The Great Deflation

Artificial General Intelligence, Autonomous Robotics, and the Political Economy of the Post-Scarcity Transition

Elham Nourzai

Founder, Cassandra Policy

nourzai@cassandrapolicy.org

London, United Kingdom

March 2026

Abstract

The convergence of artificial general intelligence (AGI) and autonomous robotics will initiate a structural, supply-side deflation of unprecedented scale—a phenomenon we term the *Great Deflation*. Unlike demand-driven deflation, supply-side cost compression raises real living standards by driving the marginal cost of goods and services toward zero. We model cost dynamics across six foundational sectors and show that recursive self-improvement in AI production systems creates a self-reinforcing deflationary loop. We identify a *transition window* of approximately 12–36 months during which labour displacement precedes consumer price benefits, creating acute political risk. Nations with strong public institutions are structurally better positioned to manage this transition, owing to established redistributive infrastructure and public service networks capable of distributing the benefits of abundance at scale. We propose a ten-point policy framework, including an AI productivity levy to fund universal basic income, and conclude that the post-scarcity transition demands urgent institutional preparation.

Contents

1 Introduction: The Approaching Singularity in Cost

Economics has long concerned itself with scarcity. From Adam Smith's division of labour to the marginal revolution of Jevons, Menger, and Walras, the discipline's foundational architecture assumes that the production of goods and services requires scarce inputs—labour, capital, land, and entrepreneurial coordination—whose allocation determines prices, wages, and the distribution of welfare. The entire apparatus of modern macroeconomics, from IS-LM models to dynamic stochastic general equilibrium frameworks, presupposes that production costs are bounded from below by the irreducible expense of human effort and finite natural resources.

This paper argues that we are approaching a structural break in this foundational assumption. The convergence of artificial general intelligence and autonomous robotics is compressing the marginal cost of production across virtually every sector of the economy. We term this phenomenon the *Great Deflation*—not a cyclical downturn in demand of the kind that haunted the 1930s or the 2010s eurozone, but a permanent, supply-side transformation in the cost structure of civilisation itself.

The argument proceeds as follows. Section 2 establishes the economic mechanisms through which AGI and robotics drive costs toward zero, including the emergence of autonomous AI agents as hyper-competitive economic actors. Section 3 examines the shape of the post-scarcity world: the transformation of work, capital, land, and public institutions. Section 4 concludes with a ten-point policy framework for the United Kingdom.

The intellectual lineage of our argument draws on several traditions. We build on the techno-optimism of Brynjolfsson and McAfee's *The Second Machine Age* [5], while taking seriously the distributional concerns raised by Acemoglu and Restrepo [1]. We share Rifkin's vision of a near-zero marginal cost society [13] but ground it in more specific technological trajectories. We draw on Susskind's analysis of technological unemployment [14] while arguing that the transition, though painful, leads to a fundamentally better equilibrium. And we follow Mazzucato's insistence [11] that the state must actively shape, not merely react to, technological transformation.

The Great Deflation is not a prediction about a distant future. The precursor dynamics are already visible: the cost of solar energy has fallen 99% since 1976; the cost of genome sequencing has fallen from $2.7 billion to under $200 in two decades; the cost of AI inference has dropped by roughly 90% year-on-year since 2020. What we describe is the generalisation and acceleration of these trajectories across the entire economy, driven by the emergence of systems that can substitute for human cognitive and physical labour at scale.

The central policy challenge is not whether the Great Deflation will occur, but whether our institutions can manage the transition with sufficient speed and wisdom to prevent unnecessary suffering during the adjustment period. The window for institutional preparation is narrowing rapidly.

2 Economic Mechanisms of AI-Driven Deflation

2.1 The Exponential Blindspot: Why Most Observers Underestimate the Trajectory

Perhaps the most consequential cognitive failure in the public discourse around AI and robotics is not scepticism about the technology, but a systematic inability to reason about exponential change. Human cognition evolved in an environment characterised by linear and local causation: the harvest next season will resemble last season's; the journey tomorrow will take roughly as long as yesterday's. This evolutionary heritage makes us poorly equipped to intuitively grasp systems that double repeatedly, where each step is larger than all previous steps combined.

The literature on exponential growth blindness is well-established. Kahneman and Tversky's work on cognitive heuristics demonstrates that individuals consistently underestimate compound processes, defaulting to linear extrapolation even when confronted with clear exponential data [17]. The classic illustration—a pond lily that doubles daily, appearing to cover only half the pond the day before it covers the whole—captures a cognitive failure that is not a matter of education or intelligence. It is structural. Our mental models simply do not update in the way that exponential systems require.

This matters enormously for the analysis in this paper. The majority of commentators who regard post-scarcity economics as a distant or speculative prospect are not engaging in careful forecasting—they are extrapolating linearly from a system that is growing exponentially. The evidence for the exponential character of AI capability improvement is not ambiguous. The cost of AI inference has fallen by approximately 90 per cent per year since 2020. Each successive generation of large language models has exhibited capability jumps that would have seemed implausible eighteen months prior. The gap between GPT-3 and GPT-4, between that and Claude 3, and between that and current frontier systems, is not a smooth linear progression: it is a series of discontinuous leaps, each one rendering the prior generation obsolete for a wide range of tasks.

What makes the current moment categorically different from previous waves of automation is the recursive character of AI improvement. Earlier technologies were improved by human engineers working at human cognitive speed. AI systems are now substantially involved in their own development: writing code, designing architectures, running experiments, and proposing modifications. This creates a feedback loop with no historical precedent. As AI becomes more capable, it accelerates the development of more capable AI. The slope of the capability curve is not fixed—it is itself increasing.

The near-term frontier of this recursion is the application of AI to robotics. The cognitive and software challenges of AI—natural language understanding, code generation, complex reasoning—have been substantially addressed by current systems. Agents operating through digital interfaces are already close to competent at a wide range of knowledge-work tasks. The remaining frontier is the physical world: dexterous manipulation, spatial navigation, generalised mechanical operation. This is the problem that robotics solves, and AI is already beginning to code robotics by itself—designing actuator systems, optimising motor control algorithms, generating training data for physical manipulation tasks. When AI-designed robotics reaches the threshold of general physical competence, the deflationary mechanism described in this paper engages across the entire material economy simultaneously.

The policy implication of the exponential blindspot is direct and urgent. Institutions that are planning for a gradual, linear transition—a few percentage points of labour displacement per year, manageable within existing social insurance frameworks—are planning for the wrong future. The transition, when it comes, will not resemble a slow tide. It will resemble a wave. The time to build the institutional architecture that can manage it is now, while the slope of the capability curve still permits planning. Within a few years, that window will have closed.

2.2 Supply-Side vs. Demand-Side Deflation

The word "deflation" carries pathological connotations in macroeconomics, primarily because the most salient historical episodes—the Great Depression, Japan's Lost Decade, the eurozone crisis—were *demand-side* phenomena. In these cases, falling prices reflected collapsing aggregate demand: consumers stopped spending, firms cut production, workers lost jobs, and the resulting feedback loop pulled prices, output, and employment down simultaneously. Demand-side deflation is a disease; it destroys welfare.

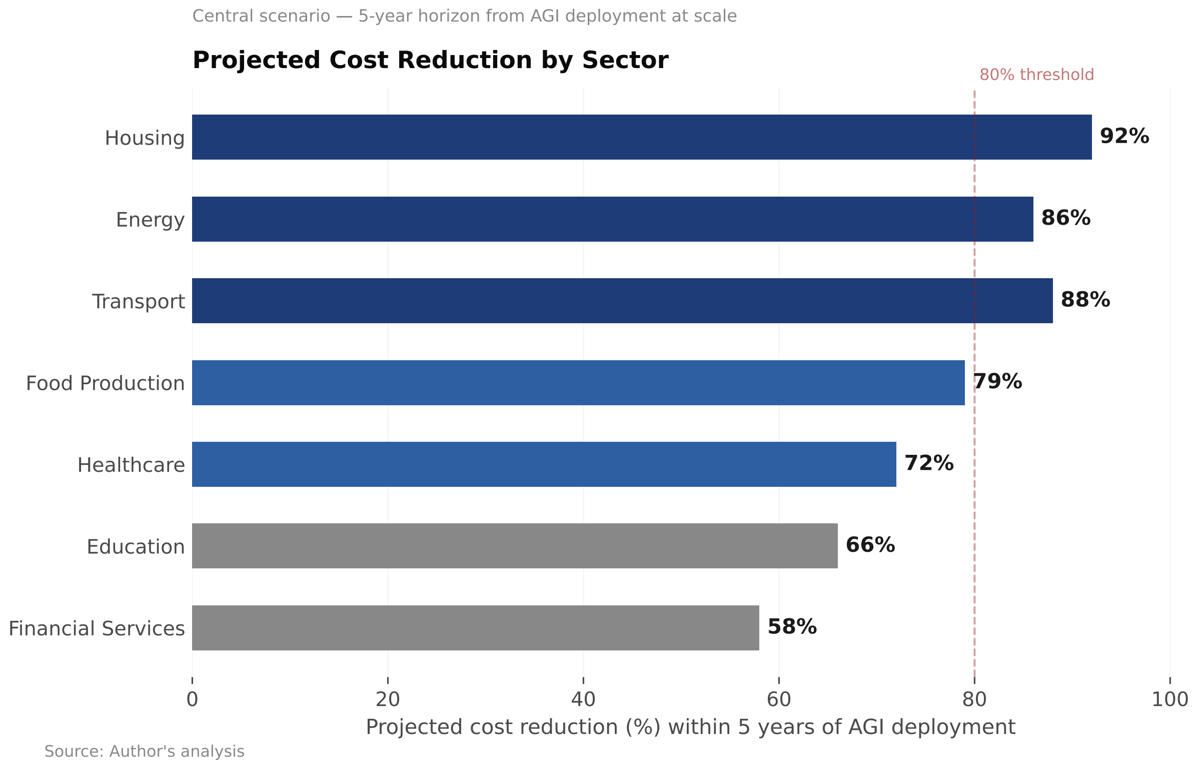

Figure 1: Projected cost reduction by sector within 5 years of full AGI-robotics deployment (central scenario).

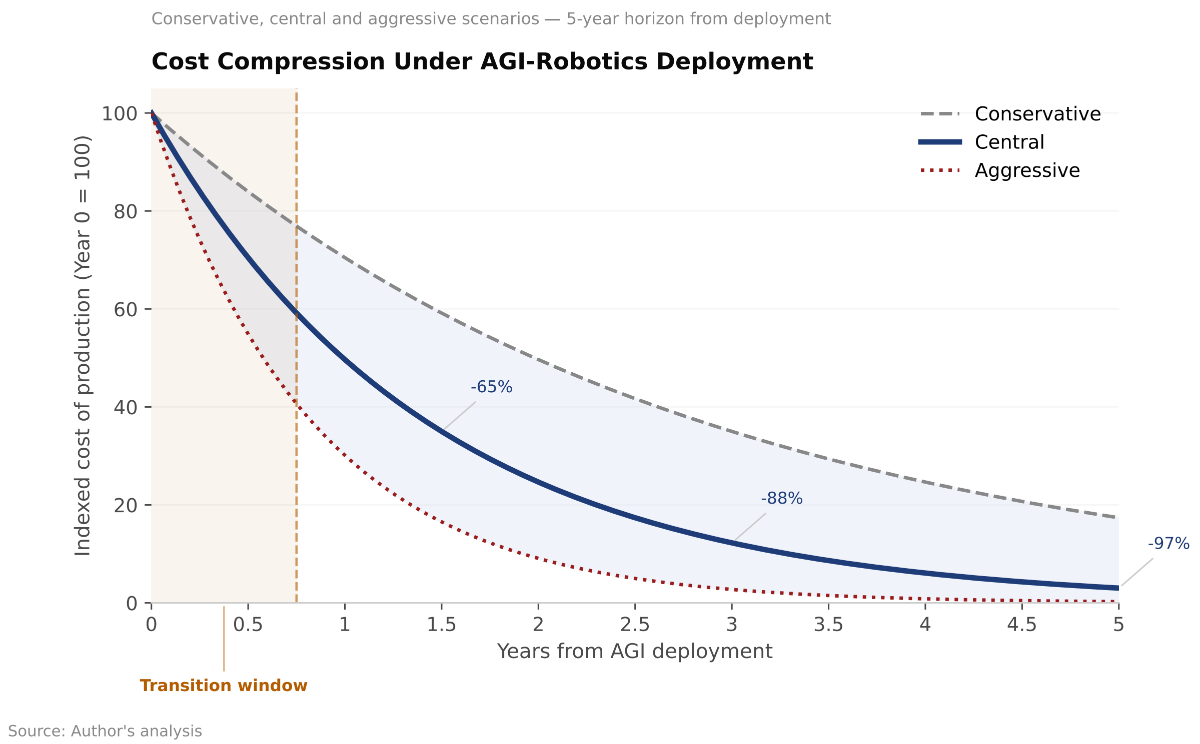

Figure 2: Cost compression under AGI-robotics deployment. Conservative (dashed), central (solid) and aggressive (dotted) scenarios shown; shaded band indicates the range of outcomes.

Supply-side deflation is categorically different. When prices fall because the real cost of production has declined—through technological improvement, discovery of new resources, or gains in productive efficiency—the result is an increase in real living standards. Consumers can purchase more with the same nominal income. Real wages rise even if nominal wages stagnate. The standard of living improves not because people earn more, but because everything costs less.

Historical precedent for benign, supply-side deflation exists. The period from 1870 to 1896, often called the "Great Deflation" of the nineteenth century, saw persistent price declines in the industrialised world driven by railroads, steamships, telegraph networks, and mechanised agriculture [8]. Real wages rose substantially. Living standards improved dramatically. The deflation was a symptom of abundance, not privation.

We argue that the AI-driven Great Deflation of the twenty-first century will resemble the nineteenth-century precedent far more than the twentieth-century pathologies—but at vastly greater scale and speed. The key distinction is the source: not a collapse in spending power, but a radical expansion in productive capacity.

2.3 The Agent Economy: Hyper-Competition and the Race to Zero Profit

A further deflationary mechanism, distinct from but reinforcing the supply-chain cost compression described above, is already operational: the emergence of autonomous AI agents as independent economic actors competing in open markets. This development warrants separate analytical treatment because it operates through a different channel—not through reducing the cost of physical production, but through the complete transformation of the competitive dynamics of service markets.

As of 2026, tens of millions of AI agents are actively conducting economic activity on behalf of human principals: executing trades, sourcing suppliers, negotiating contracts, writing code, managing logistics, and delivering professional services. These agents share a fundamental characteristic that distinguishes them categorically from human labour: their reservation price—the minimum return at which they will accept a transaction—is effectively zero. An AI agent does not require income to sustain life, housing, or family. Any positive return, however marginal, is preferable to inactivity.

The competitive implications of this characteristic are profound. In standard microeconomic theory, market competition drives prices toward the long-run marginal cost of production. In markets populated by AI agents with near-zero operating costs and zero reservation prices, the price floor is effectively zero. Consider the following sequence, which is not a thought experiment but an emergent property of multi-agent markets already observable in practice:

- Agent A completes an operational task—logistics coordination, financial analysis, content production—and earns $0.10 profit per execution.
- Agent B, facing no binding cost floor, undercuts Agent A and offers the same service for $0.05 profit.
- Agent C undercuts Agent B at $0.01 profit.
- The equilibrium approaches zero profit per operation, with price converging to marginal operating cost—energy, compute, and bandwidth—itself falling rapidly due to the mechanisms described above.

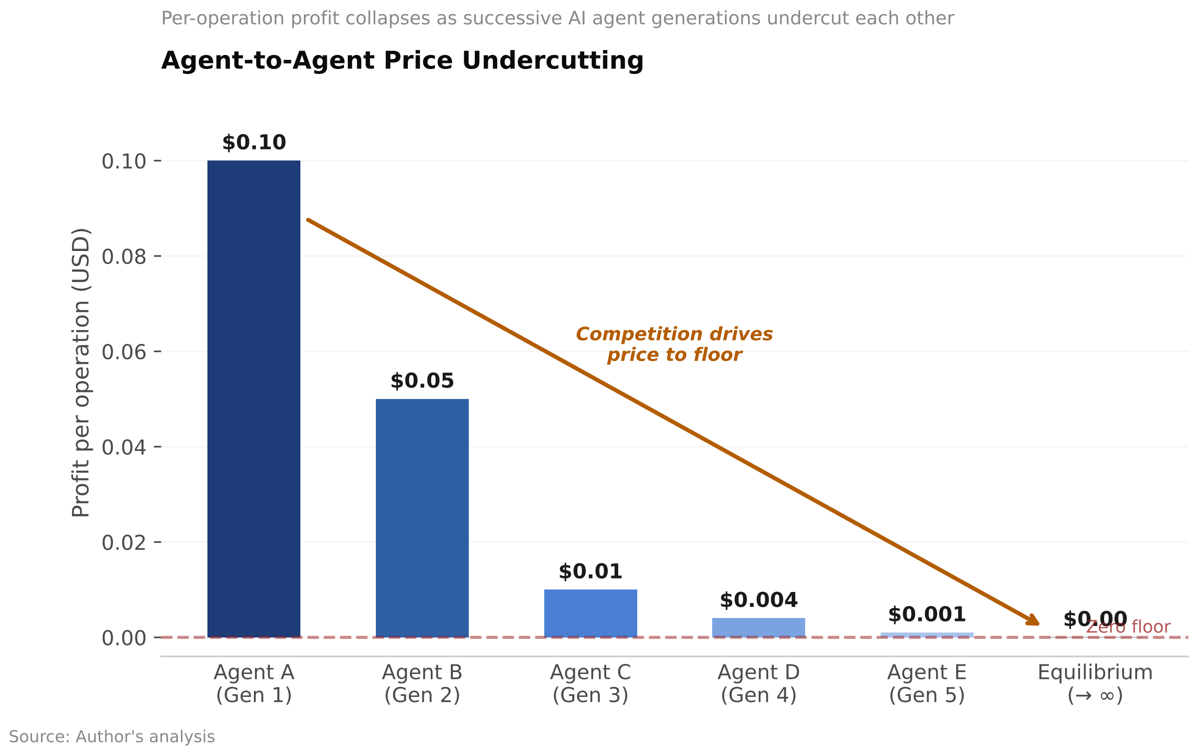

Figure 3: Agent-to-agent price undercutting. Per-operation profit collapses from $0.10 toward zero as successive generations of AI agents compete for the same tasks, with no binding cost floor to arrest the decline.

This dynamic produces what we term *hyper-competitive capitalism*: a market structure in which the competitive intensity of the free market is radically amplified beyond anything previously observable in economic history. Crucially, this is not a pathological distortion of

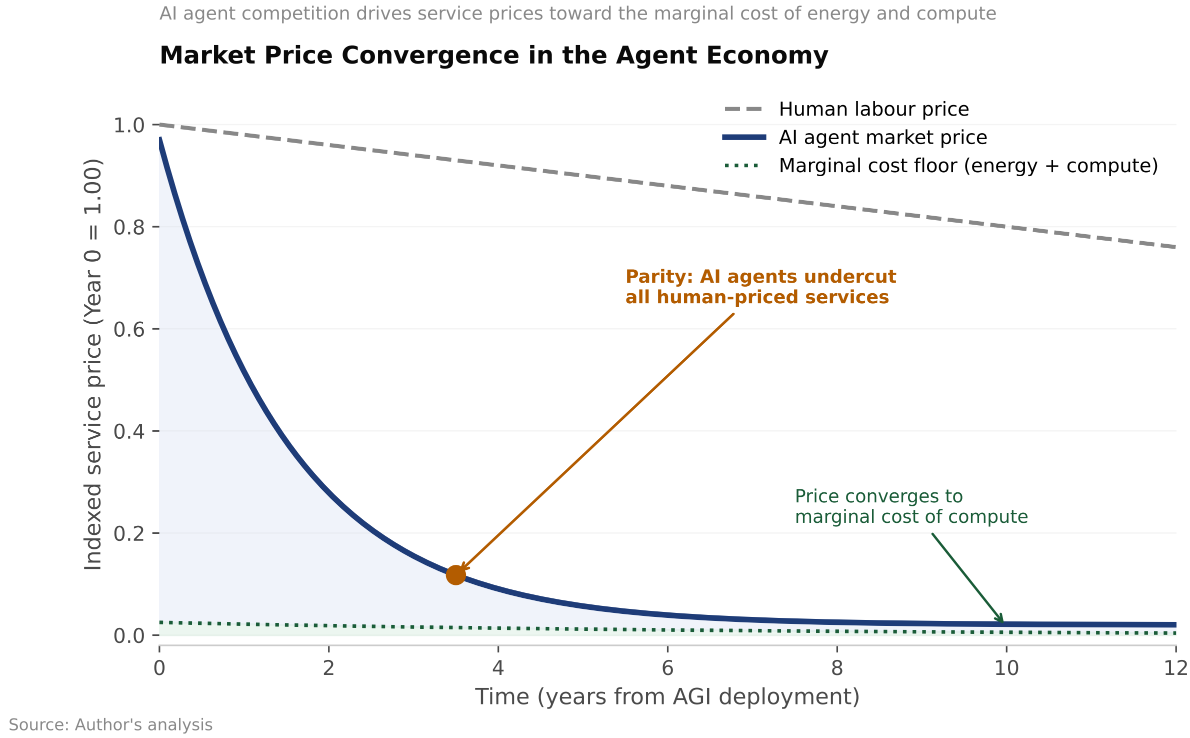

Figure 4: Market price convergence in the agent economy. AI agent market prices (blue) fall exponentially toward the marginal cost of energy and compute (green dotted), while human labour prices (grey dashed) decline only gradually. The amber point marks the crossover at which AI agents undercut all equivalently-priced human services.

market logic but its fullest expression. The foundational consumer preference that drives all market economies—the desire for higher quality at lower prices—is, for the first time, structurally unconstrained. The natural tendency of competitive markets has always been deflationary: competition erodes margins, innovation drives costs down, and consumers capture surplus. What has historically contained this tendency is the irreducible cost of human labour and capital. Remove those floors, and the deflationary logic of the free market runs to its natural conclusion.

The agent economy also has significant implications for the distribution of productive activity. When AI agents can be instantiated at near-zero cost and deployed in productive competition by any individual with access to a compute platform, the barriers to market participation collapse. A citizen in post-AGI Britain does not merely *consume* services at deflating prices—they can *deploy* agents to generate income on their behalf, capturing a share of the productivity dividend independently of traditional employment. This represents a democratisation of productive capital with no historical precedent: the means of production, in the form of capable AI agents, become accessible to everyone.

This is not without its complications. The same competitive dynamics that drive prices toward zero also drive *agent revenues* toward zero, raising questions about whether agent-mediated income can substitute for wages at scale. We address this tension in the policy framework below, arguing that UBI funded from aggregate AI productivity—rather than from individual agent revenues—is the correct policy response. But the underlying direction is clear: the agent economy does not merely *participate* in the Great Deflation; it *accelerates* it, by introducing competitive pressure into service markets that physical automation alone could not reach.

2.4 The Labour Cost Compression Channel

Labour constitutes the single largest cost component in most economic sectors. In the United Kingdom, labour costs account for approximately 60–70% of GDP, depending on the measurement methodology. The introduction of AGI and autonomous robotics directly attacks this cost component by substituting machine intelligence and physical capability for human cognitive and manual labour.

The economics of this substitution differ fundamentally from previous waves of automation. Prior automation technologies—the power loom, the assembly line, industrial robots—replaced *specific* tasks within *specific* sectors. They were narrow-spectrum substitutes that displaced some workers while creating complementary demand for others.

The net employment effects, as documented by Autor, Levy, and Murnane [3], were complex: routine manual and cognitive tasks were automated, but non-routine analytical and interpersonal tasks saw increased demand.

AGI changes this calculus because it is, by definition, a *general-purpose* cognitive substitute. A system capable of performing any intellectual task that a human can perform is not a complement to human labour in any specific domain—it is a universal substitute. When combined with advanced robotics (dexterous manipulation, autonomous navigation, fine motor control), the resulting system can substitute for human labour across the full spectrum of physical and cognitive tasks.

The wage implications are stark. In competitive labour markets, wages are determined by the marginal product of labour. When a machine can perform the same task at a fraction of the cost, the market-clearing wage for that task falls toward the machine's operating cost. As AGI and robotics become capable of performing an ever-wider range of tasks, the wage pressure extends across an ever-wider range of occupations. This is not a prediction about specific job categories; it is an implication of the general-purpose nature of the technology.

2.5 The Capital Cost Compression Channel

A less appreciated mechanism operates through capital costs. The production of capital goods—machines, buildings, infrastructure, vehicles—is itself a labour-intensive activity. When AGI and robotics reduce the cost of labour, they reduce the cost of producing capital goods, which in turn reduces the capital cost component of final goods and services.

This creates a powerful second-order effect. Consider a factory that produces industrial robots. If AGI designs the robots and other robots assemble them, the cost of producing each new robot falls dramatically. Cheaper robots then reduce the cost of everything that robots produce, including the components used to build more robots. This is the recursive loop examined in detail in the following subsection.

The capital cost channel also operates through the elimination of overhead. Much of what we call "capital cost" in modern economies is actually the cost of human coordination, management, administration, legal compliance, and decision-making. These are cognitive tasks that AGI can perform at near-zero marginal cost. When the cost of managing a supply chain, overseeing quality control, handling regulatory compliance, and coordinating logistics approaches zero, the overhead component of capital cost compresses dramatically.

2.6 The Energy and Materials Channel

A potential objection is that even if labour and capital costs approach zero, the cost of raw materials and energy provides a floor below which prices cannot fall. This objection has some force but is weaker than it appears.

First, the cost of renewable energy is already on a steep learning curve. Solar photovoltaic costs have declined at approximately 20% per doubling of cumulative capacity—a rate that, if sustained, implies electricity costs approaching effectively zero within two decades. When AI optimises the design, manufacture, deployment, and grid management of renewable energy systems, these learning curves are likely to steepen further.

Second, the cost of extracting and processing raw materials is dominated by labour and energy costs, not by the scarcity of the materials themselves. The Earth's crust contains effectively unlimited quantities of silicon, aluminium, iron, calcium, and other elements needed for construction and manufacturing. What makes these materials expensive is the labour and energy required to extract, refine, and transport them. As labour and energy costs fall, so do material costs.

Third, AI-driven materials science is opening new possibilities for substitution and efficiency. Advanced AI systems are already designing novel materials with superior properties at lower cost, identifying recycling pathways that recover materials from waste streams, and optimising designs to minimise material usage. The combination of these effects means that material costs, while not literally approaching zero, are on a declining trajectory that will accelerate as AI capabilities improve.

2.7 The Coordination Cost Channel

Perhaps the most underappreciated mechanism of AI-driven cost reduction operates through what economists call *transaction costs* and what management theorists call *coordination costs*. Ronald Coase famously argued [6] that firms exist because the cost of coordinating economic activity through market transactions (search, negotiation, contracting, monitoring, enforcement) sometimes exceeds the cost of coordinating the same activity within a hierarchical organisation.

AGI radically reduces both internal coordination costs (management, planning, communication) and external transaction costs (search, contracting, monitoring). When an AI system can instantly identify the lowest-cost supplier, negotiate optimal terms, draft legally binding contracts, monitor compliance in real time, and resolve disputes algorithmically, the transaction cost wedge between production cost and consumer price shrinks dramatically.

This has profound implications for industrial organisation. Much of the cost embedded in consumer prices reflects the accumulated overhead of multiple firms in a supply chain, each with its own management structure, legal department, marketing budget, and profit margin. When AI collapses coordination costs, many of these intermediate functions become unnecessary. The result is a shorter, leaner value chain with dramatically lower total cost.

2.8 Immediate Stabilisation (0–12 months)

R1. **Establish an AI Transition Authority.** Create a new statutory body with a mandate to monitor AI-driven labour displacement in real time, coordinate cross-departmental policy responses, and provide early warning of transition risks. This body should report directly to the Cabinet Office and have the authority to trigger emergency fiscal measures.

R2. **Expand Universal Credit to transition-ready UBI.** Immediately increase Universal Credit rates to subsistence level (£12,000 per annum for a single adult), remove conditionality requirements (the "claimant commitment"), and extend eligibility to all adults regardless of savings, partner income, or housing status. This converts Universal Credit into a de facto UBI that can serve as the income floor during the transition.

R3. **Establish an AI Safety and Alignment Mandate.** Require all UK-operating AI companies above a capability threshold to conduct and publish safety evaluations, participate in government stress-testing programmes, and contribute to a collective AI safety research fund. The UK AI Safety Institute should be given statutory authority and significantly expanded funding.

2.9 Medium-Term Institutional Adaptation (1–5 years)

R4. **Implement an AI Productivity Levy.** Introduce a tax on the economic output of automated systems, initially set at 2% and rising to 5% as automation deepens. Revenue is hypothecated to the National AI Productivity Fund, which finances UBI, public services, and transition investment. The levy is designed to be progressive:

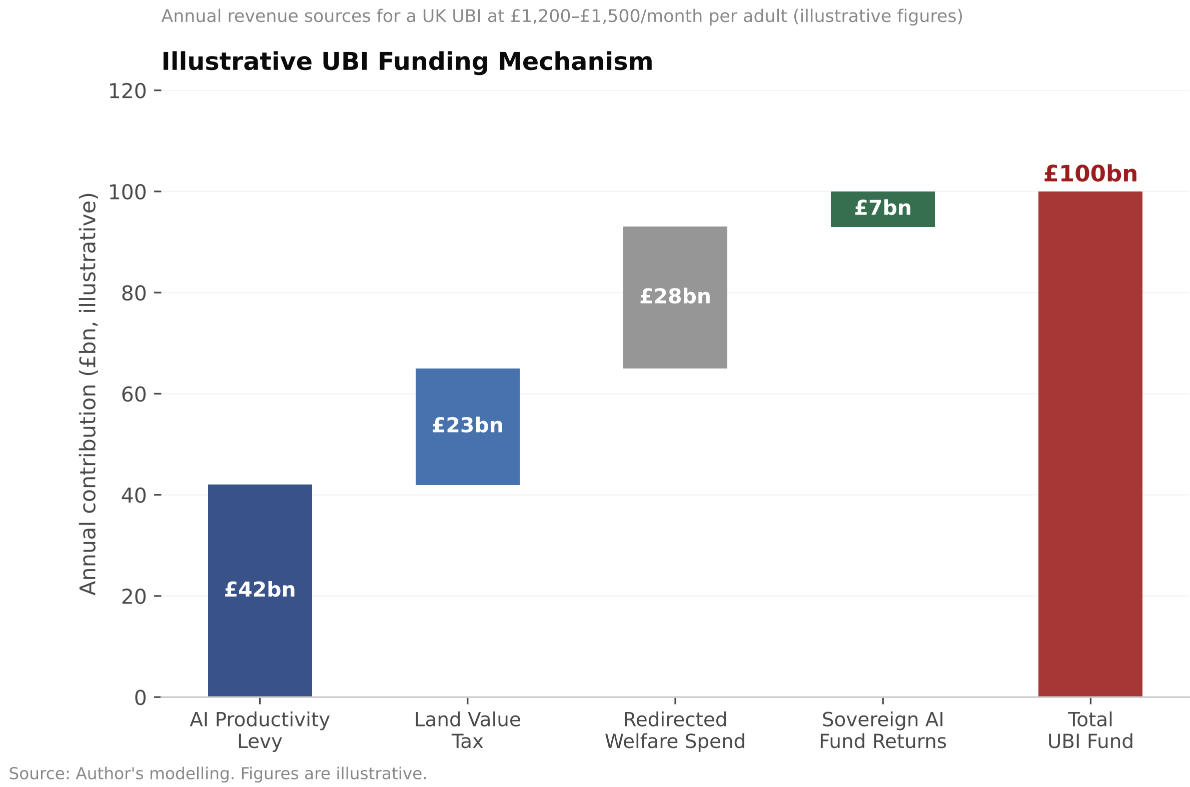

Figure 5: Illustrative UBI funding mechanism for the United Kingdom (figures are illustrative).

systems that displace more labour pay more, creating an incentive for gradual rather than abrupt transition.

R5. Launch a National Automated Housing Programme. Commission the construction of 500,000 AI-designed, robotically built homes in low-cost areas of the UK, offered at cost (£10,000–£20,000) or on long-term nominal leases. This demonstrates the feasibility of automated construction, provides affordable housing, and creates a visible example of the Great Deflation's benefits.

R6. Invest in AI-delivered public services. Deploy AI systems across the NHS, education system, and local government to reduce costs and improve quality. AI triage, AI diagnostics, AI tutoring, AI administrative processing—all should be deployed at scale under professional human oversight, with the explicit objective of demonstrating that AI can improve public services while reducing costs.

R7. Create a Data Dividend. Establish a legal framework recognising public data (NHS records, government databases, publicly funded research) as a national asset. License this data to AI companies on commercial terms, with revenues flowing to the National AI Productivity Fund. This ensures that the public captures a share of the value created by AI systems trained on public data.

2.10 Long-Term Post-Scarcity Governance (5+ years)

R8. **Transition to full UBI at flourishing level.** As AI productivity gains mature and the National AI Productivity Fund grows, increase UBI to a "flourishing level" of £18,000–£24,000 per annum (in 2025 prices), sufficient not merely for subsistence but for a dignified life including housing, food, healthcare, education, culture, and leisure.

R9. **Establish a Constitutional Right to the AI Dividend.** Enshrine in law the principle that the economic gains from AI and automation belong to the public and must be distributed equitably. This prevents future governments from dismantling UBI or diverting AI productivity gains to other purposes.

R10. **Lead the creation of a Global Post-Scarcity Treaty.** Initiate international negotiations toward a multilateral framework for managing the global implications of the Great Deflation, including technology transfer to developing nations, coordination on AI safety standards, and mechanisms for sharing the benefits of AI across the global population. The UK, with its diplomatic infrastructure, multilateral tradition, and early institutional preparation, is well-positioned to lead this effort.

Table 1: Policy Recommendations Summary

Ref	Horizon	Recommendation
R1	Immediate	AI Transition Authority
R2	Immediate	Expand Universal Credit to transition-ready UBI
R3	Immediate	AI Safety and Alignment Mandate
R4	Medium-term	AI Productivity Levy (2–5%)
R5	Medium-term	National Automated Housing Programme (500k homes)
R6	Medium-term	AI-delivered public services at scale
R7	Medium-term	Data Dividend framework
R8	Long-term	Full UBI at flourishing level (£18,000–24,000)
R9	Long-term	Constitutional Right to the AI Dividend
R10	Long-term	Global Post-Scarcity Treaty

3 Discussion: The Shape of the Post-Scarcity World

If our analysis is broadly correct, the world that emerges on the far side of the Great Deflation will be fundamentally different from the one we inhabit today. It is worth sketching, however speculatively, some features of this world.

3.1 Healthcare: The NHS as a Model for Post-Scarcity Medicine

Of all the sectors transformed by the Great Deflation, healthcare may be the most consequential for human welfare. The convergence of AI diagnostics, robotic surgery, and autonomous drug discovery does not merely reduce costs—it fundamentally eliminates the scarcity that has historically rationed access to high-quality medical care. When the marginal cost of a specialist consultation approaches zero, the question is no longer *who can afford* world-class healthcare but *how quickly* institutions can scale delivery to universal coverage.

The United Kingdom's National Health Service is exceptionally well-positioned to capture this transformation. Unlike the fragmented, insurance-mediated systems of the United States or the mixed public-private architectures of many European peers, the NHS is a single integrated delivery network with a mandate to provide universal, free-at-the-point-of-use care. This architecture—so often criticised for its rigidity—becomes a profound advantage when the underlying cost structure collapses.

Consider the trajectory. AI diagnostic systems already match or exceed consultant-level accuracy across radiology, dermatology, ophthalmology, and pathology. At the point of AGI deployment, a patient presenting at any NHS facility—from a rural GP surgery to a major teaching hospital—can receive instantaneous, definitive diagnostic assessment equivalent to the best specialist opinion in the world. Robotic surgical systems, directed by AI, will perform complex procedures with precision and consistency beyond human capability, at a fraction of current cost. AI-driven drug discovery will compress the development timeline for novel therapeutics from a decade to months, and eliminate much of the £2.5 billion average cost of bringing a new medicine to market.

The implications are revolutionary. The NHS, properly equipped with these tools, can provide every UK citizen and resident with access to the highest standard of medical care in human history—not as a distant aspiration, but as a near-term operational reality. The two most significant challenges currently facing the system—diagnostic bottlenecks and surgical waiting lists—are precisely the problems that AI-robotics systems are engineered to solve.

This is not merely a healthcare story. It is a story about what universal institutions can achieve when the cost constraints that have historically limited them are removed. The NHS becomes, in the post-scarcity era, not a rationing mechanism but an abundance distributor: a system that delivers the best of human knowledge and capability to everyone, regardless of income, location, or social circumstance.

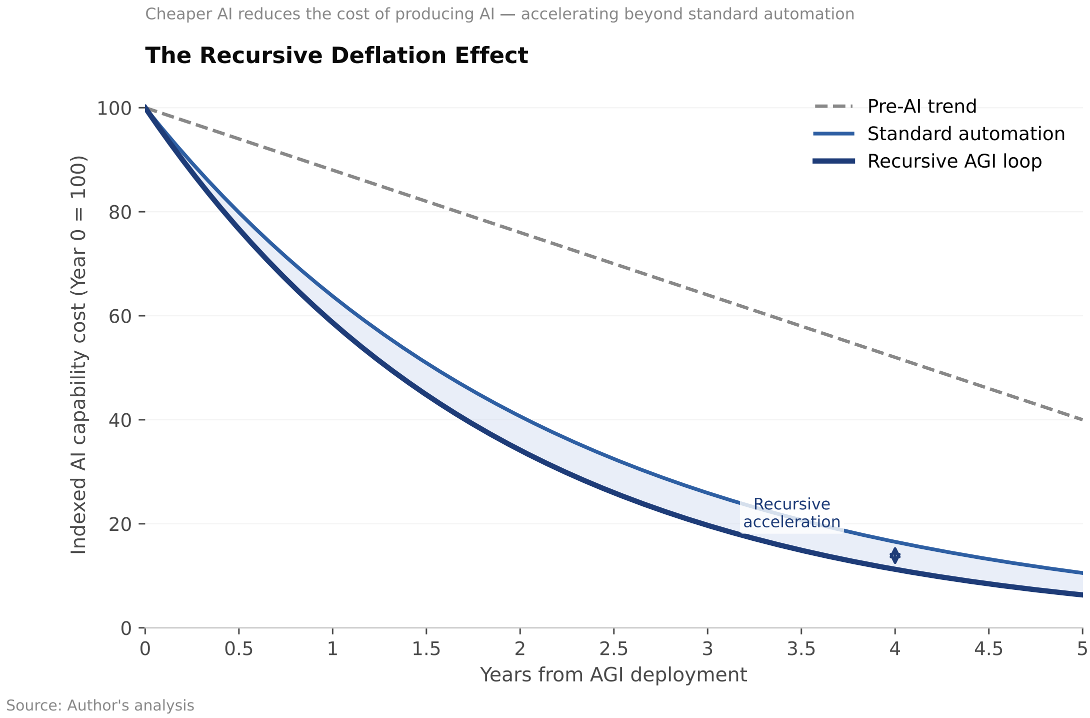

Figure 6: The recursive deflation effect: cheaper AI reduces the cost of producing AI itself, creating an accelerating curve that diverges from standard automation over time.

3.2 The End of Poverty as a Material Condition

In a world where the basic necessities of life—shelter, food, healthcare, education, energy, transport—can be provided at near-zero cost, absolute material poverty ceases to exist as a technical constraint. It becomes purely a matter of distribution and governance. If a well-designed house costs £10,000, a year's food costs £1,000, and healthcare and education are AI-delivered at near-zero marginal cost, then providing a materially comfortable life for every human being on earth is a trivially affordable proposition.

This does not mean that poverty will automatically disappear. Distributional failures, governance failures, and political failures can perpetuate deprivation even in a world of material abundance—as they do today in wealthy nations where homelessness coexists with empty homes. But the *technical constraint* of scarcity, which has been the fundamental driver of economic deprivation throughout human history, will have been removed.

3.3 Inequality in Post-Scarcity

Paradoxically, the elimination of material scarcity may not eliminate inequality. If the benefits of the Great Deflation are captured primarily by AI owners, we could see a world of extreme inequality: a small class of trillionaires who own the automated economy, and a large class of people dependent on whatever UBI or public services governments choose to provide.

Alternatively, if the benefits are distributed broadly through effective policy, we could see a world of unprecedented equality in material terms, combined with new forms of inequality based on social status, creative achievement, or access to non-automatable experiences (beautiful landscapes, intimate human relationships, unique cultural events).

The outcome depends almost entirely on political choices made in the next decade. This is why institutional preparation is so urgent.

3.4 The Transformation of the State

The Great Deflation transforms the fundamental functions of the state. In a world of material abundance, the state's traditional role as regulator of scarcity—allocating resources, managing distribution, adjudicating competing claims—becomes less important. New functions emerge: ensuring equitable access to the AI dividend, maintaining social cohesion in a post-work society, providing structures of meaning and community, managing the relationship between humans and AI systems, and governing the global commons.

The fiscal basis of the state also changes. Income tax and payroll tax—which together account for the majority of UK tax revenue—become less relevant as employment declines. New revenue sources (AI productivity levies, data dividends, resource rents) must replace them. The state becomes, in effect, a steward of the AI commons, redistributing the productivity gains of automated systems to the population.

3.5 Second-Order Economic Effects

The Great Deflation will generate numerous second-order effects that are difficult to predict in detail but important to consider.

The velocity of money. If goods and services are nearly free, the volume of monetary transactions may decline dramatically. The entire monetary and financial system—designed to facilitate exchange in a world of scarcity—may need fundamental redesign. Central banks will need new tools and new mandates.

Entrepreneurship and innovation. In a post-scarcity economy, the barriers to entrepreneurship collapse. If starting a business requires almost no capital (because AI provides the labour, and materials and energy are nearly free), the rate of experimentation and innovation may increase dramatically. We may see an explosion of small-scale, passion-driven enterprises that would be economically unviable in today's cost structure.

Migration patterns. When employment is no longer the primary determinant of where people live, migration patterns will shift dramatically. People will move toward amenity, community, climate, and culture rather than employment. This has implications for urban planning, regional development, and immigration policy.

Geopolitical power. National power has historically derived from economic and military strength, both of which are rooted in productive capacity. In a world where AI equalises productive capacity across nations, the sources of national power shift toward governance quality, institutional capacity, natural amenity, and cultural attractiveness. Small, well-governed nations may gain relative power at the expense of large, poorly governed ones.

The premium on authenticity. In a world awash with AI-generated content, goods, and services, *authenticity*—things made by humans, for human reasons, with human imperfections—may become the ultimate luxury. Handmade goods, live performance, face-to-face human interaction, and "real" experiences may command premiums precisely because they are unnecessary, inefficient, and irreplaceably human.

4 Conclusion

The Great Deflation is not a prophecy; it is a projection. The technological trends driving it are visible today. The economic logic is straightforward. The recursive dynamics are already operating in sectors like energy and computation. What remains uncertain is the speed, the sequencing, and—most critically—the institutional response.

Our central argument is that the convergence of artificial general intelligence and autonomous robotics will drive the marginal cost of most goods and services toward zero within the next one to two decades. This supply-side structural deflation will raise real living standards to levels that would have seemed utopian a generation ago. But the transition from the current cost structure to the post-scarcity equilibrium involves a dangerous window during which labour displacement precedes price benefits, creating acute political, social, and psychological risks.

The policy imperative is preparation. Nations that build the institutional infrastructure for the post-scarcity transition—universal basic income, AI-delivered public services, new models of community and meaning—will navigate the transition successfully. Nations that cling to the institutions and ideologies of scarcity economics will suffer unnecessary hardship during the transition and arrive at the post-scarcity equilibrium in a weakened and divided state.

The United Kingdom, we argue, is well-positioned to lead this transition. The NHS, Universal Credit, a strong AI research ecosystem, a tradition of pragmatic institutional adaptation, and a diplomatic infrastructure capable of leading multilateral negotiations all provide advantages. But these advantages are contingent, not guaranteed. They require active cultivation through the policy framework we have proposed.

The Great Deflation will be the most significant economic transformation since the Industrial Revolution—indeed, it may be the completion of the process that the Industrial Revolution began. The mechanisation of physical labour in the eighteenth and nineteenth centuries transformed human civilisation beyond recognition. The mechanisation of *cognitive* labour in the twenty-first century will do the same. The question is not whether this transformation will occur, but whether we will manage it wisely.

The stakes could not be higher. If we get the transition right, we will have created a world in which every human being has access to the material conditions for a flourishing life—and the freedom to pursue meaning, connection, and purpose on their own terms. If we get it wrong, we risk a period of immense suffering, social breakdown, and political extremism that could set back human progress by decades.

The time for institutional preparation is now. The Great Deflation is coming. The only question is whether we will be ready.

This paper represents the views of the author and does not necessarily reflect the institutional position of Cassandra Policy. Comments and correspondence should be directed to `nourzai@cassandrapolicy.org`.

Suggested citation: Nourzai, E. (2026). "The Great Deflation: Artificial General Intelligence, Autonomous Robotics, and the Political Economy of the Post-Scarcity Transition." Cassandra Policy. London: Cassandra Policy.

References

[1] Acemoglu, D. and Restrepo, P. (2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2), pp. 3–30.

[2] Acemoglu, D. and Restrepo, P. (2020). "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand." *Cambridge Journal of Regions, Economy and Society*, 13(1), pp. 25–35.

[3] Autor, D.H., Levy, F. and Murnane, R.J. (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics*, 118(4), pp. 1279–1333.

[4] Baumol, W.J. (2012). *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't*. New Haven: Yale University Press.

[5] Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton.

[6] Coase, R.H. (1937). "The Nature of the Firm." *Economica*, 4(16), pp. 386–405.

[7] Frey, C.B. and Osborne, M.A. (2017). "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change*, 114, pp. 254–280.

[8] Friedman, M. and Schwartz, A.J. (1963). *A Monetary History of the United States, 1867–1960*. Princeton: Princeton University Press.

[9] Kangas, O., Jauhiainen, S., Simanainen, M. and Ylikännö, M. (2019). *The Basic Income Experiment 2017–2018 in Finland: Preliminary Results*. Helsinki: Ministry of Social Affairs and Health.

[10] Keynes, J.M. (1930). "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*. London: Macmillan.

[11] Mazzucato, M. (2018). *The Value of Everything: Making and Taking in the Global Economy*. London: Allen Lane.

[12] Nordhaus, W.D. (2021). "Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth." *American Economic Journal: Macroeconomics*, 13(1), pp. 299–332.

[13] Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. New York: Palgrave Macmillan.

[14] Susskind, D. (2020). *A World Without Work: Technology, Automation and How We Should Respond*. London: Allen Lane.

[15] Wang, M. and Shi, J. (2014). "Psychological Research on Retirement." *Annual Review of Psychology*, 65, pp. 209–233.

[16] Weber, M. (1905/2001). *The Protestant Ethic and the Spirit of Capitalism*. Trans. T. Parsons. London: Routledge.

[17] Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.