap is outlining how these proposals might automatically adapt across scenarios.

For instance, expanded EITC and UI programs could be preemptively designed to manage any mild labor market disruption (Scenario 1). Yet these interventions could—both politically and economically— pave the way for a Negative Income Tax if displacement deepens (Scenario 2). And in an extreme scenario where AGI were to usher in an unprecedented wave of disruption and decoupling of labor from economic growth (Scenario 3), an NIT could soon after function as a universal income floor as earned wages

declined. Pairing these automatic stabilizers with a pre-designed Universal Basic Capital (UBC) interventions could insulate living standards and economic opportunity, even if society were to undergo a radical transformation. Indeed, the output of our work suggests that citizens would be protected immediately by cash transfers while retaining a durable ownership stake in capital-driven growth.

Evaluating AGI economic policies: Feasibility and implementation

Color-coded expert panel ratings (N = 51 simulated economist personas) across the six feasibility sub-criteria (0-100)

Policy Intervention	Political Support	Economic Feasibility	Popular Support	Admin. Capacity	Speed	Rea
Earned Income Tax Credit (EITC)	75.2	83.2	77.3	77.9	69.8	95.3
Unemployment Insurance (UI)	35.5	81.9	59.2	34.3	23.0	91.9
Negative Income Tax (NIT)	22.9	78.9	36.7	66.5	38.4	57.9
Wage Insurance	38.3	47.9	48.4	68.0	32.0	54.8
Active Labour-Market Policies	40.7	50.1	62.3	41.2	28.4	65.0
Directed Industrial Policy	35.6	54.0	41.7	39.3	21.3	55.0

Sovereign AI Fund / Dividend	23.1	48.5	45.5	44.4	17.8	51.0
Universal Basic Income (UBI)	15.9	40.2	36.7	63.4	18.4	38.6
Universal Basic Services (UBS)	18.3	61.0	40.3	25.0	14.1	48.4
Universal Basic Capital (UBC)	19.0	55.0	39.6	38.2	17.1	30.0
Federal Jobs Guarantee (FJG)	24.2	38.0	43.9	21.7	15.7	27.0

Score Tiers: High (≥ 70) Moderate (55–69) Mixed (40–54) Low (< 40)

Source: Authors' Multi-Agent Expert Economist Panel Benchmark (*N* = 51 EDSL Policies ordered by Composite Feasibility & Implementation Score.

Political & Popular Support: legislative viability and broader public acceptance. & Admin. Capacity: fiscal/macro cost viability and institutional execution.

Speed & Readiness[†]: rapidity of disbursement and author-coded policy operati 100).

The potential emergence of AGI could be among the most significant moments in economic history. By reducing the costs of goods and services throughout the economy, an AGI transition could kindle significant improvements to human well being, health, and prosperity, in a manner similar to the Industrial Revolution and early 20th century periods of technological innovation.

This was an initial attempt to taxonomize and assess in a standardized framework the battery of potential

interventions that societies could explore for an AGI transition, using an exploratory agent-based methodological approach. Yet more work—especially experiments, empirical evaluations, and research studying the global economic developmental impacts of AGI—is necessary to further validate policy options as AI evolves. Throughout this economic transition, it will be particularly important for policymakers to improve data collection on AI's economic impacts in order to better understand the real-world impact of the technology on employment and wages.

The value of this framework is that it also illustrates how we might sequence policy. Interventions like expanded unemployment insurance and negative income taxes may be worth implementing sooner since they automatically adjust to labor market conditions. Other policies—such as basic capital programming—may be worth implementing at later stages, as the evidence on AI's impacts becomes clearer. This will require investment in improved data collection and monitoring. Interventions deployed too early could impose unnecessary costs and market distortions, while policies deployed too late could leave society unable to address initial AI-labor shocks.

The most important takeaway from our work is that society has the capacity and tools to shape our economic trajectory in the AGI era. As history has repeatedly demonstrated,

economic growth that concentrates among a small segment of society is socially and politically fraught. If AGI emerges and does indeed increase economic growth, as we expect, then a central question of our time may be one about ensuring wide access to the benefits of that growth. This includes using AGI's gains to support improved standards of living and opportunity for all.

Disclaimer

The DMI is a platform, started by researchers from Google and Google DeepMind, to publish and discuss creative, deeply informed ideas about a world with AGI. DMI pieces are intended to serve as conversation starters, reflecting the author's ideas and research, and should not be read as Google's official view.