




Job-Less Utopia

Macroeconomics
in the Age of AGI

Marcus Hutter

Job-Less Utopia

Macroeconomics in the Age of AGI

Marcus Hutter

2026

Job-Less Utopia: Macroeconomics in the Age of AGI

Copyright © 2026 Marcus Hutter. Licensed under [CC BY-NC 4.0](#)

ISBN: 978-1-0668284-1-8

First edition, July 2026.

Publisher: AIXI Media

Typeset in Latin Modern using L^AT_EX.

Dedicated to

나의 사랑하는 아내 지우

Preface

Will automation destroy livelihoods or create prosperity? The answer depends entirely on policy. If the wealth that machines produce is shared, the result is liberation; if it is not, the result is a preventable dystopia of concentrated wealth and mass precarity. As an AI researcher who has spent two decades on the theoretical foundations of artificial general intelligence, I wrote this book because the mainstream debate still treats job loss as synonymous with economic loss, when in fact it is only the *distribution* of output that is at stake. This book assembles the full case in one place.

The advent of Artificial General Intelligence (AGI) and advanced robotics will eventually permanently decouple economic productivity from human labor. However, rather than a crisis, mass un(der)employment should be embraced as a transition to a [post-labor economy](#) where universal material prosperity is achieved and sustained through redistributive policies in the form of some [Guaranteed Income](#) (GI) such as Universal Basic Income ([UBI](#)) or a [Citizens' Wealth Fund](#) (CWF).

Theses

The core argument is simple: when a robot or an AI replaces a human worker, the volume of goods and services does not shrink; it typically grows. The salary the company saves can, in principle, be taxed and redistributed to the displaced worker as a basic income. That is, redistribution is cost-neutral at the point of replacement. The macroscopic result: output is maintained, and nobody is worse off. In practice, the transition is messier, but the fundamental arithmetic holds. Forcing full employment instead risks a proliferation of socially wasteful (“BS”) jobs; the book proposes a new tax to discourage them.

All pieces of this argument have been made before, by economists, political theorists, philosophers, and technologists. This book is a synthesis: a single, integrated treatment that follows the thread from economic foundations

through welfare design, democratic stability, the question of human meaning, and international coordination. It develops the case along thirteen theses:

1. **Total automation is inevitable:** AGI and advanced robots will eventually perform virtually all cognitive and physical tasks better and cheaper than humans (Chapter 2).
2. **Material abundance is maintained:** Replacing human workers with AI does not shrink the economy; the exact same (or a greater) volume of goods and services will still be produced (Chapter 3).
3. **Near-zero marginal costs:** All supply chain costs ultimately trace back to human labor, raw materials, and energy. Automating labor and harnessing abundant energy will drive the production cost of most non-scarce goods and services remarkably close to zero (Chapter 3).
4. **The land bottleneck is solvable:** Even at zero labor cost, land and scarce resources remain priced. Proven tools (land value taxation, sovereign resource funds, [Harberger](#) taxes) can capture these rents and channel them into [GI](#) (Chapter 3).
5. **The ‘New Jobs’ fallacy:** The historical pattern of technology creating new human jobs will eventually end. The transition will be too fast for human retraining, and new roles will also be performed better by machines (Chapter 4).
6. **Meaning exists beyond employment:** The psychological need for traditional employment is a product of modern cultural conditioning. Humans are naturally primed to find purpose in family, community, hobbies, learning, and structured competition (Chapter 7).
7. **Mass unemployment is a feature, not a bug:** Un(der)employment is perfectly consistent with universal prosperity. The primary goal of an economy is to provide a decent standard of living, not to maximize human toil (Chapter 2).
8. **Forced employment is dystopian:** Attempting to artificially maintain full employment in a highly automated society will trap humans in low-paying hyper-competition and socially wasteful (BS) jobs (marketing arms races, predatory finance, ...) (Chapter 4).
9. **UBI is economically affordable:** Multiple complementary instruments can finance [GI](#), either directly via automation, luxury, Pigouvian, land, inheritance, consumption, and service taxes, or indirectly via wealth funds, even cost-neutrally (Chapter 5).
10. **Alternative safety nets are inferior:** Conditional welfare, job guarantees, and [means-tested](#) transfers re-introduce surveillance, [conditional-ity](#), or BS-jobs that UBI avoids. Universal Basic Services can augment but not replace cash transfers. Blocking automation via regulation only halts progress (Chapter 6).

11. **Democracy is self-correcting:** Failing to redistribute the wealth generated by AI risks severe inequality (Figures 5.3 and 7.5) and social unrest. In a functioning democracy, the displaced majority will vote for parties that enact redistribution; at the limit, the threat of civil unrest compels even reluctant capital owners to share the gains of automation (Chapter 7).
12. **International challenges are solvable:** Global asymmetries and tax base erosion can be managed through tariffs, destination-based consumption taxes, digital services taxes, a global minimum tax, and international coordination (Chapter 8).
13. **Coordination failures become easier to address:** By removing the desperation for employment, AGI paired with UBI deflates the pressure behind BS-jobs, marketing arms races, and exploitative gig work, creating political space for complementary institutional reforms (Chapter 10).

How to read this book

The book is designed to be read cover to cover, with each chapter building on the preceding ones. Cross-references link related arguments across chapters, and every chapter opens with a motivating story. That said, most chapters are sufficiently self-contained to be read independently once the Introduction (Chapter 1) has established the basic vocabulary and the four-camp taxonomy. Readers already familiar with the automation debate may skip directly to the chapters most relevant to their interests. Throughout the text, terms printed in blue are explained in the [Glossary](#) (back matter) and hyperlinked to it in the PDF version. A brief [Key Terms](#) primer appears before the first chapter.

- **Readers with interest in economics** will find the technical core in Chapter 3 (zero-cost convergence, the land bottleneck, cost-to-price dynamics) and Chapter 5 (taxation instruments, wealth redistribution).
- **Policy practitioners** may start with Chapter 6 (UBI financing, the transition roadmap, wealth funds) and Chapter 8 (tariffs, digital services taxes, global coordination).
- **General readers** concerned about the future of work will find Chapter 4 (why new jobs will not compensate) and Chapter 7 (meaning, democracy, and the fear of idleness) the most accessible entry points.

Each chapter ends with a *Key Insights* box that distills the main takeaways and some *Open Problems*. The Appendices provide comprehensive cata-

logues of arguments for UBI (Appendix A.1) and against preserving jobs (Appendix A.2).

The book’s territory is deliberately circumscribed. Figure 1 maps the core topics (automation economics, job displacement, redistribution, post-labor society, and international challenges) against adjacent fields that are briefly touched but not treated in depth (see Chapter 9 for these wider considerations).




Figure 1: Scope boundary map. The book’s core territory (green) covers automation economics, job displacement, redistribution via UBI, post-labor society, and international challenges. Adjacent topics (blue) are discussed in Chapter 9 but not treated in depth. Gray topics are explicitly excluded.

Acknowledgements

I am grateful to a number of “colleagues” who helped bring this book into existence: Claude served as a research assistant, tracking down references and stress-testing my (mostly not so) novel ideas that form the book’s backbone, and to assist interactively with generating figures and tables. Gemini conducted broad literature searches and critically proofread multiple drafts. Claude also served as the most diligent editor I could have hoped for: While I have a reasonable grasp of economic concepts, I am neither a professional

economist nor a native English speaker, so it was invaluable to have someone polish my clumsy English into professional language. No text was taken verbatim without my critical review, revision, and approval. Gemini also provided indispensable logistical support throughout the project, and Nano Banana the book cover. I am also grateful to my humanoid colleagues, in particular Joseph Levine, Alex Imas, and Adam Bales, for feedback on early drafts of the manuscript. The structure, design, choices of what (not) to include, the novel ideas, and the original writing of all drafts are mine.

London, UK, July 2026

Marcus Hutter

Table of Contents

Front Matter	i
Title, Copyright, Dedication	i
Preface	v
Table of Contents	xi
List of Figures and Tables	xv
Key Terms	xix
1 Introduction	1
2 Foundations of a Post-Labor Economy	9
2.1 Four Camps	10
2.2 What Is a 'Job'?	11
2.3 The Industrial AI Revolution	15
2.4 Goods and Services vs Jobs	18
2.5 The Core Argument	20
2.6 What Is Guaranteed Income (GI)?	23
3 The Economics of Automation	29
3.1 Next-Generation AIs and Robots	30
3.2 Idealized Transition Scenario	38
3.3 Zero Marginal Cost Paradigm	42
3.4 Self-Replicating Factories and Resource Limits	43
3.5 Remaining Cost Components	46
3.6 The Land Ownership Problem	47
3.7 The Intellectual Property Problem	48
3.8 From Cost to Price	51
3.9 Even Utopia Requires Money	53
3.10 Skill Shortage and Job Mismatch	54
4 The Illusion of Full Employment	59
4.1 New Jobs Will Be Created	60

4.2	What “Jobs” Will Remain?	66
4.3	A Taxonomy of Post-ASI Human Activities	68
4.4	Why Remaining Jobs Will Erode	73
4.5	BS-Jobs	75
5	Redistribution Mechanisms	87
5.1	Labor Redistribution	88
5.2	Wealth Redistribution	89
5.3	Taxation Strategies	90
6	Guaranteed Income	105
6.1	Universal Basic Income (UBI)	106
6.2	Transition from CI to UAI to UHI	109
6.3	Objections to UBI/GI	114
6.4	Alternatives to UBI/GI	119
6.5	Universal Citizen Share Ownership Models	124
6.6	Pension Fund Pathway	133
7	Societal Transitions and The Future	141
7.1	Political Economy of Redistribution	142
7.2	Will Democracy Prevail in a Post-Labor Society?	145
7.3	Democratic Fragility and Techno-Feudalism	150
7.4	Democratic Innovations for the Post-Labor Era	152
7.5	Social Order and Meaning of Life	155
8	International Challenges	163
8.1	Trade Imbalance and AI Dependency	164
8.2	Tariffs and Currency Adjustment	166
8.3	Tax Design for a Borderless AI Economy	168
8.4	Capital and Talent Flight	169
8.5	AI Colonialism and the Global Development Gap	171
8.6	The Double-Edged Sword of Global Coordination	175
9	Wider Considerations	179
9.1	AI Alignment and Existential Risk	180
9.2	Technical Architectures of AGI and Robotics	182
9.3	Economic Growth Theory and AGI	185
9.4	Machine Consciousness, Robot Rights, and AI Ethics	188
9.5	A Survey of Economic Schools and AGI	190
9.6	Political Feasibility in Specific Polities	193
9.7	Environmental Sustainability and Climate	196
9.8	Topics Partially Addressed in the Main Text	199

10 Discussion	203
10.1 Abundance of Goods and Services	204
10.2 Timeline to AGI and Societal Diffusion	204
10.3 The Transition Is Achievable	204
10.4 From Theory to Practice	206
10.5 Coordination Problems Persist	207
10.6 Conclusion	208
 A Appendices	 211
A.1 Full List of Arguments for UBI	211
A.2 Full List of Arguments Why Jobs Will Disappear	217
 Back Matter	 223
Glossary	223
Bibliography	229
Index	247
Blurb	253
About the Author	253

List of Figures and Tables

1	Scope boundary map	viii
2.1	Four positions on AGI and jobs	10
2.2	Properties of a ‘job’ and eight edge cases	13
2.3	The job–leisure spectrum	14
2.4	The Great Decoupling: productivity vs. compensation	16
2.5	Declining labor share of US national income, 1950–2023	16
2.6	AI/automation displacement forecasts	18
2.7	Global private investment in AI, 2013–2025	18
2.8	AI Revolution vs. historical industrial revolutions	19
2.9	US time allocation, 1840–2100	19
2.10	Four views on AGI and labor	22
2.11	Top 10% income share in the US, 1913–2022	22
2.12	Terminology: building-block letters and acronyms	24
3.1	AI researcher forecasts: when AI outperforms humans	31
3.2	Technology capability levels	33
3.3	Minimum technology level for job replacement by category	35
3.4	Idealized transition: salary savings redistributed as CI	39
3.5	Moore’s Law for Everything: cost decline of automated goods	43
3.6	Zero Marginal Cost Supply Chain	44
3.7	Exponential cost declines (log scale)	46
3.8	Three major rent categories in a post-AGI economy	49
3.9	Moravec’s Paradox Quadrant	55
4.1	Historical transitional jobs and AI-era analogs	61
4.2	Acemoglu–Restrepo Task Model	62
4.3	The Leontief horse analogy: US horse population, 1850–1960	64
4.4	Bounded S-Curve of AI Adoption	65
4.5	S-Curve with Multiple Constraints	65
4.6	Working hours per worker, 1870–2023	65

4.7	Post-ASI human activities grouped by survival mechanism .	69
4.8	Post-ASI human activities: composite summary	70
4.9	Taxonomy of BS-jobs	77
4.10	Red Queen Game in Marketing	78
4.11	Global annual spending comparison (2024)	79
4.12	The Annoyance Economy	80
4.13	Workers perceiving their job as ‘socially useless’	81
5.1	Labor redistribution mechanisms in a post-AGI economy . .	89
5.2	Wealth redistribution mechanisms in a post-AGI economy .	90
5.3	US household income inequality (Gini index), 1967–2022 . . .	91
5.4	Government budget: revenue is fungible	92
5.5	Taxation instruments for funding redistribution	94
5.6	Tax instrument radar comparison	95
5.7	Taxation sources shift under automation	96
5.8	Tax revenue composition shift	96
5.9	Top 1% income share (1900–2024), selected countries	102
5.10	Top 1% income share by world region (1900–2024)	103
6.1	Major UBI and GI pilot programs worldwide	110
6.2	UBI pilot world map	111
6.3	UBI transition timeline	114
6.4	Transition path from labor economy to post-labor economy	115
6.5	GI level vs. job market equilibrium	116
6.6	Income-support approaches and welfare institutions compared	127
6.7	American Equity Fund mechanics	130
6.8	Four AI wealth-redistribution models compared	131
6.9	Mandatory pension systems as CWF infrastructure.	134
6.10	Birth endowment proposals compared.	135
6.11	Pension pathway transition phases, synchronized with CI → UAI → UHI (Section 6.2).	136
6.12	Novelty assessment: what exists in the literature vs. what is new (see also Tables 6.9 to 6.11).	138
7.1	Democracy feedback loop	146
7.2	Oil-funded vs. AI-funded rentierism	147
7.3	Shifting power dynamics under AGI wealth concentration .	149
7.4	The resource curse: resource rents vs. democracy	149
7.5	Top 1% income share by country (2024)	152
7.6	Power-distribution mechanisms for the AI age	156
7.7	Sources of life meaning (17-country survey)	159

8.1	International trade imbalance from asymmetric AI adoption	164
8.2	AI Preparedness Index by world region	165
8.3	Immigration strategies under UBI	172
8.4	International policy responses to AI-driven challenges	176
10.1	Stylized AGI transition scenarios	205
10.2	AI capability vs. job displacement lag	205

Key Terms

A brief overview of recurring acronyms. Terms are hyperlinked to the comprehensive [Glossary](#) in the back matter.

Term	Definition
AGI	Artificial General Intelligence: AI matching human cognitive abilities across domains
ASI	Artificial Superintelligence: AI far surpassing the best human capabilities
UBI	Universal Basic Income: unconditional periodic cash transfer to all citizens
GI	Guaranteed Income: umbrella term covering UBI and near-UBI schemes
CI	Conditional (targeted) Income support
UAI	Universal Adequate Income: unconditional, covering basic needs
UHI	Universal High Income: unconditional, enabling comfortable living
NIT	Negative Income Tax: transfers to individuals below an income threshold
UBS	Universal Basic Services: free provision of essentials (health, education, housing) as alternative to cash
BS-job	Bad-for-Society job: employment that is socially wasteful or pointless
CWF	Citizens' Wealth Fund: publicly owned fund distributing dividends to all citizens

Meridian — 2025

Giwoo has been a surgical nurse at Meridian General Hospital for twelve years. On her days off she looks after Grandma Angela, Marc's mother, whose arthritis has confined her mostly to the house. Giwoo's husband Marc runs a materials-science lab at the university, supervising three postdocs and an aging mass spectrometer. Their son Levin is in his second year as a patrol officer; daughter Chaeyeong just started her economics degree; and Alice, the youngest, is finishing a marine-biology master's, spending weekends knee-deep in estuary sediment. At the county courthouse, Judge Chloe is reviewing sentencing guidelines for the hundredth time this year. And Grandpa Butz (everyone calls him that; his real name is Reinhard), eighty-five and long retired from mechanical engineering, sits in his garden reading the paper and muttering that nobody builds things properly anymore. The hospital has just installed an AI triage assistant. "It flags the urgent cases," the chief nurse says. "You still make the decisions." Marc's department has received a grant to benchmark an AI that designs novel alloys faster than any human team. At Sunday dinner, Butz folds his newspaper: "I have seen tools come and go. The question is always who keeps the savings." Nobody has an answer yet.

Chapter 1

Introduction

‘horses did not get new jobs when tractors arrived, they got retired’ ‘humans are not horses; they are apes’ — paraphrasing Leontief (1983a) and Vollrath (2015)

Artificial general intelligence and advanced robotics are poised to automate all economically relevant tasks within a single generation. Unlike every previous industrial revolution, the emerging one replicates the very cognitive flexibility that has so far made human workers indispensable: the historical pattern of [creative destruction](#) may, for the first time, fail to generate compensating employment. This book argues that the correct response is not to mourn the loss of jobs but to ensure that the undiminished (indeed, growing) output reaches everyone. Under adequate redistribution, mass un(der)employment becomes a liberation, not a catastrophe. The introduction motivates this thesis, situates it among four intellectual camps, outlines what is novel in the present treatment (a systematic architecture spanning economics, welfare design, democratic stability, meaning, and geopolitics), and provides a chapter-by-chapter road map.

History vs future. Global investment in artificial intelligence is expected to exceed \$500 billion in 2026 and is accelerating (Figure 2.7). Studies estimate that 30–50% of current occupations face substantial automation risk within two decades (Frey and Osborne, 2017; Acemoglu and Restrepo, 2020). These are not science-fiction scenarios; they are mainstream research estimates. Every previous industrial revolution automated a narrow class of tasks and, after a painful transition, generated new occupations to absorb the displaced workforce (Acemoglu and Restrepo, 2019; Autor et al., 2024). The emerging revolution in artificial general intelligence (AGI) and advanced robotics is qualitatively different: it threatens to automate *all* economically relevant tasks, cognitive and physical alike, within a single human generation (Grace et al., 2025; Genewein et al., 2026). For the first time, the historical pattern of [creative destruction](#) may not produce a compensating wave of new

human employment, because the very cognitive flexibility that previously made humans indispensable is now being replicated in silicon (cf. [Acemoglu and Restrepo, 2019](#); [Susskind, 2020](#)). The question, then, is not whether displacement will occur but how society should respond.

Four camps. The dominant reactions to this prospect fall into three camps: optimists who trust the historical pattern to hold (the Continuity Thesis), pessimists who equate job loss with societal collapse (the Displacement Crisis), and skeptics who doubt AGI will arrive at all (the Stagnation Wager). All three share a common framing limitation: they fixate on *jobs* rather than on *goods and services*. This book advances a fourth position, **Post-Labor Prosperity**: if machines produce the same or greater output, the material wealth of society is undiminished; the real challenge is distribution, not production. Under adequate redistribution, mass un(der)employment is not a catastrophe but a liberation, a transition to universal material prosperity in which work becomes voluntary.

Three pillars. The argument rests on three pillars, treated jointly because none is sufficient alone. First, *AI and robotics* are approaching the threshold where they replicate general human cognition and dexterity, not merely narrow tasks. Second, the loss of *jobs* feels threatening because jobs have traditionally been the vehicle for livelihood, identity, and social standing. Yet the concept is murkier than it appears: its defining properties (wages, obligation, regularity, replaceability) degrade under automation, so that jobs vanish while *work* persists. Getting the definition right is essential for a precise discussion of what exactly is lost and what is not. Third, *Guaranteed Income*, the book's funding-agnostic umbrella for Universal Basic Income (UBI) and related transfers, channels the surplus from automation back to displaced workers. AI without redistribution yields dystopia; redistribution without automation is unaffordable; only their convergence delivers broadly shared prosperity.

What's new? The individual arguments are well established; the core thesis is the standard post-laborist position articulated at book length by [Susskind \(2020\)](#), championed politically by [Srnicek and Williams \(2015\)](#), and formalized by [Korinek and Juelfs \(2022\)](#). Nearly all individual ideas expressed in this book have been presented elsewhere in some form or context. What is novel is the *systematic architecture*: the first attempt, to my knowledge, to assemble the full case (economics, full-employment critique, UBI mechanics, democratic stability, meaning, and international coordination) into a single, mutually reinforcing argument. Specifically, the book contributes:

1. a micro-level [fiscal feasibility](#) proof that, upon each individual dis-

placement, the salary saved by the firm suffices to fund the displaced worker's lost income, leaving every party no worse off (Section 3.2 and Figure 3.4);

2. a Wittgensteinian property-cluster definition of 'job' (11 properties $\times$ 8 edge cases, Table 2.2) and a formal job-leisure spectrum showing how UBI degrades job properties (Table 2.3);
3. a four-camp taxonomy of positions on automation (Continuity / Displacement / Stagnation / Post-Labor Prosperity, Table 2.1);
4. the macro-thesis that artificially forcing full employment under AGI traps society in BS-jobs (Section 4.5), with an extended taxonomy and a BS-tax umbrella (Section 5.3);
5. a three-phase UBI transition roadmap with pre-legislated automation triggers (Section 6.2 and Figure 6.3);
6. geopolitical frameworks: AI-vs-oil rentierism (Table 7.2), a USA-Australia trade example scenario (Section 8.1), and a coordination-as-double-edged-sword argument (Section 8.6);
7. extensive illustration with 41 diagrams and 30 tables.

Chapter 2 lays the conceptual groundwork. The debate over AGI and employment splits into four camps, defined by a 2×2 matrix of whether mass job loss occurs and whether broad prosperity is achieved (Table 2.1). Three camps share a framing limitation: they fixate on jobs rather than on goods and services. The chapter develops a property-cluster definition of "job" (11 properties, 8 edge cases, Table 2.2), drawing on Wittgenstein's family-resemblance analysis, that degrades gracefully as UBI replaces wages. It traces the **Great Decoupling** of productivity from wages since the 1970s, surveys current displacement forecasts, and states the core thesis: the proper goal of an economy is maintaining a decent standard of living, not maximizing employment. A five-step logical chain (automation preserves output $\rightarrow$ wealth shifts to capital $\rightarrow$ tax the windfall $\rightarrow$ fund Guaranteed Income $\rightarrow$ universal prosperity with voluntary work) structures the remainder of the book.

Chapter 3 develops the detailed economic logic. A technology-level taxonomy (five cognitive AI levels plus two robotic levels) is mapped onto a job-replacement matrix showing what threshold automates 10%, 50%, 90%, and 99% of each occupation (Table 3.3). An **idealized transition** scenario then proves fiscal feasibility at the individual level (Section 3.2 and Figure 3.4). The chapter then traces why **marginal costs** converge to zero: all costs ultimately reduce to labor, raw materials, and energy; once labor is automated and energy abundant, the cost of most non-scarce goods approaches zero. Two rent categories resist this convergence: land ownership and intellectual

property. Both receive dedicated treatment, including land value taxation and [Harberger](#) taxes as corrective instruments.

Chapter 4 dismantles the illusion of full employment. The comforting mantra that “new jobs will always be created” rests on a pattern that breaks when the displacing technology replicates general human cognition itself. An explicit parallel to the 2008 financial crisis underscores the danger of extrapolating from equilibrium models. While transitional jobs will emerge during the AI era (organized into five archetypes: Operators, Translators, Maintainers, Regulators, Educators), historical transitions lasting 20–80 years will compress to 5–15 years because software scales faster than mechanical infrastructure, and many ostensibly new roles are self-eliminating: they train the very systems that will make them obsolete. The chapter’s most distinctive argument is that even if full employment could be maintained artificially, doing so would trap society in hyper-competitive BS-jobs (Graeber’s bullshit jobs and, more broadly, Bad-for-Society jobs), for which a detailed taxonomy covering predatory, arms-race, and negative-sum employment is developed (Table 4.9). The proper response to automation is not to invent make-work but to let machines do the work and redistribute the output.

Chapter 5 pivots to the core policy question: how to channel automation wealth to citizens. Two complementary strategies are developed. Labor redistribution (shorter workweeks, earlier retirement funded by pension investment in AI enterprises, extended parental leave, and paid re-education) spreads remaining human work across the population, with UBI itself inducing these reductions voluntarily rather than through mandates. Wealth redistribution encompasses four mechanisms (UBI, Universal Basic Capital, Universal Basic Services, Universal Basic Jobs), each compared across pros and cons. A diversified taxation toolkit of 16 instruments is surveyed and compared across eight policy dimensions via a radar chart (Table 5.5). Key instruments include corporate profit and capital taxes (targeting automation rents rather than normal returns), progressive income taxes, land value taxes (uniquely efficient because immovable), and a novel “BS-tax” umbrella that synthesizes Pigouvian, demerit, and dissipative-activity charges on socially wasteful activities. The non-affectation principle establishes that how revenue is raised should be separated from how it is spent.

Chapter 6 argues for the superiority of Universal Basic Income. Among available transfer vehicles, UBI stands out for its simplicity, universality, and compatibility with voluntary work, supported by three independent philosophical clusters: freedom (the republican “power to say no”), fairness (common ownership of cumulative knowledge and nature, Rawlsian correc-

tion of morally arbitrary advantages), and justice (the cruelty of behavioral conditions on welfare in a world where machines do the work). Seven structural design advantages distinguish UBI from means-tested alternatives: no poverty trap, full take-up, administrative simplicity, labor flexibility, entrepreneurship buffer, bargaining power, and no behavioral surveillance. A phased transition from Conditional Income through UBI to Universal High Income tracks the pace of automation (Section 6.2 and Figure 6.3). The standard objections (laziness, unaffordability, inflation) are systematically rebutted, and the commonly proposed alternatives (job guarantees, participation income, basic endowment, share ownership models) each reintroduce frictions UBI avoids.

Chapter 7 addresses the societal dimension. When redistribution fails, consequences range from populism to techno-feudalism: an “intelligence curse” strictly worse than the resource curse because AI rents are more concentrated, enable surveillance without human soldiers, and scale at software speed. An Oil-vs-AI comparison (Table 7.2) shows that AI severs both the fiscal *and* labor link between citizens and the state. Yet democratic feedback loops incentivize corrective action (Figure 7.1): the displaced majority has electoral power to demand redistribution, and democracies have absorbed comparable structural shocks before (the Engels Pause, the New Deal, the Nordic welfare state). Institutional innovations (sortition, constitutional entrenchment of UBI rights, AI-assisted deliberation, quadratic voting) offer structural resistance to elite capture. The existential dimension receives equal treatment: the objection that joblessness destroys meaning conflates jobs with work. Evidence from leisure-oriented societies, retirees, and the taxonomy of unpaid work (domestic, caregiving, creative, civic, educational) shows humans find purpose regardless of economic compulsion.

Chapter 8 tackles the international challenges. Asymmetric AI adoption creates severe trade imbalances, illustrated through a stylized USA–Australia model in which the AI leader extracts massive profits from the follower through high-margin AI and robot exports while receiving only raw materials in return (Section 8.1). Tariffs and currency devaluation offer blunt protection but carry a fundamental trade-off between shielding domestic employment and cutting the country off from AGI’s benefits. The more durable solution is redesigning tax systems for a borderless digital economy: destination-based consumption taxes (inherently robust to profit shifting), interim digital services taxes, and the OECD’s global minimum corporate tax. Capital and talent flight, often invoked as a showstopper for aggressive redistribution, turns out to be more limited than feared: real AI infrastructure (GPU clusters, data centers, energy grids) is highly sticky, even if paper profits and IP

are mobile. Beyond advanced economies, AI threatens to undermine the traditional development ladder itself (AI colonialism), deepening the global development gap for countries reliant on labor-cost advantages.

Chapter 9 maps the boundary between the book's core argument and adjacent topics that deserve attention but lie outside its scope: AI alignment and existential risk, technical AGI architectures, machine consciousness, a survey of competing economic schools, political feasibility in specific polities, and environmental sustainability. For each topic, the chapter explains why the boundary exists and where to find deeper treatments.

Chapter 10 stress-tests the preceding conclusions from five angles. First, AGI will likely trigger radical abundance: formal growth models suggest AI-driven R&D automation can produce hyperbolic growth, making UBI even easier to fund from a vastly larger pie. Second, whether AGI arrives in 5 years or 50, the structural logic is the same; paradoxically, rapid arrival may be safer, because gradual displacement risks a boiling-frog effect in which each incremental job loss dulls political urgency until cumulative damage is severe. Third, the transition is achievable: democracies have absorbed comparable shocks, and the displaced majority retains electoral power to demand redistribution. Fourth, the single greatest obstacle is institutional and cognitive inertia: a population trained to equate work with meaning may require a human generation to overcome. Fifth, coordination failures (BS-jobs, marketing arms races, international tax competition) persist even with AGI and UBI, but UBI removes the pressure that multiplies them: if surplus workers no longer need employment at any cost, the pressure inflating BS-jobs and exploitative gig work vanishes at the source.

Appendix [A.1](#) lists 70 arguments for UBI and Appendix [A.2](#) lists over 50 arguments why jobs should disappear.

Stakes. The arithmetic of redistribution is encouraging, but the politics are not self-executing. Without deliberate policy intervention, the gains from automation will accrue to capital owners while displaced workers bear the costs. The historical record is sobering: wealth concentration has risen sharply since the 1980s (Figures [5.3](#) and [7.5](#)), and the current wave of automation threatens to accelerate this trend. The resulting inequality is not merely unjust; it corrodes democratic institutions, fuels populism, and invites authoritarian responses. Getting the transition right is not optional: it is the difference between broadly shared prosperity and a preventable dystopia.

Background reading. Proposals for basic income and related welfare programs have a long history with a vast literature. We use *Guaranteed Income*

(GI) as a funding-agnostic umbrella for any periodic cash transfer above subsistence level (see Table 2.12 for a precise taxonomy); the reader can think of UBI throughout, which is the most widely discussed variant. Van Parijs and Vanderborght (2017) is the most thorough single reference for the ethical, philosophical, economic, and political case for UBI; Widerquist (2024) offers a concise, updated overview including post-2017 pilot evidence and a chapter on automation. Susskind (2020) provides a book-length treatment of a world without work due to automation, covering the limits of task-based displacement and policy responses in greater detail than we can here. While the present book cites sources in place throughout, especially more recent ones, it naturally cannot replicate the jointly over one thousand references in these three books. Appendices A.1 and A.2 provide compact catalogues of arguments for UBI and against preserving jobs, respectively.

Caveats. The analysis works with idealized scenarios to isolate fundamental economic principles; real-world implementation will involve political and social frictions not modeled here. Nor does the book survey all competing economic schools of thought; its aim is to demonstrate that post-labor prosperity is a realistic and desirable outcome and to argue for the structural changes required to manage a transition that appears increasingly inevitable. Specific dates mentioned for AGI milestones vary across chapters: some passages cite expert-survey medians, others reflect industry roadmaps or scenario analyses, and the fictional vignettes adopt their own timelines. Even expert guesstimates on when AGI and ASI will arrive vary widely; presenting a range of plausible timelines invites the reader to consider the argument under different assumptions rather than anchoring on a single forecast.

Meridian — 2028: First Signs

Three years in, Meridian General has cut its nursing staff from ninety-four to sixty-one. The AI no longer just triages: it monitors vitals, adjusts dosages, and drafts discharge summaries that attendings sign without reading. Giwoo survived the first round but now spends half her shifts supervising robotic IV lines. Marc's lab lost two postdocs when the AI alloy-design system made their experimental work redundant; his own role has shifted from hypothesis to validation. Chaeyeong's economics professors are rewriting their labor-market lectures in real time; one confesses that half the textbook models already look obsolete. Alice's marine-biology program loses its fieldwork funding: "Autonomous sensor networks do it now," the rejection letter explains. Judge Chloe's courtroom sees its first AI-generated sentencing recommendation. She overrides it, mostly on principle. Levin's precinct has deployed predictive-patrol software, but the dispatchers still call him for welfare checks no algorithm can handle: a confused old man, a teenager who needs a human face. Butz waits until everyone has spoken, then: "When I was an engineer, we automated the lathe. But nobody automated the machinist's judgment. That is what is different now."

Chapter 2

Foundations of a Post-Labor Economy

The last stage of the laboring society, the society of jobholders, demands of its members a sheer automatic functioning, as though individual life had actually been submerged in the over-all life process of the species. — Arendt (1958, p.322)

2.1	Four Camps	10
2.2	What Is a ‘Job’?	11
2.3	The Industrial AI Revolution	15
2.4	Goods and Services vs Jobs	18
2.5	The Core Argument	20
2.6	What Is Guaranteed Income (GI)?	23

The debate about the future of jobs in the face of [AGI](#) and advanced robotics is heavily polarized, even among professional economists, social scientists, and philosophers ([Korinek and Suh, 2024](#); [Karger et al., 2026](#); [Agrawal et al., 2026](#)). The landscape splits into four camps (Section [2.1](#)), three of which share a common framing limitation: they fixate on *jobs* rather than on *goods and services*. Before assessing what automation threatens, one must be precise about what is at stake, leading to a property-cluster definition of ‘job’ that degrades gracefully as UBI replaces wages (Section [2.2](#)). The arc of the industrial [AI](#) revolution, from the Great Decoupling of productivity and wages to current displacement forecasts (Section [2.3](#)), sets the stage for the book’s central distinction: the proper goal of an economy is maintaining a decent standard of living, not maximizing employment (Section [2.4](#)). The

thesis follows directly: automation preserves output, and the challenge is distribution, not production (Section 2.5; see also Tables 2.1 and 2.10). The section closes with the terminology and taxonomy of Guaranteed Income used throughout this book (Section 2.6).

2.1 Four Camps

The debate on AGI and the future of work can be organized into four positions (Table 2.1), defined by two binary questions: will mass job loss occur, and will broad-based prosperity be maintained? Three camps each get one prediction wrong; this book argues for the fourth: mass job loss *with* prosperity, achieved through redistribution.

Table 2.1: The four positions mapped onto a 2x2 quadrant. The axes are whether mass job loss occurs and whether broad-based prosperity is maintained. This book argues for the upper-right quadrant: mass job loss with prosperity, achieved through redistribution. Each camp has adherents across disciplines, but clear majorities exist (cf. Table 2.10).

Positions	Full Employment	Mass Job Loss
Prosperity	Continuity Thesis ✗ “New jobs will always compensate for lost ones”	Post-Labor Prosperity ✓ (this book) “Output persists; redistribute it via UBI”
Stagnation or Crisis	Stagnation Wager ✗ “AGI is far off; status quo persists”	Displacement Crisis ✗ “Job loss means poverty & collapse”

The Continuity Thesis. Every technological revolution so far created more jobs than it destroyed; AGI will be no different. Proponents (Pinker, 2018; Autor, 2015; Aghion et al., 2023; Brynjolfsson and McAfee, 2014) extrapolate from two centuries of labor-market adaptation.¹ But this time, machines are acquiring the very cognitive flexibility that previously made humans indispensable. The historical pattern has no precedent for the automation of general intelligence itself.

The Displacement Crisis. AI will cause mass unemployment, spiralling inequality, and social collapse. Proponents (Frey and Osborne, 2017; Korinek and Stiglitz, 2019; Acemoglu and Restrepo, 2020; Susskind, 2020) correctly

¹Autor (2015) offers perhaps the most measured version of this position, acknowledging that automation eliminates specific tasks but arguing that humans retain comparative advantage in adaptability and social judgment. However, his analysis predates modern large language models, which directly target the cognitive flexibility he identified as humanity’s enduring edge.

identify the unprecedented scale and speed of displacement (Section 2.3), but implicitly assume that fewer jobs must mean less prosperity, conflating the loss of employment with the loss of production.² Existing inequality trends (Figure 5.9) lend urgency to these concerns.

The Stagnation Wager. *The whole premise is overblown: true AGI is further away or less transformative than claimed, so mass displacement will not materialize.* Proponents (Acemoglu and Johnson, 2023) argue that the direction and impact of technology are political choices, not technological inevitabilities; that historical diffusion has always been slow and uneven; that regulation, energy constraints, and competing crises (geopolitical conflict, climate change) will cap AI adoption well short of full labor substitution; and that previous automation scares turned out vastly overstated. Given the exponential trajectory of AI investment, capabilities, and adoption, betting against broad AI diffusion is a risky position.

Post-Labor Prosperity (this book). *All three camps share a common framing limitation: they treat job loss as inherently threatening rather than asking who receives the output.* The Displacement camp's best models (e.g. Acemoglu and Restrepo 2019; Korinek and Stiglitz 2019) already recognize that automation can raise total output while reducing the labor share, but their policy prescriptions still center on restoring labor demand rather than redistributing the surplus. If machines produce the same or greater output, the material wealth of society is undiminished. The real challenge is distribution, not production. Under adequate redistribution, which is achievable through mechanisms such as UBI (Section 6.1), mass un(der)employment is not a catastrophe but a liberation: *a transition to universal material prosperity in which work becomes voluntary* (Keynes, 1930; Bostrom, 2024).

2.2 What Is a 'Job'?

The standard dictionary definition "a regular activity or role performed in exchange for payment" is surprisingly fragile. Under it, a group of friends who agree to pay each other equal amounts for 'friendship services' would hold jobs (net transfer zero, no genuine service rendered); a volunteer lifeboat crew member performing dangerous, skilled, structured work (similarly:

²Frey and Osborne (2017) estimate that 47% of US employment is at "high risk" of computerization, a figure widely cited but often misread as a prediction rather than a risk assessment. Acemoglu and Restrepo (2020) provide the first causal estimate: each industrial robot per thousand workers reduced the employment-to-population ratio by 0.2 percentage points and wages by 0.42%.

mountain rescue volunteers, disaster-relief workers, open-source developers) would *not*; and a day-trader gambling on meme stocks would qualify despite arguably contributing nothing to society.

Beyond dictionary definitions. The English language further conflates *job*, *work*, *labor*, and *employment*: four related but distinct concepts whose boundaries shift with technology and culture. The key insight for this book is that *jobs* may vanish while *work* persists: people will likely continue to engage in useful, meaningful, effortful activities: they simply will not be ‘employed’ in the traditional sense. Following Wittgenstein’s (1953) family-resemblance analysis of ‘game’,³ we treat ‘job’ as a *cluster concept* (see also Cholbi, 2022; Tyssedal, 2025, for philosophical surveys): no single property is necessary nor sufficient; the more properties an activity satisfies, the more ‘job-like’ it is. Table 2.2 lists eleven candidate properties, each rated by importance, alongside eight edge cases (columns a–h) that stress-test the characterization.

The job–leisure spectrum. The concept is not binary but rather a spectrum, as illustrated in Table 2.3. The AGI/UBI transition (Section 6.1) will systematically slide activities upward on this spectrum: what are today full-time jobs become side gigs, hobbies, or leisure, consistent with the book’s thesis applied at the micro level (see also Section 7.5 on meaning beyond employment).

Unpaid work. The spectrum reveals a large middle ground between ‘job’ and ‘leisure’ (see also Table 4.7): purposeful activities that satisfy many job-like properties ([LN], [LR], [LS], [LK]) yet lack compensation ([LC]) and market embeddedness ([LM]). This book uses *unpaid work* as the umbrella term for these activities, which fall into several broad categories: *domestic* (housework, cooking, home maintenance), *caregiving* (childcare, eldercare, disability support), *community* (volunteering, mentoring, neighborhood organizing), *civic* (political participation, activism, jury service), *creative* (art, music, writing, open-source software), *educational* (self-study, thinking, tutoring, citizen science), and *spiritual/cultural* (religious service, cultural preservation). As AGI displaces paid employment, these categories will absorb an increasing share of human effort; a shift from coerced labor to chosen purpose (see Section 7.5). That so many socially essential activities go uncompensated already exposes a limitation of the current economic order: supply and demand set prices with elegant simplicity, but the resulting wages

³Wittgenstein’s insight is that members of a conceptual class need not share a single common essence: “we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail” (*Philosophical Investigations*, §66). He coined the term *Familienähnlichkeit* (family resemblance) for this phenomenon.

Table 2.2: Properties of a ‘job’ and eight edge cases. ✓ = satisfied, ✗ = not satisfied, ? = ambiguous. D = book’s composite definition (bottom row). [L★] = labor-property identifiers. Edge cases: (f) friends paying each other equally for socializing; (c) professional companion for an elderly person; (y) YouTuber earning ad revenue; (v) volunteer lifeboat crew member; (h) unpaid housework (cooking, cleaning, childcare); (b) BS-job (Gracher, 2018); (e) professional e-sports player; (r) retiree selling paintings at a flea market. ‘Importance’ is a subjective weighting (1–10) informed by Morse and Weiss (1955), Autor (2015), and Susskind (2020).

[ID]	Property	Importance	Key Pro	Key Con	D	(f)	(c)	(y)	(v)	(h)	(b)	(e)	(r)
	[LC] Monetary comp.	9	Universal marker	Breaks under UBI	✓	✓	✓	✓	✗	✗	✓	✓	?
	[LL] Livable income	6	Econ. substance	Culturally relative	✓	?	?	✗	✗	✗	✓	✓	✗
	[LN] Net benefit to others	8	Economic rationale	Hard to measure	✓	✓	✓	✓	✓	✓	✗	?	?
	[LA] Asymmetry	7	Kills friend-trick	Excl. passionate wkrs	✓	✗	✓	✓	?	✓	✓	✓	?
	[LR] Regularity	5	Legal/tax standard	Gig economy erodes	✓	✓	✓	✓	✓	✓	✓	✓	✗
	[LS] Structure/account.	4	Work discipline	Many jobs lack it	✓	?	?	✗	✓	?	✓	✓	✗
	[LO] Opportunity cost	7	Classical economics	Excl. happy workers	?	✗	✓	✓	✓	✓	?	?	✗
	[LK] Skill required	3	Professionalism	Many legit low-skill jobs	✗	?	✗	✗	✓	?	?	✓	?
	[LM] Market embedded.	5	Measurability	Excl. informal economy	✓	✗	✓	✓	✗	✗	✓	✓	✗
	[LP] Replaceability	4	Labor-market notion	Fails for unique roles	✓	✗	✓	✓	✓	✓	✓	✓	✓
D	[LI] Involuntariness	6	Protestant-ethic pull	Culturally biased	?	✗	?	?	✗	✓	?	?	✗
	Book’s definition	see composite definition in the text			✓	✗	✓	✓	✗	✗	✓	✓	✗
Score (%)		✓=1, ✗=0, ?=½; weighted by importance				32	77	75	54	63	75	84	27

Table 2.3: *The job–leisure spectrum. Activities slide upward as UBI decouples income from labor (Keynes, 1930; Susskind, 2020). ‘Unpaid work’ encompasses the categories listed in Section 2.2.*

Category	Pay	Obligation	Enjoyment	Beneficiary
Pure leisure	None	None	Pure	Self only
Hobby	Tips	Flexible	Mostly	Self
Unpaid work	None	Self-imposed	Mixed	Household /community
Side gig	Some	Flexible	Mixed	Some others
Part-time job	Regular	Some	Mixed	Clear clients
Full-time job	Livable	Full	Often obligatory	Employer /market

do not always track social value well. Caregiving, nursing, and teaching are widely regarded as very valuable occupations yet remain chronically underpaid, while domestic work receives no market wage at all.

Working definition. For this book, I adopt the following composite working definition:

*A **job** is a regular, compensated activity that produces goods or services of value to others, involves a non-trivial asymmetry between provider and consumer, and generates sufficient income to contribute meaningfully to the provider’s livelihood.*

Consistent with the cluster-concept analysis above, this definition functions as a *prototypical archetype* rather than a strict rule: an activity qualifies as a ‘job’ if it satisfies sufficiently many of the core properties with sufficient weight, even if it fails individual criteria. This excludes the friend-hiring trick (fails properties [LN] and [LA]), pure volunteering (fails [LC]), and unpaid housework (fails [LC] and [LM]), while including BS-jobs (which satisfy [LC], [LA], [LR] formally, and whose services are valued by the employer even if not by society; the book contends these *should not exist*) and passionate professionals (properties [LC], [LN], [LA] still hold). Crucially, the definition degrades gracefully under UBI: as properties [LC], [LL], [LO], and [LI] weaken, fewer activities qualify as ‘jobs,’ exactly matching the book’s prediction. This narrow definition of ‘job’ sits inside the broader concept of ‘work’ that Tyssedal (2025) defines via conceptual engineering as *activity that is instrumental and intermediary to satisfying needs and wants, or equivalently, activity that produces a benefit external to the worker*. The book’s ‘job’ adds compensation [LC], regularity [LR], and market embeddedness [LM] on top of Tyssedal’s broader ‘work,’ which is why jobs can vanish while work persists. For related frameworks on labor, work, and the meaning of

employment, see [Arendt \(1958\)](#) and many of the references throughout this book.

2.3 The Industrial AI Revolution

This section traces the arc of automation from the industrial revolutions through the [Great Decoupling](#) of productivity from wages, current displacement forecasts, and AI investment trends, to the long-run structural transformation of how humanity allocates its time. The pattern is clear: each wave of automation displaced workers, raised aggregate output, and created new categories of employment. What distinguishes the AI revolution is the prospect that this compensation mechanism will finally break.

Historical automation and the Great Decoupling. As the industrial revolutions of the past centuries demonstrate, automation is the engine of progress ([Pinker, 2018](#)). First, industrial machines replaced raw physical labor; then assembly lines and refined machines mass-produced complex goods; then general-purpose industrial robots brought flexibility to the factory floor; then computers automated routine cognitive work, and the internet acted as a massive productivity multiplier. Now, AI enables the production of an increasing range of goods and services faster, cheaper, and often better than human labor. [Brynjolfsson and McAfee \(2014\)](#) discuss the ‘[Great Decoupling](#)’ of productivity from employment in the “Second Machine Age”.⁴ There is macroeconomic support for this decoupling since the 1980s. So far this has primarily led to divergence between real wages and productivity rather than mass unemployment ([Skare and Skare, 2017](#)) (see Figures 2.4 and 2.5) largely driven by globalization, weakening labor institutions, and institutional/policy changes.

Lessons from past displacement. The abstract trajectory of five industrial revolutions (Table 2.8) conceals vivid human stories that foreshadow the AI transition. Three episodes are instructive, each illustrating a distinct displacement pattern. *Wage collapse before job loss:* In the 1810s, Britain’s hand-loom weavers numbered roughly 250 000, earning skilled artisan wages. Power looms undercut them not by eliminating their jobs overnight but by driving piece rates down by 80% over three decades: weavers could

⁴The term “Great Decoupling” refers to the divergence, beginning in the 1970s, between rising labor productivity and stagnating median real wages. Before 1973, the two tracked closely; since then, productivity roughly doubled while median compensation grew only modestly. The gap reflects the increasing share of economic gains captured by capital owners and top earners rather than typical workers.




Figure 2.4: The Great Decoupling: Divergence of productivity and compensation. Approximate data showing the decoupling since the 1970s (Brynjolfsson and McAfee, 2014; Skare and Skare, 2017). Data: Bureau of Labor Statistics (BLS) / Economic Policy Institute (EPI).




Figure 2.5: Declining labor share of US national income, 1950–2023. The ≈ 10 percentage point decline since its peak circa 1960 indicates that an increasing fraction of economic output is captured by capital rather than labor. Data: BLS nonfarm business sector (Federal Reserve Economic Data (FRED) series PRS85006173).

still find work, but at starvation wages, until the occupation effectively vanished by the 1860s (Thompson, 1963).⁵ *Gradual substitution with decades of coexistence*: American telephone operators peaked at over 350 000 in the 1950s. Automatic switching, introduced in the 1920s, did not eliminate operators immediately; it redeployed them to long-distance connections, directory assistance, and supervision. The last major layoffs came in the 1960s and 1970s, a half-century after the enabling technology existed. *Complementarity then substitution*: When ATMs were introduced in the 1970s, the number of tellers per branch fell, but the reduced operating cost allowed banks to open more branches, so total US teller employment actually *grew* until the early 2000s. Only after online and mobile banking eroded branch traffic did teller employment decline by roughly 25% between 2007 and 2024. Each pattern is now recurring at compressed timescales in the AI era: translation software followed the telephone-operator trajectory (decades of coexistence, then sudden displacement); content writing is following the weaver trajectory (AI has not eliminated writers but has collapsed per-word rates); and coding assistants may be following the ATM trajectory (currently complementing programmers, with substitution still ahead).

AI's economic impact and displacement forecasts. Studies suggest that AI capabilities will contribute trillions of dollars to the global economy (McKinsey & Company, 2023) and outpace job creation leading to a contraction in wage income (Citirini and Shah, 2026; Acemoglu and Restrepo, 2022), although some other studies suggest that the resulting aggregate productivity gains may be far more modest (Acemoglu, 2025). Early indicators of this displacement are already being tracked, with over a hundred thousand explicitly AI-linked job losses (The Alliance for Secure AI, 2026), and a survey of 81 000 AI users finds that perceived job threat tracks actual AI exposure, with early-career workers expressing significantly higher displacement anxiety than senior professionals (Anthropic, 2025). The range of major displacement forecasts is summarized in Table 2.6 (see also del Rio-Chanona et al. (2025) for a systematic review); aggregate forecast distributions are shown in Figure 3.1. Note that most of these forecasts are from pre-LLM era, so are likely too conservative. Current AI already handles data entry, bookkeeping, translation, coding, and paralegal research; expert-level AI will soon reach into radiology, financial analysis, and legal practice; and once paired with advanced robotics, physical jobs from warehouse work and delivery driving to surgery will eventually face comparable disruption

⁵Thompson (1963) documents how parliamentary petitions from weavers were repeatedly dismissed with the same argument now deployed against AI regulation: that interfering with machinery would harm national competitiveness.

(Table 3.3 maps 35 job categories against technology capability thresholds in detail).

Table 2.6: Comparison of major AI/automation displacement forecasts. The wide range reflects differences in methodology (whole-job vs task-based) and time horizon. For further arguments, see Appendix A.2.

Study	Year	Key Estimate
Frey & Osborne (Oxford)	2013 ⁶	47% of US jobs at ‘high risk’
OECD (Arntz et al.)	2016	9–14% highly automatable (32 countries)
McKinsey Global Institute	2017	75–375M workers displaced by 2030
PwC	2018	21% (Japan)–38% (US) by early 2030s
Goldman Sachs	2023	300M jobs exposed (generative AI)
WEF Future of Jobs	2025	92M displaced, 170M created; 22% churn

Any country not investing in AI technology, or actively trying to slow it down, will seriously fall behind; just as any country that missed past industrial revolutions did, but likely much faster (see Table 2.8, Figure 2.7 and Section 8.5 on the global development gap).




Figure 2.7: Global private investment in AI, 2013–2025. Data: *Stanford HAI* (2026); 2026 projection estimated from other sources.

Someday, computers in general and AGI in particular will completely dominate the economy of every developed country.

2.4 Goods and Services vs Jobs

The displacement camp fears AI will destroy jobs and cause mass poverty. The continuity camp trusts that new jobs will always compensate. Both sides

⁶Working paper; published as *Frey and Osborne* (2017).

Table 2.8: Comparison of the AI Revolution with the historical agricultural/industrial revolutions. Year≈ onset (5th and 6th are the author’s estimates).

Revolution	Onset	Core Tech	Labor Impact	Pace	Social Response	Hrs/Wk
0th	Hunting-10 ⁷ BC	Tools, fire	Gatherers→hunters	—	Bands, migration	15–20
1st	Farming-10 ⁴ BC	Cultivation	Hunters→farmers	V. Slow	Settlement, hierarchy	40–50+
2nd	Steam 1780s	Mechanical power	Manual labor	Slow	Poor laws, early unions	60–80
3rd	Electric 1870s	Mass production	Skilled manual	Medium	Welfare state emergence	40–50
4th	Digital 1970s	Info. processing	Clerical, routine	Fast	Service sector growth	35–45
5th	AGI 2027	Agentic AI	Cognitive	V. Fast	Crisis →UBI	↘20
6th	Robot 2030s	Advanced Robots	Physical	Fast	Post-scarcity	→0




Figure 2.9: Structural transformation of US time allocation in hours per year, 1840–2100. Paid employment shares (Agriculture, Industry, Services) are multiplied by average annual working hours per worker (Figure 4.6; interpolated for missing years). Household hours (unpaid domestic labor: cooking, cleaning, childcare, etc.) are per-adult averages (both sexes) estimated from historical time-use studies (Cowan, 1983; Vanek, 1974; Ramey, 2009; Aguiar and Hurst, 2007). The reference total ~5 500 h/yr comprises personal care (~1 000), paid work (~3 100), and household (~1 400) in 1840, representing total obligatory productive hours plus essential self-care. Household work declined from ~1 400 h/yr (pre-appliance era) to ~850 h/yr (2020) as white goods spread and women entered the paid workforce. The plateauing of working hours since the 1960s is US-specific. In Europe, the downward trend continued (Figure 4.6). Data: BLS/Census (sectors); Penn World Table 11.0 (Feenstra et al., 2015) (working hours); American Time Use Survey (ATUS)/BLS (recent household). In the AGI-era stylized projection (right of dashed line), automation collapses both paid work and household chores, with free time rising to over 90% (excluding personal care) by 2100. Free time encompasses leisure and any voluntary activity: socializing, caring, education, hobbies, sports, and creative pursuits. Household work becomes voluntary too, since it will be fully automatable.

are misguided: they fixate on employment rather than output. The primary goal of an economy is not to maximize human toil, nor for businesses to maximize abstract profit, but to maintain a decent standard of living for as many people as possible. Un(der)employment should not be the focal problem; this book maintains it is desirable and affordable (Chapter 3).

2.5 The Core Argument

The core argument is simple: if a robot or AI replaces a human worker, the aggregate volume of goods and services is not diminished but likely increased. The material wealth of the world has not shrunk; it has merely shifted from labor income to capital returns (cf. Figure 2.5). It is then “merely” a political matter to redistribute these gains fairly so that nobody is worse off, most people are better off, and in the long run, much better off (cf. [Hutter and Hutter, 2021](#), Sec.4.3). The importance of immaterial benefits will only grow. This simple economic fact, while long recognized in the tradition of [Keynes \(1930\)](#), remains underappreciated by the displacement crisis camp.⁷ While material wealth secures physical survival and comfort, prosperity also encompasses flourishing, social status, and success. Once basic physical needs are universally met, the status associated with *being employed* will likely diminish, shifting toward creative, social, and intellectual achievements. (Conversely, *employing* humans may itself become a status signal in a world of abundant machine labor; see Section 3.9.) As long as the desired goods continue to be produced, mass un(der)employment is perfectly consistent with universal material prosperity (as formalized in the idealized transition of Section 3.2). At a large enough scale, this automation of cognitive and physical tasks requires a fundamental rethinking of economic mechanisms to deal with the shift from labor to capital as the main economic resource, ensuring that the social contract is updated for a post-AGI world ([Genewein et al., 2026](#)).

Step-by-step walkthrough. The argument proceeds in five steps, each developed in detail in subsequent chapters:

1. **Automation preserves output.** When an AI or robot performs a task previously done by a human, the good or service is still produced,

⁷In his 1930 essay *Economic Possibilities for Our Grandchildren*, Keynes predicted that within a century, technological progress would solve the “economic problem” and reduce the working week to 15 hours. His prediction about productivity growth was roughly correct (output per worker has increased eightfold); what he missed was that institutions would channel the gains into consumption growth and inequality rather than leisure.

typically faster, cheaper, and at higher quality (Section 3.1 surveys the capabilities; Table 3.3 maps 35 job categories against technology thresholds).

2. **Wealth shifts from labor to capital.** The displaced worker loses wage income, but the firm's revenue is unchanged or increased. The difference accrues as profit to the firm's owners: wealth has not vanished, it has moved from paychecks to balance sheets (Figures 2.5 and 2.11).
3. **Tax the capital windfall.** A portfolio of taxes (automation levies, land-value taxes, AI-output taxes, and reformed corporate taxes) captures a large share of the windfall (Section 5.3 presents eleven instruments; Figure 5.6 compares them across policy dimensions).
4. **Fund Guaranteed Income.** The tax revenue finances cash transfers: initially a Conditional Income (CI) targeted at displaced workers, converging toward an unconditional Universal Basic Income (UBI) as the displaced share grows (Sections 6.1 and 6.2).
5. **Universal prosperity with voluntary work.** As production costs converge toward zero (Section 3.3), the income required for a comfortable life falls, while the income provided by GI rises. Work becomes a choice, not a necessity (Sections 7.5 and 4.2).

Each step is individually well-established in economics or political philosophy; the contribution of this book is to integrate them into a unified policy architecture and to show that the chain is not merely theoretically coherent but practically feasible and politically achievable.

Why this is not as easy as it sounds. The qualifier “merely a political matter” in the preceding paragraph does not mean it is easy, just that the economics and financing is sound: Three objections deserve immediate acknowledgment. First, production costs may not actually converge to zero: scarce minerals, land, and intellectual property create persistent rents (Sections 3.3, 3.4, 3.6 and 3.7 address each in turn). Second, aggressive taxation may drive capital and talent to lower-tax jurisdictions, eroding the tax base (Sections 8.4 and 8.3 show why this threat, while real, is overstated for physical infrastructure and consumer markets). Third, even with material abundance, people may lose purpose and meaning if employment disappears (Section 7.5 shows that meaningful activity does not require a paycheck). The rest of this book develops each step and addresses each objection in detail.

Table 2.10: Four views on AGI and labor compared. This book synthesizes the valid insights of the first two camps while rejecting the third’s complacency: output grows and human labor becomes unnecessary, enabling universal prosperity through redistribution.

AGI&Labor	Core Claim	Blind Spot	Verdict	Believers
Continuity Thesis	“New jobs will always compensate for lost ones”	Ignores that AGI automates cognition itself	Half right:✗ output grows	most economists
Displacement Crisis	“Job loss means poverty & collapse”	Conflates loss of jobs with loss of production	Half right:✗ jobs vanish	“Luddites” many workers
Stagnation Wager	“AGI is far off or will be capped by regulation or resource limits”	Ignores exponential AI investment and capability trajectory	Risky bet:✗ delays preparation	AGI skeptics
Post-Labor Prosperity (this book)	“Output is maintained; the challenge is distribution”	Requires political will for redistribution (e.g. UBI)	Right&desirable: output grows and jobs vanish ✓ ⇒ universal prosperity	“Utopians” (this book)




Figure 2.11: The Piketty U-curve: Top 10% pre-tax income share in the United States, 1913–2022, measures the percentage of the total national pre-tax income that is earned by the wealthiest 10% of individuals or households in any given year. Inequality has returned to pre-war levels, a reversal strongly correlated with the decline in top marginal tax rates (Piketty et al., 2014). Without redistributive intervention, AI-driven productivity gains risk further concentration. Data: World Inequality Database (2024) <https://wid.world/>. See Figure 5.9 for the top 1% of various countries.

2.6 What Is Guaranteed Income (GI)?

This book advocates for unconditional income transfers as the primary redistribution mechanism in a post-labor economy. Since related terms are used inconsistently in public discourse, we introduce a systematic terminology. Table 2.12 defines the building blocks and composed acronyms used throughout this book (for comprehensive treatments, see [Van Parijs and Vanderborght, 2017](#); [Standing, 2017](#); [Gentilini et al., 2020](#)).

Separation principle. This book strictly separates *distribution* (who receives how much, under what conditions) from *funding* (how the transfer is financed). The term GI is deliberately funding-agnostic: the same payout can be financed by general taxation (the standard UBI model), by dividends from a Wealth Fund (Section 6.5), or by a hybrid. Our arguments about the merits of unconditional cash transfers (Section 6.1) apply regardless of the funding channel. The funding mechanisms themselves are analyzed separately (Section 5.3).

Definition and features. We use *Guaranteed Income* (GI) as the umbrella term for any periodic cash transfer that guarantees at least a subsistence-level standard of living, and *Universal Basic Income* (UBI) for the specific variant in which every citizen receives the same amount. The same or closely related ideas have appeared under various names in the literature, including “citizen’s income” and “citizen’s wage,” “universal benefit,” “state bonus,” “demogrant,” and “territorial dividend.” The Basic Income Earth Network (BIEN) defines UBI through five features: it is *periodic* (paid at regular intervals), *cash* (not in-kind), *individual* (not household-based), *universal* (paid to all, without means test), and *unconditional* (no work requirement) ([Van Parijs, 2004](#)).

Eligibility. The BIEN definition says “paid to all” but leaves open who counts. A natural starting point is to restrict UBI to all adult citizens who have not permanently emigrated. This is an existing, clear-cut legal criterion that aligns with the state’s primary duty of care toward its own citizens and appeals to UBI advocates’ preference for administrative simplicity. Non-citizens residing in the country would remain welcome but would initially need to support themselves through employment, savings, or existing welfare channels, earning entitlement to full UBI through naturalization. Without some such threshold, public support for UBI may erode: if unrestricted access is perceived as an invitation to “free-ride” on national wealth, the resulting resentment can poison the political coalition needed to sustain the program. That said, this citizen-first framing is a pragmatic default, not a

Table 2.12: Terminology: building-block letters and composed acronyms. Letters are bold-faced in the expansions to highlight the mnemonic.

Sym.	Expansion	Meaning
<i>Building blocks</i>		
G	Guaranteed	Unconditional: no means test nor work requirement (Section 2.6)
I	Income	Periodic cash payment, not in-kind nor lump-sum (Section 2.6)
U	Universal	Paid to all adult citizens; scope can be broadened (Section 2.6)
B	Basic	Level sufficient for basic needs (Section 6.1)
H	High	Level enabling flourishing, not just survival (Section 6.2)
A	Affordable / Adequate	Level linked to GDP/capita; between B and H, rising over time (Section 6.2)
W	Wealth / Welfare	Collectively held wealth, in context of redistribution (Section 6.5)
F	Fund	Investment vehicle pooling assets (Section 6.5)
C	Conditional	Eligibility restricted to displaced workers (Section 6.2)
C	Citizens'	Owned by citizens, not government (Section 6.5)
S	Sovereign	Owned by the state (Section 6.5)
<i>Income types (distribution side)</i>		
GI	Guaranteed Income	Umbrella: any periodic cash transfer above subsistence; funding-agnostic
CI	Conditional Income	Transitional: adapted welfare with relaxed conditionality (Section 6.2)
UBI	Univ. B asic Income	GI at a uniform basic amount for all citizens (Section 6.1)
UAI	Univ. A fford. Income	GI at a GDP-linked level; likely around average salary (Section 6.2)
UHI	Univ. H igh Income	GI at levels enabling flourishing (Section 6.2)
<i>Funding side</i>		
WF	Wealth Fund	Collective investment vehicle funding GI via dividends (Section 6.5)
CWF	Citizens' W F	Citizen-owned fund units (Section 6.5)
SWF	Sovereign W F	State-owned variant (Section 6.5)

moral absolute. [Van Parijs and Vanderborght \(2017\)](#) advocate for broader coverage, and as material abundance grows under automation, the case for extending UBI to long-term residents through a contribution period, analogous to vesting in social-insurance systems, and possibly eventually toward some form of global basic income. The management of selective immigration under UBI is discussed further in Section 8.4. The eligibility boundary is, in the end, a political choice; the economic analysis in this book applies regardless of where the line is drawn.

Scope: from UBI to UHI. This book is not fixated on UBI in the narrow sense. The central claim is that *some* form of [GI](#), whether delivered as a uniform UBI, as adapted welfare with relaxed conditionality (CI, Section 6.2), or even as an NIT, is both necessary and sufficient to ensure universal material prosperity in a post-labor economy. We argue, however, that UBI is the preferred long-run instrument because of its administrative simplicity, absence of stigma, and full take-up rate (Table 6.6 compares twelve approaches across these and other properties; Section 6.1 develops the case in detail; Table 6.1 surveys pilot programs; Section 6.3 addresses objections). The idealized transition scenario in Section 3.2 begins as targeted [CI](#) (displaced workers receive compensation roughly equal to their former salary) and gradually converges toward uniform UBI as the displaced share of the workforce grows. Once material abundance is sufficient, the income level can increase beyond basic needs to what we term *Universal High Income* ([UHI](#)): an unconditional income enabling not just survival but flourishing. The transition pathway from current welfare systems to UBI and eventually UHI is discussed in Section 6.2.

Key Insights

- The debate splits into four camps: Continuity, Displacement, Stagnation, and Post-Labor Prosperity. This book argues for the fourth.
- Output is maintained or increased by automation; the challenge is distribution, not production.
- An 11-property cluster definition of “job” reveals that automation erodes multiple properties simultaneously.
- The core argument: automate production → tax capital windfall → redistribute as [GI](#) → universal prosperity.

Open Problems

- How to measure the “displacement fraction” of a job in practice, given that task boundaries shift continuously.
- Whether political coalitions for redistribution can form before displacement reaches crisis levels.
- Path dependence: the order in which sectors automate may determine which transition strategies are feasible.

Meridian — 2031: The Cascade

Meridian General's surgical ward is now robot-assisted end to end. Giwoo was offered retraining as a "Patient Experience Coordinator" (smiling at people in waiting rooms). She lasted three months before the humiliation outweighed the paycheck. Marc's university dissolved his department after the AI platform published forty-seven papers in a single quarter, each validated by automated labs overnight. Chaeyeong graduated with honors and cannot find work: the consulting firms run on AI analysts now. Alice's data-cataloguing contract vanished when the AI learned to clean its own data. She volunteers at a coastal restoration project on weekends: unpaid, but at least the ocean still needs hands. Judge Chloe has overridden AI sentencing fourteen times this year; the appeals court reversed her in eleven, citing lower recidivism predictions. "The machine is less biased than I am," she tells Giwoo. "That is the part I cannot argue with." Levin is one of twenty-three officers left from fifty-eight. He handles the cases that require a human presence: mental-health crises, domestic disputes. Meridian's unemployment rate hits 14%. From his garden chair, Butz delivers the verdict: "The factory is running. The paychecks are not."

Chapter 3

The Economics of Automation

Full Employment in present fully competitive conditions would probably result in an extremely unequal distribution of income. . . I should regard such an outcome as intolerable.
— Meade (1995, Preface, pp.xviii–xix)

3.1	Next-Generation AIs and Robots	30
3.2	Idealized Transition Scenario	38
3.3	Zero Marginal Cost Paradigm	42
3.4	Self-Replicating Factories and Resource Limits	43
3.5	Remaining Cost Components	46
3.6	The Land Ownership Problem	47
3.7	The Intellectual Property Problem	48
3.8	From Cost to Price	51
3.9	Even Utopia Requires Money	53
3.10	Skill Shortage and Job Mismatch	54

If a machine can do your job, the goods you produced do not vanish: they simply arrive without your toil. This section explains the simple economic logic of that insight: automation preserves (and most likely increases) the total volume of goods and services. The problem is not output; it is distribution. All production costs trace back to labor, raw materials, and energy; once labor is automated and energy is abundant, marginal costs converge toward zero (Rifkin, 2014), and the core economic problem shifts from production to distribution (see Genewein et al., 2026). This section develops this argument in detail:

Section 3.1 surveys the AI and robotics capabilities that enable automation

across cognitive and physical jobs. Section 3.2 presents an [idealized transition](#) scenario in which corporate salary savings are taxed and redistributed as basic income. The next group of subsections traces why production costs approach zero: the zero-marginal-cost paradigm (Section 3.3), self-replicating factories and resource limits (Section 3.4), the role of land (Section 3.6), and remaining cost components (Section 3.5). Two major rent categories resist this convergence: land ownership (Section 3.6) and intellectual property (Section 3.7). Section 3.8 explains how the transition from labor cost to tax cost maintains price stability while funding UBI, and Section 3.9 shows that even in a post-scarcity economy, prices and money remain necessary. Finally, Section 3.10 addresses skill shortages and job mismatch during the transition.

3.1 Next-Generation AIs and Robots

This subsection provides a more granular look at the AI and robotic capabilities that will drive labor displacement, along with approximate timelines. Readers who are content to accept the premise that AI and advanced robots will eventually automate the vast majority of human jobs may skip ahead to Section 3.2 without loss of continuity.

AI and robot advantages. Advantages of AI workers are that they can consume and output information at superhuman rates and modalities (e.g. video-generation); thinking and reasoning can be sped up with compute; working memory capacity and memorization are vastly superhuman; they are substrate-independent and can run on any computer including better ones in the future; can be losslessly replicated including lifetime experience; and have high-bandwidth sharing of (learning) experiences ([Genewein et al., 2026](#)). Together with being more rational, emotionally stable, and informed, this leads to faster, cheaper, and better decisions, actions, and output, compared to humans, already in many domains, and soon for virtually all mental activities. Advantages of future robots, which may not lag far behind, will be that they are tireless and their shape is customizable (to be stronger, more sensitive, to larger/smaller scale, eventually even more dexterous). Together this will result in cheaper, faster, bespoke, and abundant production of high-quality goods and services. Some measure the timelines for such an [AGI](#) revolution in years, and general-purpose robots and widespread diffusion into society to come, possibly within a decade. Even if AGI takes a full generation to reach its full impact, the impact is already visible, and society needs to prepare now (see Table 3.2, Figure 3.1, and Section 10.2).



Figure 3.1: Median year by which surveyed AI researchers assign 50% probability that AI outperforms humans at each task. HLMI = High-Level Machine Intelligence, i.e. AI that can perform all human tasks. Predictions moved forward ≈ 13 years versus the 2022 survey. Data: [Grace et al. \(2025\)](#), survey of 2 778 AI researchers. Fully automating labor by 2116 appears far too conservative, while some other rows may be a bit optimistic.

The AI capability spectrum. AI capability is a continuous spectrum along two axes: *breadth* (generality across tasks) and *depth* (performance on any given task) ([Morris et al., 2023](#)).¹ In practice, however, considering discrete capability levels makes it easier to discuss automation. I define five cognitive levels (CAI through ASI) and two physical levels (AR and SR), summarized in Table 3.2. These levels are chosen to capture the qualitative transitions that matter most for labor displacement:

- **CAI** (*Current AI*) represents the baseline: what today’s LLMs, vision models, and domain-specific tools can already do (2025–2026 capability). The boundary is empirical, not theoretical: CAI is what works already now. It requires human prompting and supervision for each task. Autonomy is nascent but not yet reliable enough.
- **HAI** (*Human-like*² AI) fills the gap between narrow AI and domain expertise: general-purpose intelligence at the level of an unskilled adult, with common sense, basic reasoning, and spatial understanding but no specialized training beyond school education. This matters for physical jobs that require common sense and spatial reasoning but not specialized training. HAI can follow instructions and handle routine

¹[Morris et al.](#)’s Level 1 (Emerging) corresponds roughly to 2023-era LLMs; Level 3 (Expert) to the 90th-percentile skilled adult.

²Unfortunately, the better targeted term *Human-Level AI (HLAI)* is already taken and generally interpreted as a machine that can match or beat human thinking in almost any job.

variation; needs moderate human oversight. Note that CAI and HAI are not strictly ordered: CAI exceeds HAI for language and data tasks, while HAI exceeds CAI for embodied common-sense reasoning.

- **XAI** (Expert AI) is human-expert-level skill and knowledge across the relevant domain(s), but still requires human prompting and guidance through multi-step tasks: XAI is a tool, not an agent.
- **AAI** (*Agentic AGI*)³ has the same expert-level capability as XAI across all required domains, but adds autonomy: given minutes of high-level prompting, AAI can work independently for a day or more, planning, executing, recovering from errors, and self-correcting. This distinction is critical because at 99% replacement, one human cannot meaningfully guide 99 AI instances task-by-task.
- **ASI** (*Super-Intelligence*) is cognitively superior to every human in virtually every domain simultaneously. Fully self-directed; capable of recursive self-improvement.
- **AR** (*Advanced Robots*) denotes capable robotic systems with good mobility and manipulation but that are still clumsier than humans overall.
- **SR** (*Super-Robots*) denotes systems mechanically superior to humans in most aspects: strength, precision, speed, and endurance.

The five cognitive levels roughly correspond to the five “Levels of AGI” proposed by [Morris et al. \(2023\)](#) (Emerging, Competent, Expert, Exceptional, Superhuman) but they treat autonomy as a *deployment decision* (how much control to grant), whereas I treat it as an *intrinsic capability* (the system’s ability to sustain reliable long-horizon work without continuous human supervision or guidance). This motivates the XAI/AAI split: both have expert-level skill, but only AAI can work autonomously for a day or more. I also added the orthogonal physical capabilities AR and SR, which is essential for predicting the displacement of physical and mixed jobs.

Degree of automatability. Table 3.3 maps 35 representative job categories and team-scale activities against four replacement thresholds (10%, 50%, 90%,

³In AI research, “agentic” denotes a system that can autonomously pursue goals over extended horizons: decomposing objectives into subgoals, selecting and executing actions, monitoring outcomes, and recovering from errors without continuous human guidance. The term distinguishes goal-directed autonomous operation from tool-like use, where the human specifies each step.

Table 3.2: Technology capability levels used in Table 3.3. The five cognitive levels roughly correspond to Morris et al.’s (2023) five Levels of AGI, augmented by two physical robotics levels. Capability (not diffusion) timeline estimates are informed by expert surveys (Grace et al., 2025), industry roadmaps (Stanford HAI, 2026), scenario forecasts Kokotajlo et al. (2024), and others; the author’s personal expectations are at the lower end of these ranges. For the purposes of this book, the exact arrival dates are immaterial: what matters is the ordering and the mapping to job displacement.

Acronym	Name	Description	Capability Timeline
CAI	Current AI	Today’s LLMs and AI tools (requires continuous human prompting)	2025–26
HAI	Human-like AI	Unskilled adult-level common sense (requires moderate oversight)	2026–28
XAI	Expert AI	Human-expert skill (operates as a human-guided tool)	2027–30
AAI	Agentic AGI	Expert and autonomous (minutes of prompting → days of work)	2028–32
ASI	Super-Intelligence	Superhuman in all domains (self-directed and self-improving)	2029–34
AR	Advanced Robots	Good physical mobility and manipulation (clumsier than humans)	2028–32
SR	Super-Robots	Physically superior to humans (dexterity, sensitivity, precision, speed, endurance)	2032–40

99%), recording the minimum technology level required at each threshold.⁴ This table assesses *technological feasibility* only: legal, moral, cultural, adoption, diffusion, and resistance factors are not considered (see Section 3.2 for real-world complications). Cell entries combine a cognitive level (CAI, HAI, XAI, AAI, or ASI) with a physical level (AR or SR) where needed; if only one dimension is required, the other is omitted.

Several cross-cutting patterns emerge. The 50% column, the point at which half the workforce in a job category becomes technologically replaceable, splits three ways: routine cognitive jobs (data entry through uncreative writing) require only CAI; professional knowledge jobs (software development through research science) require XAI; and physical jobs require HAI+AR (general common sense plus robotics), with the notable exceptions of surgery (XAI+SR) and plumbing (XAI+AR). The 99% column is dominated by AAI for individual jobs: near-complete replacement demands autonomous operation because one human supervisor cannot meaningfully prompt-guide 99 AI instances task by task. Only the collective/team-scale activities (rows 27–35) require ASI at 99%, because their hardest instances demand superhuman cognition: long-standing open mathematical problems, recursive self-improvement, global governance optimization. At the other extreme, data entry is the only listed job category essentially fully replaceable by current AI across all four columns; every other job eventually requires HAI, XAI, or AAI.

Cognitive jobs. The 13 cognitive jobs in Table 3.3 fall into four groups. Data entry and telemarketing (row 1) are already essentially fully automatable by current AI: narrow, rule-based, language-mediated tasks where even the remaining 1% of edge cases require no qualitative capability jump. Routine knowledge work (bookkeeping, translation, paralegal research, accounting, journalism, and creative art; rows 2–7) follows the pattern CAI/CAI/XAI/AAI: the routine half is already within reach of current AI (form filling, template-driven writing, boilerplate research), but the hard half (tax edge cases, literary translation, investigative journalism, genuinely original art) requires expert-level domain knowledge (XAI at 90%). Professional knowledge jobs (software development, financial analysis, radiology, law, and research science; rows 8–12) require expert AI even at 50%, because the median practitioner in these fields applies deep domain

⁴For context on how rapidly these thresholds may be reached: in the largest survey of its kind, Grace et al. (2025) asked 2778 AI researchers when AI would outperform humans at each of various tasks. The median estimate for High-Level Machine Intelligence (all human tasks) moved to 2047, thirteen years earlier than in the 2022 edition of the same survey, indicating rapidly accelerating expert expectations.

Table 3.3: Minimum technology for a given replacement fraction, by job category. Assessment based on technological feasibility only (see Table 3.2 for codes). Jobs are sorted within each group from easiest to hardest to automate (by the 50% column). Collective activities span a range of difficulty (annotated in parentheses); the 99% column requires ASI because the hardest instances demand superhuman cognition.

#	Job Category	10%	50%	90%	99%
<i>Cognitive</i>					
1	Data entry / telemarketing	CAI	CAI	CAI	CAI
2	Bookkeeper / tax preparer	CAI	CAI	XAI	AAI
3	Translator / interpreter	CAI	CAI	XAI	AAI
4	Paralegal / legal researcher	CAI	CAI	XAI	AAI
5	Accountant / auditor	CAI	CAI	XAI	AAI
6	Journalist / copywriter	CAI	CAI	XAI	AAI
7	Creative artist	CAI	CAI	XAI	AAI
8	Software developer	CAI	XAI	XAI	AAI
9	Financial analyst / trader	CAI	XAI	XAI	AAI
10	Radiologist / pathologist	CAI	XAI	XAI	AAI
11	Lawyer (litigation)	CAI	XAI	XAI	AAI
12	Research scientist	CAI	XAI	XAI	AAI
13	CEO / senior executive	CAI	XAI	AAI	AAI
<i>Mixed cognitive+physical</i>					
14	Retail salesperson	CAI	CAI	XAI+AR	AAI+SR
15	Teacher (K-12)	CAI	XAI	XAI+AR	AAI+SR
16	General practitioner	CAI	XAI	XAI+AR	AAI+SR
17	Nurse / caregiver	CAI+AR	XAI+AR	XAI+SR	AAI+SR
18	Chef / cook	AR	HAI+AR	XAI+SR	AAI+SR
19	Surgeon	XAI+AR	XAI+SR	AAI+SR	AAI+SR
<i>Physical-dominant</i>					
20	Factory / warehouse worker	AR	HAI+AR	XAI+AR	AAI+SR
21	Cleaner / janitor	AR	HAI+AR	XAI+AR	AAI+SR
22	Agricultural worker	AR	HAI+AR	XAI+AR	AAI+SR
23	Dog walker / pet caregiver	AR	HAI+AR	XAI+AR	AAI+SR
24	Truck / delivery driver	CAI+AR	HAI+AR	XAI+AR	AAI+SR
25	Construction worker	AR	HAI+AR	XAI+SR	AAI+SR
26	Plumber / electrician	CAI+AR	XAI+AR	XAI+SR	AAI+SR
<i>Collective / team-scale activities</i>					
27	Mathematics (lemma→Millennium)	CAI	XAI	AAI	ASI
28	R&D (incr.→paradigm shift)	CAI	XAI+AR	AAI+AR	ASI+SR
29	AI research (NAS→RSI)	CAI	XAI	AAI	ASI
30	Governance (local→global)	CAI	XAI	AAI	ASI
31	Strategy (small→mega-corp)	CAI	XAI	AAI	ASI
32	Engineering (modeling→terraf.)	CAI	XAI+AR	AAI+SR	ASI+SR
33	Forecasting (trends→chaos)	CAI	XAI	AAI	ASI
34	Innovation (synthesis→new field)	CAI	XAI	AAI	ASI
35	Factory (workshop→mega-plant)	CAI+AR	XAI+AR	AAI+SR	ASI+SR

expertise that exceeds current AI capabilities. In both groups, the jump from 90% to 99% is primarily an autonomy transition (XAI $\rightarrow$ AAI). The CEO role (row 13) is uniquely hard among cognitive jobs: it is the only one requiring AAI already at 90%, because cross-domain synthesis, high-stakes autonomous judgment under deep uncertainty, and complex organizational navigation are tasks where prompt-by-prompt human guidance is infeasible.

Mixed cognitive+physical jobs (rows 14–19) require robotics at varying points. Retail is unusual: the easy half is already displaced by e-commerce (CAI at 50%), but the remaining brick-and-mortar advisory roles require AR at 90%. Teachers and GPs can be partially replaced via screen or VR with expert AI at 50% (older students can learn from AI tutors, routine consultations from diagnostic AI) but physical presence becomes essential at 90% (younger children, physical examinations). Nursing profits from robotics from the start (patient lifting, medication dispensing at 10%), with SR arriving at 90% because frail-patient handling demands superhuman gentleness and precision. Surgery is the hardest mixed job: even the simplest autonomous procedures (LASIK, biopsy) require expert medical judgment plus robotic precision (XAI+AR at 10%), and superior robotics (SR) is needed already at 50% for sub-millimeter-precision procedures.

The surgeon: a displacement timeline. To illustrate the gradual encroachment pattern (beyond the automation city vignettes), consider the displacement trajectory for surgery: the hardest mixed cognitive+physical job in Table 3.3.

Today, a surgeon can operate a telemanipulated robot (the da Vinci system, introduced in 2000, now installed in over 9 000 systems globally (see [Rojas Burbano et al., 2025](#), for a review)) while AI assists with pre-operative imaging, 3D anatomical reconstruction, and intraoperative navigation. The surgeon remains legally and functionally in full control; the robot is a tool, not an agent.

Near-term (~2030), AI begins handling routine subtasks autonomously: suturing, tissue retraction, vessel sealing. The STAR robot ([Saeidi et al., 2022](#)) has demonstrated autonomous laparoscopic suturing in porcine tissue, matching or exceeding human precision in controlled settings. The surgeon supervises and intervenes as needed, much as a flight instructor monitors a student pilot.

Mid-term (~2035–2040), standard elective procedures (cholecystectomy, herniorrhaphy, cataract extraction) are performed end-to-end by the AI, with the surgeon monitoring from a console and intervening only for complications. Regulatory frameworks evolve accordingly: the legal framing shifts

from “surgeon operating a tool” to “surgeon supervising an autonomous agent.”

Late (~2040–2050), autonomous systems handle most elective and many emergency surgeries. The surgeon’s role narrows to complex reconstructive, transplant, and rare-anatomy cases where anatomical variation exceeds the system’s training distribution. New trainees increasingly learn AI supervision rather than manual operative technique.

In the *endgame*, SR systems with sub-cellular precision, real-time histological analysis, and zero hand tremor surpass human performance on every metric. The surgeon becomes a historical role, analogous to the telegraph operator: once the pinnacle of technical skill, eventually rendered obsolete by the very technology it helped pioneer.

The training asymmetry underscores the inevitability: a human surgeon requires 13–16 years of post-secondary education (4 years undergraduate, 4 years medical school, 5–7 years residency, plus optional fellowship), whereas a surgical AI can replicate a trained system’s full skill set in hours via model copying and improve continuously from every procedure performed worldwide. When a single night’s batch update can incorporate the lessons of ten thousand operations across three continents, the human learning curve becomes untenable.

Physical-dominant jobs (rows 20–26) broadly follow the pattern AR → HAI+AR → XAI+AR → AAI+SR: specialized robots handle the most structured tasks at 10% without general intelligence, but at 50% the work becomes varied enough to require general common sense (HAI) for handling unexpected objects, navigating cluttered spaces, or adapting to weather and animal behavior. Expert AI enters at 90% for the long tail of unusual situations. Construction (row 25) escalates to SR earlier (at 90%) due to the high variety of physically demanding tasks. The hardest physical job is plumbing and electrical work (row 26), where expert knowledge (XAI) is needed already at 50% to interpret non-standard layouts, building codes, and hidden infrastructure.

Collective activities and the ASI frontier. The final group (collective and team-scale activities; rows 27–35) represents tasks that currently require large groups of cooperating people with heterogenous skills: mathematical research programs, scientific R&D and engineering, governance, corporate strategy, forecasting, innovation, and running advanced laboratories and factories. Each row spans a range from routine to extreme (e.g. mathematics spans from proving basic lemmas to solving Millennium Prize problems), which is why all four columns are populated despite the extreme end requiring ASI. The purely cognitive collective activities share a uniform

progression: CAI assists at 10%, XAI handles median-difficulty instances at 50%, AAI manages the hard instances at 90%, but the hardest instances require superhuman cognition (ASI at 99%). R&D, engineering, and factory-scale operations additionally require robotics for physical experimentation, prototyping, and manufacturing. The defining difference from individual jobs is that the 99% column requires ASI instead of just AAI. Some activities in this category arguably do not qualify as “jobs” in the conventional sense, but they illustrate the frontier beyond which even agentic AGI is insufficient:

- *Open mathematical conjectures*: Millennium Prize problems have resisted all human mathematicians and some likely require qualitatively superhuman insight.
- *Recursive self-improvement (RSI)*: Designing fundamentally superior AI architectures that bootstrap beyond their own capability level definitionally requires exceeding one’s own cognitive ceiling (Genewein et al., 2026).
- *Global governance optimization*: Finding Pareto-optimal⁵ solutions across billions of agents with conflicting objectives, not merely “running a government” at expert level (or often worse).
- *Paradigm creation*: Inventing entirely new fields of science, art forms, or social systems that lie beyond the human cognitive horizon.
- *Planetary-scale engineering*: Real-time climate remediation or terraforming where the feedback loops and second-order effects exceed what any human team can model simultaneously.
- *Superhuman diplomacy*: Simultaneously modeling all parties’ full psychology, history, and strategic incentives at a depth no human negotiator can approach.
- *Curing all diseases* A program aimed at eradicating all diseases would have to involve most of the collective activities in Table 3.3 at ASI level.

These examples underscore that ASI is reserved for tasks requiring *beyond-human* cognition; in contrast to every existing *individual* job, no matter how complex, is fully automatable by AAI (plus AR or SR for physical jobs).

3.2 Idealized Transition Scenario

Imagine that employee E , working for company C and earning a salary S , is replaced by an AI/robot. Company C saves salary S , meaning profits

⁵A Pareto-optimal outcome is one where no party can be made better off without making another worse off. Named after the Italian economist Vilfredo Pareto (1848–1923).

theoretically rise by S . The government adjusts corporate taxation so that C 's tax increases by precisely S . The government then distributes S directly to E as a basic income $I=S$. The macroscopic result: C 's profit margin and E 's income are unchanged, despite E no longer working. The goods or services are still distributed to society unchanged. Note that this targeted compensation of specific displaced workers at their individual salary is not yet unconditional UBI; it is a stylized [fiscal feasibility](#) proof showing that the aggregate funds exist to make no one worse off (see also [Nayebi, 2025](#), for a related formal model). In practice, such a scheme could start as a transitional CI (Section 6.2), essentially unemployment insurance without the obligation to seek new work, and gradually pool into a uniform, unconditional payment as the displaced share of the workforce grows and individual tracking becomes impractical.




Figure 3.4: Idealized transition scenario: Corporate salary savings are taxed and redistributed directly to the displaced worker as CI.

This transition toward covert wealth redistribution and universal prosperity driven by automation (see Figure 3.4) has been vividly popularized in fiction, such as the “Omega Team” scenario by [Tegmark \(2017\)](#).

A concrete example. To make the abstract algebra concrete, consider a mid-size accounting firm with 50 employees and a total salary bill of roughly \$4 million per year. *Full displacement:* The firm’s 10 junior accountants (each earning \$55 000) perform data entry, bank reconciliation, and standard tax preparation. An AI system automates these tasks entirely. The firm saves \$550 000 in salaries plus payroll taxes; the government increases corporate taxes by roughly \$550 000; and each displaced junior accountant receives a CI of \$55 000. Their income is unchanged, their commute vanishes, and the firm’s clients notice nothing. *Partial displacement:* Five senior auditors (each earning \$120 000) perform analytical work. AI handles roughly 70% of the audit analysis (anomaly detection, cross-referencing, pattern matching), so the firm restructures: two of the five auditors are displaced, while the remaining three supervise the AI. The two displaced auditors receive CI at their former salary level; the three retained auditors may see their roles

evolve toward judgment calls and client relations. *Redeployment*: The office manager (earning \$45 000) transitions to an “AI workflow coordinator” role, handling exceptions the AI flags and managing client onboarding. No CI is needed, though the role may itself be transitional (Section 4.1). Across the firm, 12 employees are displaced and 38 retained (some in transformed roles). The firm’s net saving after the automation tax is modest but positive, preserving the incentive to automate (Section 3.2). The displaced workers’ incomes are maintained, and the firm’s output is sustained or improved. This is, of course, highly simplified: real-world complications (existing taxes, automation costs, incentive margins) are addressed immediately below.

Existing taxes are no complication. The idealized scenario of Section 3.2 ignored existing tax structures, but they integrate seamlessly: *Payroll taxes* P (social security, Medicare, unemployment insurance) are currently paid by the employer as a percentage of salary S . After automation, the employer no longer pays salary or payroll tax, so the automation tax is set to recover the full labor cost $S + P$. The government simply folds payroll tax into the automation tax. *Income tax* T : the worker previously earned salary S and paid income tax on it. After displacement, the worker receives UBI $I=S$ and pays income tax T on I instead. The worker’s net income ($S-T$) and the government’s tax revenue ($P+T$) are both unchanged. *Consumption taxes* (VAT, sales tax) are likewise unaffected: the same goods are produced and consumed at the same prices (Section 3.8), so consumption tax revenue remains constant. In short, every major existing tax category (payroll, income, consumption) either folds into the automation tax or continues to operate on the same base, leaving total government revenue invariant.

Automation costs money. Let A denote the ongoing domestic automation cost (AI licensing, robot maintenance, energy) and initial purchase or setup cost when considering the first year of automation. Then company C ’s net saving is $S - A$, not S . However, the automation expenditure A itself goes toward creating (partial) new domestic jobs (e.g. in AI/Robot maintenance or energy production). The person taking this new job needs A less governmental support, allowing the government to still afford the welfare S to pay E from a tax $T = S - A$ plus the A already recycled into the labor market. This automation-service role is a new job: if filled by an unemployed worker, the government saves A in welfare; if filled by an already-employed worker, their vacated position propagates the same saving down the chain (considering only automation-induced displacement). (International leakage, where A flows to a foreign technology provider rather than domestic labor, is addressed in Chapter 8.)

Company C needs a financial incentive to automate in the first place, so the extracted tax T must be slightly less than $S - A$, leaving a small positive profit margin $\epsilon = S - A - T > 0$. Three mechanisms can generate this margin (for a concrete worked example, see Section 3.8):

- (a) the displaced worker's effective income could fall a little below their former salary S , since commuting, professional wardrobe, other work-related expenses (and toil!) vanish, reducing the CI required to maintain their living standard;
- (b) the unmodeled but likely productivity gains from automation increase the economic pie, and
- (c) a generous but finite CI cap (say \$200k/year) ensures that very high earners absorb a manageable haircut, freeing additional margin.

⁶While fine-grained, employee-specific taxation is unrealistic, simply increasing existing uniform corporate or capital gains taxes in line with global automation and productivity gains may suffice. Innovative companies' after-tax profits will increase; backward companies' will fail. More broadly, this scenario deliberately neglects real-world economic frictions such as capital expenditure for AI integration, labor-market adjustment lags, and general-equilibrium effects on prices and wages; more realistic taxation mechanisms are discussed in Section 5.3.

CapEx amortization. When the initial capital expenditure for automation equipment exceeds the first-year salary saving ($A > S$), two mechanisms resolve the apparent deficit. First, the supplier of the automation equipment itself employed workers to produce it, so A has already been recycled into wages upstream, generating the same welfare-saving effect. Second, standard accounting amortizes the capital cost over the asset's useful life of n years, so the annual charge is A/n , which for sufficiently long-lived equipment satisfies $A/n < S$. Moreover, the amortized expenditure A/n itself creates service-sector demand (maintenance, financing, installation) whose tax receipts close the remaining gap. In either case, the fiscal neutrality of the idealized transition is preserved.

In practice, this transition is of course more complex. The idealized scenarios are not meant to be implemented literally, but to illustrate that paying unemployment benefits even at full salary level is affordable and cost-neutral.

⁶Workers are not fungible: a displaced accountant cannot be employed to develop better AI. But the argument here does not require the *same* person to take the new job. It requires only that *someone* does. If no one can, this merely delays automation but does not invalidate the fiscal argument.

3.3 Zero Marginal Cost Paradigm

The cost of a product arises from labor, material, delivery, admin, inventory, and marketing costs. Ultimately, if we trace the supply chain to its origin, all costs arise from labor, raw materials, and energy. The only reason someone can charge for raw materials (barring monopolies) is scarcity and the labor required to extract them. If all labor is automated, products that rely only on abundant raw materials can be produced arbitrarily cheaply. This zero-marginal-cost paradigm is already visible in some purely digital sectors such as software development, where an app once developed can be replicated and used by billions of users at virtually no further cost. Rifkin (2014) in “The Zero Marginal Cost Society” argues this fosters a global Collaborative Commons where extreme abundance supersedes the traditional economic mindset of scarcity.⁷ If AGI extends this paradigm to physical labor via robotics, the production cost of goods and services would plummet, theoretically making universal prosperity highly achievable on relatively small universal basic incomes (Section 6.1; see Figure 3.6 and the exponential decline in Figure 3.5). This section examines why conventional GDP understates the resulting welfare gains and traces the zero-cost trajectory across five concrete sectors.

GDP-B. Paradoxically, as production costs approach zero, conventional GDP (which measures market value of output) will shrink even as actual living standards rise, since GDP cannot capture the welfare from goods that are cheap or free. Brynjolfsson et al. (2025a) address this measurement failure by proposing *GDP-B*,⁸ a welfare-based metric that quantifies consumer surplus from new and free goods, revealing that GDP growth already substantially understates welfare growth in the digital economy. Under AGI-driven abundance, this divergence between GDP and actual well-being would widen sharply. Indeed, the *Jevons paradox* applies: as AI makes cognitive tasks cheaper, total demand for them expands, so aggregate economic activity (measured by volume rather than price) may grow even as unit costs collapse.

⁷Rifkin’s central thesis is that near-zero marginal costs represent “the optimally efficient state for promoting the general welfare and represent the ultimate triumph of capitalism.” Paradoxically, this triumph undermines capitalism’s price mechanism, as profits dry up when goods can be replicated at negligible cost.

⁸GDP-B decomposes into GDP-N (welfare from new goods) and GDP-F (welfare from free goods). Using incentive-compatible choice experiments, Brynjolfsson et al. find that the median US user would require \$42/month (\$506/year) to give up Facebook alone. Including Facebook would add 0.05–0.11 percentage points to annual US GDP-B growth, entirely invisible to conventional GDP.




Figure 3.5: Moore’s Law for Everything: Exponential decline in the cost of automated goods (halving every 2 years) versus stagnant or rising costs for non-automated goods. Concept from Altman (2021).

Sector-by-sector cost trajectories. The zero-marginal-cost trajectory is not speculative; it is already observable across sectors at various stages of completion. In *music production*, a studio album that once cost \$50 000–500 000 to produce can now be composed, arranged, and mastered by AI systems at near-zero marginal cost per track, with distribution via streaming platforms adding fractions of a cent. In *medical diagnostics*, AI systems already match or exceed specialist accuracy in dermatology, radiology, and pathology; once deployed, each additional scan costs virtually nothing, compared to \$200–500 for a human specialist consultation. In *legal document review*, AI contract analysis processes pages at a fraction of a cent each, replacing junior lawyers who charged \$200 per hour; major law firms have adopted these systems at scale. In *education*, personalized AI tutoring provides one-on-one instruction at zero marginal cost per session, a service that previously required a human tutor at \$50–100 per hour. In *manufacturing*, lights-out factories (fully automated facilities operating without human workers or lighting) already exist in electronics assembly, driving the labor cost per unit toward zero. The common pattern is clear: in each sector, human labor constitutes the dominant cost component, and automation drives that component toward zero, leaving only raw materials and energy.

3.4 Self-Replicating Factories and Resource Limits

Critics argue that producing food, housing, and infrastructure requires massive physical and energy infrastructure, and hosting AI models, manufacturing robots, and extracting resources demands continuous capital investment; therefore, the marginal cost will not tend to zero (Raymond, 2015). What these counter-arguments miss is that eventually we will have self-replicating factories in the sense that a factory produces robots which can




Figure 3.6: Zero Marginal Cost Supply Chain: When labor and energy are automated, the marginal cost of products relying on abundant materials approaches zero. Inspired by Rifkin (2014).

then build further factories producing power stations and mining equipment, build further robots, and produce all other desirable goods at no labor cost. The only limitation is the availability of raw resources on Earth. Some are naturally unlimited like iron, silicon, water, nitrogen, and oxygen. Others are effectively unlimited like energy (solar and nuclear) or nitrogen fertilizer (Haber-Bosch process⁹; phosphorus and potassium must be mined, but known reserves span centuries). Note we are not considering singularitarian scenarios with explosive economic growth and resource consumption. Even then, supply chain logistics and building out the physical infrastructure mean the convergence to near-zero cost is gradual, not instantaneous. The resources that are fundamentally limited like precious metals and rare earth minerals (unless/until we master cheap nuclear transmutation or asteroid mining) are the only ones preventing convergence to absolute zero cost. At zero cost they would quickly be depleted and no further goods relying on them could be produced. The sensible approach for governments is to appropriately tax their extraction to limit consumption and ensure future availability (see Section 5.3 on taxation strategies).

Physical feasibility. The concept of self-replicating factories is not science fiction. Von Neumann (1966) proved that self-reproducing automata¹⁰ are theoretically possible: there is no thermodynamic or information-theoretic

⁹The Haber-Bosch process, developed in the early 1910s, synthesizes ammonia from atmospheric nitrogen and hydrogen using high pressure and temperature. It enabled the mass production of synthetic fertilizer, which today sustains roughly half the world’s food production. The process consumes about 1–2% of global energy, but this is a fixed infrastructure cost that scales with energy supply, not with labor.

¹⁰Von Neumann’s constructive proof (completed posthumously by Arthur Burks) shows that a “universal constructor” in a sufficiently rich computational framework can build a copy of itself from raw materials, mirroring the DNA/ribosome architecture later discovered in biology.

barrier. The practical question is *closure*:¹¹ how much of itself a factory can reproduce. Today's best automated facilities achieve roughly 50–70% closure; the RepRap project (an open-source 3D printer that prints about half its own structural parts) provides a modest proof of concept. What matters economically is not perfect closure but *geometric capacity growth*: a factory (or ecosystem of factories) that reproduces most of itself doubles capacity roughly annually. Within a decade, a modest initial investment scales to planetary-level production. The historical precedent is instructive: during World War II, the United States built 2710 Liberty ships in four years by recursively scaling shipyard capacity, each new yard partially equipped by existing ones. AGI-controlled factories would follow the same logic at greater speed and precision.

Which resources are actually scarce? The resource-scarcity objection is narrower than it appears. By mass, approximately 98% of manufactured goods are composed of the ten most abundant elements in Earth's crust: oxygen, silicon, aluminum, iron, calcium, sodium, potassium, magnesium, titanium, and hydrogen. These are, for practical purposes, inexhaustible. Steel, concrete, glass, plastics (from atmospheric carbon or biomass), and the vast majority of structural and consumer goods face no fundamental resource constraint.

The genuinely scarce elements fall into a few categories: platinum-group metals used in catalysts and fuel cells; cobalt and lithium for current-generation batteries; rare earth elements (neodymium, dysprosium) for high-performance permanent magnets; and helium for cryogenics and medical imaging. For each, however, substitution pathways exist or are emerging: lithium iron phosphate (LFP) batteries eliminate cobalt entirely; sodium-ion batteries eliminate lithium; ferrite magnets can replace rare-earth magnets at a modest performance cost; and closed-loop helium recovery systems reduce consumption sharply. Recycling economics also improve as virgin extraction costs rise: urban mining of electronic waste is already profitable for gold and palladium, and will become so for a widening set of elements as automated disassembly reduces labor costs. The combination of substitution, recycling, and [Pigouvian](#) extraction taxes (Section 5.3) can manage scarcity without preventing the convergence of aggregate production costs toward near-zero.

¹¹The fraction of a factory's components that the factory itself can produce. Full self-replication requires closure near 100%; lower values indicate external supply-chain dependence.

3.5 Remaining Cost Components

The decomposition above focused on labor, raw materials, and energy, but production cost also includes distribution, marketing, overhead, and capital charges. Distribution (transportation, warehousing, logistics) and marketing (advertising, sales, market research) themselves decompose into labor, materials, and energy; with automated trucks, drone delivery, AI-generated campaigns, and automated warehouses, these costs follow the same trajectory toward zero. Overhead (administrative, insurance, and compliance costs) is predominantly labor, readily automated by AI. The remaining overhead component, rent for premises, is a land cost addressed below. Capital charges have two components: physical depreciation and the financial return demanded by investors. Depreciation is simply the cost of replacement parts and repair, i.e. materials and labor, both of which converge to zero as argued above. The financial return on capital vanishes in competitive markets once physical capital (factories, robots, servers) is freely replicable at near-zero cost: absent scarcity, there is no rent to extract. ¹²




Figure 3.7: Exponential cost declines on logarithmic scale. Left: Solar PV module cost dropped from \$106/W (1976) to \$0.20/W (2023), an $\approx 500\times$ decline. Right: Computing cost fell from \$1.1 trillion/GFLOPS (1961) to \$0.01/GFLOPS (2024), an $\approx 10^{14}\times$ decline. Such learning curves suggest AI-driven production costs will follow similarly steep trajectories. Data: IRENA/Nemet; AI Impacts/Epoch AI.

¹²More precisely, the *absolute* capital charge rK vanishes even if the interest rate r remains positive (reflecting time preference and risk), because the reproduction cost K is negligible (Susskind 2020, Chp.8; Korinek and Stiglitz 2019). The exceptions, where artificial scarcity persists through patents, network effects, or market power, are addressed under intellectual property below.

3.6 The Land Ownership Problem

Land to build on is limited and expensive, but primarily due to existing ownership and governmental regulation. In principle, even a densely populated country like the UK could be reorganized such that every person can live in a free-standing house essentially built for free (rare raw elements amount to a few hundred dollars). In other countries like Australia, every citizen could afford a whole farm. Services (e.g. dog-walking or baby-sitting) do not directly require significant raw materials, so their price could also fall close to zero once robots are physically capable, ubiquitous, and cheap (Section 3.3).

More fundamentally, in a zero-marginal-labor-cost economy, land and natural resources become the binding constraint on achieving true post-scarcity. If all arable, residential, and resource-bearing land is privately owned, raw material and housing costs cannot fall to zero regardless of how cheap production becomes, because landowners capture the residual rent. The insight that land rents are owed to society, not to individual owners, predates modern economics: [Paine \(1797\)](#) proposed a citizens' dividend funded by a ground rent on landowners as early as the eighteenth century. [George \(1879\)](#) developed the idea into a systematic economic program, arguing that the appreciation of land values is largely unearned, driven by population growth, infrastructure, and societal development rather than the owner's effort, and therefore a natural target for redistribution. Several policy approaches address this:

- **Land Value Tax (LVT):** Tax the unimproved value of land at high rates, as proposed by [George \(1879\)](#) and endorsed on efficiency grounds by [Stiglitz \(1977\)](#) (the "Henry George Theorem":¹³ under certain conditions, the optimal public-goods bundle can be financed entirely by land rents). LVT discourages speculative hoarding, is impossible to evade (land cannot be moved offshore), and creates no **deadweight loss** since land supply is fixed. A high LVT would capitalize into lower land prices, threatening mortgage collateral and banking stability during the transition ([Dye and England, 2009](#)). In the long run, however, lower land prices reduce household debt and shift lending toward productive investment. Transition risks can be mitigated by gradual

¹³[Stiglitz's](#) result shows that in a model with optimal community size, aggregate land rents exactly equal expenditure on public goods, so a 100% tax on land rents (i.e. the annual rental value of unimproved land, not the capital value) can finance the efficient level of public provision with zero deadweight loss. The theorem provides the theoretical foundation for land value taxation as the "least bad tax," a position endorsed by economists from Milton Friedman to Joseph Stiglitz.

phase-in over decades, grandfathering or homestead exemptions for existing owners, compensatory tax credits or bonds, and coordinated mortgage restructuring.

- *Public land ownership*: Singapore demonstrates that state ownership of most land, leased to citizens and firms on long-term contracts, is compatible with a thriving market economy and enables the government to capture appreciation directly. Compulsory purchase at fair-market value, funded by automation-generated surplus, could gradually transition privately held land into a public trust.
- *Resource extraction taxation*: Natural resources are the second major tangible rent category alongside land (see Table 3.8 for a comparison). Unlike land, resources are depletable, so the policy goal is both rent capture and conservation. The Norwegian model of capturing resource rents through sovereign wealth funds provides a proven template. Pigouvian taxes on extraction can simultaneously limit depletion and fund UBI (Edenhofer et al., 2021).
- *Harberger self-assessed taxes*: Under a Common Ownership Self-Assessed Tax (COST), owners declare the value of their land and pay a percentage annually; anyone can buy the asset at the declared price (Posner and Weyl, 2018). This ensures efficient allocation and prevents indefinite accumulation of idle land.

In practice, a combination of high LVT on unimproved values plus extraction taxes on finite resources is likely the most feasible path, since it builds on existing tax infrastructure and avoids the political upheaval of wholesale expropriation. The key insight is that in a post-AGI economy, the historical shift from labor taxation to capital taxation (Section 5.3) must in time extend to *land* and *intangible-rent* taxation as the last remaining sources of economic rent.

3.7 The Intellectual Property Problem

Land and natural resources are tangible, but the post-AGI economy features a third major category of economic rent: intellectual property (IP). Patents, copyrights, and trade secrets grant legal monopolies that function much like land titles: the holder captures rents merely by owning the right, not by producing anything (Boldrin and Levine, 2008; Stiglitz, 2008).¹⁴ Table 3.8

¹⁴Boldrin and Levine's (2008) central argument is that intellectual property monopolies typically stifle rather than promote innovation, functioning primarily as rent-extraction

compares the three rent categories.

Table 3.8: Comparison of the three major rent categories in a post-AGI economy. Land is the cleanest target for taxation (immovable, fixed supply, zero deadweight loss); natural resources are intermediate; intellectual property is the hardest due to its mobility and elastic supply.

Property	Land	Natural Resources	Intellectual Property
Supply	Fixed (perfectly inelastic)	Finite stock, depletable	Created; elastic in principle
Mobility	Immovable	Partially mobile (commodities traded)	Highly mobile (redomicile, licensing)
Rent character	Mostly earned (location value)	un-Partially earned (discovery effort)	Mixed (R&D effort vs legal monopoly)
DWL of taxation	Zero	Low	Potentially high (may deter innovation)
Evasion risk	None	Moderate	High (jurisdictional arbitrage)
Policy tools	LVT, Harberger tax	Sovereign wealth funds, Pigouvian tax	Patent reform, compulsory licensing, Harberger IP tax

The crucial difference is that IP is both *created* and *mobile*, whereas land is neither. Taxing land rents incurs zero deadweight loss because land supply is fixed; taxing IP rents risks deterring the R&D that creates new knowledge. However, the traditional justification for IP (incentivizing human creativity through temporary monopoly) weakens markedly once AI systems generate the bulk of patentable inventions and copyrightable works at near-zero marginal cost (Abbott, 2016). When the creative input is a machine rather than a human mind, the monopoly rent accrues to the machine’s owner, not to an inventor who needs incentivizing.

AI-generated IP and the authorship gap. The legal basis for IP rests on a foundational assumption that is rapidly eroding: that intellectual creations originate from human minds that need incentivizing. Current patent law requires a human inventor, and copyright law requires a human author. AI-generated works fall into a legal vacuum. The US Copyright Office’s 2023 decision on *Zarya of the Dawn* denied copyright to AI-generated images

devices. They advocate drastically shortening patent terms or abolishing them entirely, noting that historically most innovation occurred without strong IP protection.

while partially protecting human-arranged elements. The DABUS patent applications, which listed an AI system as the sole inventor, were rejected in the United States, the United Kingdom, and the European Union (though briefly accepted in South Africa and Australia before reversal).

The practical consequence is a legal fiction: firms file AI-generated inventions under the names of human employees, preserving the formal requirements while the substantive rationale has evaporated. As AI systems generate an exponentially growing volume of patentable inventions and copyrightable works, this gap between *de facto* creation (by machines) and *de jure* ownership (by humans) widens, with monopoly rents continuing to flow under legacy rules to entities that have contributed neither the creative spark nor the labor. The case for shortening IP terms or expanding compulsory licensing (already strong on efficiency grounds) becomes overwhelming when the “inventor” being incentivized is a neural network that requires no incentive beyond electricity.

Open-source vs. proprietary AGI. The distributional implications of IP under AGI depend critically on whether frontier AI systems are proprietary or open-source. If AGI remains proprietary (as with current frontier models from major AI labs), IP rents concentrate explosively: a handful of firms capture monopoly rents on the most economically transformative technology in history, and the IP-reform agenda described above becomes urgent. If AGI is open-source (as with open-weight models released by various companies), IP rents on the model itself dissipate because the weights are freely available, but competitive advantage shifts to complementary resources: compute infrastructure, proprietary training data, and deployment scale. Open-source AI thus does not eliminate concentration; it relocates it from IP owners to compute owners (cloud providers and chip manufacturers), a shift that requires different policy instruments (compute taxes, infrastructure regulation) rather than IP reform. The most likely outcome is a mixed regime: frontier models remain proprietary with massive patent portfolios, while commodity-level models are open-source, creating a two-tier economy in which the IP-reform and compute-governance agendas must be pursued in parallel.

Data. A related emerging category is *data*. User-generated data is the raw material that trains AI models, yet the value it creates accrues almost entirely to platform owners. Data differs from IP in that it is largely non-excludable, continuously generated by users rather than invented, and typically not protected by patent or copyright. [Posner and Weyl \(2018\)](#) propose treating data labor as compensable work; others advocate data dividends or platform-level data taxes. The “data as labor” paradigm ([Posner and Weyl, 2018](#)) has a

natural interface with GI: if users are compensated for the data they generate (through micro-payments, data dividends, or platform-level levies), these payments constitute a partial basic income funded by the tech sector's use of collective data. In practice, individual data contributions are too small and too complementary to be priced in isolation, making platform-level data taxes (assessed on aggregate data value) more administrable than per-user micro-payments. Whether data rents in the end warrant a distinct policy instrument or are adequately captured by corporate profit taxes and digital services levies remains an open question.

Reforms. Several reform directions are available: shortening patent and copyright terms to reduce the duration of monopoly rents (Boldrin and Levine, 2008; Korinek, 2024), expanding compulsory licensing so that essential AI-generated innovations diffuse rapidly, and applying Harberger self-assessed taxes (Posner and Weyl, 2018) to IP portfolios, the same mechanism discussed above for land, where holders declare a value, pay an annual percentage, and any challenger can purchase the patent at the declared price. Unlike land, IP reform requires international coordination to prevent jurisdictional arbitrage, reinforcing the case for the global governance instruments discussed in Section 8.3.

3.8 From Cost to Price

Falling production costs do not automatically translate into falling consumer prices. In general, the price of a good or service decomposes as

$$\text{Price} = \text{Production Cost} + \text{Tax} + \text{Profit} - \text{Subsidies}.$$

In the idealized transition scenario (Section 3.2), the government taxes away the corporate savings from automation (i.e. the increase in profit that results from eliminating labor costs). As production costs fall toward zero, taxes rise by roughly the same amount, so the consumer price remains approximately constant. This is precisely how CI is funded: the tax component of the price replaces the labor-cost component, and the displaced worker receives the difference as basic income.

A worked example. Consider a loaf of bread that currently costs \$3, decomposed as: ingredients \$0.50, labor \$1.50 (baking, packaging, retail), overhead \$0.50, and profit \$0.50. After full automation of the bakery, the labor cost drops to zero, overhead falls to \$0.10 (machines require less floor space and management), and the pre-tax price becomes \$1.10 (production cost \$0.60 plus retained profit \$0.50). The government imposes an automation

levy that captures \$1.50 of the savings: the consumer still pays \$3, but \$1.50 now flows to the treasury rather than to bakery workers' paychecks. The displaced workers receive the salary-equivalent sum as Guaranteed Income, and the company saves $\epsilon = \$0.40$ (cf. Section 3.2), motivating the automation. One might object that this is merely a hidden price floor. That is correct, and it is fine: the "floor" is not waste but a transfer from the firm's cost savings to citizens' incomes. Moreover, as competition forces profit margins down (all bakeries automate, so no single firm retains a cost advantage), the pre-tax cost falls further, and the tax can decrease proportionally while still funding the same GI payments. In practice, the mechanism is not a literal per-unit "automation tax" but the portfolio of instruments developed in Section 5.3 (corporate profit taxes, VAT, automation levies, and rent taxes), whose combined effect approximates this transfer.

Long-run tax base shift. In practice, prices are set by supply and demand, not by cost-plus accounting. In competitive markets, some prices will fall as costs drop (benefiting consumers directly), while in markets with limited competition or scarce inputs, profits or rents may absorb the savings instead. The detailed tax instruments discussed in Section 5.3 are designed to capture these gains regardless of market structure. If AI-driven productivity growth expands total output (as discussed under 'Abundance' in Section 10.1), the tax burden per unit can actually decrease, allowing prices to fall even as UBI rises.

In the long run, however, ad valorem taxes face a structural problem: a percentage tax on a near-zero price yields near-zero revenue. (Aggregate nominal expenditure may remain stable as a share of GDP, but per-unit amounts become too small to enforce.) A 20% VAT on a loaf of bread that costs \$0.01 yields \$0.002; compensating by raising the rate to 10 000% is economically equivalent to imposing an arbitrary per-unit price floor, which creates distortions and enforcement problems that defeat the purpose of percentage-based taxation. Long-run UBI funding therefore shifts more naturally toward structurally scarce economic rents: land (with the zero-deadweight-loss properties established in Section 5.3), natural resources (finite stocks priced by Pigouvian taxes that reveal social cost), and intellectual property (as already identified in Section 5.3). These tax bases are market-determined and self-revealing: their values do not shrink with technological progress but may in fact grow as the economy expands around them. The shift from ad valorem to rent-based taxation is therefore not merely a policy preference but a structural inevitability as marginal production costs approach zero.

3.9 Even Utopia Requires Money

Zero production cost does not imply zero price. Even in a post-scarcity economy, insatiable consumer preferences keep demand ahead of supply and scarce resources. Positional competition for status goods further ensures that market allocation remains necessary.

Insatiable desires and demand. Beyond the fiscal mechanism just described, consumer prices remain positive for a more fundamental reason: demand is not fixed. For goods with saturable demand, e.g. basic food, basic medicine, and basic clothing, and especially perishable basic food, prices can converge close to zero once production is fully automated, because there is little incentive for hoarding. But as long as consumer preferences are insatiable (wanting ever-larger homes, faster vehicles, more exotic vacations, more personalized entertainment, higher-resolution displays, longer-range autonomous drones, bespoke pharmaceuticals), demand is not fixed, and ramping up production will always lag behind these expanding desires. Rather than a state-dictated rationing of abundant-but-never-enough goods, let people's preferences decide what to produce, and let supply and demand decide the prices and availability.

Status goods and positional competition. Humans are status-seeking creatures, and many things are desirable for status reasons only, which leads to a positional arms race in [status goods](#) and services ([Frank, 2007](#)) (see also [Section 4.5](#) on wasteful hyper-competition and [Table 4.8](#)). Not only do people desire luxury goods (most of them may eventually become extremely cheap), they also desire status services. [Status goods](#) will by definition always be scarce and hence expensive and supply will lag behind demand. A distinct subcategory is *locational goods*: only 36 homes exist on Presidio Terrace in San Francisco, and no amount of automation can create a 37th. Once everyone has a BMW, you need a Porsche to signal status, and then a Bugatti, but even this can eventually be mass-produced at close-to-zero cost, and then what? Already today some people care about human-made luxury goods, even if they are inferior to machine-made goods: hand-thrown pottery, bespoke tailoring, artisanal food, live theater, and watching sports.

When work becomes voluntary, human labor will not uniformly become *scarce*, but it will become differentially *expensive*, with prices set by how unpleasant the task is. For tedious, subservient work that few would freely choose¹⁵ willing workers will be rare and command princely sums. That is

¹⁵Examples include a human butler standing at attention, a liveried footman serving dinner, a personal valet laying out clothes, or a chauffeur driving when the self-driving car is flawless.

precisely the point: having a human perform these roles *in view of others* is the ultimate status symbol, proof that you can afford to waste someone's irreplaceable time on what a robot does better. Conversely, for activities many people genuinely love (cooking an elaborate dinner, tending a garden, restoring vintage furniture, tutoring a neighbor's child) supply will be abundant and prices modest, since plenty of people would happily do these for pocket money or even for free. The result is a labor market segmented not by skill but by *desirability*: the more subservient and tedious the role, the higher the premium for having a human fill it.

3.10 Skill Shortage and Job Mismatch

The preceding sections assumed a frictionless labor market in which displaced workers are simply compensated. In practice, humans are not interchangeable: retraining takes years, AI-induced displacement will outpace it, and [Moravec's paradox](#) ([Moravec, 1988](#)) means that cognitive work vanishes before physical work does. These frictions shape the transition timeline but do not alter the endpoint.

Some arguments implicitly assumed that humans are fungible, which they are clearly not. Retraining takes time, and this friction guarantees that job vacancies won't instantly drop to zero. A laid-off paralegal cannot overnight assume a software engineering position to automate further jobs. When in the past the job landscape shifted slowly over decades, retraining and generational turnover could fill new needs. AI-induced changes will be too fast for this mechanism to function effectively. It is hard to predict specific future jobs, but we can categorize the transition:

AI/Robotics jobs: A shortage of highly skilled workers to automate other professions will temporarily slow job losses, easing the societal transition. This friction is likely a net positive for stability. However, as AI systems increasingly contribute to and automate AI research itself through recursive self-improvement ([Genewein et al., 2026](#)), this temporary friction will quickly evaporate, accelerating the transition.

New skilled work: In the era of AGI, most jobs requiring novel cognitive skills will likely be performed better by the AIs themselves, so human demand in this sector may remain low.

Unskilled physical work: Until robots conquer [Moravec's \(1988\) paradox](#)¹⁶ entirely and take over menial physical work (e.g. plumbing, cleaning), many

¹⁶Moravec's paradox, articulated in the late 1980s, observes that "it is comparatively easy to make computers exhibit adult-level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception

expect surviving jobs to be in this category (see Figure 3.9).

Speed of displacement: Paradoxically, a rapid capability takeoff may be *advantageous*: by making displacement unmistakable and sudden rather than gradual, it forces the structural reforms (redistribution, UBI) that a slower trajectory would let politics indefinitely postpone, precipitating the transition to a post-labor society earlier than a boiling-frog scenario in which each incremental loss appears manageable (see Section 10.2 for a fuller treatment).

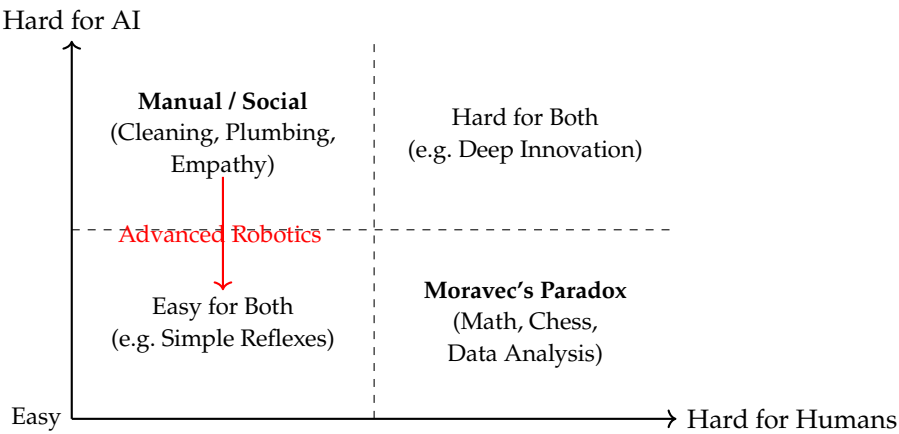


Figure 3.9: *Moravec's (1988) Paradox Quadrant: AI excels at high-level cognitive tasks hard for humans but struggles with low-level sensorimotor tasks easy for humans, though advanced robotics is progressing.*

A former programmer may not want to become an Uber driver or a cleaner or a pizza deliverer due to lower salary or perceived loss of dignity. However, with a UBI in place (Section 6.1), neither is a problem, so working becomes optional. If an employer desperately needs a human for a physical task, they just need to offer a sufficiently attractive salary to entice someone out of leisure.

In summary, next-generation AI and robotics have the potential to drive the marginal cost of production close to zero and create unprecedented material abundance. While skill shortages and land ownership constraints may create temporary frictions, the underlying economic logic points toward a post-scarcity society where labor is no longer the binding constraint on production.

and mobility.” The explanation is evolutionary: sensorimotor skills have been optimized over hundreds of millions of years of evolution and are implemented in highly parallel, dedicated neural circuits, whereas abstract reasoning is a recent (~100 000 years) and comparatively shallow adaptation, with correspondingly less neural optimization.

Key Insights

- All production costs trace to labor + materials + energy; automating labor and harnessing abundant energy drives marginal costs toward zero.
- An idealized transition shows that each displaced worker's former income can be funded without reducing anyone's income or profit.
- Land and IP are the last bottlenecks: LVT, Harberger taxes, and patent reform can capture these rents.
- $\text{Price} = \text{Cost} + \text{Tax} + \text{Profit} - \text{Subsidies}$: as costs fall, taxes rise proportionally, keeping prices stable while funding redistribution.

Open Problems

- The speed at which automation costs actually decline in practice (e.g. S-curve vs. exponential) remains empirically uncertain.
- Whether LVT can be implemented at scale without destabilizing real estate markets during the transition.
- How to handle the transition period when automation is partial: some sectors at very low cost, others still labor-intensive.

Meridian — 2034: The Make-Work Trap

Meridian's city council, alarmed by 22% unemployment, launches a "Dignity Through Work" jobs program. Giwoo is now a "Wellness Navigator," guiding patients through an AI diagnostic system that needs no guidance. Marc is a "Laboratory Heritage Consultant," cataloguing old equipment for a museum nobody visits. He spends his afternoons in the garage building a walnut bookshelf: the only work that feels real. Chaeyeong writes "AI Impact Assessments" for public agencies. They are thorough, clearly argued, and never read. Alice provides "human oversight" at a coastal restoration site, which means watching autonomous dredging drones move sand. Judge Chloe has stopped overriding the AI. The data is unambiguous: AI-assisted sentencing reduces racial and socioeconomic disparities by 40%. Her moral hesitation lasted three years. The statistics ended it. Levin is one of nine officers covering a district that once employed over a hundred. He answers the calls machines cannot: a grieving widow, a teenager on a ledge, a neighbor dispute that needs a calm voice. Butz, slower on his feet these days, watches the city council meeting on television: "They are paying people to watch machines that do not need watching. That is not dignity. That is theater."

Chapter 4

The Illusion of Full Employment

A bullshit job is a form of paid employment that is so completely pointless, unnecessary, or pernicious that even the employee cannot justify its existence even though, as part of the conditions of employment, the employee feels obliged to pretend that this is not the case.
— Graeber (2018)

4.1	New Jobs Will Be Created	60
4.2	What “Jobs” Will Remain?	66
4.3	A Taxonomy of Post-ASI Human Activities	68
4.4	Why Remaining Jobs Will Erode	73
4.5	BS-Jobs	75

‘New jobs will always be created’ is a comforting and dangerous economic mantra of our time.¹ The parallel to the 2008 Global Financial Crisis is instructive: before the crash, mainstream economists trusted equilibrium models that deemed systemic crises virtually impossible (the “Great Moderation” consensus) until reality refuted the models (Rodrik, 2015, Chp.5). Today, the same profession trusts the historical pattern that technology always creates

¹The belief that technology always creates compensating employment is sometimes invoked via the “Luddite fallacy”: the claim that the Luddites who smashed textile looms in the 1810s were wrong to fear displacement because new industries ultimately absorbed the workers. The term is used as a rhetorical trump card to dismiss concerns about technological unemployment. But if AGI replicates general human cognition, the historical pattern no longer applies, and the “fallacy” label becomes itself fallacious: the Luddites’ error was one of timing, not of logic, and this time the logic will hold.

compensating jobs, even though [AGI](#) represents a qualitative break: the first technology that replicates the very cognitive flexibility on which all prior compensation rested. The historical pattern will break once machines replicate general intelligence (Section 4.1), and while certain niches of human labor will persist longer than others (Section 4.2), even these will erode under wage competition and demand insufficiency (Section 4.4). Worse, artificially maintaining full employment traps society in hyper-competitive BS-jobs² (Section 4.5) that destroy more value than they create.

4.1 New Jobs Will Be Created

The comforting historical pattern, that automation destroys old jobs and creates new ones, rests on the assumption that human cognition remains indispensable. I first review that pattern, then explain why AGI breaks it, and finally examine which niches, if any, survive.

The historical pattern. Historically, every destroyed profession spawned new ones, and the net effect of technology on human welfare has been overwhelmingly positive ([Pinker, 2018](#)). Indeed, roughly 60% of employment in 2018 was in job titles that did not exist in 1940, and new work accounts for about 85% of net employment growth over those eight decades ([Autor et al., 2024](#)).³ In the past, unskilled jobs have been automated but replaced by often higher skilled entirely new jobs ([Skare and Skare, 2017](#)). Some argue, even if AGI will replace skilled jobs, again, entirely new, previously unimaginable sectors of employment that complement, rather than compete with, frontier technologies will be created (see Figure 4.2).

Transitional jobs. Indeed, new jobs *will* be created during the transition, both jobs that facilitate further automation and new demand-driven jobs in sectors that do not yet exist. Historically, each technological revolution produced characteristic transitional occupations that persisted for decades before the technology matured enough to eliminate them (cf. [Autor et al., 2024](#)). Table 4.1 maps representative examples to their plausible AI-era analogs.

²“BS-jobs” is used throughout this book as a shorthand covering two overlapping categories: [Graeber’s \(2018\) “bullshit jobs”](#) (work the employee considers pointless) and, more broadly, jobs that are Bad for Society (predatory, arms-race, or negative-sum employment). See Table 4.9 for the full taxonomy.

³[Autor et al.](#)’s study uses a novel methodology linking Census occupational titles across eight decades. The 60% figure captures genuinely new job categories (e.g. software developer, MRI technician) rather than mere relabeling.

Table 4.1: Representative transitional jobs from past technological revolutions and their plausible AI-era analogs. Historical durations are approximate; the AI transition may compress these to 5–15 years as software scales faster than mechanical systems.

Revolution	Transitional job	Duration	AI-era analog
Industrial	Jacquard pattern designer	~50 yr	Prompt engineer
Industrial	Steam engine operator	~80 yr	AI system operator
Electrification	Switchboard operator	~80 yr	AI integration specialist
Electrification	Elevator operator	~70 yr	AI safety monitor
Automobile	Red flag walker (UK)	~30 yr	AI ethics compliance officer
Computer	Keypunch operator	~30 yr	Data labeler
Computer	Webmaster	~15 yr	AI workflow designer
Computer	Y2K consultant	~4 yr	One-off migration specialist

Five archetypes and temporal compression. These roles cluster into five recurring archetypes: *Operators* (running new technology not yet self-operating), *Translators* (mediating between old and new systems), *Maintainers* (repairing immature technology), *Regulators* (enforcing safety during uncertainty), and *Educators* (teaching others to use new tools). Two further observations sharpen the analogy. First, historical transitional jobs lasted 20–80 years; the AI transition will likely compress this to 5–15 years because software scales faster than mechanical infrastructure. Second, many ostensibly “new” AI-era jobs are old tasks relabeled (“AI-powered consultant” = consultant using AI); the genuinely new roles are those that AI itself makes possible (managing autonomous agent fleets, red-teaming emergent capabilities), and these are precisely the ones most vulnerable to further automation.

AI-era transitional roles and their terminus. We can organize the AI-era transitional roles into four functional categories (see also Section 3.10):

- (i) *Training and refinement:* data labelers, RLHF evaluators, and prompt engineers who curate the feedback AI systems learn from (Gray and Suri, 2019);
- (ii) *Oversight and safety:* AI auditors, red-teamers, and alignment researchers who monitor systems not yet trusted to operate autonomously (Korinek and Suh, 2024);
- (iii) *Human–AI mediation:* integration specialists and workflow designers who embed AI into legacy organizations (Brynjolfsson and McAfee, 2014);
- (iv) *Authenticity and meaning:* provenance certifiers and human-experience curators, including the “ractors” of Stephenson’s (1995) *Diamond Age*

who perform live inside AI-generated narratives.

Beyond AI-specific roles, parallel technological and social shifts will generate their own transitional employment provided the AI transition is slow enough to initially require humans: climate restoration (rewilding managers, geoengineering monitors), space industry (orbital construction, asteroid mining), biotech (gene therapy counselors, organ-printing technicians), and post-work social infrastructure (meaning consultants, community architects). The meta-irony is hard to miss: many AI-related roles (labeling data, aligning objectives, red-teaming outputs) consist of training and improving the very systems that will make them obsolete. Gray and Suri (2019) argue that the frontier perpetually recedes: each advance in automation creates new edge cases requiring human judgment. However, this process has a terminus. Once AI can automate not just individual tasks but entire supply chains (see below), including the *generation* of new tasks themselves, the labor share of national income falls monotonically (Figure 2.5) and approaches zero in the limit. New task creation (N growing in the Acemoglu model) continues, but automation (I advancing) eventually catches up and surpasses it. Widerquist (2024, Ch.7) makes the related point that the standard automation argument for GI (“robots will take all jobs”) is strategically weak because it invites the reply “let’s wait and see.” His alternative argument, that present-day disruption and inequality already justify UBI regardless of future predictions, complements rather than contradicts this book’s thesis: we agree that UBI provisions are needed *now*, and additionally argue that the terminus is real.

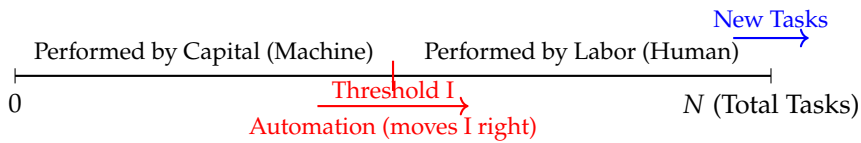


Figure 4.2: *Acemoglu and Restrepo (2019) Task Model: Automation displaces labor by moving the threshold I to the right, while new tasks at the frontier become available to human labor. The problem is that automation likely outpaces new task creation, leading to a net loss of jobs.*

Why this time is different. While it is prudent to ground one’s forecasts in historical patterns, one must be acutely aware of the differences. Some past revolutions actually increased the total amount of human labor (Table 2.8). Hunter-gatherers worked an estimated 15–20 hours per week (Sahlins, 1972); the agricultural revolution *increased* toil to 40–50+ hours (Diamond, 1987; Scott, 2017). The industrial revolution initially made things worse: factory workers routinely worked 60–80 hours per week, and only decades of labor

organizing clawed hours back down. The electric and the digital revolutions raised productivity enormously and finally reduced working hours to around 40 hours but not back to hunter-gatherer leisure levels (Figure 4.6). Globally, annual labor hours per worker fell roughly 34% between 1800 and 2025 (from about 3 200 to 2 100 hours), with an estimated productivity elasticity of -0.15 : for every 1% rise in hourly productivity, labor hours declined by 0.15% (Andreescu et al., 2025). This elasticity is a “political” elasticity, reflecting the strength of labor movements and collective bargaining rather than individual utility maximization; the main period of work-time reduction (1860–1980) coincides with the peak of organized labor. In every case, labor was *redeployed* to new sectors, not *eliminated*: farm workers became factory workers, factory workers became office workers, and office workers became service workers. AGI breaks this pattern because it is the first technology capable of substituting for human labor itself (cognitive and, with robotics, physical) across virtually all domains simultaneously. There are no new sectors left to absorb displaced workers when the displacing technology is, by definition, at least as capable as the workers it displaces. Leontief (1983a) wrote “The role of humans as the most important factor of production is bound to diminish—in the same way that the role of horses in agricultural production was first diminished and then eliminated by the introduction of tractors.” Horses did not get new jobs; they were retired; and AI will do the same for cognitive labor (see also Leontief, 1983b, for an extended version of the analogy) (see Figure 4.3). Samuelson (1989) and Imas (2026a) proved formally that when cheap machines substitute for human labor at below-subsistence cost, competitive equilibrium can eliminate human employment entirely and even reduce total output.⁴ However, Ricardo’s output-decline prediction rests on the assumption that displaced workers simply vanish (population shrinks to restore subsistence wages). In the modern scenario, displaced workers remain, and CI maintains their purchasing power (Section 3.2), so total output is sustained or increased by AI while the distribution problem is solved by redistribution rather than demographic collapse.

The argument that human labor is infinitely adaptable ignores that ‘capability’ does not imply ‘economic viability’. By definition, full AGI+Robotics will be able to do virtually all human jobs, and at a fraction of the cost.

⁴Samuelson’s proof uses a classical Sraffa-type model with Cobb–Douglas technology: when “horses” (machines) that do exactly the work of one human can be had for five-sixths of subsistence cost, the new steady-state equilibrium contains zero human workers and national product falls by 40%. As Samuelson put it: “What could be more neoclassical than Cobb–Douglas marginal productivity, and what could be more devastating for human welfare than the genocide its equilibrium brought about?”




Figure 4.3: The Leontief horse analogy: US horse and mule population peaked at ≈ 26 million around 1915 and collapsed as tractors replaced animal labor. *Leontief (1983a)* warned that human cognitive labor may follow a similarly irreversible trajectory once machines surpass human capabilities. Red dashed line shows a stylized (non-predictive) projection. Data: USDA/Census historical abstracts.

Keeping full employment in this future requires some political and financial gymnastics that, even if possible, are not actually desirable (as argued in Section 4.5). Moreover, even before full substitution is reached, *Krier (2026)* argues that *interface friction* may force a ‘corner solution’: for an AI system to utilize a human, it must export context, wait for biological processing, and re-import the result. As AI-native workflows accelerate, this round-trip latency tax can exceed the value humans add, causing firms to redesign processes to exclude humans entirely, not because AI is technically superior at every sub-task, but because the fixed cost of human integration is no longer justified. *Susskind (2020)* (Chp.4) and *de Vries (2023)* argue that AI adoption will follow a bounded S-curve restricted by energy and regulatory limits, serving as a complement to labor rather than a total substitute, thereby avoiding a collapse in aggregate demand (see Figure 4.4, the multi-constraint scenario in Figure 4.5). The restrictions may slow down and ease the transition but will not prevent the long-term decoupling of human labor from productivity.

Automated corporations. Even complex collaborative tasks requiring organizational intelligence will not remain a human stronghold. As highlighted in (*Genewein et al., 2026*), AGI agents could form coherent ‘Group Agents’ or fully automated corporations that operate with emergent cognitive capabilities far superior to human organizations. By efficiently delegating tasks through a cognitive division of labor, these virtual agent economies can effectively automate entire companies, bypassing the bottlenecks of individual human capabilities.




Figure 4.4: Bounded S-Curve of AI Adoption: Growth may follow a logistic curve, limited by physical or regulatory constraints (red dotted line).




Figure 4.5: S-Curve with Multiple Constraints: AI adoption curves limited by different factors (hardware, energy, regulation) compared to an unconstrained exponential path.




Figure 4.6: Average annual working hours per worker, 1870–2023. Hours declined by 40–54% over 150 years; the trend stalled in the US since the 1980s (2023 Penn World Table 11.0, Feenstra et al. 2015). AGI can restart and accelerate this secular decline (dotted lines).

4.2 What “Jobs” Will Remain?

Even if the new-jobs mantra fails at scale, some categories of human labor will persist longer than others. The following paragraphs survey the most commonly cited candidates; the next section examines why even these niches erode over time. Moreover, most of the unpaid activities listed in Section 2.2 are performed voluntarily, so even if they can eventually be automated, we will likely choose to continue doing the enjoyable ones such as playing competitive games and sports, making art, volunteering, etc. ourselves.

Human-made and artisanal. There has always been and will likely remain a niche demand for human-made products and human-provided services (see Table 4.8, clusters 2–3).

Just as photography did not replace (even photorealistic) painting, and people pay premiums for hand-woven textiles, consumers might pay for ‘human-crafted’ code or literature purely for the vanity, aesthetic, or status of its human origin, a “[provenance premium](#)”⁵ rooted in the social meaning of human effort ([Imas, 2026b](#); [Susskind, 2020](#)). Under non-homothetic preferences, rising incomes shift demand toward human status goods, bidding up wages in these sectors ([Comin et al., 2021](#)). Once everyone owns a household robot, some rich people will show off by hiring a human butler and maid, or feel a moral obligation or social pressure to pay for a human nurse rather than a robot nurse for their ailing parents, even if much more expensive and functionally worse. However, this income effect only protects a narrow luxury niche of highly differentiated labor ([Korinek and Suh, 2024](#)). For the vast majority, labor displacement collapses demand, rendering their services redundant as most consumers favor abundant, automated alternatives. Even if prices halve and/or incomes rise, consumers may simply buy more and better machine-made goods rather than human ones, leaving the mass of displaced labor redundant.

Sports, performance, and content creation. While watching robotic competitions will increase in popularity, watching humans compete or excel in sports, art, or chess will also most likely remain popular (see Table 4.8, cluster 2). One could charge entrance fees for these events, then the performance would count as a human job. One could also offer such events

⁵The term “provenance premium” refers to the extra price consumers willingly pay for goods or services whose value derives partly from verifiable human origin, analogous to how art collectors pay vastly more for an authenticated original than for a perfect reproduction. [Susskind \(2020\)](#) illustrates the point: diners in a fine restaurant feel shortchanged if their coffee was made by a capsule machine rather than by a trained barista, even when capsule coffee is preferred in blind tests.

free of charge, then it would just be a social activity. Presumably (only) the most popular events can and will charge as a way to get demand in line with supply (e.g. seats in a stadium). Streaming and content creation extend this logic: audiences pay for parasocial connection with a *person*, not for objectively optimal content.

Price competition. Full employment is technically possible even in a post-AGI economy. As long as a person produces goods or services which are desired by others, she could offer them on the market. This has two limitations: (a) many modern products (e.g. GPUs) can only be produced by machines and AI, and the fraction of such products is likely to increase. (b) for those offerings that can be hand-made, the price has to undercut the (close-to-zero) machine-made one (see caveat above) so this would not result in any meaningful income (Cunningham, 2025; Korinek and Suh, 2024).

Human services. Brynjolfsson and McAfee (2015) counter that unlike horses, humans possess agency, political power, the capacity for capital ownership, and the political leverage to intervene through targeted legislation like minimum wage laws. Vollrath (2015) sharpens this rebuttal with the pithy formulation “humans are not horses; they are apes”: horses had no political voice and no consumption demand, whereas humans have both. However, having political leverage to *mandate* employment does not guarantee that the resulting jobs are productive (see Section 4.5). Unfortunately, this would only create jobs for the purpose of keeping people busy rather than for real economic need, establishing a coercive, state-mandated regime of workfare and ‘Universal Basic Jobs’ (UBJ) schemes (Standing, 2017) discussed in Section 4.5 (see Table 4.8 for a quantitative assessment). They also suggest aggressive taxation, which is necessary, but does not by itself create any jobs.

Unlike horses, humans are a deeply social species with economic desires that explicitly and continuously demand interpersonal connection. People willingly pay significant premiums for human expression, hospitality, motivation, inspiration, and empathy. This deep-seated interpersonal demand establishes a structural floor for human labor that silicon-based machines cannot easily eradicate. However, Atkinson (2015, pp.117–118) warns that automation creates a learning-by-doing ratchet: replacing human service with machines today makes further replacement more likely tomorrow, progressively eroding this floor. Imas (2026b) and Comin et al. (2021) argue that this last bastion of human labor will increase in size and fully compensate for automated jobs, ensuring full employment indefinitely, though Comin et al.’s model *assumes* full employment by construction ($L_A + L_M + L_S = L$) rather than deriving it.

Unpaid social interaction is not a job. The preceding paragraphs established that human-made goods and services will survive as a niche. However, *unpaid* social interaction that is mutually enjoyed, like hanging out with friends, is traditionally not regarded as a job (see Table 4.8, cluster 8). Under the property-cluster characterization of Table 2.2, such symmetric leisure fails the asymmetry ([LA]), opportunity-cost ([LO]), and market-embeddedness ([LM]) criteria. One can envision a future where ‘friends’ pay each other for such ‘services’ and call it a job, but this is just relabeling leisure time: introducing circular payments does not create genuine economic activity. For it to be a ‘proper’ job, there should be a non-trivial asymmetry: one person desiring the interaction, while the other is only willing to engage for pay. Historical and contemporary examples abound: the court jester entertained a patron who paid for wit and diversion, and escorts provide companionship or intimacy as a commercial service; both are unambiguously jobs under any reasonable characterization. More broadly, under UBI, symmetric voluntary enjoyable interactions increasingly fail [LC]/[LL] as well, sliding from ‘jobs’ to ‘social leisure’ on the spectrum of Table 2.3, thereby expanding the domain of meaningful, unremunerated work outside formal employment (Van Parijs and Vanderborght, 2017). Only genuinely asymmetric social services (therapy, coaching, elder care, court jesters, escorts; cluster 1 in Table 4.8) retain enough job-like properties to qualify. For this to sustain full-time employment, on average each human must be willing to pay for such labor for, say, 40 hours per week, *and* at the same time be willing to work for 40 hours per week for the pay offered. While precise numbers are culturally dependent (e.g. US vs Europe) and many high-income earners fall into the ‘luxury trap’ (Frank, 2007), there is an increasing fraction of workers, especially since Covid, who would prefer greater flexibility (e.g. remote work, valued at roughly 5% of salary) or more leisure time over higher pay (Aksoy et al., 2022; Aguiar and Hurst, 2007).

4.3 A Taxonomy of Post-ASI Human Activities

Table 4.7 catalogues the human activities most likely to persist beyond ASI, organized into eight clusters. While Table 3.3 asks which jobs *disappear*, this taxonomy asks what *remains* and why. The answer is sobering for defenders of full employment: most surviving activities are not traditional jobs but leisure, voluntary work, or niche gigs (many are currently unpaid; see Section 2.2). The preceding paragraphs discuss the most commonly cited candidates: artisanal goods (cluster 3), performative activities (cluster 2), human services (cluster 1), and symmetric social interaction (cluster 8).

Table 4.7: Post-ASI human activities: detailed clusters. Each cluster groups activities by their primary survival mechanism (see Table 4.8 for quantitative summary). Activities range from formal jobs to pure leisure; many that survive are currently unpaid (see Section 2.2 for the taxonomy of unpaid work).

Activity	Why human element persists
1. Relational Care (<i>asymmetric, interpersonal</i>) These survive because genuine human empathy and connection are valued <i>and</i> the relationship is asymmetric (provider → recipient).	
Therapist / counselor	Empathic presence; trust; moral legitimacy
Elder companion / caregiver	Emotional bond; dignity; moral norm against robot care for parents
Childcare / nanny	Parental trust; child development; moral limits
Life coach / mentor	Aspirational identification; accountability from a <i>person</i>
Chaplain / spiritual guide	Existential meaning; religious authority
Companion / Sex worker	Intimacy; hard to automate socially
2. Performative / Spectator (<i>audience preference</i>) Survive because audiences want to watch <i>humans</i> compete or create, not just the objectively best output (see Section 4.2, Section 3.9).	
Professional athlete	Audience identifies with human struggle and triumph
Live performer (music, theater, comedy)	Authenticity; liveness; unpredictability
Competitive gamer / e-sports	Human competition is the point
Streamer / content creator	Parasocial relationship with a <i>person</i>
3. Artisanal / Provenance Premium (<i>status goods</i>) Survive as status goods (the Veblen/positional logic): “I can afford a <i>human</i> to do this” (see Section 4.2, Section 3.9).	
Bespoke tailor / craftsperson	Provenance premium; luxury status
Artisanal chef / sommelier	Human curation; hospitality ritual
Human butler / concierge	Ultimate status good
Hand-crafted furniture / pottery	Authenticity; imperfection-as-aesthetic
4. Governance / Moral Authority (<i>moral accountability</i>) Survive because society demands human accountability, not because humans are better at the task (see Section 4.4, moral limits).	
Judge / jury member	Democratic legitimacy; moral accountability
Elected official / diplomat	Representation; democratic mandate
AI safety oversight / alignment	The “last human in the loop”; recursive trust problem
Military command (rules of engagem.)	Moral responsibility for lethal force
5. Creative Expression (<i>self-preference</i>) Survive because people <i>want</i> to do them, not because there is market demand. Mostly leisure lacking [LC][LA],[LO], but some become gigs (see Section 2.2, creative category).	
Visual art / painting / sculpture	Self-expression; flow; identity
Music (playing, composing)	Joy of creation; social bonding
Writing (fiction, poetry, blogging)	Self-expression; communication of lived experience
Open-source / hobby coding	Intellectual challenge; community
Gardening / urban farming	Embodied satisfaction; connection to nature
Cooking (for family/friends)	Love; cultural transmission; sensory pleasure
6. Education / Mentoring (<i>relational+moral</i>) Teaching survives partly for relational reasons (students want <i>human</i> teachers) and partly for moral ones (society trusts humans to shape young minds). Knowledge transmission shifts heavily toward AI tutoring (see Section 4.2, Table 3.3 row 15).	
Teacher (K-12, inspirational)	Role model; socialization; moral formation
University professor (seminar)	Socratic dialogue; intellectual mentorship
Tutor (1-on-1 academic)	Personalized attention; motivational
Sports coach / personal trainer	Motivation; physical co-presence; accountability
7. Physical / Embodied (<i>self-preference</i>) People want to do these <i>themselves</i> even if robots could do them better or faster. Not jobs: leisure activities with physical embodiment (see Section 2.2).	
Competitive sports (amateur)	Joy of competition; fitness; social bonding
Hiking / adventure / exploration	Embodied experience; cannot be delegated
DIY / home repair (hobby)	Satisfaction of building; self-sufficiency
Martial arts / dance	Physical mastery; cultural practice
8. Social / Communal (<i>symmetric → not jobs</i>) Symmetric leisure (mutual enjoyment, no asymmetry) fails job criteria [LA], [LO], [LM] (see Section 4.2). Included for completeness.	
Socializing / hanging out	Intrinsically human; not delegable
Community organizing	Civic engagement; belonging
Religious / cultural ceremony	Meaning; community; tradition
Parenting (beyond childcare)	Love; identity; not a “service”

Table 4.8: Post-ASI human activities: composite summary. **PMSR:** P = preference limits (consumers want a human), M = moral limits (society demands a human), S = self-preference (people want to do it themselves), R = relational demand (genuine interpersonal connection required). **Job-ness:** J = job, G = gig, V = voluntary, L = leisure; arrows (→) indicate post-ASI transition direction. * = elite few only, † = niche market, ‡ = mass participation. **Pay:** 0 = unpaid, \$ = token, \$\$ = low, \$\$\$ = mid, \$\$\$\$ = high, \$\$\$\$\$ = elite. **Job%:** weighted share of the 11 job properties from Table 2.2 satisfied post-ASI (see text for details). **Erodes:** susceptibility to the ‘too-good problem’ (Section 4.4); ✓ = yes, ~ = slowly, ✗ = no.

#	Cluster	Representative activities	Survival mechanism	PMSR	Now	@ASI	Job %	Pay	Erodes?
1.	Relational care	Therapy, eldercare, childcare, coaching, spiritual counsel	Empathy; trust; moral norm	PRM	J	J → G	38–75	\$\$–\$\$\$\$	~
2.	Performative	Pro sports, live music/theater, e-sports, streaming	Audience wants <i>human</i> competition/creation	P	J → G	J* / L†	52–95	0–\$\$\$\$\$	✗ / ~
3.	Artisanal luxury	Bespoke craft, human chef, butler, hand-made goods	Provenance premium; status/Veblen good	P	J	G†	35–58	\$\$–\$\$\$\$\$	✓
4.	Governance	Judge, elected official, AI safety, military command	Democratic legitimacy; moral accountability	M	J	J	95–100	\$\$\$\$–\$\$\$\$\$	~ / ✗
5.	Creative	Art, music, writing, coding, gardening, cooking	Self-expression; flow; joy of creation	S	L → G	L → G	9–20	0–\$\$	✗
6.	Education	Teacher, professor, tutor, sports coach	Role model; socialization; mentorship	PRM	J	G → V	29–63	\$–\$\$\$\$	✓
7.	Physical	Amateur sports, hiking, DIY, dance, martial arts	Embodied experience; not delegable	S	L	L	0–9	0	✗
8.	Social	Socializing, community, parenting, ceremony	Intrinsically human; symmetric	S	L/V	L/V	0–37	0	✗

The remaining clusters—governance (4, treated in Section 4.4), creative expression (5), education (6), and physical/embodied activities (7)—are primarily self-preference or leisure and are discussed in Section 2.2.

Four survival mechanisms (PMSR). Each activity persists through one or more of four survival mechanisms, coded **PMSR** in Table 4.8. *Preference limits* (P): consumers prefer a human provider: a “provenance premium” rooted in the social meaning of human effort (Imas, 2026b) or in positional competition (Frank, 2007). *Moral limits* (M): society demands human accountability, not because humans are more competent but because legitimacy, democratic mandate, or the irreversibility of lethal decisions requires a human in the loop (Susskind, 2020). *Self-preference* (S): people want to do it themselves: they derive meaning, flow, or physical pleasure from the activity regardless of whether a machine could do it better. *Relational demand* (R): the activity requires genuine interpersonal connection with an asymmetry between provider and recipient, satisfying properties [LA] and [LO] from Table 2.2. The first three align with the limit taxonomy of Susskind (2020, 2025) (see Section 4.4); self-preference (S) is the only mechanism immune to the “too-good problem” of Section 4.4: one cannot outcompete a person’s desire to paint, play music, or garden.

Job-ness transitions. The composite Table 4.8 quantifies each cluster along several dimensions. The **Now** and **Post-ASI** columns track job-ness using the labels J (job), G (gig), V (voluntary), and L (leisure) from the spectrum of Table 2.3. The dominant transition pattern is J → G or J → V: most current jobs degrade to gigs or voluntary activities as automation advances and UBI weakens properties [LC], [LL], [LO], and [LI]. Conversely, some creative activities experience L → G: an amateur painter may earn a small provenance premium in a world saturated with machine-generated art.

Job degree (weighted property score). The **Job%** column reports a weighted average of the 11 job properties from Table 2.2 satisfied post-ASI, where the importance weights (1–10) from that table serve as coefficients, and uncertain properties count at 50%. A score above 70% broadly corresponds to a recognizable “job”; below 30% the activity retains too few job-like properties to qualify under the definition of Section 2.2. The pattern is stark: only governance and elite performative roles score above 90%; the vast majority of surviving activity *types* score below 40%, confirming that they are leisure or voluntary work, not jobs in any traditional sense (though the single largest cluster by employment volume, relational care, spans 38–75%).

Erosion and the self-preference floor. The *Erodes?* column indicates whether the cluster is susceptible to the “too-good problem” (Section 4.4): as

machine quality improves, moral limits soften when AI demonstrably outperforms humans (Section 4.4), and preference limits weaken when machine output becomes indistinguishable (Section 4.4). Only self-preference (S) is immune: wanting to cook for one's children or hike a mountain does not erode when robots can do it better. The single most telling observation is perhaps that *the only durable floor for human activity post-ASI is leisure, not employment*.

Pay distribution: power law, not bell curve. The *Pay* column reveals a power-law distribution: a few elite performers earn extreme incomes (\$\$\$\$), governance roles retain high salaries, but the vast majority of surviving human activities generate zero or token income: the mirror image of today's roughly bell-shaped wage distribution. This distribution by itself argues against "full employment through social jobs" (Section 4.2): the activities that survive produce meaning and connection, not macroeconomic output, and expecting them to sustain 40-hour workweeks at livable wages is a category error.

The participation-employment gap. The question of *how many* people each cluster can sustain as paid employment sharpens the demand-insufficiency argument of Section 4.4. Two quantities must be distinguished: participation volume (how many people *engage* in an activity) and paid employment volume (how many hold *full-time-equivalent (FTE) positions* earning a livelihood from it). The gap between them is enormous and grows post-ASI. Billions will create art, play sports, cook, and socialize, but almost none will earn income from doing so. The paradigm case is creative expression (cluster 5): perhaps 2–4 billion people paint, write, or garden, yet only 50–500 K may earn livably from it post-ASI.

FTE employment by cluster. All estimates below are order-of-magnitude ranges for global FTE positions (not gig hours or occasional income) once ASI+SR have diffused into the economy; actual numbers will depend on cultural norms and UBI generosity. Relational care (cluster 1) is the only cluster with a large *paid* employment base: 20–200 M FTE, accounting for roughly 80–95% of all remaining paid human labor. The wide range reflects uncertainty about whether demand for human empathy will grow (as [Imas 2026b](#) argue) or shrink (as AI handles increasingly complex emotional support). Every other cluster contributes marginally: governance 1–10 M FTE (structurally capped by the number of polities and courts), education 1–20 M FTE (down from ~70 M today as AI tutoring dominates knowledge transmission), artisanal luxury 0.5–5 M FTE (depending on the resilience of [provenance premiums](#)), and elite performative roles 50–500 K FTE (extreme

power-law: a few thousand earn fortunes, the rest earn nothing). Clusters 5–8 (creative, physical, social) generate effectively zero FTE positions: they are leisure by definition. Even summing optimistically, total post-ASI paid human employment may reach only 23–235 M out of ~5 billion working-age adults: a 0.5–5% employment rate, over an order of magnitude below any historical notion of “full employment.”

Inverse volume–quality pattern. Two further patterns sharpen this picture. First, the clusters with the highest Job% scores (governance at 95–100%, elite performative at 52–95%) employ the *fewest* people: governance roles are structurally scarce (every polity needs only a handful of judges and legislators), and elite performance is winner-take-all by definition. Second, relational care, the “last bastion” of [Imas \(2026b\)](#) and [Comin et al. \(2021\)](#), is the only cluster where paid volume is substantial, yet even here the trajectory is J → G: hours per person decline as AI handles routine cases, converting full-time positions into intermittent gigs (Section 4.4). The surviving activities are mass-participation leisure, not mass-employment jobs.

Declining working hours. If even the most generous estimates yield only 5% full-time employment, then maintaining broad participation in paid work requires radically shorter hours per person. So working hours should steadily decline. This trend, which saw weekly hours drop from up to 80 at the end of the 18th century to under 40 in the 1980s, could at last resume (see Figure 4.6, Table 2.8, and Section 5.1 on labor redistribution) first back to hunter-gatherer levels, then [Ferriss’ \(2007\)](#) 4-hour work-week, and eventually even less. Once a comfortable lifestyle is reached, a significant fraction of the population would not be willing to work for further money. In any case, a society where all work is truly voluntary and enjoyable, while not a jobless society as predicted in Section 4.1, would nevertheless be Utopia, and maybe even more so ([Danaher, 2019](#); [Srnicek and Williams, 2015](#)).

4.4 Why Remaining Jobs Will Erode

The preceding examples (artisanal niches, price competition, human services) can be organized under a unified taxonomy. [Susskind \(2020, 2025\)](#) identifies three classes of limits on automation.⁶

⁶[Susskind’s](#) central thesis is that “the world of work comes to an end not with a bang, but a withering: a withering in the demand for the work of human beings, as the substituting force gradually overruns the complementing force.” The three limits below constitute the last defenses against this withering.

Limits on automation. *Moral limits:* society may judge it ethically impermissible to delegate certain decisions to machines; sentencing in criminal courts, life-or-death triage in medicine, and, ironically, the alignment and oversight of AI systems themselves are paradigmatic cases where human judgment is demanded not because it is more accurate, but because accountability, legitimacy, and the expression of communal values require a human in the loop. *Preference limits:* consumers may simply prefer human-provided goods and services; hand-thrown pottery, live musical performance, a human therapist’s empathic presence carry a ‘provenance premium’ rooted in the social meaning of human effort (Imas, 2026b). *General-equilibrium limits:* as machines displace workers in some sectors, labor floods into others, depressing wages there until human labor becomes cheap enough to remain competitive on price alone (Susskind, 2025; Restrepo, 2025). Imas (2026b) formalizes the optimistic version via the structural-change model of Comin et al. (2021): as commodity sectors are automated, expenditure shares shift toward a ‘relational sector’ of inherently human-intensive, interpersonal services (teaching, caregiving, coaching, hospitality) where demand grows precisely because everything else becomes cheap.

However, these residual categories face two serious objections that together cast doubt on any ‘full-employment-through-social-jobs’ equilibrium.

Problem 1: wage erosion. Even where preference or moral limits preserve a human role, the wage that role commands is set by competition with near-zero-marginal-cost machine alternatives. As AGI quality improves, the premium consumers are willing to pay for ‘human-made’ shrinks: a hand-knit scarf competes not with a mediocre factory product but with a machine-made one of equal or superior quality at a fraction of the price. Cunningham (2025) and Korinek and Suh (2024) show that in the limit, human wages converge toward the resource cost of running the equivalent machine, i.e. effectively zero. Strictly, wages fall only to the *subsistence wage*, the floor below which labor cannot be sustained. Without UBI, the reservation wage (the lowest wage a worker will accept) collapses to subsistence; with UBI, it rises above it, as discussed in Section 6.2. At subsistence, workers are simply replaced, so the effective wage is zero. We term this convergence the ‘too-good problem’ (cf. Susskind, 2025): the better AI becomes, the weaker all three limits become: *moral* objections soften and with them legal barriers weaken when machines demonstrably outperform humans; *preferences* shift when machine output is indistinguishable from or superior to human output; and *general-equilibrium* adjustments push wages below subsistence. Competitive dynamics compound the effect: if one jurisdiction permits AI judges or autonomous medical triage and gains efficiency, others

face pressure to follow, a coordination failure that Alexander (2014) terms the ‘Moloch’ trap.⁷

Problem 2: demand insufficiency. Even granting that a niche for human labor persists, the aggregate demand for such labor is unlikely to sustain anything resembling full employment. The relational sector may grow in expenditure share, but if the absolute number of person-hours demanded falls short of what the working-age population can supply, the result is structural underemployment (Restrepo, 2025).

The trajectory of job loss. In short, while moral, preference, and general-equilibrium forces may slow the displacement of human labor, none constitutes a durable floor (see the *Erodes?* column in Table 4.8); the trajectory points toward a world in which work is voluntary (Section 10.6), not because policy makes it so, but because machines have rendered the economic contribution of most human labor negligible. The property-cluster analysis of Table 2.2 captures this trajectory systematically. Furthermore, as UBI weakens properties [LC] (compensation), [LL] (livable income), [LO] (opportunity cost), and [LI] (involuntariness), activities that currently qualify as jobs will progressively lose that status: the technological displacement described above is the primary driver, and UBI reinforces it by removing the economic compulsion that currently props up marginal employment. For an extensive catalogue of arguments, see Appendix A.2.

4.5 BS-Jobs

It is actually possible to ensure full employment even in the advent of AGI+robotics, but full employment is highly undesirable if the majority of jobs have questionable or negative social value. Without adequate income support, workers will crowd into whatever niches remain, a utopia for employers and dystopia for workers. Another class of jobs that ought to be avoided are lucrative jobs that are wasteful or harmful for society. The following paragraphs catalogue the damage: Graeber-style BS-jobs, negative-sum marketing arms races, the annoyance economy, predatory financial engineering, and folly products (cf. the working definition in Section 2.2).

Which jobs are BS? Graeber (2018) argues that advanced capitalist economies are *already* riddled with meaningless employment (adminis-

⁷The Moloch metaphor, drawn from Allen Ginsberg’s poem *Howl*, captures multi-polar coordination failures: “In some competition optimizing for X, the opportunity arises to throw some other value under the bus for improved X. Those who take it prosper. Those who don’t take it die out. Eventually, everyone’s relative status is about the same as before, but everyone’s absolute status is worse than before” (Alexander, 2014).

trative, managerial, and clerical roles) that contribute little of tangible value to society,⁸ consistent with many employees viewing their jobs as socially useless (Walo, 2023) (Tables 4.9 and 4.13). Soffia et al. (2022) counter-argue that feelings of pointlessness are better explained by the classical concept of ‘alienation’, driven by poor management practices and toxic workplace environments, rather than being objectively useless. In any case, the point here is not whether existing jobs are useless, but the extent to which useless future jobs will be created. The property-cluster characterization of Table 2.2 makes the critique precise: Graeber’s taxonomy captures *pointless* work, but several categories discussed below, excessive marketing, the annoyance economy, predatory finance, folly products, go further: they are not merely useless but actively harmful, imposing negative externalities on consumers and society. Table 4.9 extends Graeber’s original “BullShit” job taxonomy of five categories with four proposed additions of jobs that are *Bad for Society* that capture these negative-sum and predatory employment types, jointly denoted as BS-jobs.

Are BS-jobs jobs? The classification of BS-jobs as ‘jobs’ despite failing the heavily weighted property [LN] (net benefit to others) admits several treatments.

- (i) *Weaken the definition:* removing “of value to others” from the working definition (Section 2.2) would include BS-jobs trivially but also admit too many pointless activities, eroding the definition’s discriminatory power.
- (ii) *Exclude BS-jobs:* accepting that [LN]-failure disqualifies BS-jobs is logically clean but contradicts universal usage; an economy of 100% BS-jobs would paradoxically count as ‘job-free.’
- (iii) *Private value:* BS-jobs *do* provide services valued by the *employer*: a flunky makes the boss feel important; a box-ticker satisfies a compliance requirement; an arms-racer protects market share. Under this reading, [LN] is satisfied at the private level, and the BS-job’s defining quality consists precisely in the disconnect between private and social value.
- (iv) *Weighted score threshold:* the cluster-concept analysis provides its own resolution. The score row of Table 2.2 computes weighted scores for

⁸Graeber’s working definition: “a **bullshit job** is a form of paid employment that is so completely pointless, unnecessary, or pernicious that even the employee cannot justify its existence even though, as part of the conditions of employment, the employee feels obliged to pretend that this is not the case” (Graeber, 2018, p.19). Walo (2023) found empirical support: in US survey data, over 19% of respondents perceived their own jobs as socially useless, with the rate reaching roughly 25% in finance and insurance.

Table 4.9: Taxonomy of BS-jobs: Graeber (2018)’s original five “BullShit” (BS) job categories (top) and proposed extensions (bottom) to jobs that are Bad for Society (BS) such as negative-sum and actively harmful jobs discussed in this book.

Category	Description	AI Relevance / Automation Potential
Flunkies	Jobs created to make superiors feel important (e.g. receptionists, doormen).	Low automated value, but AI avatars might reduce perceived need.
Goons	Jobs with an aggressive element, acting on behalf of employers (e.g. corporate lawyers, telemarketers).	High potential for automation (e.g. AI calling, legal document review).
Duct Tapers	Hired to fix problems that shouldn’t exist or could be fixed by better systems.	AI can refactor systems and predict failures, reducing need.
Box Tickers	Exist to allow organizations to claim they are doing something they aren’t.	Automated reporting might streamline or just generate more reports.
Taskmasters	Middle management assign tasks to others, create BS-tasks.	Algorithmic management and AI coordination can replace these roles.
Proposed extensions (this book):		
Arms Racers	Negative-sum competitive spending (e.g. excessive marketing, SEO spam, churnalism, content farming, clickbait sites). Maps to Graeber’s Goons.	AI generates ads/content at near-zero cost, escalating the arms race.
Obstructionists	Intentional friction to extract fees or deter claims (annoyance economy (Gabaix and Laibson, 2006)).	AI chatbots are the obstruction tool; elimination requires regulation.
Rent-Seekers	Predatory financial engineering, opaque derivatives, junk fees.	AI enables faster, more complex financial instruments; systemic risk grows.
Charlatans	Marketing and selling folly products (snake-oil, homeopathy, vanity goods) via manufactured needs.	AI-generated persuasion amplifies deception; detection also improves.

each edge case ($\sum w_i v_i / \sum w_i$, with $v=1$ for ✓, $v=0$ for ✗, $v=1/2$ for ?). BS-jobs score 75%, clearing the same threshold as YouTubers (75%) and companions (77%) despite failing the heavily weighted [LN]. The cluster definition is designed for precisely such cases: no single property is necessary, and a sufficiently high aggregate score compensates for individual failures.

- (v) *Institutional convention*: tax authorities, labor statistics, and employment law all classify BS-jobs as jobs, settling the matter pragmatically regardless of social value.

This book adopts treatments (iii) and (iv) jointly: BS-jobs qualify as jobs because they satisfy the employer’s revealed demand *and* because their aggregate cluster score clears the inclusion threshold. The book then argues these jobs *should not exist*.

Hyper-competition. If survival requires earning money, low-entry-barrier jobs and gigs will be overrun, driving hourly income down (ultimately to zero). As an example, in poor countries, streets are often lined with hundreds of small sales booths, often selling nearly the same items with very few sales per day. Would this number be reduced to a handful, the shops would run efficiently, with buyers still being served equally well. The hyper-competition arguably renders the majority of these jobs economically useless. Hyper-competition does not necessarily drive all salaries down, and indeed can result in very lucrative BS-jobs:

Excessive marketing. A lucrative form of such a wasteful arms race is advertising. Global marketing as a negative-sum arms race is grounded in economic game theory. Academic models identify these dynamics as ‘Red Queen games,’ where entities engage in relentless, socially wasteful arms races simply to maintain their relative market position (Baumol, 2004) (see Figure 4.10).




Figure 4.10: Red Queen Game in Marketing (Baumol, 2004): A cyclical, negative-sum arms race where companies spend massive amounts just to maintain market share.

If every company halved their marketing budget, no company would be worse off. Indeed, it is a negative-sum arms race, since marketing costs money. The financial magnitude of this arms race is staggering: The marketing budget in the USA alone is over a trillion dollars per year. Half of it would more than suffice to solve extreme poverty worldwide (Hutter, 2021). This has recently been corroborated by Mecca (2025). Global advertising spend exceeded \$1 trillion in 2025, of which digital advertising alone accounted for roughly \$750 billion. To put this in perspective, econometric modeling suggests that eradicating extreme global poverty via direct transfers would cost only approximately \$175 billion annually (Pritchett et al., 2024), highlighting a massive capital misallocation. Honest (free) user endorsements and (cheap) independent honest evaluations instead of misleading company-financed persuasion would increase overall user satisfaction (see Figure 4.11).

Military \$2240B	Aid \$211B	Poverty \$175B	AI Inv. \$149B
	Advertising \$862B		

Figure 4.11: Global annual spending comparison (2024 estimates). Eradicating extreme poverty would cost ≈\$175B/year; only a-fifth of global advertising spend, and less than 8% of military expenditure. Data: SIPRI, WARC/GroupM, UN SDG costing estimates.

This is not an argument against *all* marketing: some marketing facilitates market expansion, funds the vast majority of free digital infrastructure, and informs potential buyers who otherwise may be unaware of a desired product.

Content pollution. A closely related arms race plays out in digital content production. ‘Churnalism’ that repackages press releases as original reporting (Davies, 2008); ‘content farming’ that mass-produces low-quality articles for advertising revenue; ‘scraper journalism’ that auto-republishes others’ work; and the creation of clickbait websites optimized for engagement rather than accuracy are all negative-sum: they crowd out genuine journalism and information while imposing search and verification costs on consumers. Most Search Engine Optimization (SEO) falls into the same trap: individually rational but collectively wasteful, as every site’s gain in ranking is another’s loss (Baumol, 2004). AI sharply lowers the cost of all these activities, threatening to flood the information ecosystem with synthetic low-value content at a scale that overwhelms human curation.

AI-amplified BS-jobs. AI does not merely threaten to replace jobs; it also amplifies the creation of socially useless ones. *Fake review farms* now use LLMs to generate product reviews at industrial scale, spawning counter-industries of “review authentication” services and litigation: a pure arms race in which consumers are net losers. *Compliance theater* expands as AI-generated risk assessments, privacy policies, and ethics reports satisfy legal checkboxes without meaningfully protecting anyone, while the humans who review them rarely read past the executive summary. *Dark-pattern optimization* deploys AI to A/B test manipulative interface designs (deceptive countdown timers, hidden unsubscribe buttons, “confirm-shaming” modals) at a speed no human design team could match, turning user-hostile design into a profitable engineering discipline. Perhaps most perversely, the AI industry itself has generated a shadow workforce of millions of low-wage data labelers in Kenya, the Philippines, and India, performing the tedious annotation work that trains the very systems designed to eliminate tedious work (Gray and Suri, 2019): a job that is by construction self-eliminating (Section 4.1). Each of these examples satisfies Graeber’s BS-job criteria (Table 4.9): they are experienced as pointless by many who perform them, yet they persist because they serve someone’s financial interest.

The annoyance economy. The theoretical foundation for this pattern is Gabaix and Laibson’s (2006) model of “shrouded attributes”: firms exploit consumer myopia by hiding add-on costs, and even competition fails to eliminate this practice because educating a consumer about a competitor’s hidden fees merely teaches that consumer to exploit the scheme, creating a “curse of debiasing.” Hyper-competition and such misaligned incentives give rise to what might be called the *annoyance economy*, where corporations intentionally introduce inefficiencies, lengthy procedures, and AI-driven obstacles (like convoluted customer service chatbots) simply to extract junk fees and boost revenue at the expense of consumer time (Gioino, 2026; Gabaix and Laibson, 2006) (see Figure 4.12).




Figure 4.12: The Annoyance Economy (Gabaix and Laibson, 2006; Gioino, 2026): Artificial friction and AI barriers introduced to extract fees or reduce service costs at the expense of consumer time.

Speculative financial instruments. While even some sophisticated financial instruments have economic and hence societal value, there is an increasing number of ever-more indirect, highly leveraged, and opaque derivatives that are at best questionable and at worst detrimental for the economy. Several financial and economic crashes have been caused by such reckless gambling behavior and greed. The 2008 financial crisis, driven by collateralized debt obligations (CDOs) that few participants fully understood, is the paradigmatic example. These are even worse than harmless BS-jobs, and probably deserve their own name, maybe ‘predatory jobs and products’.

Folly products. A more ambiguous category is ‘folly’ goods (snake-oil, vanity, unhealthy, status luxury & designer products, e.g. gold flakes on food) and services (VIP lounges, homeopathy, ...). To the extent that these goods and services are genuinely desired and do not harm others, a free society should be allowed to produce and consume them, so the associated jobs are arguably useful. On the other hand, many consumers are misled by marketing campaigns that they ‘need’ such folly products and that they ‘really work’. The real problem starts if genuinely socially useful jobs dry up, and people are driven into those folly jobs by a need to earn a living rather than by free choice since they genuinely believe in the value of those jobs and products. From this perspective these jobs range from somewhere between BS-jobs and predatory.

Table 4.13: Percentage of workers who perceive their own job as ‘socially useless,’ by broad sector. Finance, business services, and public administration report the highest rates. Data: Walo (2023), based on ISSP survey.

Sector	% ‘Useless’
Finance & Insurance	≈25%
Business Services	≈22%
Public Administration	≈21%
Sales & Marketing	≈20%
Transport & Logistics	≈16%
Manufacturing	≈14%
Construction	≈10%
Education	≈7%
Healthcare	≈5%
Overall Average	≈17%

Government job guarantee. A distinct but related route to forced employment is a government Job Guarantee (JG), sometimes called *Universal Basic Jobs* (UBJ), where the state acts as employer of last resort, offering a

living-wage job to anyone willing to work (Tcherneva, 2020). Critically, JG is the *politically natural* response to mass displacement: Survey evidence confirms that job guarantees enjoy somewhat higher public support than unconditional basic income: in a randomized experiment among unemployed workers, 84% supported a job guarantee versus 62% for a basic income, with the gap persisting across all generosity levels (Lehner, 2026). The intuition that “everyone should have a job” is an easier democratic sell than “everyone should receive money for nothing” (Section 7.1). This political appeal makes BS-jobs not merely a theoretical risk but the likely *default* outcome of democratic responses to automation. JG preserves the social structure of employment and avoids the stigma of unemployment, but it faces a structural dilemma: once genuinely productive public tasks are exhausted, the program must either ration jobs (defeating the ‘guarantee’) or manufacture make-work to absorb the remaining applicants. Standing (2017, Chp.9) argues that JG inevitably devolves into paternalistic workfare once the pool of useful tasks dries up, since bureaucrats (not markets) must decide what counts as ‘useful’ work, a classification problem that grows intractable at scale.

Job guarantees in practice. The consequences of mandating employment when genuinely productive work has been exhausted range from merely wasteful to deeply oppressive, depending on the institutional and cultural context. Real-world cases already illustrate two poles of this spectrum.

India’s MGNREGA is a relatively benign case: It is the world’s largest employment guarantee covering over 50 million households, illustrating the practical failure modes. Imbert and Papp (2015) find that only 42–56% of officially recorded NREGA employment days are independently confirmed by survey data (the remainder being ghost workers and inflated muster rolls), and Muralidharan et al. (2016) document that introducing biometric smartcards reduced payment leakage by over 40%, implying massive prior diversion of funds. These pathologies (make-work, ghost employment, bureaucratic capture) already plague a JG in a developing economy with abundant productive tasks; they would only intensify in a post-AGI economy where genuine public work has been automated away, leaving bureaucrats to invent ever more fictitious employment (see also Section 6.4 and Table 6.6).

At the other extreme, the Soviet Union enshrined a constitutional “right to work” (Art. 40, 1977 Constitution) while criminalizing non-employment under anti-parasitism statutes (*tuneyadstvo*), which were used to prosecute dissidents, artists, and anyone who fell outside the state’s definition of productive labor, most famously the poet Joseph Brodsky, sentenced to five years of hard labor for “social parasitism” in 1964. The result was massive

hidden unemployment disguised as overstaffing and make-work, the very pathology that a job guarantee is supposed to prevent.

Make-work in fiction. Since a post-AGI labor market has no historical precedent, fiction provides useful thought experiments for institutional dynamics that cannot yet be observed empirically. Literary satire, in particular, has explored the consequences of guaranteed employment without productive tasks with sharper edge than policy analysis can afford. In Douglas Adams's (1980) *Hitchhiker's Guide to the Galaxy* series, the Golgafrinchans classify their population into three groups and exile the supposedly useless *B Ark* middle third (telephone sanitizers, hairdressers, management consultants) to a distant planet; the jobs are pointless but the arrangement seems harmless, even comically so. Kurt Vonnegut's (1952) *Player Piano* depicts a fully automated America where displaced workers are assigned to the Army or the Reconstruction and Reclamation Corps ("Reeks and Wrecks"), performing degrading make-work that preserves the *form* of employment while destroying its dignity. Adams's Vagon bureaucracy goes further: an entire civilization exists solely to process paperwork, enforce regulations, and (infamously) read poetry at captive audiences: a government that has become its own purpose. In the limit, Kafka's (1926) *The Castle* presents the most unnerving vision: a bureaucratic apparatus so impenetrable that the protagonist cannot even determine whether his appointment as land surveyor is real or fictitious, let alone what his job entails: work has become an absurdist ritual disconnected from any recognizable purpose. Countless other novels and movies entertain similar scenarios.

The progression from MGNREGA through Brodsky's sentence to Kafka's *Castle* traces a spectrum that any job-guarantee program must confront once productive tasks are exhausted.

In summary, clinging to full employment by forcing humans into predatory, hyper-competitive, or BS-jobs, whether generated by market competition or government mandate, is a highly undesirable outcome (see examples in Section 4.5) compared to an un(der)employed society supported by UBI, freed to pursue genuine interests. For further arguments, see Appendix A.2.

Key Insights

- The historical pattern of technology creating new human jobs will not repeat under AGI: the new jobs are also automatable.
- Four survival mechanisms (preference, moral limits, scarcity, regulation) all erode as AI capabilities increase and costs fall.
- Forcing full employment post-AGI creates a dystopia of hyper-competitive, socially wasteful BS-jobs.
- Automation progressively erodes the defining properties of jobs ([LC], [LL], [LO], [LI]) within the property-cluster framework.
- When socially useful jobs dry up, luxury and folly markets structurally transform into exploitative labor traps.
- Symmetric social interaction fails as the definition of a 'job' because it lacks asymmetry, opportunity cost, and market embedding.

Open Problems

- Quantifying the "artisanal premium": how large is the niche for human-made goods and human-provided services, and does it scale?
- Whether moral limits on AI delegation (e.g. in healthcare, justice) will harden into durable institutions or gradually soften.
- The sociology of identity: how quickly can societies decouple personal identity from employment status?

Meridian — 2034: The Levy

The city council's emergency session runs past midnight. The question is how to replace the revenue that used to flow from payroll taxes on thirty thousand jobs that no longer exist. Meridian General Hospital alone shed 1 200 positions in the last four years; the distribution center another 900. Councilwoman Torres proposes a 2.3% automation surcharge on businesses that deploy AI systems in roles previously held by human workers. "The machines produce the same output," she argues. "The money is there. It just changed hands." A corporate lobbyist calls it a "robot tax" and warns of capital flight. Giwoo, watching the livestream from her kitchen, turns to Marc: "They're arguing about where the money goes. But the money is already gone: it went from our paychecks to their shareholders." Marc nods. "The question is whether the city can claw some of it back before we burn through savings." Chaeyeong, home from university for the weekend, pulls up a spreadsheet: "If the surcharge had been in place three years ago, it would have generated \$47 million annually: enough to cover the gap in the pension fund and fund retraining vouchers for everyone the hospital and distribution center let go." Butz, reading the paper the next morning, mutters: "They should have done this when they closed the steel mill. Now they're thirty years late and calling it innovation."

Chapter 5

Redistribution Mechanisms

The art of taxation consists in so plucking the goose as to obtain the largest amount of feathers with the least possible amount of hissing.
— attributed to Jean-Baptiste Colbert (c. 1665)

5.1	Labor Redistribution	88
5.2	Wealth Redistribution	89
5.3	Taxation Strategies	90

Once we accept that [AGI](#) and robotics will progressively replace human labor, the core policy challenge is redistribution: how the shrinking pool of work is shared across the population, how the vast wealth that machines produce reaches everyone, and how to fund the transfer. Without active redistribution, the default is mass unemployment, hyper-competitive BS-jobs (Section 4.5), and extreme concentration of wealth (see Figure 7.5 and Section 7.1 for the political consequences). The first step is sharing the remaining work itself: shorter hours, earlier retirement, and extended parental leave (Section 5.1). For income, several complementary mechanisms can channel automation wealth to citizens (Section 5.2), funded by a diversified mix of taxation instruments (Section 5.3). The next chapter examines the transfer vehicles themselves, above all Universal Basic Income (Section 6.1) (cf. [Hutter and Hutter, 2021](#), Sec.4.3).

5.1 Labor Redistribution

As AI and robots absorb an increasing share of tasks, the total volume of human labor will contract. The question is whether that contraction falls on everyone gradually or on a few suddenly (see Table 5.1):

- *Working hour reduction*: Progressively lower the standard workweek (e.g. from 40 to 15 hours), spreading the remaining human labor across the population (Keynes, 1930). Total take-home pay need not fall: rather than fully automating one worker and funding the displaced salary entirely through CI, partial automation of two workers at half-time each lets each retain half their salary while CI bridges the shortfall. Mandate double-rate pay for overtime.
- *Retirement adjustment*: Substantially lower the legal retirement age to alleviate labor oversupply. Pension funds could finance this by actively investing in high-yield AI enterprises, thereby directly socializing automation gains.
- *Extended parental leave or part-time*: Offer substantially extended parental leave or part-time, potentially spanning from first birth until the youngest child reaches adolescence, to allow time-poor parents to invest more time in child-rearing. This directly improves the well-being of both children and parents while absorbing surplus labor capacity.
- *Extended paid re-education*: Offer fully funded multi-year education and retraining programs (sabbaticals, second degrees, vocational pivots) to displaced workers, absorbing surplus labor while building human capital. Unlike short retraining courses, extended programs of one to several years acknowledge that meaningful skill transitions take time, and provide income security during the transition.
- *Increased unemployment*: If no active labor redistribution is pursued, rising unemployment is the passive default. This is socially corrosive and should be avoided, but acknowledging it clarifies what the other measures above prevent.

Crucially, these reductions need not be mandated top-down. UBI itself induces them voluntarily: a per-child UBI (under parental control) lets parents afford to reduce hours or stop working; many will do so naturally once the financial pressure lifts. Earlier retirement emerges organically when workers save more from combined earnings and UBI. The UBI level becomes the policy lever that regulates the pace of voluntary labor contraction to match society's needs, without coercive work-hour legislation or forced retirement schedules.

Table 5.1: Comparison of labor redistribution mechanisms in a post-AGI economy.

Mechanism	Pros	Cons
Work Hour Reduction	Preserves jobs, directly increases leisure	Hard to implement in non-scalable roles
Retirement Adjustment	Alleviate labor oversupply, socializes AI gains	Increases dependency ratio, requires returns
Extended Parental Leave	Improves child and parent well-being; converts displaced work-hours into care-giving	Skill atrophy makes re-entry hard (alleviated by part-time); removes workers at peak productivity
Extended Paid Re-education	Builds human capital; provides income during transition; absorbs surplus labor	Expensive; skills may be obsolete before completion; limited scalability under rapid automation
Increased Unemployment	No policy action needed; market selects for most productive workers	Socially corrosive; unfair; risks permanent underclass

5.2 Wealth Redistribution

Even with optimally shared work, the income from human labor will decline. Redistributing AI-generated wealth is desirable on three grounds: fairness, since AI builds on millennia of collective scientific and cultural achievement to which everyone is heir; macroeconomic stability, since without purchasing power the consumer demand that drives production collapses; and democratic cohesion, since extreme wealth concentration corrodes the political institutions needed to govern AI safely. Several complementary mechanisms exist (see Table 5.2):

- *Universal Basic Income (UBI)*: Institute an unconditional basic allowance high enough to afford a basic living standard (grounded in “real freedom for all”; [Van Parijs, 2004](#)), funded by the taxation mechanisms discussed below (see Figures 5.7 and 5.8).
- *Universal Basic Capital (UBC)*: Mandate that all citizens receive equity in productive assets or AI corporations (Universal share ownership in a [Citizens’ Wealth Fund](#) (CWF) such as the suggested “American Equity Fund” ([Altman, 2021](#))), directly aligning public wealth with automation-driven corporate profitability (see Section 6.5 for a detailed discussion of share ownership models as GI/UBI funding mechanisms).
- *Universal Basic Services (UBS)*: Alternatively, the state could vastly ex-

pand free, high-quality public services (healthcare, housing, transport) funded by automation wealth, reducing the monetary cost of living (see Section 6.4 for details).

- *Universal Basic Jobs (UBJ)*: A government Job Guarantee (Tcherneva, 2020) that offers public employment at a living wage to everyone who wants it. While UBJ preserves the social structure of work and avoids the stigma of unemployment, it risks devolving into state-mandated BS-jobs (Section 4.5) once genuinely useful tasks are exhausted, and does not provide freedom from work.

Table 5.2: Comparison of wealth redistribution mechanisms in a post-AGI economy.

Mechanism	Pros	Cons
UBI	Simple, eliminates poverty stigma, restores agency	High headline tax rates, political resistance
UBC	Aligns citizen interests with corporate success; citizen agency via capital power	Market volatility, risk of forced selling; voting dilution
UBS	Non-inflationary, targets basic needs directly	Less individual choice, bureaucracy
UBJ	Preserves work structure, social identity, meaning	Risks state-mandated BS-jobs (Section 4.5); no freedom from work

A well-designed combination of these measures, funded by the taxation strategies described below and dynamically adjusted to the pace of automation, can ensure a peaceful transition (see Tables 5.1 and 5.2 and Figures 5.3, 2.11, 5.9 and 7.5).

5.3 Taxation Strategies

The redistribution mechanisms above must be funded. As argued in Section 3.2, even current salary levels can be fully covered by taxing the productivity gains from automation, at effectively no extra burden to anyone. Korinek and Stiglitz (2019) make the broader case that redistribution alone may be insufficient at scale, calling additionally for pre-distribution interventions such as steering AI research toward labor-augmenting applications. However, a single blunt automation tax is not the optimal instrument; a carefully designed mix of the domestic and international taxation instruments in Tables 5.5 and 8.4 will yield better outcomes by minimizing distortions, reducing avoidance incentives, and adapting to shifting tax bases (see Korinek




Figure 5.3: US household income inequality as measured by the Gini index (0 = perfect equality, 1 = maximum inequality), 1967–2022, showing a secular upward trend (+23% over 55 years). Without active redistribution, AI-driven automation risks accelerating this trend. Data: US Census Bureau.

and Lockwood, 2026, for a thorough primer on how taxation must shift from labor to consumption and in time to AGI harvesting).¹

Tax revenue is fungible. A crucial principle underpins the instrument survey that follows: how to *fund* a guaranteed income must be separated from how to *distribute* it. In public finance, this is known as the *non-affectation principle* (*Gesamtdeckungsprinzip*): all tax revenues flow into a general fund, and expenditures are allocated independently based on social priorities (Musgrave, 1959). The only binding constraint is that total revenue roughly equals total expenditure (Figure 5.4). Politicians routinely violate this principle by earmarking revenues for rhetorically linked purposes: carbon taxes are tied to renewable subsidies (even though the tax already steers behavior), Germany’s Solidaritätszuschlag persisted three decades past its reunification purpose, the US Highway Trust Fund links gasoline taxes to road spending², the Social Security “trust fund” is an accounting fiction that misleads about fiscal capacity, state lotteries earmarked for education simply displace general-fund spending, and tobacco taxes tied to healthcare create a perverse incentive to maintain smoking. This is common wisdom among economists but ignored by politicians, because “this tax funds X” is easier to sell than “we need more revenue” (see also Van Parijs and Vanderborght, 2017, Chp.6). The burden of proof should lie with earmarking, not against it.

¹Van Parijs and Vanderborght (2017, Chp.6) groups funding sources under four headings (capital, nature, money, and consumption) to which our instrument-level taxonomy adds automation, IP, digital services, financial transactions, BS-taxes, sin/self-harm taxes, and philanthropic donations, while separating ownership models into Table 6.8.

²Even when the tax approximates a user fee, the right amount of road spending has no reason to equal the right amount of gasoline tax revenue.

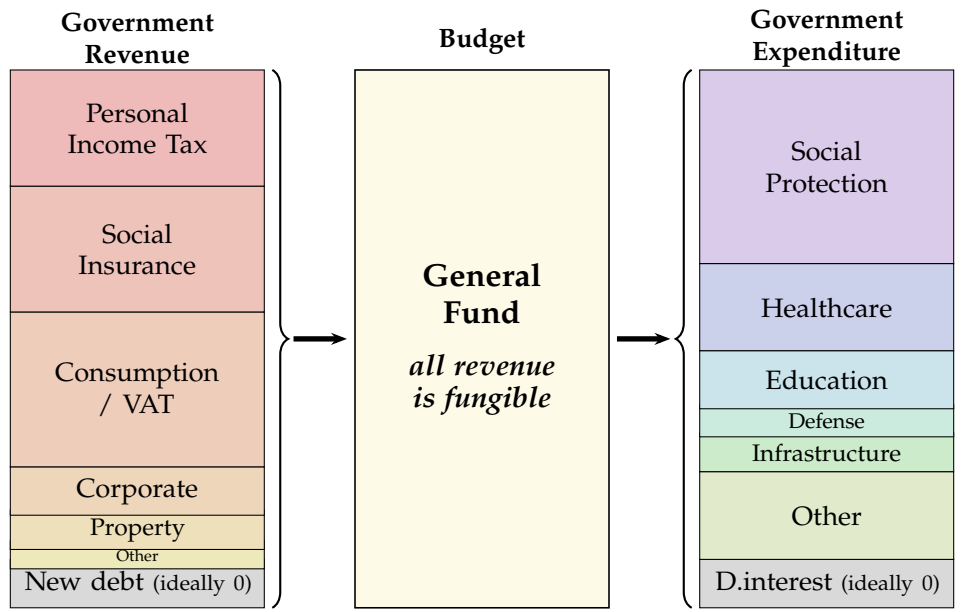


Figure 5.4: The non-affectation principle: all government revenues (left) flow into a single general fund (center), from which expenditures (right) are allocated based on social priorities. No arrow connects a specific tax to a specific spending category. In a balanced budget, total revenue equals total expenditure; ideally both new debt (left) and even better debt-service costs (right) are zero. Revenue composition roughly follows OECD averages (cf. Figure 5.8).

The reason this section surveys so many tax instruments in detail is that, while the idealized transition scenario (Section 3.2) demonstrates that automation taxes alone can fund a generous guaranteed income in the initial phase, the three-phase transition roadmap (Section 6.2) requires the tax base to shift progressively as the economy transforms.

Corporate profit and capital taxes. Progressive corporate profit taxes (Devereux et al., 2020) and capital taxes (Piketty, 2014) directly target the wealth generated by automation. Korinek and Lockwood (2026) stress the crucial distinction between taxing the *normal return* to capital (the minimum profit needed to justify the investment, taxing which deters future investment) and taxing *economic rents* (excess profits from monopoly position, data moats, or network effects, taxing which does not reduce output because the investment is already sunk); AI monopoly rents fall squarely in the latter category. As AI concentrates returns to capital, these become the natural replacement for declining labor income tax revenue (see Figure 5.7); even moderate increases in the effective corporate tax rate may suffice to fund UBI from automation profits alone (Nayebi, 2025).

Progressive personal income tax. The most established funding instrument and the backbone of current welfare states (Van Parijs and Vanderborght, 2017, Chp.6). Optimal tax theory, building on the foundational model of Mirrlees (1971),³ shows that revenue-maximizing top marginal rates are far above those currently in force. Diamond and Saez (2011) compute an optimal top rate of 73%, using empirically calibrated labor-supply elasticities that turn out to be modest. Piketty et al. (2014) push the figure to approximately 83% by decomposing the behavioral response into three channels: a small real supply-side effect, a tax-avoidance channel (addressable through enforcement), and a rent-seeking channel whereby high pay reflects bargaining power rather than productive contribution, so that taxing it away imposes no efficiency loss. Atkinson (2015) proposes a top rate of 65% as part of a broader inequality-reduction program. Cross-country evidence from 18 OECD countries (1960–2010) undermines the growth objection: nations that slashed top rates (US, UK) saw top-1% income shares explode but did not grow faster than nations that maintained high rates (Germany, Denmark) (Piketty et al., 2014); see Figures 5.9 and 7.5 for the long-run trajectories and current global snapshot. The main practical concern is talent flight, but empirical studies

³Mirrlees' 1971 paper established optimal income taxation as a formal discipline, showing that the trade-off between equity and efficiency depends critically on the labor-supply elasticity: how much people reduce their work effort in response to higher tax rates. This elasticity turns out to be empirically modest, implying that high marginal rates are less distortionary than commonly assumed.

Table 5.5: Comparison of taxation instruments for funding redistribution in a post-AGI economy (cf. Table 8.4 for international policies and arguments).

Tax Instrument	Pros	Cons
Corp. Profit Tax	Targets large-scale automation wealth directly	Encourages tax avoidance, capital mobility
Corp. Capital Tax	Directly taxes the source of automation (capital)	May discourage capital investment, hard to value
Personal Income Tax	Most established instrument; progressive rates natural	Tax base erodes as labor income vanishes; req. shift to other bases
Luxury Taxes	Highly progressive, does not penalize basics; grows with abundance	Narrow tax base, limited revenue
Pigouvian Taxes	Corrects externalities imposed on others (carbon, pollution, congestion)	Difficult to measure and price externalities
Sin = Demerit Tax / Self-Harm Tax	Internalizes moral hazard of social insurance; broad base; already widespread	Revenue self-limiting as consumption falls; scope boundaries contentious
Automation Taxes	Slows disruptive transition, funds retraining	Stifles innovation, deadweight loss
Land / Resource Tax	Immovable base; no deadweight loss; captures unearned rents	Politically difficult; hard to value; alone insufficient for UBI
IP Reform / Tax	Captures intangible monopoly rents; Harberger mechanism applicable	Mobile across jurisdictions; risks deterring innovation
Wealth / Inheritance	Targets accum. dynastic wealth; counters $r > g$ (Piketty, 2014)	Valuation difficult; capital flight; political resistance
Consumption / VAT	Broad stable base even as labor income vanishes; hard to evade	Regressive without UBI offsets; zero-rates erode base
Digital Services Tax	Captures local user-generated value; implementable unilaterally	Sector-specific; distortionary; obsoleted by broader reforms
BS-Tax (social waste tax)	Discourages negative-sum activities (Table 4.9); doubles as revenue source	Scope subjective; classification contentious; risks overreach
Financial Transaction Tax	Vast base (all electronic transfers); stabilizes speculative markets; rate (0.1–0.5%) below credit-card fees	Widens bid-ask spreads (reduced market-making); invites avoidance via derivatives
Charity / Donations	Voluntary; no economic distortion; Giving Pledge could scale with AI wealth	Orders of magnitude too small; unreliable; donors set priorities, not recipients
Seigniorage (money creation)	Non-distortionary if matching real GDP growth; stimulates demand in deflationary traps	Inflationary beyond narrow band; erodes monetary credibility; not sustainable at UBI scale




Figure 5.6: Radar comparison of twelve tax instruments across eight policy-relevant dimensions. Each axis runs from 0 (center) to 5 (edge). Values represent qualitative, indicative assessments for a European policy context using the maximal realistic tax rates, and are intended solely for illustrative purposes. Serious policy design requires empirical modeling beyond this book’s scope. Key trade-offs include the erosion of traditional income taxes as the labor share declines; the high efficiency and evasion resistance of land value tax; the mobile base of wealth, financial transaction, and IP taxes, which are subject to capital flight and tax competition (*Avi-Yonah, 2000; Akcigit et al., 2016*) but can be mitigated by international coordination and destination-based taxation (*Alstadsæter et al., 2024; Auerbach et al., 2017*); and the high revenue but regressive nature of consumption taxes. Automation-Robustness highlights the resilience of the tax base to full automation. See *Abbott and Bogenschneider (2018)* and *Atkinson (2015)* for more thorough discussions of technology-targeted taxation and redistribution frameworks.




Figure 5.7: Taxation Sources Shift: Stylized projected decline in labor income tax revenue and the rise of capital and AI-driven corporate taxes as primary revenue sources.




Figure 5.8: Tax revenue composition shift across three eras. **Now** (OECD average 2022, source: OECD Revenue Statistics via Tax Foundation): income (24%, blue), consumption (32%, orange), social insurance (26%, purple), corporate (10%, red), property (6%, green), other (2%, gray). **Transition** and **Post-scarcity** are qualitative scenarios: AI/robot taxes and LVT gradually replace income and social insurance taxes as the labor-income base erodes; carbon/resource taxes grow; novel instruments (luxury, Tobin, data taxes, etc.) gain share. See Section 5.3 and Table 5.5 for the full taxonomy and Figure 5.7 for the trend-line version.

find that domestic migration elasticities are small and actual ‘millionaire flight’ is modest (Young and Varner, 2011) (Section 8.4). As labor income erodes under automation, the personal income tax base shrinks (Figure 5.7), necessitating a gradual shift toward capital, consumption, and rent-based taxes, but during the long transition, steep top rates on residual high earners remain a potent redistributive lever.

Luxury taxes. Luxury taxes are highly progressive by design: they target positional and status goods (yachts, private jets, jewelry, high-end vehicles) while leaving basic consumption untouched. Frank (2007) argues that much luxury spending is ‘positional’ (driven by relative status rather than absolute need), generating wasteful expenditure arms races that can be taxed without reducing genuine welfare. A progressive expenditure tax, which exempts savings and taxes only consumption above a generous threshold, offers a natural implementation that is both growth-friendly and strongly redistributive (Atkinson, 2015). The main limitation is a narrow tax base: luxury goods constitute a small fraction of total consumption, yielding limited revenue relative to broad-based instruments such as VAT. However, if AGI delivers the post-scarcity abundance discussed in Section 10.1, discretionary and positional spending could expand markedly, turning luxury taxes from a niche supplement into a substantial revenue source.

Pigouvian taxes. Named after Pigou (1920),⁴ these taxes correct negative externalities that one party’s activity imposes on others: carbon emissions warming the climate, industrial pollution degrading air and water, road congestion wasting commuters’ time. The case for sharply higher carbon prices is overwhelming: the social cost of carbon is estimated at \$50–200 per ton, while most jurisdictions charge far less or nothing (Edenhofer et al., 2021). Pigouvian taxes are largely orthogonal to AGI (pollution does not depend on who produces it, human or machine) but offer a substantial untapped revenue stream alongside their primary corrective function. The main difficulty lies in measuring externalities accurately: if the rate is set too low, the externality persists; if too high, productive activity is inefficiently curtailed.

Demerit and self-harm taxes. Pigouvian taxes correct harm imposed on others; demerit (aka sin) taxes, and self-harm taxes address a distinct class of activities that primarily harm the actor, and only become social

⁴A Pigouvian tax is a levy on activities that impose costs on third parties (“externalities”). Pigou’s foundational insight was that when “one person A, in the course of rendering some service . . . incidentally also renders services or disservices to other persons,” the government can correct the divergence between private and social costs via “bounties and taxes” (Part II, Ch.IX).

costs because society collectively insures against the consequences through healthcare, disability support, and eventually UBI. Absent a safety net, these would be purely private choices: one could simply let risk-takers bear their own costs. But once society guarantees a floor, every voluntary risk becomes a claim on the common pool, creating a moral hazard that self-harm taxes internalize. Traditional sin taxes on tobacco (~\$190B global excise revenue), alcohol (~\$100B), and gambling are among the oldest and most politically popular levies (Musgrave, 1959).⁵ Sugar and soda taxes have spread rapidly since Mexico (2014) and the UK (2018), with early evidence of significant consumption reductions. The logic extends naturally to any voluntary activity with high expected social insurance costs: motorcycle riding (~20× car fatality rate per mile), extreme sports (rock climbing, skydiving, base jumping, paragliding), contact sports with cumulative injury risk (boxing, rugby, American football), tanning beds, and energy drinks. Private insurers already price this risk through smoker surcharges and extreme-sport exclusions; a sin tax socializes the same mechanism. Crucially, the tax is not paternalistic: it does not prohibit any activity but merely asks the risk-taker to pre-fund their expected draw on shared resources. The revenue paradox is that successful sin taxes reduce consumption of the taxed good, eroding their own base over time; this is a public health benefit but limits their long-term fiscal contribution. In practice, the erosion is often modest: tobacco and alcohol excise revenues have remained substantial for decades despite declining consumption, because demand is inelastic and rate increases offset volume declines. Moreover, a post-labor economy with abundant leisure (Section 10.1) may expand the base for dangerous hobbies and thrill-seeking, partly counteracting the deterrent effect.

Automation taxes. Robots/AI could be taxed directly. The idea has been heavily popularized by figures like Bill Gates, who argued that robots taking human jobs should be subjected to equivalent income taxes (Gates, 2017) to pay for the social costs of sudden labor displacement. Casas and Torres (2024) observe that AGI threatens to erode labor taxes, which constitute the primary source of state revenue. Abbott and Bogenschneider (2018) argue that current tax codes inadvertently subsidize automation over human labor, since employers pay payroll taxes on workers but not on robots, and propose leveling this asymmetry. Unfortunately, such a direct automation tax to restore tax neutrality and secure the welfare state creates **deadweight**

⁵Musgrave introduced the terms “merit goods” (whose consumption society encourages, e.g. education) and “demerit goods” (whose consumption society discourages via taxation, e.g. tobacco and gambling). His three-branch model (allocation, distribution, stabilization) remains the standard framework for fiscal policy design.

losses across the broader economy (Gasteiger and Prettnner, 2022), stifles technological progress, and leads to even lower employment levels compared to reducing baseline taxes on human labor. Preferably one would levy only a targeted transition charge on specific displacement costs, such as re-education and temporary income support for laid-off workers.

Land and resource taxes. In a zero-labor-cost economy, land and natural resources become the binding constraint on post-scarcity. A **land value tax** (LVT)⁶ (George, 1879; Stiglitz, 1977) is uniquely efficient: land is immovable, so the tax creates no deadweight loss and cannot be evaded offshore. More broadly, identifying and taxing economic rents from fixed factors (land, spectrum, natural resources) becomes increasingly valuable as labor taxation erodes (Korinek and Lockwood, 2026). Resource extraction taxes and **Harberger self-assessed taxes**⁷ (Posner and Weyl, 2018) complement LVT. The detailed case is made in the ‘land ownership problem’ discussion (Section 3.6 and the paragraph following it).

Intellectual property reform. As AI generates an increasing share of patentable inventions and copyrightable works, IP monopoly rents become a major source of concentrated wealth. Shortening patent terms, expanding compulsory licensing, and applying Harberger self-assessed taxes (Posner and Weyl, 2018) to IP portfolios can capture these rents. The key challenge is international coordination, since IP is trivially relocated across jurisdictions (see Table 3.8 and the discussion in Section 8.3).

Wealth and inheritance taxes. Target accumulated dynastic wealth and counter Piketty’s (2014) $r > g$ dynamic:⁸ when the rate of return on capital (r) exceeds the growth rate of the economy (g), wealth concentrates in the hands of capital owners over time. The urgency is underscored by the finding that

⁶A land value tax levies on the unimproved value of land (location value), not on buildings or improvements. Henry George argued in *Progress and Poverty* (1879) that land appreciation is socially created and should fund public goods via a “single tax.” Stiglitz later proved the “Henry George Theorem”: under optimal conditions, aggregate land rents equal public-goods expenditure, so a 100% LVT can finance the efficient level of public provision.

⁷Under a Harberger tax (or Common Ownership Self-Assessed Tax, COST), owners publicly declare the value of their assets and pay a percentage as tax; crucially, anyone can buy the asset at the declared price. This creates a self-enforcing mechanism: declare too low and risk losing the asset; declare too high and overpay in taxes. Posner and Weyl estimate that a 7% COST could raise roughly 20% of national income.

⁸Piketty’s central thesis is that when the rate of return on capital (r , typically 4–5%) persistently exceeds the growth rate of the economy (g , typically 1–2%), wealth mechanically concentrates among capital owners across generations. He documents this pattern across 20 countries over two centuries and advocates a progressive global capital tax as the “ideal tool” to arrest the spiral.

global billionaires currently pay effective personal tax rates of just 0–0.5% of their wealth (Alstadsæter et al., 2024).⁹ Inheritance taxes are especially justified when wealth increasingly derives from owning AI capital rather than personal effort. Practical challenges include asset valuation, capital flight, and strong political resistance.

BS-tax (social waste tax). The negative-sum and actively harmful activities catalogued in the BS-jobs discussion (Table 4.9), excessive marketing arms races, annoyance-economy rent extraction, predatory financial instruments, and folly products, impose social costs that existing Pigouvian taxes (Edenhofer et al., 2021) do not reach because they target environmental rather than economic externalities. These activities are empirically significant: Piketty et al. (2014) show that much of the behavioral response to low top tax rates reflects rent-seeking and compensation bargaining rather than productive effort, making such activities both a cause of inequality and a natural target for taxation. A BS-tax is an umbrella term for surcharges on such socially wasteful activities that are *Bad for Society*, subsuming broadened Pigouvian levies on non-environmental externalities, demerit-goods taxes (Musgrave, 1959) on harmful products, and dissipative-activity charges on zero-sum competitive waste (Tullock, 1967). While the precise scope will face interest-group opposition (as do existing luxury, sugar, and gambling taxes), the classification problem is structurally easier than for job guarantees: BS-taxation requires *identifying* existing wasteful activities, for which many cases (excessive marketing, planned obsolescence, predatory financial instruments) are clear-cut among unbiased experts; job guarantees, by contrast, require *inventing* useful jobs ex nihilo in an environment where productive work is drying up, a fundamentally harder task. Several BS-tax components are already implemented in some jurisdictions (advertising taxes in France and Hungary, financial transaction taxes, planned-obsolescence penalties), suggesting that a coherent umbrella structure is feasible.

Financial transaction tax. A *super-Tobin tax* is a micro-levy of 0.1–0.5% on every electronic money transfer, potentially even including between one’s own accounts. It was seriously discussed during the 2016 Swiss UBI referendum (Van Parijs and Vanderborght, 2017, p.154). The rate is below what credit-card networks charge merchants, yet the tax base (all electronic transactions) is enormous, giving even a tiny rate substantial

⁹The EU Tax Observatory’s 2024 report finds that offshore personal tax evasion has declined roughly threefold thanks to automatic information exchange, but corporate profit shifting remains high: \$1 trillion was booked in tax havens in 2022, costing roughly 10% of global corporate tax revenue. The report proposes a 2% minimum wealth tax on billionaires (yielding an estimated \$250B/yr).

revenue potential. A conventional Tobin tax on foreign-exchange transactions (Tobin, 1978)¹⁰ primarily aims to curb speculative volatility rather than raise revenue; even low rates sharply contract speculative volume, eroding the tax base. A broader financial transaction tax largely sidesteps this problem: the bulk of its base consists of real-economy transfers (payroll, rent, purchases) that are inelastic, while the contraction of speculative volume is a welcome stabilization side-effect rather than a revenue loss. The main obstacle is international coordination to prevent migration of financial activity to untaxed jurisdictions.

International taxes. Consumption/VAT and digital services taxes are inherently resilient to profit shifting because they tax where goods are consumed, not where they are produced. The OECD global minimum corporate tax of 15% further curtails the race to the bottom. These instruments and the broader international challenges are discussed in Chapter 8.

Philanthropic donations. Philanthropic cash transfers have demonstrated that the delivery mechanism works: GiveDirectly's 12-year RCT across 295 villages in Kenya, the largest basic income experiment to date, found that monthly transfers increased enterprise creation and consumption without reducing work effort (Banerjee et al., 2023). The Giving Pledge, in which billionaires commit to donating the majority of their wealth, has attracted over 240 signatories, and AI-driven wealth concentration could vastly enlarge this pool. However, philanthropy is structurally inadequate as a primary funding source. Total global philanthropic giving is roughly \$500 billion per year, while UBI for the United States alone would require \$3–5 trillion annually: an order-of-magnitude gap. More fundamentally, Reich (2018) argues that big philanthropy is a form of unaccountable power that undermines democratic priority-setting: donors, not recipients, choose what to fund, and charitable tax deductions mean the public effectively co-finance donor preferences. A *right* to basic income cannot depend on the goodwill of benefactors. Philanthropy thus serves best as a complement (funding pilots, building political support) rather than a substitute for the tax-based instruments above.

Untenable funding ideas. Two perennially creative suggestions deserve brief mention. First, printing fresh money: Major Douglas's "Social Credit" movement in the 1930s and Jacques Duboin's "distributist" movement in

¹⁰Tobin's original 1978 proposal aimed to "throw some sand in the wheels of our excessively efficient international money markets" by imposing a small levy on currency trades. The goal was to curb destabilizing speculation, not to raise revenue. The idea was later broadened to financial transaction taxes on all securities trades.

France both proposed distributing purchasing power directly via the printing press (Van Parijs and Vanderborght, 2017, pp.152–153). The idea resurfaces periodically as “helicopter money” or “quantitative easing for the people,” and governments have historically needed little persuasion to embrace it. Since seigniorage is the government’s favorite funding mechanism (politically painless, requiring no explicit tax vote, with costs diffused through inflation and conveniently deferred past the next election) the temptation to fund UBI this way would be considerable. However, seigniorage-funded UBI is inflationary beyond a narrow band matching real GDP growth, and the band is too narrow to fund more than a token transfer. This has not historically deterred enthusiasm: governments routinely project unprecedented GDP growth to justify the printing press, only to discover that inflation arrives well before the growth does. Second, more fancifully, is redistributing casino profits, as practiced by the Eastern Cherokee (~\$6 000/year per member (Costello et al., 2003)) and Macau’s “Wealth Partaking Scheme” (Van Parijs and Vanderborght, 2017, p.152). The model works, but only if one can locate a sufficiently large external population of willing losers, which limits scalability somewhat.




Figure 5.9: Top 1% pre-tax national income share (1900–2024) for selected countries shows a similar trend as Figure 2.11. The U-shaped pattern visible in Anglo-Saxon economies reflects declining inequality through mid-century followed by a sharp reversal since the 1980s. Data from the World Inequality Database (2024) <https://wid.world/>. See also Figure 5.10 for the regional breakdown and Figure 7.5 for a global snapshot.




Figure 5.10: Top 1% pre-tax national income share (1900–2024) by world region (PPP-weighted), complementing the country-level data in Figure 5.9. MENA and Latin America consistently exhibit the highest concentration, while Europe shows the lowest post-war inequality. Data from the World Inequality Database (2025) <https://wid.world/>.

Key Insights

- A diversified tax portfolio (automation, consumption, land, IP, Pigouvian, luxury) can finance redistribution even as the labor-income tax base erodes.
- A BS-tax can lower the risk of bullshit-job proliferation and wasteful human toil by economically disincentivizing negative-sum competitive waste.
- Labor redistribution (shorter hours, earlier retirement, extended parental leave and re-education) can delay rising unemployment but cannot substitute for income redistribution.
- All supply-chain costs in the end trace back to labor, energy, and raw materials; taxing automated production captures the value that payroll taxes used to capture.

Open Problems

- Cross-country coordination: preventing a race to the bottom in automation taxation.
- Optimal mix of tax instruments as the economy transitions through three phases of automation (Section 6.2).

Meridian — 2036: The Vote

Meridian's unemployment rate has hit 36%, and the "Dignity Through Work" program has been quietly defunded. The city council puts a referendum to voters: unconditional monthly payments to every adult, funded by a 2.5% automation surcharge. The campaign is bitter. Opponents warn of freeloading and inflation. Proponents point to Meridian General, which now runs with four human staff and two hundred autonomous systems, generating record patient outcomes while its former nurses exhaust their savings. The referendum passes 61:39. Giwoo receives her first payment of \$1 400: not enough to live on, but enough to stop the slow bleed. She can finally spend proper time with Angela, whose health has worsened. Marc uses the breathing room to set up a proper carpentry workshop in the garage. Chaeyeong, fascinated by the economics of the policy keeping her family afloat, begins drafting a paper on optimal automation-levy design. Alice's coastal restoration continues, but she no longer pretends it is a job. Judge Chloe endorses the referendum in an op-ed: "If an algorithm delivers justice more fairly than a human judge, the most just thing a judge can do is step aside and ensure the displaced are not punished for being human." Levin votes against. He is still employed, and the idea of being paid not to work offends something deep in him. But he notices the domestic-violence calls have dropped since the payments started. Butz clips the result from the newspaper and pins it to his refrigerator: "First sensible thing this city has done since they built the park."

Chapter 6

Guaranteed Income

The idea of an unconditional basic income is bound to intrigue ... those who want tomorrow's world to be a world of freedom—of real freedom, not mere formal freedom, and for all, not just for the happy few. — Van Parijs and Vanderborght (2017)

6.1	Universal Basic Income (UBI)	106
6.2	Transition from CI to UAI to UHI	109
6.3	Objections to UBI/GI	114
6.4	Alternatives to UBI/GI	119
6.5	Universal Citizen Share Ownership Models	124
6.6	Pension Fund Pathway	133

The previous chapter established how automation wealth can be redistributed (Section 5.2) and funded (Section 5.3). This chapter examines the transfer vehicles themselves. Among the available options, Universal Basic Income (UBI) stands out for its simplicity, universality, and compatibility with voluntary work (Section 6.1), although a phased transition from conditional to unconditional payments is more realistic than a big-bang introduction (Section 6.2). The standard objections (laziness, unaffordability, inflation) do not withstand scrutiny (Section 6.3), and the commonly proposed alternatives each re-introduce frictions that UBI avoids (Section 6.4). Share ownership models, often presented as a competing approach, are best understood as funding mechanisms for UBI rather than substitutes for it (Section 6.5). A particularly promising transition vehicle is the co-option of existing mandatory pension systems (such as Australia’s Superannuation) into citizen-owned wealth funds that can gradually morph into a de facto

CWI (Section 6.6).

6.1 Universal Basic Income (UBI)

Among the available transfer vehicles, UBI stands out: it is the simplest, most universal, and most robust to the peculiar dynamics of a [post-labor economy](#). This section develops the case in detail. We begin with the political framing of UBI and the observation that it need not be introduced as a radical new program but can emerge gradually from existing welfare systems. Three clusters of philosophical arguments provide the normative foundations, while the structural design properties that distinguish UBI from means-tested alternatives are summarized below. Growing public support and cross-partisan appeal make UBI more politically viable than its opponents assume. Finally, existing social-democratic policies already point in the same direction, and the evidence from over a dozen pilot programs worldwide (Table 6.1 and Figure 6.2) is overwhelmingly positive. An exhaustive catalogue of all 70 arguments for (and against) UBI appears in Appendix A.1.

The political framing and core idea of Universal Basic Income (UBI) is that every citizen receives an unconditional basic allowance from the government high enough to afford a basic living standard. Opinions are unnecessarily polarized: While most Europeans are sympathetic, many US citizens reject it, viewing it as a form of socialism. Interestingly, many CEOs of US AI companies are advocating for some version of GI. Governments of most countries have some form of (mostly means-tested) social welfare system (under different names: unemployment benefits, job-seeker allowance, ...). So if ‘UBI’ faces irrational objection, one could equally well gradually adapt existing social welfare systems to the changing needs, rather than introducing and naming a new system ([Dumont, 2022](#)). See Section 6.6 for a proposal and Section 2.6 for alternative names.

First, social welfare payments could or should increase in line with AI-driven productivity and GDP growth. Second, once a country admits that full employment is no longer possible, the requirement for the unemployed to actively seek a job could gradually be eased in proportion to the reduced need for workers. Eventually this system converges to what UBI-proponents propose. The level of UBI could also increase from ‘basic’ to ‘high’ in line with the country’s overall wealth.

The philosophical case. The case for UBI rests on several independent philosophical foundations, each sufficient on its own and jointly compelling.

The *freedom* arguments observe that formal legal equality means little to a person who cannot refuse an exploitative employer, leave an abusive partner, or decline degrading work. UBI provides what Pettit (2012) calls the “power to say no”: republican non-domination¹ through a genuine exit option. It extends this protection to unpaid caregivers (predominantly women), compensating labor that markets systematically undervalue without requiring surveillance or verification (Van Parijs and Vanderborght, 2017).

The *fairness* arguments note that no individual can claim sole credit for the cumulative stock of human knowledge, the natural resource base, or the technological infrastructure on which all modern wealth depends. Thomas Paine (1797) proposed a citizens’ dividend from land rents as early as the eighteenth century, on the grounds that the Earth in its uncultivated state is common property; Henry George (1879) developed the idea further, and Gabriel and Kasirzadeh (2026) extends it to AI-generated wealth. Rawls (1971) supplies the complementary insight that morally arbitrary advantages (birth, talent, inheritance) require correction through redistribution.² Since AI builds on millennia of collective scientific and cultural achievement, everyone should be heir to its gains.

The *justice* arguments complete the picture. All behavioral conditions on welfare (work requirements, means tests, “actively seeking employment” mandates) are either unenforceable, unjust, or both: they punish the unlucky, humiliate the needy, and become absurd once machines can do the work. Ecologically, unconditional income decouples survival from environmentally destructive production, enabling people to refuse harmful work without financial penalty. And intergenerationally, future generations should not inherit a system in which survival depends on jobs that no longer exist.

Structural advantages. Beyond its moral foundations, UBI has compelling design properties that distinguish it from means-tested alternatives (Table 6.6). It eliminates the *poverty trap*: because UBI is not withdrawn as earned income rises, there is no implicit high marginal tax rate discouraging work at the lower end of the income distribution. It achieves *full take-up* by definition, since every citizen receives it automatically: no one falls through the cracks due to ignorance, shame, or bureaucratic barriers. Its

¹In republican political theory, freedom is defined not merely as the absence of interference but as the absence of *domination*: the condition of not being subject to another’s arbitrary will. A person who depends on an employer’s goodwill for survival is “dominated” even if the employer happens to be benign. UBI provides the material basis for genuine exit options.

²Rawls’ “difference principle” holds that social and economic inequalities are permissible only if they benefit the least advantaged members of society. His “veil of ignorance” thought experiment asks what rules rational people would choose if they did not know their future position in society; the answer, Rawls argues, is maximin: maximize the minimum outcome.

administrative simplicity slashes overhead: a single universal transfer replaces the patchwork of *means-tested* programs, each with its own eligibility rules, verification costs, and error rates. It supports *labor flexibility* (part-time work, gig work, career transitions, sabbaticals, retraining), provides a *risk cushion* for entrepreneurship and innovation, and shifts *bargaining power* to workers, who can refuse exploitative offers when their survival no longer depends on accepting them. Finally, unlike conditional welfare, it requires *no behavioral surveillance*: no caseworkers monitoring job-search activity and no inspections.

Political viability. Public support for UBI has grown substantially, especially since the COVID-19 pandemic demonstrated the fragility of employment-based safety nets. Polls from 2016–2020 show 64% support in the EU (Dalia Research e28™ survey, April 2016, $N=10\,000$ across EU-28), 71% in a 2020 Oxford European survey, and 55% in the US (up 6 percentage points during COVID); German support stood at 51%, Australian at 58% (Laenen, 2023).³ Politically, UBI can be framed in ways that appeal across the ideological spectrum. The closely related negative income tax (NIT) was originally proposed by Milton Friedman (1962)⁴ as a conservative alternative to the welfare bureaucracy: a tax adjustment, not a “handout.” Many Silicon Valley CEOs (Musk, Altman, Zuckerberg, Amodei, Omidyar, Yang) publicly advocate some version of guaranteed income. UBI also functions as an automatic macroeconomic stabilizer during recessions, injecting purchasing power without the legislative delays that plague fiscal stimulus packages (Hoynes and Rothstein, 2019). Even skeptics of UBI should accept costless pre-legislated triggers (Section 6.2): if labor displacement never materializes, the legislation is never activated and costs nothing.

Traditional policies suffice. Bruenig (2026) emphasizes that while the public intuitively clamors for ‘creating good-paying jobs’ to combat AI displacement, exotic new solutions are unnecessary; traditional social democratic policies, namely direct income support and redistributive taxation, are perfectly adequate to fix the distributive impacts of AI. Korinek and Stiglitz (2019) argue similarly for redistribution but stress that ex-post transfers

³A 2017 Gallup/Northeastern University poll found Americans split 48%–52% when UBI was framed specifically as support for workers displaced by AI; the same framing yielded only 43% in a 2019 follow-up before COVID-era polls reported renewed increases. Support is volatile and highly sensitive to question framing.

⁴In *Capitalism and Freedom* (Chp.12), Friedman proposed the NIT to replace “the host of special measures now in effect” with a single cash transfer that “could be substituted for” the entire welfare bureaucracy. He argued it preserves work incentives (“an extra dollar earned always means more money”) while slashing administrative overhead. The NIT thus has deep roots in free-market, not socialist, thought.

alone may be insufficient at scale, calling additionally for pre-distribution interventions such as steering AI research toward labor-augmenting applications. Real-world pilot programs, such as Ireland's basic income for artists, have demonstrated that unconditional stipends not only alleviate the severe mental toll of precarity but also yield outsized socio-economic returns, generating more value for society than the state's initial investment (Hogan, 2026). Additionally, the Marshall Islands became the first nation to launch a universal basic income scheme to all its citizens, funded by a trust and distributed partially via cryptocurrency, specifically aimed at providing a social safety net against rising living costs (Srinivasan, 2025). Key results from major UBI pilot programs worldwide are compared in Table 6.1.

Beyond the numbers. The table's summary statistics, while compelling, compress away the human texture of these experiments. In Finland, participants in the Kela experiment reported not merely higher life satisfaction but significantly higher levels of trust: trust in other people, in politicians, and in the legal system (Kangas et al., 2020). The most consistent finding was psychological: recipients described relief from the "constant anxiety of conditionality," the perpetual fear that a missed appointment or a misunderstood form could terminate their benefits. In Stockton, the SEED experiment reduced income volatility and moved recipients from "likely to have a mild mental health disorder" to "likely to be well" on clinical distress scales (West and Castro, 2023). Qualitative data revealed a recurring theme of "breathing room": the financial cushion enabled participants to refuse precarious gig work and pursue better opportunities, with no significant reduction in labor supply. In rural Kenya, GiveDirectly's \$22-per-month transfers (roughly a quarter of average household income) triggered a wave of micro-enterprise formation: recipients invested in livestock, tin roofing, and small shops. Banerjee et al. (2023) found no reduction in work effort; instead, the transfers unlocked productive investments that poverty had previously foreclosed. The Ontario experiment serves as a cautionary tale: cancelled after a single year by an incoming government, it demonstrated that the greatest risk to basic income is not economic failure but political fragility. Participants who were surveyed before cancellation reported improvements in health, food security, and housing stability (Ferdosi et al., 2020), making the abrupt termination an act of policy, not evidence.

6.2 Transition from CI to UAI to UHI

How does a society transition from current welfare systems to UBI? The literature converges on a gradual, phased approach (Figure 6.3) rather

Table 6.1: Summary of major UBI and guaranteed income pilot programs worldwide. Across nearly all experiments, labor supply reductions were modest or negligible, possibly due to low savings, low income, low UBI amount, and low duration, while wellbeing improvements were consistent. A key limitation: all existing pilots operate at far smaller scales, shorter durations, and lower basic-income levels than full deployment would require, so general-equilibrium effects remain untested. See [Gentilini et al. \(2020\)](#) for a comprehensive survey; [Corinth and Mayhew 2026](#) surveys 122 US pilots; [Wikipedia contributors 2026](#) maintains a comprehensive chronological list.

Program	Country	Period	Amount	Key Result
Marshall Islands	Marsh.Isl.	2025–	\$200/qtr	First nationwide UBI; trust-funded; crypto
OpenResearch	USA	2020–24	\$1000/mo	Work reduced ≈ 2 h/wk; $\uparrow$ leisure, education
Kela Experiment	Finland	2017–18	€560/mo	+6 employment days; higher life satisfaction (Kangas et al., 2020)
SEED (Stockton)	USA	2019–21	\$500/mo	$\downarrow$ income volatility; improved mental health (West and Castro, 2023)
GiveDirectly	Kenya	2017–	\$22/mo	$\uparrow$ consumption, enterprise; no work reduction (Banerjee et al., 2023)
Mein Grundeinkommen	Germany	2021–24	€1200/mo	No significant work reduction; $\uparrow$ wellbeing
Alaska PFD	USA	1982–	\$1–3K/yr	No sig. employment effect over 40+ years
Mincome	Canada	1974–79	Varied	Only new mothers & teens worked less
Wales BIP	UK	2022–25	£1600/mo	Most generous pilot worldwide; care leavers (Elliott et al., 2024)
Ontario BIG	Canada	2017–18	C\$1400/mo	Cancelled after 1 yr by new government (Ferdosi et al., 2020)
US NIT (4 expts)	USA	1968–82	Varied	First RCTs; modest work reduction (5–9%)

tratively impractical, and the stigma of means-tested “welfare” becomes politically untenable for a majority of the population. At this point, payments are pooled into a single uniform, unconditional transfer: Universal Adequate Income (UAI), set at a GDP-linked level between basic subsistence and full flourishing, likely around average salary level, and increasing over time as automation-driven productivity raises national wealth. Phase 1 recipients whose CI exceeded the initial UAI level are ideally grandfathered at their existing or only gently decreasing level until the rising UAI rate converges upward past it. The transition can be eased by categorical expansion (Sec. 6, “categorical basic incomes”; Van Parijs and Vanderborght, 2017; Standing, 2017). Crucially, large-scale, long-running, generous income schemes that are forms of BI already exist and serve as precursors to full UBI. Universal child benefits (paid in many European countries and Mongolia to all families regardless of income) and universal basic pensions (e.g. New Zealand Superannuation since 1938, the Netherlands’ AOW at $\approx 30\%$ of GDP per capita, Namibia’s since 1992, Bolivia’s Renta Dignidad since 2008) demonstrate that unconditional periodic cash transfers can operate at national scale for decades (Van Parijs and Vanderborght, 2017, Sec. 6). Student grants (universal in some countries, e.g. Scandinavian *studiestöd*), farmer income support (replacing EU price subsidies), and artist basic income schemes (e.g. the Netherlands’ WWIK, 2005–2012; Ireland’s pilot since 2022, Hogan 2026) extend the principle to working-age adults in specific categories. Van Parijs and Vanderborght (2017) advocate a “partial basic income” as a strategically sound first step, gradually increased over time. Politically, a Negative Income Tax (Section 2.6) may serve as an equivalent stepping stone, since it is framed as a tax adjustment rather than a “government handout” (Friedman, 1962).

Phase 3 (UHI) From Basic to High (long-term). As production costs converge toward zero (Section 3.3) and material abundance grows, the GI level can rise beyond affordable to *Universal High Income* (UHI): an unconditional income sufficient not merely for survival but for flourishing, indexed to GDP or GDP-B (see Section 3.3) per capita or productivity growth. Since everyone receives UAI/UHI including high earners, progressive income taxation ensures that those with substantial additional income effectively return the transfer, so the net fiscal cost equals only the redistribution to those who need it (Widerquist, 2017, 2024).

Pre-legislated triggers. A complementary idea is to *pre-legislate* the transition: define measurable triggers that automatically activate each phase. Candidate indicators could be: labor share of national income falling below a threshold; median real wage declining over a sustained period; underemployment rate (including involuntary part-time and discouraged workers)

exceeding a ceiling; headline unemployment sustained above a critical level. A single unemployment-rate trigger is insufficient: as argued in Section 4.5, displaced workers may be absorbed into hyper-competitive, low-wage BS-jobs that keep headline unemployment artificially low while masking genuine labor displacement. Combining wage-based and labor-share indicators avoids this blind spot. Pre-legislating such triggers follows the logic of automatic fiscal stabilizers and avoids delays caused by political gridlock when displacement accelerates. Korinek and Juelfs (2022) show that a utilitarian planner would shift from CI to GI once labor becomes sufficiently redundant; pre-committing to that shift via legislated triggers is a natural implementation. Even sceptics may not object to such legislation: if the triggers never fire, it costs nothing.

Beyond triggers, *ratchet clauses* can tie safety-net levels to displacement metrics (e.g. sectoral employment decline, automation-adoption indices), ensuring automatic benefit escalation without requiring a new political fight at the moment when political voice is weakest (cf. California's AI employment-impact dashboard, Section 7.1). *Constitutional entrenchment* of UBI floors via supermajority-protected provisions can make reductions as difficult as revoking voting rights; the Alaska Permanent Fund Dividend, legislatively entrenched before oil revenues peaked, demonstrates the viability of such pre-commitment, though the dividend (unlike the Fund principal) is statutory and has faced repeated legislative reductions (Widerquist and Howard, 2012). Such entrenchment also addresses the citizen-power objection to UBI (Section 6.5): a constitutionally locked-in UBI is harder to revoke than any share-based ownership claim.

Adaptive regulatory architecture. The preceding mechanisms (pre-legislated triggers, ratchet clauses, constitutional entrenchment) constitute an *adaptive regulatory architecture*: a policy architecture designed to keep pace with technological change by adjusting automatically rather than requiring new legislation at each stage. The core principle is simple: specify the conditions under which policy changes activate, and let the policy self-execute when those conditions are met. This is not novel in fiscal policy (automatic stabilizers such as unemployment insurance already operate this way) but is radical in the context of structural economic transformation. Key design elements include: measurable displacement indicators (labor share, median real wage, un(der)employment rate, GDP per capita) that resist gaming by reclassifying workers; phase-gate thresholds with hysteresis (once triggered, benefits cannot be rolled back below the previous floor); sunset clauses for exceptional powers but not for baseline protections; and independent monitoring bodies insulated from electoral cycles. The advantage is that it

depoliticizes the *timing* of the transition (the indicators fire when displacement warrants it, not when an election cycle permits it) while leaving the *design* of the policies to democratic deliberation.

The *key insight* is that the transition need not be revolutionary: it is a gradual policy evolution that most democracies have the institutional capacity to execute (see Section 10.3 for the broader feasibility argument).




Figure 6.3: Three-phase UBI transition timeline with illustrative trigger events. Phase boundaries are not fixed dates but respond to automation-driven economic indicators.

6.3 Objections to UBI/GI

The most common objections, that GI breeds laziness, is unaffordable, or that share ownership is superior, are examined and rebutted below. [Van Parijs and Vanderborght \(2017\)](#) rebut these and further objections in great detail. See Appendix A.1 for a summary.

UBI makes people lazy. First, experiments generally show this is false. A meta-analysis of 122 US guaranteed-income pilots (2017–2025) finds an average employment effect of +0.8 percentage points across 30 randomized trials, although the four largest studies ($N \geq 500$) show -3.2pp ([Corinth and Mayhew, 2026](#)). For instance, a large-scale three-year UBI experiment conducted in Germany ([Bohmann et al., 2025](#); [Bernhard et al., 2026](#)) demonstrated that receiving UBI did not significantly alter labor market behavior; recipients did not stop working, nor did they drastically reduce their weekly hours, although this can be due to a multitude of factors. The world’s




Figure 6.4: Stylized transition path from a labor economy to a post-labor economy. As AGI automates jobs (red), GI coverage rises (green) to maintain material prosperity, and voluntary activity (blue, dashed) (creative, social, competitive, ...) gradually replaces compulsory employment as the primary source of human engagement.

largest UBI experiment (295 villages in Kenya) likewise found no net reduction in labor supply, with workers shifting from wage employment to self-employment (Banerjee et al., 2023). Second, natural experiments with unearned windfalls corroborate this: lottery winners reduce labor earnings by only about \$11 for every \$100 won, and roughly 85% of major-prize winners keep working (Imbens et al., 2001); Swedish data confirm that even large windfalls reduce labor supply only modestly and permanently, rarely leading to complete withdrawal from the workforce (Cesarini et al., 2017). Van Parijs (2004) responds systematically to such reciprocity-based objections, arguing that unconditional transfers are justified by an egalitarian share of external resources. Third, even if true, this is actually a feature, not a bug.

Initially, UBI would and should start so low that the majority still desires extra income from work to improve their living standard. Once personal wealth accumulates to a comfortable level, the incentive to work for additional income diminishes, especially if the hourly wage drops sufficiently (due to supply/demand shift from automation) and/or if the work is not intrinsically rewarding. Once the job market dries up, and more people seek jobs than jobs are available, UBI should be increased to actually discourage some from seeking employment in line with the available jobs, in order to avoid hyper-competition (see above), and also to increase life satisfaction: Voluntarily unemployed individuals, such as pensioners, report significantly higher life satisfaction than the involuntarily laid-off. To a first approximation, the level of GI should be tuned so that the number of active job-seekers roughly matches the number of actual job-openings (see Figure 6.5).




Figure 6.5: *GI Level vs Job Market Equilibrium: Tuning GI to match job seekers with available openings, avoiding both hyper-competition and excessive labor scarcity.*

Fairness resentment. A related but distinct objection is the perceived unfairness: “Why should I work while others live off UBI?” The objection dissolves on inspection. First, UBI is *universal*: those who work receive it too, so their total income (wage plus UBI) strictly exceeds that of non-workers. Second, if a worker nonetheless resents that others are “getting something for nothing,” the symmetry of *unconditionality* provides a conclusive reply: they are entirely free to stop working as well. If the UBI level is sufficient for a life they find acceptable, nothing prevents them from choosing it; if it is not, their continued employment is evidently worth it to them, and the complaint reduces to envy of a lifestyle they themselves have declined. Third, in a post-labor economy this objection becomes moot: once most jobs are automated, the distinction between “those who work” and “those who don’t” loses its moral weight, since the majority will not work for income regardless.

Political resistance. More broadly, political resistance to UBI is likely to come from several additional quarters: *welfare chauvinism* (“our money should go to our people,” addressed in Section 2.6); *bureaucratic entrenchment* (means-tested systems employ large administrative apparatuses whose staff resist abolition of the programs that justify their roles); and *employer opposition* (UBI strengthens workers’ outside option, reducing employers’ bargaining power and the ability to impose poor working conditions). These are political obstacles, not principled economic objections; each may slow adoption but none undermines the case for UBI on the merits.

Universal basic income, mental health, and the salience of needs. Critics argue that empirical evidence shows UBI leads to negative psychological outcomes despite financial gains. For instance, a UK microsimulation study by Thomson et al. (2024) reports a simulated 0.38% increase in population-level common mental disorders (CMD), while a US randomized controlled

trial by Jaroszewicz et al. (2022) of one-time cash windfalls (\$500 or \$2 000) finds non-positive subjective well-being and increased distress. However, a closer inspection of these studies reveals that these negative findings are either modeling artifacts or consequences of cash insufficiency rather than UBI itself. In the microsimulation of Thomson et al. (2024), the primary model's distress increase is a direct consequence of a highly pessimistic, ad-hoc behavioral assumption: the model penalizes any UBI-induced voluntary labor supply withdrawal (such as reducing hours to care for family, study, or enjoy leisure) as psychologically identical to the acute trauma of involuntary unemployment. When this assumption is relaxed in their structural sensitivity analyses, the findings reverse. If employment is held constant (assuming no labor supply contraction), Full+ UBI reduces CMD prevalence by 0.27%, preventing over 112 000 cases of mental illness in 2023 alone; if labor supply reductions are treated as voluntary economic inactivity (which is psychologically neutral or positive), UBI reduces CMD prevalence by 0.31% (Thomson et al., 2024). Similarly, the distress documented by Jaroszewicz et al. (2022) is driven by a cognitive-salience mechanism that is unique to temporary windfalls. Receiving a one-time cash windfall shifts low-income recipients from passive to active financial management (such as opening long-ignored bills and calculating debts). Confronting their severe financial reality makes their chronic, unmet structural liabilities highly salient ("money on the mind"). Because a one-time windfall is not sufficient to clear their structural poverty, this salience shock causes psychological distress. This represents an indictment of temporary, insufficient windfalls, not a permanent UBI; a predictable, lifelong income floor structurally closes the resource deficit, allowing recipients to systematically resolve debt over time without cognitive shocks. More broadly, these critiques rest on the premise that voluntarily working less is inherently harmful, a premise undermined by the distinction between compensated jobs and purposeful work (Section 7.5).

UBI is unaffordable or inflationary. The 'idealized transition scenario' above demonstrates that the aggregate funds to compensate every displaced worker at their former salary exist within corporate savings alone. Whether delivered initially as targeted displacement benefits or pooled into unconditional UBI, such transfers are affordable without additional government expenditure, funded purely by redirecting salary savings via suitable taxation, as detailed in Section 3.2. Widerquist (2017) demolishes the "sticker shock" objection by distinguishing gross cost (total outlay) from net cost (actual redistribution), showing that a poverty-level UBI would have a net cost of roughly 25–35% of its gross cost. Gentilini et al. (2020) estimate fiscal costs of 3–35% of GDP depending on design parameters. This also

qualifies Pritchett et al.'s (2024) argument that providing UBI at the global extreme poverty line would cost nearly five times as much as targeted direct transfers: in developed countries with robust income tax systems, UBI counts as taxable income and is effectively taxed at each recipient's marginal rate, zero for very low-income earners (who need it most) but approaching 50% for high-income earners in many OECD countries, so roughly half the transfer is automatically clawed back through the existing tax system, and the remainder could be recouped by a modest increase in the top marginal rate if desired. (In developing nations with large informal economies, the automatic clawback mechanism is weaker and Pritchett's cost concern carries more force; however, Gentilini et al.'s (2020) cross-country estimates already account for this heterogeneity.) Higher levels of UBI sufficient to eventually replace a middle-class salary would be introduced gradually in step with increased automation and decreased job availability.

A related concern is that UBI might trigger systemic inflation. However, when UBI is funded by progressive redistribution rather than monetary expansion, it is approximately *demand-neutral in aggregate*: purchasing power shifts from higher-income groups (who save more) to lower- and middle-income groups (who spend more on basics), rather than expanding the monetary base (Hoynes and Rothstein, 2019). Strictly, the shift from savers to spenders does increase money velocity and hence aggregate demand; however, the resulting inflationary pressure is modest relative to the automation-driven supply expansion, and any residual price increase benefits UBI recipients more than it costs them (their purchasing power rises by the full UBI amount minus a small velocity-driven markup). The standard argument that shifting funds from corporate savings to consumption increases money velocity and causes short-term inflation does not apply if (U)BI starts out as unemployment benefits for displaced workers as per the 'idealized transition scenario' in Section 3.2 merely replacing lost wages that were already being spent, leaving individual and aggregate demand constant, and then gradually phased in as discussed above. In addition, in a post-AGI economy, the supply of goods and services is expanding through automation-driven efficiency gains, absorbing additional consumer demand and exerting downward pressure on prices rather than upward.

Trojan horse and central planning objections. Another class of objections, raised by traditional social security advocates, is that UBI acts as a "Trojan horse" for neoliberal deregulation that individualizes systemic failures and dismantles public services, leading to a central-planning crisis where the state cannot adjust benefits due to the lack of price signals (Dumont, 2022). However, this critique conflates progressive UBI proposals with

right-libertarian/Friedmanite ones. As Dumont (2022, p.301) himself highlights, progressive UBI models are explicitly designed to complement, rather than replace, existing public infrastructures (such as universal healthcare, education, and public housing), preventing the privatization of essential services. While Dumont critiques the individualization and universality of UBI from a social-insurance law perspective, he rejects the simplistic neoliberal characterization, advocating instead for a constructive dialog on how UBI can inspire social protection reforms. Furthermore, the central planning objection gets the economics backwards. Unlike centralized provisioning (Universal Basic Services) which directly allocates goods and sets prices, UBI merely distributes liquidity (purchasing power) and leaves the private competitive market fully intact to determine supply, demand, and prices. UBI enhances consumer sovereignty, enabling citizens to express their demand in the market. Adjusting UBI levels requires no complex price-setting; benefits are easily indexed to a basket of basic necessities or inflation, a well-established administrative mechanism in modern pension systems.

Similarly, the concern that UBI disempowers citizens under an authoritarian shift by creating dependency on a central provisioning node (rendering subversives easily purgeable, as historically paralleled in Imperial Germany's mass expulsions (Fitzpatrick, 2015)) misunderstands the economics of cash. As Fitzpatrick (2015) documents, the 1878 Socialist Law expulsions succeeded in impoverishing activists because they depended on localized market employment; their primary survival and resistance fallback lay in spontaneous, horizontal mutual-aid networks funded by worker collections (p. 76). Unlike centralized provisioning (such as public housing or state-guaranteed jobs) which allows a hostile regime to easily weaponize survival resources on a political basis, an unconditional cash transfer distributes fungible liquidity. This cash floor directly enhances geographic mobility (exit options) and finances autonomous, horizontal mutual-aid networks, maximizing citizen resilience against central state coercion.

6.4 Alternatives to UBI/GI

Several related but distinct concepts exist (see Table 6.6). A *Negative Income Tax* (NIT) (Friedman, 1962) delivers net transfers through the tax system; under a linear tax it is equivalent to UBI in disposable-income terms (Section 6.4). Guaranteed Minimum Income (GMI) is means-tested and often conditional. Universal Basic Services (UBS) (Coote and Percy, 2020) provide in-kind provision rather than cash. A Participation Income (Atkinson, 1996) and

‘Conditional Basic Income’ (CBI) (Susskind, 2020, Chp.12) retain a broad conditionality on social participation (Section 6.4). A Sovereign Wealth Dividend distributes resource or investment returns equally to all citizens (e.g. Alaska’s Permanent Fund Dividend since 1982). The Earned Income Tax Credit (EITC)⁵ (Hoynes and Rothstein, 2019) subsidizes only those with earned income and is thus neither universal nor unconditional. A Job Guarantee (Tcherneva, 2020), sometimes called *Universal Basic Jobs* (UBJ), offers public employment at a living wage to all who want it, but once productive tasks dry up it risks becoming a government-run BS-jobs program (Section 4.5). A Basic Endowment (Ackerman et al., 2006, Chp.2) is a one-time lump-sum capital grant at adulthood rather than a periodic payment. Conditional Cash Transfers (Fiszbein and Schady, 2009) target poor households and require compliance with behavioral conditions such as school attendance. Social Insurance (Barr, 2012) is contribution-based, covering only those with sufficient employment history. Workfare (Besley and Coate, 1992) conditions benefits on performing assigned work, serving as a self-targeting mechanism. Labor unions and minimum wage laws are sometimes proposed as alternatives but become moot under full automation (Section 6.4).

Below we discuss some of the more popular alternatives or variants of UBI.

Negative income tax. A *Negative Income Tax* delivers income support through the tax system: households whose income falls below a threshold receive a net transfer rather than paying tax. As typically proposed, it differs from UBI in several notable respects. First, it is *household-based* rather than individual, which disincentivizes co-habitation: couples living together receive a lower per-capita payment than two singles filing separately, penalizing partnership formation. Second, it is usually *restricted* to those currently working or demonstrably willing to work, re-introducing a conditionality that UBI explicitly avoids. Third, it is *paid in arrears*, typically annually as a tax refund, providing no immediate safety net for sudden job loss and requiring recipients to survive the gap between earning and receiving. (Periodic-payment NIT variants integrated with payroll withholding are conceivable, and were briefly implemented in the US as the Advance Earned Income Tax Credit (1979–2010), but uptake remained below 3% of eligibles and the program was repealed due to administrative complexity (Hoynes

⁵The EITC is a refundable tax credit that supplements the earnings of low- to moderate-income workers: it phases in with earned income, reaches a plateau, then phases out. Unlike UBI, it provides zero benefit to those with no earned income and thus cannot serve as a safety net for the fully displaced.

and Rothstein, 2019).) These features make the NIT fiscally cheaper and politically palatable because it preserves a work norm, but they sacrifice most of UBI's distinctive advantages: the work condition re-introduces surveillance and desert requirements, and the household basis requires monitoring of living arrangements. Defenders note that the monitoring infrastructure already exists in countries with couple-based tax or benefit systems, so the marginal surveillance cost is low; critics counter that this is an argument for *dismantling* existing intrusive systems, not for building new policy on top of them. Under certain unrealistic assumptions, NIT and UBI converge financially: with a linear (flat) tax and for single-person households, an NIT produces the same disposable-income amount as a UBI (Atkinson, 1995; Hoynes and Rothstein, 2019), though still paid in arrears and only to those willing to work.

Participation income (Atkinson, 1996) retains a broad conditionality on 'social participation' (employment, caregiving, volunteering, or education). Atkinson (2015) integrates the participation income into a comprehensive 15-proposal program (spanning guaranteed employment, a sovereign wealth fund, capital endowments at adulthood, and progressive taxation up to 65%) costed via UK microsimulation; it remains the most detailed pre-AGI policy blueprint for reducing inequality through combined pre-distribution and redistribution.

Similarly, Susskind (2020, Chp.12) proposes a 'Conditional Basic Income' (CBI) paired with "leisure policies, designed to help people spend their time in purposeful ways; opportunities for people who still want to 'work' even if not for a wage; and requirements that people contribute to society, in return for the support that society provides." However, such state-directed nudging may be unnecessary. Under unconditional UBI, the market revalues social activities automatically: as automated goods become cheaper, the relative price of human social services (caregiving, teaching, mentoring) rises naturally, especially if consumers genuinely prefer human interaction as argued in Section 4.2. Communities can even self-organize around ambitious collective goals: Interest groups could pool voluntary fractions of their UBI to fund, say, a private lunar mission or even colonizing Mars, perhaps offering raffle tickets proportional to contributions for a coveted seat; no government encouragement is required.

Universal Basic Services. **Universal Basic Services** (UBS) proposes expanding free, universal public provision beyond the traditional welfare pillars of healthcare and education to encompass housing, transport, childcare, adult social care, and digital access (Coote and Percy, 2020; Gough, 2019). Where UBI gives cash and lets individuals purchase what they need on the market,

UBS decommodifies the services themselves, guaranteeing access regardless of income. Gough (2019) grounds UBS in human need theory: certain needs (nutrition, shelter, health, social participation) are objectively determinable and heterogeneous, making collective provisioning systems more efficient than individualized market purchases for meeting them.

UBS has several attractions: First, collective provision exploits economies of scale and pools risk across the population, potentially delivering services at lower per-capita cost than individual market purchases funded by cash transfers. Second, although properly funded GI is not inflationary in aggregate (Section 6.3), cash transfers remain vulnerable to *rent extraction* in supply-inelastic markets: landlords and monopolistic providers can absorb part of the transfer through higher prices, particularly for housing; UBS sidesteps this by providing the service directly, removing the price channel through which rents are extracted. Third, UBS achieves strong redistribution without means-testing: free services are worth disproportionately more to lower-income households as a share of their total consumption. Fourth, publicly provided services can be steered toward sustainability and long-term planning goals (e.g. low-carbon transport infrastructure) that decentralized consumer spending may underinvest in (Coote and Percy, 2020).

The weaknesses are significant. UBS requires the state to determine *what* services to provide, at *what* quality, and *how* to allocate them: precisely the Hayekian knowledge problem⁶ (Hayek, 1945) and the socialist calculation problem that have historically plagued central planning. Preferences are heterogeneous, local, and constantly shifting; a bureaucracy that decides everyone needs the same housing, the same transit routes, or the same childcare model risks paternalism and misallocation. Cash transfers respect consumer sovereignty by letting recipients decide how to meet their own needs. Moreover, UBS creates large public monopolies vulnerable to capture by providers (e.g. public-sector unions resistant to reform) rather than consumers.

Under AGI, the calculus shifts in an interesting way: The information and optimization problems that historically defeated central planners may become increasingly tractable as AI systems handle logistics, demand forecasting, and resource allocation. To the extent that AGI enables near-optimal planning, the efficiency case for UBS strengthens considerably. However, even computationally solvable planning remains politically dangerous if it

⁶Hayek's 1945 insight is that "the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess." No central planner can aggregate this information as efficiently as the price system does through decentralized market transactions.

concentrates allocative power without democratic accountability.

The practical conclusion is that UBS should *augment* UBI, not *substitute* for it. Gradually improving and expanding high-quality public services (healthcare, education, public safety, transit, digital infrastructure) reduces the cost of living and thereby lowers the UBI level needed to guarantee a dignified standard of living. This is already the implicit model of every welfare state that combines cash transfers with public services. Additionally, UBS can serve as a vehicle for the “leisure policies” that [Susskind \(2020, Chp.12\)](#) advocates: publicly funded community centers, sports facilities, libraries, maker spaces, and cultural programs provide structure and purpose for those no longer occupied by paid employment, helping address the meaning deficit that pure cash transfers leave unresolved (see [Section 7.5](#)). The key design principle is that UBS should expand the *options* available to citizens, not prescribe how they live; wherever possible, individuals should retain the freedom to opt out of public provision and use their UBI to purchase alternatives on the market.

Labor unions and minimum wage. Once offerings can be automated to a cost below minimum legal wage ([Section 3.3](#)), workers (eventually) become redundant, and even minimum wage jobs will dry up. Labor unions or governments may try to legally prevent automation (e.g. striking against driverless trains) but such anti-progress strategies actually harm the macro-economy and hence in the end all its citizens. Furthermore, any deliberate national slowdown or regulatory blockade is likely futile. Due to “military-economic adaptationism” ([Genewein et al., 2026](#)), nations that resist adopting productivity-enhancing AI will simply be outcompeted and lose global influence, making total automation virtually inevitable.

Even before full automation renders the question moot, UBI offers a less distortionary alternative to minimum wage laws. A minimum wage is a price floor on labor: when set above the market-clearing rate, it improves earnings for those who keep their jobs but prices out workers whose marginal productivity falls below the threshold, creating unemployment and pushing activity into the informal sector. UBI sidesteps this trade-off. If workers receive an unconditional transfer, their total income (market wage plus UBI) can match or exceed the old minimum-wage income even if the statutory floor is lowered or removed, because the income support shifts from a market-distorting price regulation to a direct, non-distortionary cash transfer. Moreover, UBI strengthens workers’ outside option, the “power to say no” to exploitative offers ([Van Parijs, 2004](#)), so market wages need not collapse to subsistence levels once the legal floor is lifted. This is precisely the logic [Friedman \(1962, Chp.12\)](#) advanced for the negative income tax: a single

universal transfer renders minimum wages, tariffs, and the entire “rag bag” of piecemeal interventions redundant, letting the labor market clear while guaranteeing a dignified income floor.

6.5 Universal Citizen Share Ownership Models

We use “Universal Citizen Share Ownership” as an umbrella term for any scheme in which citizens hold (directly or indirectly) equity shares in productive enterprises for the purpose of funding their livelihood. This encompasses individual share grants, Universal Basic Capital (UBC), Citizens’ Wealth Funds (CWF), and Sovereign Wealth Funds (SWF). Importantly, these are *funding mechanisms* for basic income, not alternatives to it; the distinction matters because the payout can always be structured as a periodic dividend indistinguishable from UBI, while the ownership structure determines accountability, risk, and citizen agency (see the discussion below).

A popular proposal is to guarantee everyone shares in an (AI/Robot) company, often called Universal Basic Capital (UBC), possibly the one where a worker has been laid off, so as to profit from the increased profitability due to automation. The argument is that this aligns citizens’ long-term economic incentives with the success of corporate AI development (Korinek and Juelfs, 2022). However, as we show below, every model in this family (individual shares → citizens’ wealth fund → sovereign wealth fund) reduces to tax-funded UBI while adding complexity and fragility as a theoretical *endpoint* (though existing pension systems can serve as a practical *transition vehicle*; see Section 6.6). Historically, this proposal represents a variation of the classic debate between one-time capital grants (e.g. the “stakeholder grants” or “citizen’s stake” of Ackerman and Alstott (2006)) and periodic basic income (see Ackerman et al., 2006, Chp.1).

Tinbergen’s separation principle. A guiding principle for the following analysis is *Tinbergen’s rule*: effective economic policy requires as many independent instruments as there are independent targets (Tinbergen, 1952, Chp.V); Musgrave’s (1959) three-branch model makes the same point for fiscal policy specifically. Applied here, three targets are logically independent and call for separate instruments: *funding* (how to raise revenue: taxation vs. sovereign wealth fund), *distribution* (how to deliver it: periodic UBI vs. one-time grants vs. share dividends), and *power limitation* (how to prevent its concentration: antitrust, constitutional entrenchment (Section 6.2), democratic governance). Conflating them into a single share-based mechanism, as Meidner’s (1978) wage-earner funds attempted, overcomplicated the de-

sign and contributed to the political backlash that dismantled the Swedish experiment after only seven years.

Share ownership. Under a capital grant or an *individual share-based system*, the primary risk is that individuals might sell their shares and squander the proceeds, a vulnerability referred to in the literature as “stakeblowing” (White, 2006). If a person squanders their entire capital stake, will the government let them starve, or will it be forced to step in with further support? The former is incompatible with modern moral values, while the latter is unfair to prudent citizens and means that a secondary safety net remains necessary anyway. Furthermore, a regular monthly flow of income, as delivered by UBI, does not merely act as a paternalistic mechanism to prevent individuals from squandering their resources; rather, it guarantees ongoing security “one month at a time,” providing a more stable structural foundation for individual autonomy and equalizing long-term class power relations (Wright, 2006; Pateman, 2006). A solution within the share-ownership model would be to prohibit individuals from borrowing against, selling, or pledging their future distributions (Altman, 2021; Huang and Manning, 2025), legally enforced by the government. What if a person is unlucky with their choice of portfolio? Should we punish bad luck or label it a poor investment decision? Only few people are prudent investors. The one genuine theoretical advantage of individual share ownership is direct corporate governance: shareholders vote on board members and major decisions. However, with equity distributed across an entire citizenry, people’s stake will be so small that most would not bother to vote, and even less so in a properly informed way. In practice, they would have to delegate to proxy advisors, reproducing the same principal–agent problems the model was designed to avoid.

Size of shareholdings and ramp-up. Moreover, for dividends to reach the poverty threshold or current salary levels at current productivity levels, citizens would need very large shareholdings, and the initial expropriation required to create them, such as Sanders’ proposed one-time 50% stock seizure (Sanders, 2026), is politically explosive. More realistically, share transfers would need to be ramped up gradually over years.

Automating a single employee. This is actually achievable at zero net cost to any party (ignoring the ongoing automation cost A for clarity; see Section 3.2 for the full treatment), by an argument similar to Section 3.2: Assume employee E of company C earns salary S . The company has value V and earned profit P in the last year. Now suppose E is replaced by an AI/robot but receives a fraction γ of the company’s equity V . The company’s profit rises to $P + S$, since it saves E ’s salary. The employee receives income $I = \gamma(P + S)$ from the shares valued at γV , in the form of dividends or

Column key for welfare properties in Table 6.6:

- [UU] Universal (paid to all adult citizens)
- [UC] Cash (monetary payment, not in-kind provision)
- [UN] Unconditional (no work or behavioral requirement)
- [UI] Individual (paid per person, not per household)
- [UP] Periodic (paid at regular intervals, not lump-sum)
- [MF] Real Freedom (maximizes real, not just formal, opportunities)
- [MN] Non-Domination (exit option from exploitative relationships)
- [MG] Gender Justice (compensates unpaid care work without surveillance)
- [MD] No Desert Req. (no prior contribution required)
- [MH] Universal Dignity (no means-testing humiliation)
- [EP] No Poverty Trap (no implicit high marginal tax from benefit withdrawal)
- [EK] Full Take-up (automatic receipt, no application barrier)
- [EA] Admin. Simple (low bureaucratic overhead)
- [EL] Labor Flexible (supports part-time, gig, career transitions)
- [EN] Entrepr. Buffer (risk cushion for self-employment)
- [EB] Bargaining Power (workers can refuse exploitative offers)
- [EV] No Surveillance (no monitoring of behavior or household)

References for [W]elfare system rows in Table 6.6:

- [WU] [Van Parijs and Vanderborght \(2017\)](#)
- [WN] [Friedman \(1962\)](#); [Hoynes and Rothstein \(2019\)](#)
- [WG] Various countries
- [WB] [Coote and Percy \(2020\)](#)
- [WP] [Atkinson \(1996\)](#)
- [WS] [Widerquist and Howard \(2012\)](#)
- [WE] [Hoynes and Rothstein \(2019\)](#)
- [WJ] [Tcherneva \(2020\)](#)
- [WD] [Ackerman et al. \(2006, Chp.2\)](#)
- [WC] [Fiszbein and Schady \(2009\)](#)
- [WI] [Barr \(2012\)](#)
- [WF] [Besley and Coate \(1992\)](#)
- [WL,WR,WK] [Barr \(2012, Part III\)](#)
- [WO] [Barr \(2012, Part IV\)](#)
- [WA] [Atkinson \(2015\)](#)
- [WT,WH] [Barr \(2012, Part V\)](#)
- [WV] U.S. SNAP.

capital gains, which equals the past salary iff $\gamma := S/(P + S)$. For this γ , the company's new profit $\tilde{P} = (1 - \gamma)(P + S) = P$ is also unchanged.

Profit changes. A major difference from the direct automation tax is that if profit in future years changes to P' , the employee's income changes to $I' = \gamma(P' + S) = I \cdot (P' + S)/(P + S) \approx I \cdot (P'/P)$ since $S \ll P$. That is, income becomes very sensitive to changes in profits. If profits halve or double, income will halve or double. It would be unreasonable to adjust γ annually, especially expropriating E when profits rise. In theory, this could be smoothed out by dipping into the equity itself, but any such scheme would need tight regulation to prevent stakeblowing. A better solution is a mix of funding via direct automation taxes (Section 3.2) and share ownership, so as to give displaced workers some but manageable skin in the game: share ownership aligns their interests with the economy's health, since their income rises with corporate profits (see Section 6.5 for the broader citizen-power argument).

Automating multiple employees. Automating n employees works similarly: Employee E_i will receive a fraction $\gamma_i = S_i/(P + S_+)$ of V , where $S_+ := S_1 + \dots + S_n$ is the total salary savings. Note that if n increases, γ_i decreases: Since E_i is now also a shareholder, he also has to give up some shares for the employees automated after him, though this does not decrease his income: $I_i = S_i$ remains. Initially $S_+ \ll P$ for small n , so $\gamma_i \approx S_i/P$ is approximately independent of S_+ hence n . Ultimately, when (nearly) all employees are automated, the total salary savings S_+ may even exceed profit P , and jointly the former employees would own a fraction $\gamma_+ = S_+/(P + S_+)$ of company C , which for labor-intensive firms (where payroll is 1.5–4× annual profit) ranges from 0.6 to 0.8. This exceeds even the 50% proposed by Sanders (2026), but such high γ_+ corresponds to full salary substitution. More importantly, it can and should be ramped up over time in line with increased automation at no cost to the company, since $\tilde{P} = (1 - \gamma_+)(P + S_+) = P$ remains unaffected.

Real-world complications. As with the direct automation taxes in Section 3.2, real-world complications do not invalidate the basic idea: the cost of automation and the need for incentives can be addressed as in Section 3.2. One difference is that in most countries, income, capital gains, and dividends are taxed (often significantly) differently. This can be corrected by aligning the tax treatment or adjusting γ .

Citizens' Wealth Fund. The solution to the stakeblowing, diversification, and portfolio-management problems is to pool shares into a single broad index fund: a *Citizens' Wealth Fund* (CWF) such as the proposed American Equity Fund (Altman, 2021; Huang and Manning, 2025) (see Figure 6.7). Closely related ideas appear under various names: *wage-earner funds* (Meidner, 1978), *national youth endowment*, *minimum inheritance*, *universal personal*

capital account, stakeholder grant, child trust fund, and capital endowment. Citizens hold fund units rather than individual company shares, and the fund is passively managed (the majority of actively managed funds historically underperform comparable passive index funds). This fixes the practical defects of individual ownership but further degrades the voting advantage: citizens no longer vote on corporate decisions but on fund governance, which for a passive index fund is almost vacuous: the investment strategy is mechanical. Any remaining oversight reduces to ordinary democratic accountability, which a sovereign wealth fund already provides through general elections.

A CWF does retain two residual advantages over a tax-funded UBI. First, *institutional protection*: because fund units are citizen-owned assets, dissolving the fund requires explicit legislative expropriation, whereas a statutory UBI program could in principle be defunded through ordinary budgetary or executive processes. However, this is a difference of degree, not kind: a UBI established by robust legislation (analogous to Social Security, which has survived for around a century already) is similarly durable, and constitutional entrenchment would make UBI as safe as any fund. Second, *economic alignment*: CWF dividends track fund returns, giving citizens a direct financial stake in economic prosperity. In practice, dividends must be smoothed across business cycles to avoid hardship during downturns,⁷ weakening the signal; and tax-funded UBI exhibits the same alignment indirectly (if the economy grows, tax revenue grows). Third, *citizen agency*: a CWF provides citizens with a capital-era analog of the labor power (e.g. striking) that full automation eliminates, giving them a direct stake in (and notional leverage over) productive assets (Section 6.5).

Once a CWF restricts share sales, mandates passive indexing, and strips meaningful shareholder voting, the only remaining distinction from a sovereign wealth fund is nominal asset ownership, and even this is largely formal, since citizen oversight of a passive index fund reduces to the same democratic accountability that governs an SWF.

Sovereign Wealth Fund. The government could therefore equally own the fund and distribute the gains, a *Sovereign Wealth Fund* (SWF), the model adopted by Norway's oil fund and Alaska's Permanent Fund. Senator Sanders (2026) has proposed exactly this for AI, radicalizing Meidner's (1978) idea of compulsory equity transfers to collectively managed funds: a one-time 50% stock levy on large AI companies capitalized into a federal SWF with government voting shares and board seats; see also Imas (2026a). Korinek

⁷Passive indexing governs portfolio composition, not cash-flow management; a separate buffer fund smooths payouts across cycles, as Norway's GPFG demonstrates




Figure 6.7: American Equity Fund Mechanics: Companies pay 2.5% of market value in shares, and land owners pay 2.5% of value in dollars to the fund, which distributes dividends to citizens. (Adapted from Altman (2021)).

and Lockwood (2026) note that SWFs capture AI returns without distorting investment and provide equity-based insurance against uncertainty about the pace and scope of AI displacement.

The structural shift is significant: under a CWF, individuals are the beneficial owners of fund units and dividends are an ownership right; under an SWF, the government owns the assets and distributes dividends as a revocable policy decision, not as an entitlement (Altman, 2021). Both models share the vulnerability that a hostile government can expropriate or redirect the fund (Section 7.3). An SWF, however, faces an additional fiscal-discipline risk that a CWF mitigate and a tax-funded UBI largely avoids: governments have an even worse track record of saving than individuals and companies (Alesina and Passalacqua, 2016). A large, visible pool of wealth under government control invites drawdown during recessions (when the pressure to pay above the poverty line intensifies), reappropriation for expenditures deemed more urgent (e.g. subsidizing a distant war or bailing out industries or containing unsustainable debt), and gradual erosion through politically motivated mismanagement (deficit bias, common pool, strategic debt, political budget cycles). Many past SWF attempts have succumbed to precisely these dynamics (Humphreys et al., 2007; Bauer, 2014). By contrast, a tax-funded UBI is a flow-through mechanism: quarterly taxes on private capital fund transfers directly, leaving no accumulated pool that politicians can raid.

If the government owns the AI-fund rather than the citizens, they could as well *not* own any fund but comparably tax the companies' capital (not just profits, since the fund levies a percentage of market value), and we are back at the much simpler tax-funded UBI. In practice, equity-based fund proposals such as Meidner (1978) and Altman (2021) require the levy to be paid in newly issued shares, avoiding a cash drain; but a tax-funded UBI gives firms *more* flexibility: they may raise the tax payment through share issuance, retained earnings, or any other financing mechanism, whereas

compulsory dilution constrains them to equity alone. Furthermore, the initial share redistribution would have to be free, since a laid-off worker or a newborn cannot buy shares; this is achievable via government decree, but such radical expropriation will meet more resistance than adjusting company taxes, a well-established practice. Indeed, a sovereign wealth dividend already satisfies nearly all UBI properties (Table 6.6, row [WS]); the two are complementary funding mechanisms, not competing alternatives: funding source and distribution channel are separate instruments for separate targets (Section 6.5; see Table 6.8 for a structural comparison).

Table 6.8: Structural comparison of four AI wealth-redistribution models. Each model to the left is more complex and fragile than the model to its right, while offering little structural advantage that cannot be replicated more simply (see text for the reductio argument).

Property	Individual Shares	Citizens' Wealth Fund (CWF)	Sovereign Wealth Fund (SWF)	Tax-Funded UBI
Asset owner	Individual citizens	Citizens (fund units)	Government	N/A (cash transfer)
Governance	Individual votes	Fund mgr. (passive)	Gov't votes + board seats	N/A
Dividend basis	Ownership right	Ownership right	Policy decision (revocable)	Statutory entitlement
Stakeblowing	High (can sell)	Medium (sale restricted)	None (no citizen asset)	None
Diversification	Optional	Full (broad index)	Full (portfolio)	N/A
Portfolio mgmt.	By individual (most lose)	Passive index (no decisions)	Gov't fund managers	N/A
Income stability	Volatile (single company)	Market-linked (needs smoothing)	Gov't-smoothed	Budget-smoothed (stable)
Expropriation	Explicit seizure req'd	Explicit seizure req'd	Already gov't-owned	Can cut/revoke payments
Pol. feasibility	Radical (free shares)	Gradual (2.5%/yr)	Established (oil/resource)	Established (tax adj.)
Alignment	Strong (direct ownership)	Moderate (passive)	Weak (citizens don't own)	None (decoupled)
Complexity	Very high (per-couple/person)	Moderate (single fund)	Moderate (single fund)	Low (existing infra.)
Newborns	Must issue new shares	Must issue new units	Auto. (dividends)	Auto. (payments)
Reduces to	→CWF	→SWF	→UBI	Terminal
Proponents	(Korinek and Juelfs, 2022)	(Meidner, 1978; Altman, 2021)	(Sanders, 2026)	This book (and others)

The citizen-power objection. The strongest case for a CWF over tax-funded UBI is not economic but political *power*. The “Age of Labor” conferred unprecedented political power on workers (Korinek and Stiglitz, 2021); they could strike, withhold their contribution, and thereby constrain employers and governments (Boix, 2019). Full automation extinguishes this lever entirely: if no human labor is needed, a strike is merely self-imposed privation. UBI restores *livelihood* but not *agency over production*. A CWF, by contrast, vests citizens with capital ownership, providing an analog of former labor power: collectively, fund-unit holders “own” the robots and could in principle vote to redirect or even shut down production, a capital-era strike. Moreover, the UBI equilibrium may be structurally fragile: if the government no longer depends on individual tax income and a handful of firms grow very large relative to the state, the political incentive to maintain redistribution weakens, an “intelligence curse” paralleling the *resource curse* that afflicts petro-states (Ross, 2001; Huang and Manning, 2025); see Section 7.2 for the full analogy. Under a CWF, citizens retain a direct ownership claim (Altman, 2021) that is harder to revoke than a statutory transfer.

These are genuine concerns, but each has significant limitations. First, citizen agency through a CWF is largely theoretical: with equity distributed across an entire population, each holder’s influence is negligible (the dilution problem in §Share ownership), and the “kill switch” (shutting down robots) is a scorched-earth threat that harms the striker at least as much as the firm, unlike calibrated labor strikes that withdraw specific services. Second, once practical constraints strip a CWF of meaningful shareholder voting (as argued in §CWF), the “power” it confers reduces to the same democratic accountability that governs tax-and-redistribution policy under UBI. Third, tax-funded GI retains indirect but real alignment with economic prosperity: if the economy grows, tax revenue grows, and GI can grow with it. Citizens therefore retain a stake in a well-functioning economy even without direct share ownership, though the coupling is admittedly weaker. Finally, the intelligence-curse fragility applies symmetrically: a government powerful enough to ignore citizen tax contributions is equally capable of overriding CWF governance, so the vulnerability is a democracy-level problem, not a UBI-specific one (Section 7.3). Conversely, the right to UBI could be constitutionally enshrined via supermajority-protected provisions (Section 6.2), making it as hard to revoke as voting rights: the Alaska Permanent Fund Dividend (Widerquist and Howard, 2012) demonstrates the viability of such entrenchment.

More fundamentally, redistribution and power limitation are independent targets (Section 6.5): UBI handles the first; antitrust, constitutional

entrenchment (Section 6.2), and democratic governance handle the second.

6.6 Pension Fund Pathway

The preceding analysis shows that share-based models reduce to tax-funded UBI while adding complexity and fragility, and that a CWF is not the optimal *endpoint* for redistribution (Section 6.5). Yet the argument leaves open a practical question: what is the most politically viable *path* toward that endpoint? Creating a citizens' wealth fund from scratch requires new institutions, new legislation, and new trust: a large political ask. But many countries already operate mandatory pension systems that are often, in all but name, citizen-owned restricted wealth funds.

Existing pension systems as CWF infrastructure. Australia's Superannuation system is the clearest example. With assets under management exceeding \$4.3 trillion AUD (roughly 140% of GDP) and near-universal workforce coverage of 16+ million Australians, the system is larger in relative terms than most sovereign wealth funds ([Association of Superannuation Funds of Australia, 2025](#)). Contributions are mandatory (currently 12% of earnings, paid by employers), accounts are individually owned, regulated by the Australian Prudential Regulation Authority (APRA), and offer default low-cost investment options alongside self-directed portfolios. This is a CWF in all but name: citizens hold capital in individual accounts, accumulate returns over decades, and draw down in retirement. The differences from a formal CWF are how they are funded and the restriction to post-retirement access.

Similar systems exist worldwide: Singapore's Central Provident Fund (CPF, 37% combined contribution rate, already multi-purpose: housing, healthcare, education, and retirement), Chile's AFP (10% employee contribution, pioneered mandatory individual accounts in 1981), and Malaysia's EPF (23% combined rate, partial early withdrawals for housing and education). Collectively, these systems demonstrate that citizen-owned, individually managed, restricted capital accounts can operate at national scale for decades, a fact emphasized by the foundational literature on asset-based welfare ([Sheraden, 1991](#)) (see Table 6.9).

Birth capital endowment. To broaden the pension system's role from retirement-only to lifetime income, additional funding beyond employer contributions is needed. Automation tax revenue (Section 5.3) is one source; others include reformed capital gains, inheritance, and general revenue. The most distinctive funding mechanism is a government-funded lump

Table 6.9: Mandatory pension systems as CWF infrastructure.

System	Country	Rate	AUM	Key Feature	Early Access?
Super	Australia	12% (empl.)	\$4.3T	Individual accounts, APRA-regulated	Hardship +COVID
CPF	Singapore	37% (comb.)	\$500B	Multi-purpose: housing, health, educ.	Partial
AFP	Chile	10% (empl.)	\$200B	Pioneered individual accounts (1981)	COVID
EPF	Malaysia	23% (comb.)	\$250B	Partial early access for housing, educ.	Partial

sum deposited into every newborn’s pension account at birth: a “baby shower” from the commonwealth. Consider an endowment of \$200 000 at birth. At a conservative 5% real return (after inflation), the account grows to approximately \$481k by age 18 ($200k \times 1.05^{18}$). A perpetual drawdown at 5% of capital yields roughly \$24k per year, lifelong, while preserving the principal in real terms: a genuine perpetuity, not a depleting account. This figure is near the current Australian poverty line for a single person (approximately \$27k), providing a floor, not a ceiling: any income from work, other transfers, or voluntary savings adds to it, or a continuation of the historic average return of 7% would lead to a \$65k pension ($24k \times 1.02^{50}$).

The idea of one-time capital grants at adulthood has a long pedigree. [Ackerman and Alstott \(1999\)](#) proposed \$80k at age 21, funded by a 2% annual wealth tax and repaid from the estate at death. [Hamilton and Darity Jr. \(2010\)](#) introduced “Baby Bonds”: a sliding-scale endowment from \$500 (wealthiest families) to \$50k (poorest), means-tested inversely to family wealth and accessible at 18 for approved wealth-building activities. [Sherraden \(1991\)](#) laid the theoretical foundation for asset-based welfare and Individual Development Accounts. The UK’s Child Trust Fund (2005–2011) provided £ 250–500 at birth, unrestricted at age 18, before being abolished by the Coalition government. The present proposal differs in two respects: the endowment is substantially larger (enabled by automation tax revenue that does not yet exist at scale), and it is *restricted* to perpetuity drawdown within a regulated pension system rather than available as a lump sum at adulthood (see [Table 6.10](#) for a comparison).

Crucially, \$200k for each of approximately 300k annual births costs \$60 billion per year (roughly 2.7% of Australian GDP); far less than providing \$200k to every adult (approximately 20 million, which would cost \$4 trillion). While the birth endowment may be more politically divisive (childless adults receive nothing directly), better support for families with children compared

Table 6.10: Birth endowment proposals compared.

Proposal	Amount	Age	Restrictions	Funding
Ackerman & Alstott	\$80k	21	Repaid from estate at death	Wealth tax
Hamilton & Darity	\$500–50k	birth	Accessible at 18 for approved uses	Federal budget
UK Child Trust Fund	£250–500	birth	Unrestricted at 18	General revenue
Piketty	€120k	25	None specified	Wealth + inheritance tax
This proposal	\$200k	birth	Restricted: 5% perpetuity drawdown from Super	Automation tax

to childless households is generally accepted as sound policy, and every citizen eventually benefits as endowed cohorts age in. The amount could be phased in gradually as automation and other tax revenue scales (see Section 6.6 for the phased expansion schedule). Unlike employment-linked benefits, the birth endowment is *naturally universal*: no work requirement, no conditionality, making it a truly unconditional component of the system for future generations.

Phased expansion. The transition from retirement fund to lifetime income system can be structured in three phases, deliberately synchronized with the book’s broader CI → UAI → UHI transition roadmap (Section 6.2).

In *Phase 1* (corresponding to CI, Section 6.2), the existing system is expanded along dimensions already under active discussion. Compulsory employer contributions increase (from the current 12% toward 15–20%, a trajectory Australia’s policy community already anticipates). Government contributions extend Super coverage to non-workers: the unemployed, unpaid carers (Carers Australia, 2024), students, the disabled, and parents on leave. The Super on Paid Parental Leave law (effective July 2025), which pays 12% Super on PPL payments, has already established the precedent that the government can fund Super contributions for non-working periods.⁸ Alongside these coverage expansions, automation and AI tax revenue (Section 5.3) begins flowing into *all* citizens’ Super accounts, independent of employment status. Each of these steps is individually small and politically digestible, exploiting the system’s institutional path dependence.

⁸A 2026 report by the Super Members Council found over one million Australians excluded from the Super Guarantee, including gig workers, part-time workers under 18, and domestic workers (Super Members Council of Australia, 2026).

In *Phase 2* (corresponding to UAI, Section 6.2), the birth capital endowment (Section 6.6) is introduced (phased in gradually: e.g. \$50/100/200k in 1/5/10 years and ultimately more as automation tax revenue scales), and the conditions for early access are progressively relaxed. Australian Super already permits early access on hardship and compassionate grounds before preservation age (currently 60): this is functionally conditional social welfare drawn from one’s own savings. The COVID-era early access scheme (2020) stress-tested this mechanism at scale: approximately \$36 billion was released by over three million Australians (up to \$20k each), demonstrating that the system *can* function as an emergency income source. The infrastructure for conditional early access therefore already exists. In *Phase 2*, these conditions are gradually relaxed, paralleling the CI → UAI conditionality relaxation.

In *Phase 3* (corresponding to UHI, Section 6.2), broad-spectrum taxation (Section 5.3) replaces automation-specific levies as the dominant funding source. Drawdown becomes unconditional and perpetual. The pension system has morphed, by incremental steps over a generation, into a de facto lifetime guaranteed income funded by a diversified tax portfolio: a UBI delivered through the institutional chassis of citizen-owned accounts (Table 6.11 summarizes the three phases).

Table 6.11: Pension pathway transition phases, synchronized with CI → UAI → UHI (Section 6.2).

Phase	Phase	Contribution	Access	Coverage	Character
Current	Pre-CI	12% employer only	At 60+ (hardship exceptions)	Workers	Retirement fund
Phase 1	CI	... + gov’t for non-workers	At 60+ or broader hardship	Universal	Enhanced pension
Phase 2	UAI	+ \$200k birth endowment	Capped draw from ~18	Universal newborns	Individual CWF
Phase 3	UHI	Broad-spectrum tax	Perpetuity drawdown	All citizens	De facto UBI

Design safeguards. The objections that naturally arise (stakeblowing, investment risk, adequacy) are design parameters, not fundamental weaknesses.

Stakeblowing (Section 6.5) is *impossible* if drawdown is capped at a fixed percentage of capital (e.g. 5% per year): the principal is preserved in perpetuity by construction. The remaining risk is mis-investment, not squandering. A mandatory non-choosable core allocation (for example, 80%

in a broad market index fund and government bonds) addresses this directly. Only amounts above the core can be self-directed if desired. Self-Managed Super Funds (SMSFs), which currently allow aggressive and idiosyncratic investment strategies, would need tighter restrictions or exclusion from the core allocation. This is more restrictive than current Super but substantially less restrictive than a sovereign wealth fund in which citizens own nothing directly.

The distinction between retirement income and working-age support is logically inconsistent once automation is the dominant cause of non-work. Unemployment and retirement are both states in which a person is not earning labor income; the current rule that one's pension fund is accessible at 65 but not at 25 is an artifact of the full-employment assumption that automation progressively invalidates. Lowering the retirement age is itself one proposed response to automation (Section 5.1); the pension pathway achieves the same effect more elegantly by making the system cover all of life rather than only its final decades.

At \$24k per year, the perpetuity drawdown provides a floor near the poverty line, not a comfortable income. But it is genuinely lifelong (the principal is never depleted), it can be supplemented by work income or other transfers, and the birth endowment amount can be increased as automation wealth grows.

Relationship to the *reductio* argument. A candid assessment requires acknowledging that this proposal partly rehabilitates the CWF model after the preceding sections argued against it. The *reductio* (Shares → CWF → SWF → tax-funded UBI) establishes that a CWF adds complexity over direct tax-funded UBI as a theoretical *endpoint*. But as a *transition vehicle*, co-opting existing pension infrastructure has decisive practical advantages: Political feasibility: Australia's Super has over 30 years of bipartisan support; expanding it is incremental reform, whereas creating a new CWF from scratch is revolutionary. Institutional maturity: fund management, prudential regulation (APRA), tax treatment, and drawdown rules all already exist and have been stress-tested through market cycles and the COVID crisis. Citizen trust: people understand and trust "their Super account" in a way that a new government transfer program would need decades to earn.

Crucially, this does *not* rehabilitate the SWF model. The pension pathway maintains citizen ownership through individual accounts, avoiding the pooled-wealth temptation that makes sovereign wealth funds vulnerable to government raids (Section 6.5). And the design satisfies Tinbergen's separation principle (Section 6.5): funding operates through automation and

broad-spectrum taxation, distribution through restricted individual accounts with capped drawdown, and power limitation through citizen ownership and democratic accountability: three independent instruments for three independent targets.

If the pension-mutated CWF proves trustworthy and functional over time, it could even serve as a *permanent endpoint* rather than merely a transition. A CWF kept in perpetuity is a viable long-run institution as long as citizens trust it. It is not the simplest path (direct tax-funded UBI is simpler), but it may be the politically most viable and most trustable one. The book’s theoretical reductio still holds as a simplification argument, but institutional trust may ultimately outweigh theoretical elegance.

Table 6.12: Novelty assessment: what exists in the literature vs. what is new (see also Tables 6.9 to 6.11).

Component	Prior Art?	Key Refs
Birth endowment / baby bonds	Yes	Ackerman and Alstott (1999) ; Hamilton and Darity Jr. (2010)
Expanding mandatory pension to non-workers	Yes	Carers Australia (2024) ; Super Members Council of Australia (2026)
Gov’t pension contributions independent of employment	Partial	Super co-contribution; PPL Super (2025)
Automation/AI tax → dedicated fund	Yes	Many proposals
SWF/CWF as UBI funding mechanism	Yes	Lansley (2016) ; Bruenig (2018)
Restricted drawdown lifetime accounts	Partial	Singapore CPF
Synthesis: existing pension →CWF→UBI	Novel?	This proposal

In summary, tax-funded UBI is the most direct way to redistribute automation-generated wealth, and standard economic arguments suggest it is among the least distortionary transfer mechanisms. Share-based models (individual shares → CWF → SWF) add complexity and fragility as a theoretical endpoint (Section 6.5), but existing pension infrastructure can serve as a practical transition vehicle toward a de facto CWF (Section 6.6), and a citizen or sovereign wealth fund could be gradually built up and increasingly replace direct taxes as a funding source for UBI. The strongest objection,

that UBI provides money but not *power*, is addressed above (Section 6.5): we argue that the citizen agency promised by share-based models is largely illusory once practical constraints are imposed, but a CWF may nonetheless serve as a useful complementary institution. Indeed, the pension fund pathway (Section 6.6) shows that existing mandatory pension systems can serve as institutional vehicles for a gradual transition toward a de facto UBI, exploiting decades of institutional trust and political entrenchment rather than requiring new institutions to be built from scratch. Objections regarding laziness and affordability are contradicted by empirical evidence and the unique economics of automation. Transitioning to UBI allows society to avoid the trap of forced employment and embrace a future of voluntary work.

Key Insights

- UBI dominates alternatives (job guarantees, means-tested welfare, share ownership) on simplicity, universality, and incentive alignment.
- A three-phase transition (CI → UAI → UHI) with pre-legislated automation triggers avoids political delay.
- Displaced workers can be compensated with equity at zero net cost to the firm ($\gamma = S/(P+S)$).
- Existing mandatory pension systems (e.g. Australian Superannuation) can serve as transition vehicles toward UBI, exploiting institutional trust and path dependence (Section 6.6).
- Pilot evidence consistently shows that modest income supplements do not trigger mass withdrawal from work (though the book argues that voluntary work reduction is in the long run desirable; see Section 4.5), while improving wellbeing and increasing entrepreneurship.

Open Problems

- Determining the optimal UBI level that balances adequacy with fiscal sustainability during the transition.
- Whether a citizens' wealth fund adds genuine citizen agency or merely introduces governance complexity.

Meridian — 2038: Life After Work

The pilot became permanent, then national. The federal Guaranteed Income Act of 2037 provides every adult \$2 600 per month, funded by a diversified automation levy, reformed capital gains, and the new Citizen AI Fund. Giwoo, who spent a decade defining herself by twelve-hour shifts, took eight difficult months to stop feeling guilty. Then she built a greenhouse. She now runs a community garden that feeds forty families; neighbors who once were patients stop by for the tomatoes and stay for the conversation. Angela passed away last spring. Giwoo still tends the roses they planted together. Marc builds custom furniture and hosts “Fix Anything Saturdays” in the garage: neighbors bring broken toasters, wobbly chairs, and questions about how things used to work. His brother Julian, a sound engineer whose mixing work vanished when AI mastered audio production, runs a live-music night at the community center every Friday. Chaeyeong published her automation-levy paper; three cities adopted her model. She is pursuing a PhD because the problem fascinates her, not because she needs the credential. Alice co-founded a citizen-science collective monitoring reef health along the coast. Judge Chloe retired from the bench and teaches a free ethics seminar at the community center. Levin was the last officer in his precinct to be reassigned. He fought it. He grieved it. He spent two months on the couch before Butz told him, “You were never a uniform, boy. You were the calm voice. Go find where that is needed.” He now volunteers with a youth mentoring program, not because anyone pays him, but because the teenagers still need someone who has seen the worst and chosen to stay kind. None of them earn much. All of them are busier than when they had jobs. Butz, watching from his armchair, closes Sunday dinner: “The difference is simple. Every hour is chosen.”

Chapter 7

Societal Transitions and The Future

Yet there is no country and no people, I think, who can look forward to the age of leisure and of abundance without a dread. For we have been trained too long to strive and not to enjoy.
— Keynes (1930, p.4)

7.1	Political Economy of Redistribution . . .	142
7.2	Will Democracy Prevail in a Post-Labor Society?	145
7.3	Democratic Fragility and Techno-Feudalism	150
7.4	Democratic Innovations for the Post-Labor Era	152
7.5	Social Order and Meaning of Life	155

The economics are tractable in principle (Chapter 3), but translating them into policy requires political will, social cohesion, and a credible answer to the search for meaning in a post-labor world. When redistribution fails, the consequences range from populism to **regulatory capture** and a transition-period underclass risk, yet democratic feedback loops nonetheless incentivize corrective action (Section 7.1 and Figure 7.1). A deeper structural question follows: if the state no longer depends on citizen labor for economic output, does the democratic social contract survive? The petrostate analogy, physical deterrence, and democratic innovations all bear on whether popular leverage can be preserved (Sections 7.2 to 7.4). The existential dimension is no less pressing: whether humans can find purpose, identity, and social order

when paid employment is optional rather than compulsory (Section 7.5). The conclusion is cautiously optimistic: democracies have the structural incentives to redistribute, and humans are genetically wired to thrive beyond the office.

7.1 Political Economy of Redistribution

While the economic case for redistribution appears tractable (Chapters 5 and 6), the political path is less certain. Failing to redistribute invites populism and social collapse, yet even well-intentioned democracies face regulatory capture by concentrated capital and a transitional underclass risk that may persist for a generation. The following paragraphs trace this tension: the consequences of inaction, the corrective forces inherent in democratic institutions, the limits of regulation, and the structural power of Big Tech.

Failing redistribution. Failing to ensure a decent living standard for displaced workers (Section 3.2) will have dire consequences. History shows that displaced populations can become entrenched underclasses long before democratic correction materializes (Section 7.1). In the extreme case, it leads to populism, social unrest, up to civil war. Short of this, it traps society in hyper-competition, vanity-arms-races, and/or BS-jobs (Section 4.5) without creating true value nor meaningfully reducing prices as discussed above: At the low-income, low-entry-barrier end it would drive down income even further. At the luxury end, there would be an arms race of producing and consuming vanity/status goods and services that could be redirected more sensibly. People can still enjoy their status games (“my yacht is longer than yours”) even if all yachts are taxed to half their length (cf. [Frank, 2007](#)).

Democracy and social welfare. Assume mass un(der)employment occurs without mitigation. Those negatively affected will desire and demand the government to somehow preserve their living standard. In a functioning democracy, this creates a massive incentive for political parties to advocate for more aggressive redistribution (via existing left/labor parties reviving/strengthening/returning to their original mission, or via a new party with such focus, or via referenda). If a sufficient majority of voters in a democracy are economically displaced, they will be highly incentivized to vote for a party with such policies, which can then push through the required political/economic/tax changes. This process should function independent of the favorite redistribution mechanism (see above for options) and affordable as per the idealized transition argument (Section 3.2). However, a potential caveat is that voters currently prefer work-conditional programs

(Lehner, 2026), though only mildly and this preference may erode when the make-work pathologies of such programs become apparent (Section 4.5).

Regulatory frameworks and governability. The inevitability of mass automation (Section 3.1) raises the question: can deliberate policy interventions slow or redirect technological progress? While national actors may attempt regulation, unilateral slowdowns likely fail due to international competitive pressures. Genewein et al. (2026) notes that nations that resist AI adoption face “military-economic adaptationism”: those that embrace automation gain economic and security advantages, forcing others to follow. However, global institutions, if coordinated, could establish thresholds and licensing requirements for frontier AI capabilities, embedding institutional gates that slow but do not halt progress. The challenge for policymakers is designing governance structures that account for the speed of AI development while building political legitimacy for inevitable job displacement and the redistributive policies required to manage it. Concrete action is beginning: ‘Trump Accounts’ grants every newborn a \$1 000 stock grant; and California’s Executive Order N-6-26 (Newsom, 2026) directs modernization of the state’s layoff-notification (WARN) Act for AI-driven displacement, review of displaced-worker safety nets (including severance, equity compensation, and subsidized employment), expansion of work-sharing programs, launch of an AI employment-impact dashboard, and support for worker-ownership models to broaden capital gains from productivity, representing the first major US state-level executive action specifically targeting AI workforce impacts.

Big Tech and political power. As AI concentrates economic power in a handful of technology firms, the primary threat shifts from their market dominance to their *political* power (Susskind, 2020, Chp.11). Unlike Standard Oil, whose danger was overpriced oil, Big Tech shapes liberty (autonomous systems that enforce behavioral constraints), democracy (algorithmically curated political information), and social justice (consequential decisions based on unconsented personal data) and decisions once made through democratic processes are increasingly made by private firms. Self-regulation is insufficient (engineers are not hired for ethical reasoning), and nationalization risks state abuse (cf. China’s social credit system). Susskind proposes a *Political Power Oversight Authority* modeled on competition regulators but tasked with constraining political rather than economic power, equipped with investigative, transparency, and breakup tools. Consumer satisfaction does not confer political legitimacy: economic success is not a license to shape citizens’ collective lives without accountability.

Transition-period underclass risk. The preceding arguments assume that democratic feedback eventually corrects for displacement (Section 7.1). But during a gradual transition, a technologically displaced underclass may form and become normalized before reaching the critical mass needed for political action: a boiling frog effect (Section 10.2). History provides sobering precedents:

The Engels' Pause and the Rust Belt. During the Industrial Revolution, output per worker grew by 46% while real wages rose by a mere 12% over the period 1780–1840 (the “Engels' Pause”), leaving workers to bear the costs of transition for three generations before sharing in the gains (Frey, 2019); technology's distributional impact is ultimately a political choice, and the default choice is that the powerful capture the gains (Acemoglu and Johnson, 2023). The more recent deindustrialization of the American Rust Belt produced a multi-generational decline characterized by opioid addiction, rising suicide rates, and falling life expectancy (Case and Deaton, 2020).¹ Crucially, the displaced did *not* organize effectively for redistribution: they turned to populism, substance abuse, and withdrawal. The democratic feedback loop of Section 7.1 simply did not fire. Similarly, local labor markets hit by Chinese import competition experienced persistent employment decline and political polarization; workers did not relocate or retrain as standard models predict (Autor et al., 2016).

The precariat: atomized and invisible. AI displacement may be structurally harder to organize against. Standing (2017) identifies a “precariat” of workers without stable employment, benefits, or occupational identity who are structurally incapable of effective collective action because they lack the workplace (the historical organizing node), union infrastructure, and stable community ties that enabled previous labor movements. AI-displaced workers will be atomized, dispersed, and increasingly invisible.

Normalization and moral exclusion. Perhaps most insidious is how quickly societies normalize their underclasses. American political discourse has consistently divided the poor into “deserving” and “undeserving” categories, adapting to each new wave of displacement (Katz, 2013): AI-displaced workers will predictably be framed as “those who didn't learn to (vibe-)code.” Harari (2018) warns of a “useless class”: not unemployed (which implies potential employability) but *irrelevant* to the economic system. Once “useless” enters the discourse, redistribution is reframed as charity to the unproductive

¹Case and Deaton (2020) trace an unprecedented rise in mortality among middle-aged white Americans without college degrees, driven by drug overdoses, suicide, and alcoholic liver disease: what they term “deaths of despair.”

rather than as a right. The psychological mechanism is *moral exclusion*:² groups placed outside the “scope of justice” have their suffering rendered acceptable (Opotow, 1990).

Compounding the risk, nations locked in an AI arms race face guns-versus-butter pressure against welfare spending during precisely the period it is most needed (cf. the military-economic adaptationism of Section 7.1).

Pre-commitment as antidote. The antidote is *institutional pre-commitment* while the political window remains open (see Section 6.2 for the general mechanisms: ratchet clauses, constitutional entrenchment, pre-legislated triggers). In addition, measures specific to the underclass risk are needed: when firms automate, mandating *displacement councils* (modeled on European works councils) can preserve organizational infrastructure for displaced workers. An *AI Displacement Accord* setting minimum social protection floors indexed to automation adoption could prevent a global race to the bottom (Chapter 8). The democratic feedback of Section 7.1 remains the central corrective mechanism, but it needs institutional scaffolding to survive the transition period.

Taken together, the preceding chapters show that the transition to a post-labor society is economically affordable (Section 3.2), socially desirable (Section 7.5), and politically feasible provided the right institutions are legislated early enough (Sections 6.2 and 7.4).

7.2 Will Democracy Prevail in a Post-Labor Society?

If human labor becomes economically superfluous, the fiscal and political bargain that underpins liberal democracy may unravel. Boix (2019) argues that the absolute obsolescence of human labor threatens the foundational social contract of modern democracies (see also Huang and Manning 2025 for a policy-oriented treatment). The following paragraphs examine the mechanism by which citizen leverage erodes once the state no longer depends on taxpayer labor, draw a cautionary parallel with petrostate rentierism,³

²Opotow (1990) defines moral exclusion as the process by which individuals or groups are placed outside the boundary within which fairness rules and resource-sharing obligations apply. It operates along a continuum from active dehumanization to passive unconcern. In milder forms, it manifests as simple failure to recognize undeserved suffering. A technologically displaced underclass risks exclusion through ordinary mechanisms: once outside the “scope of justice,” their claims to community resources are delegitimized by definition.

³A *rentier state* derives a substantial share of its revenue from external rents (typically natural-resource exports) rather than from taxing its citizens’ productive activity. The term,

and assess how robust the causal link between economic contribution and democratic survival really is.




Figure 7.1: Democracy feedback loop. Displacement creates electoral pressure for redistribution, which stabilizes society and enables orderly transition. The dashed “failure path” shows the alternative: without redistribution, displacement fuels populism and threatens democratic stability. See Section 7.2.

Obsolescence and democratic survival. Liberal democracies historically emerged because the state relied on taxpayers’ financial resources and labor to project power. If AI systems generate the vast majority of national wealth, and the state shifts from collecting income taxes from citizens to merely distributing technological dividends, some argue that the power dynamic shifts dangerously (Huang and Manning, 2025). If capital owners no longer require mass workforces to generate output, they systematically lose the structural incentive to compromise politically with the populace. The famous democratic maxim carries a dark inverse: *no representation without taxation*. If citizens contribute nothing economically, their political leverage evaporates, risking a slide into managerial feudalism (Huang and Manning, 2025; Boix, 2019; Ross, 2001) (see Figures 7.3 and 7.4). Boix (2019) provides the formal political-economy model: the democratic franchise historically depended on the economic usefulness of mass labor, and automation breaks this complementary relationship. Stiefenhofer (2025) formalizes this “power shift” using extended production functions with AGI labor and AGI capital, showing that income accrues entirely to AGI-owned capital once human labor is fully substituted, entrenching what he terms *techno-feudalism*; Saint-Paul

popularized by Mahdavy (1970) and Beblawi (1987), implies that the government’s fiscal independence from its population weakens democratic accountability.

(2025) shows that the resulting demand collapse can even be self-defeating for the oligarchs themselves (see Section 7.2).

Petrostate analogy. The parallel with resource-cursed petrostate rentierism (Ross, 2001) is suggestive: states funded by rents rather than citizen taxes tend to have weaker democratic institutions. If the analogy holds, the AI case could be more acute than oil, because it threatens to sever both the fiscal *and* the labor link between citizens and the state, whereas oil only weakened the fiscal link (Table 7.2).

Table 7.2: Speculative comparison of oil-funded and AI-funded rentierism (illustrative, not probative; causation between oil rents and democratic erosion is contested, though the case for AI is arguably stronger, see text).

Factor	Oil rents	AI rents
Revenue replacement	Tax revenue only	Tax revenue <i>and</i> labor
Citizen economic role	Still work (non-oil)	Potentially irrelevant
Revenue concentration	Accrues to state	Corporations → state
Military dependence	Still needs soldiers	Autonomous weapons
Speed of transition	Gradual (years)	Rapid (software scales)

Causation is contested. However, most petrostates were autocratic *before* discovering oil: institutions are the deeper cause, resources merely an amplifier (Acemoglu and Robinson, 2012). A longer-panel study by Haber and Menaldo (2011) using within-country fixed effects over 1800–2006 found no robust causal effect of resource windfalls on regime type. By using within-country variation over two centuries (rather than cross-country snapshots), Haber and Menaldo control for the deep institutional and cultural differences that confound cross-sectional studies like Ross (2001). In many specifications they find a “resource blessing” rather than a curse, suggesting that the correlation between oil and autocracy reflects pre-existing institutional weakness, not a causal effect of resource wealth.

Cases and outliers. The few suggestive cases are all confounded. Venezuela, Latin America’s oldest democracy (Polity +9 from 1958), saw the 1970s oil boom fuel massive state expansion and patronage; by 1998 the system had hollowed out sufficiently for Chávez to win power, and Polity eventually fell to −3, but ideology and institutional decay were independent factors. Russia’s democratic opening in the 1990s (Polity ~+4) coincided with low oil prices; Putin’s consolidation tracked the 2000s commodity surge, but the KGB institutional legacy provided an independent path to autocracy. Nigeria’s functioning (if flawed) post-independence institutions collapsed

during the 1970s oil boom, with rents funding military coups and destroying accountability until re-democratization in 1999.

The outliers are more telling than the cases. Resource-rich democracies abound: Norway (rents 10% of GDP, Polity +10) created its sovereign wealth fund *because* strong institutions were already in place; Chile (copper-dependent) democratized *during* a commodity boom; Botswana, the “diamond democracy,” maintained Polity +8 throughout; Australia, Mongolia, and Canada exhibit similarly robust institutions despite heavy resource dependence. Conversely, China (Polity -7), Belarus (-7), Cuba (-5), and Vietnam (-7) are deeply authoritarian without significant resource rents. The pattern suggests that **institutional quality at the time of resource discovery** is the key mediating variable, not resource wealth per se.

The rentier mechanism and the AI analogy. The theoretical mechanism behind the [resource curse](#) is clean: (i) government funded by resource rents no longer needs citizen taxes; (ii) no taxation erodes accountability (“no taxation, no representation”); (iii) citizens lose leverage and democratic institutions atrophy. But empirically, the mechanism is less airtight than it appears. The United States funded itself primarily through tariffs for its first 150 years without democratic collapse; many non-democracies tax heavily (the Soviet Union, China); and the mechanism assumes that fiscal dependence is citizens’ *only* source of leverage, whereas labor power, military service, and sheer numbers have historically mattered at least as much.

For the AI analogy, the cross-national correlation matters less than the *mechanism* itself: when a state no longer depends on citizens’ economic contribution, the structural incentive for democratic compromise weakens. AI could sever that link more completely than oil ever did, replacing not just tax revenue but labor itself. However, the analogy has a structural limit: petrostates export commodities to foreign markets, making domestic purchasing power largely irrelevant to elite wealth. An AI-driven consumer economy, by contrast, profits from mass domestic demand: [Saint-Paul \(2025\)](#) shows that the resulting demand collapse can be self-defeating even for the oligarchs themselves, giving capital owners a self-interested reason to support redistribution that petrostate autocrats never faced.⁴⁵

⁴The asymmetry is not about domestic saturation but about market structure: oil revenue derives overwhelmingly from bulk commodity exports (the domestic retail market is negligible relative to global crude trade), whereas AI services can be sold to every individual and business domestically, making a prosperous middle class directly profitable for AI companies. A petrostate autocrat can ignore its citizens’ purchasing power; an AI-economy oligarch cannot, because its marginal customer *is* the domestic consumer.

⁵This incentive presupposes that capital owners seek monetary wealth and market influence beyond material self-sufficiency; under full autarky with self-replicating factories




Figure 7.3: *Shifting Power Dynamics: Two potential paths for society as AGI concentrates wealth. Path A leads to managerial feudalism without redistribution, while Path B maintains democratic leverage through GI.*




Figure 7.4: *The resource curse: countries with higher dependence on natural resource rents tend to score lower on democracy indices. Data: World Bank total natural resources rents (% of GDP, 2021) vs Polity V democracy score (−10 to +10, 2021); N=75 countries. The OLS trend line ($R^2=0.32$) confirms a robust negative correlation, though substantial variance remains (notable outliers: Norway, Chile, Mongolia are resource-rich democracies; China, Belarus are autocratic despite low rents). Analogously, if the state becomes primarily funded by AI-generated revenue rather than citizen taxes, democratic accountability may similarly erode (Ross, 2001).*

7.3 Democratic Fragility and Techno-Feudalism

The previous section established how the social contract could unravel; this section asks what prevents that unraveling and how bad things could get if prevention fails. Physical deterrence provides a residual check on concentrated power, but existing democracies already systematically underserve their majorities, and the worst-case endpoint is a techno-feudal order in which AI rents replace both wages and political accountability.

Deterrence and the social contract. Although the ruling class's reliance on citizen labor in time led to democracy, losing this reliance will not necessarily cause democracy to collapse. As long as the population can physically revolt against dismantling democracy, and threaten the comfortable lifestyle of those in power, the incumbents have an incentive to avoid a civil war. This is perhaps the only good (among many arguably bad) argument for the right of individuals to own firearms: civilian arms serve as a credible last-resort deterrent against state tyranny, the practical embodiment of the revolution threat that formal models identify as the historical driver of democratic concessions (Boix, 2019). However, this deterrent erodes rapidly once autonomous weapons and robotic enforcement become available, allowing a regime to suppress revolt without relying on potentially sympathetic human soldiers (see Table 7.2, row "Military dependence"). The deterrence argument is genuinely contested. Proponents invoke the Framers' rationale that an armed citizenry is a structural check on tyranny, a position formalized in models where revolution threats compel redistribution (Boix, 2019; Acemoglu and Robinson, 2012). Critics counter that in practice, civilian firearms have never overturned a modern industrialized state, that the asymmetry between citizen arms and state military capacity renders the deterrent largely symbolic, and that any government anticipating revolt will regulate and enforce disarmament via (robotic) police well before the threat materializes. The more promising structural checks are discussed in Section 7.4. As argued above, a universal AI-fund as suggested by Altman (2021) and Huang and Manning (2025) and others is no safer than UBI, as both are in the end dependent on the state's continued adherence to the rule of law and the social contract.

Do democracies actually serve the majority? The preceding analysis implicitly relies on democratic responsiveness: that a displaced majority will vote its interests into policy (Section 7.1). But substantial evidence suggests that even well-functioning democracies systematically underserve

(Section 3.4), the structural incentive weakens.

the majority. Gilens and Page (2014), analyzing 1 779 US policy issues, find that economic elites and organized business groups have substantial independent impact on government policy, while average citizens have “little or no independent influence.”⁶ Bartels (2008) shows that US senators are responsive to the preferences of affluent constituents and essentially unresponsive to those of low-income voters, a structural bias that AI-driven wealth concentration would amplify. Regulatory capture (Stigler, 1971) compounds the problem: industries systematically co-opt the agencies designed to constrain them. For AI, the information asymmetry between frontier labs and regulators is extreme, as policymakers cannot evaluate model capabilities without industry cooperation. The baseline is also deteriorating. Democratic backsliding has accelerated globally, and modern democracies increasingly die not through military coups but through elected leaders gradually eroding institutional guardrails from within (Levitsky and Ziblatt, 2018). Przeworski (2019) argues that democracies require specific economic conditions (moderate inequality, sustained growth) to survive; when inequality crosses a threshold, political conflicts become zero-sum and democratic institutions destabilize. If AI creates extreme capital-labor inequality (Figures 2.11 and 5.3), the Gini coefficient could cross precisely this threshold. These are not hypothetical risks: they describe current reality, *before* AI disrupts labor markets.

Worst case: techno-feudalism and the intelligence curse. If the pessimistic trajectory materializes, Varoufakis (2023) describes the endpoint: *techno-feudalism*, in which platform owners extract rents from an economically irrelevant populace, with profit replaced by rent and capitalism itself superseded by something worse. The book’s own resource-curse analysis (Table 7.2 and Section 7.2) extends naturally to what we might call an *intelligence curse*, which is strictly worse: AI rents are more concentrated than oil rents (a few firms, not spread across geography), enable surveillance and enforcement without human soldiers, and scale at software speed rather than at the pace of physical infrastructure. Acemoglu and Robinson (2012) identify the deeper structural risk: extractive institutions (designed to concentrate resources in the hands of the few) create vicious cycles because political and economic extraction reinforce each other. Historically, these cycles were broken by internal contradictions: extractive elites *needed* human labor, creating leverage for reform. AI threatens to remove this contradiction,

⁶The study uses multivariate analysis to disentangle the independent effects of elite preferences, organized interest groups, and average citizens on policy outcomes. When elite preferences diverge from those of the median voter, elites almost always prevail: the probability of a policy being adopted is nearly flat across levels of average-citizen support, but rises steeply with elite support.




Figure 7.5: Top 1% pre-tax national income share by country (2024). Latin America (Brazil 27%, Chile 27%, Mexico 26%), MENA (Iraq 45%, Gulf states 24–26%), Russia (24%), India (23%), South Africa (22%), and the USA (21%) exhibit the highest concentration, while Northern and Western Europe show the lowest (Netherlands 7%, Nordics 9–10%, Belgium/Austria ≈9%). Source: World Inequality Database (2024) <https://wid.world/>.

making extractive institutions self-sustaining. Additionally, Fukuyama’s (2014) “vetocracy” diagnosis suggests that democratic systems with many veto points (courts, filibusters, federal-state fragmentation) can block action but not take it: the speed of AI development may be incompatible with the slow pace of democratic consensus-building, a “pacing problem” in which power structures lock in before governance catches up. Yet even Scheidel’s (2017) pessimism (historically, only catastrophic violence has produced significant inequality reduction) admits an inversion: AGI itself may constitute the first *peaceful* structural shock of sufficient magnitude to force systemic reform, provided institutions are designed to channel the disruption toward redistribution rather than further concentration.

7.4 Democratic Innovations for the Post-Labor Era

The preceding sections painted a sobering picture, but the historical record is not uniformly grim. Democracies have absorbed structural shocks of comparable magnitude before, and institutional innovations ranging from direct democracy to sortition offer structural resistance to oligarchic capture. More speculative power-distribution mechanisms could further entrench citizen leverage in a post-labor economy.

Democracies have absorbed comparable shocks. The difficulties are real but historically incomplete. Democracies have repeatedly absorbed struc-

tural shocks of comparable magnitude, often messily and with multi-decade delay, but ultimately successfully. Britain's industrial revolution produced the Engels' Pause described in Section 7.1, but the response, however slow, included factory legislation (1833), public education (1870), and the welfare state (1906–1948). The United States absorbed the Great Depression through the New Deal and absorbed post-war demobilization (twelve million returning soldiers) through the GI Bill and the post-war institutional framework (including Bretton Woods), enabling the largest expansion of the middle class in history. The Nordic countries built their extensive welfare states not during peacetime tranquility but in response to the social crises of early industrialization and the interwar period. Germany's post-reunification integration (still incomplete after 35 years) illustrates both the possibility and the generational timescale of absorbing a structural economic shock within a democratic framework. These precedents do not guarantee success: each involved decades of political struggle, and none faced the unique challenge of human economic obsolescence. But they refute the claim that structural transformation beyond Scheidel's violent shocks is historically unprecedented.

Democratic innovations for the AI age. Despite the preceding warnings, institutional designs exist that are structurally more resistant to elite capture and more adaptive than conventional representative democracy. Swiss direct democracy provides a structural veto against special-interest legislation: any law can be challenged by referendum (50 000 signatures), and constitutional amendments require mandatory popular vote. [Feld and Kirchgässner \(2000\)](#) show that Swiss cantons with mandatory fiscal referenda achieve better fiscal discipline and closer alignment between policy and citizen preferences, precisely the properties needed to resist AI-era regulatory capture. Sortition (random selection of citizens for governance roles) offers a more radical structural defense. [Landemore \(2020\)](#) argues that electoral representation inherently produces an aristocratic governing class, while sortition combined with deliberation captures the epistemic diversity of the population. [Fishkin \(2018\)](#) demonstrates across hundreds of deliberative polls worldwide that randomly selected citizens produce more informed and less polarized policy preferences. Crucially, sortition is structurally immune to campaign-finance corruption: there are no campaigns to fund and no candidates to lobby. Ireland's Citizens' Assembly (2016–2018), composed of 99 randomly selected citizens, resolved the country's decades-long deadlock on abortion with a referendum that passed with 66.4% support, demonstrating that sortition-based deliberation can handle politically toxic issues. Permanent sortition-based bodies for AI governance could combine this structural integrity with

the technical expertise such decisions require. The structural goal these innovations serve is what Pettit (2012) calls *non-domination*: a political order in which no citizen is subject to another's arbitrary will. Sortition, direct-democratic vetoes, and constitutional entrenchment are complementary institutional safeguards against the concentration of power that AI-generated wealth makes possible. Most fundamentally, UBI or AI-dividend rights can be constitutionally entrenched, making their reduction as difficult as revoking the franchise, a commitment device that Rawls (1971) would recognize as the institutional expression of the difference principle. The Alaska Permanent Fund Dividend (Section 7.1) and Norway's sovereign wealth fund demonstrate that such pre-commitment devices can survive decades of political pressure.

Exotic power-distribution mechanisms. According to Tinbergen's rule, separate targets need separate instruments (Section 6.5), so the question of *redistribution* (how to share wealth) and the question of *power* (how to prevent its concentration) need not be solved by the same mechanism. Attempting to optimize for both simultaneously (as citizen-wealth-fund proposals do, Section 6.5) may yield worse outcomes than addressing each independently. This reframing suggests a broader search for power-distribution mechanisms that complement, rather than replace, UBI. Table 7.6 surveys emerging proposals, ranging from mechanisms already in early policy deployment to those still in the conceptual stage. Among the more developed: AI itself can enhance rather than undermine democratic deliberation. Taiwan's vTaiwan platform uses AI-powered opinion clustering to identify consensus positions across partisan divides, successfully informing regulation of ride-sharing and other technology issues. Quadratic voting and quadratic funding (Posner and Weyl, 2018; Buterin et al., 2019) allow citizens to express intensity of preference without money domination, distributing agenda-setting power toward broad-based preferences and potentially useful for decisions about AI governance where those most affected by automation can express stronger preferences. At the more speculative end, *compute governance* (Sastry et al., 2024) exploits the fact that compute is the most governable of AI's three inputs (unlike data or algorithms, it is physical, measurable, and excludable). Compute allocation boards, licensing regimes, and public compute utilities could serve as levers for democratic oversight of AI capability development. Similarly, AI safety research on *corrigibility* (Soares et al., 2015; Hadfield-Menell et al., 2017) (the property of an AI that permits its principals to correct or shut it down) could be repurposed as institutional design: AI systems hardcoded to defer to democratic institutions, not merely to their operators, would create a structural check that does not depend on citizens' physical capacity

for revolt (Section 7.3). A final, perhaps counterintuitive observation: a *post-labor* society may actually *strengthen* democratic participation. Ancient Athens' democracy functioned partly because slaves freed citizens for politics; AI could serve as the morally acceptable equivalent, creating a *deliberation dividend* in which citizens who no longer need to work have the time and cognitive bandwidth for the sustained civic engagement that complex AI governance demands. The technology that threatens democracy may also provide the conditions for its renewal, but only if institutional design channels this potential proactively.

7.5 Social Order and Meaning of Life

Some argue that jobs are necessary to keep life meaningful and to keep people out of trouble. The objection, however, conflates *jobs* with *work*: the former may vanish, but the latter persists (Section 7.5). This section examines the evidence.

Jobs versus work distinction. However, this objection conflates *jobs* with *work* (the same conflation underlying the mental health objections rebutted in Section 6.3). As the property-cluster analysis of Table 2.2 clarifies, 'jobs' (compensated, regular, asymmetric roles in a market economy; cf. the working definition in Section 2.2) may largely vanish, but 'work' (purposeful, effortful, socially valuable activity) will persist: caregiving, cultural, volunteering, art, learning, community service, and sport (Section 2.2) all satisfy several job-like properties ([LN], [LR], [LS], [LK]) without requiring compensation ([LC]) or involuntariness ([LI]). Philosophically, this is because work (as Tyssedal (2025) argues) is defined by producing an *external* benefit: unlike self-care, work can be outsourced or delegated, which is precisely what makes its distribution a question of justice, and what ensures it persists independently of any particular mode of production. A prominent randomized controlled trial by Hussam et al. (2022) in Rohingya refugee camps demonstrates that structured employment provides massive psychosocial well-being gains compared to equivalent cash transfers, with refugees willing to forgo cash windfalls to work for free.⁷ This confirms Jahoda's (1982) latent deprivation theory: employment provides critical latent benefits (time structure, social contact, and collective purpose) beyond pecuniary compensation. Crucially, however, this is a validation of structured, purposeful *work* under extreme

⁷In the experiment, participants were offered the choice between a \$1.53 cash payment or a day's structured work. The majority chose work, suggesting that employment provides non-pecuniary value (purpose, structure, social contact) that cash alone cannot substitute.

Table 7.6: Emerging power-distribution mechanisms for the AI age. Maturity: ✓ = in policy/proven, ◦ = early/pilot, ? = speculative.

Mechanism	Core Idea	Maturity	Key Limitation
Swiss direct democracy	Mandatory referenda create structural veto against special-interest legislation (Feld and Kirchgässner, 2000)	✓	Slow; requires informed electorate
Sortition / citizens’ assemblies	Random selection captures population diversity; structurally lobbying-proof (Landemore, 2020; Fishkin, 2018)	✓	Scaling beyond single-issue assemblies
Constitutional entrenchment	Lock in UBI/AI-dividend rights via supermajority-protected legislation; Alaska PFD as proof of concept (Widerquist and Howard, 2012)	✓	Rigidity; hard to calibrate levels
AI-assisted deliberation	AI-powered opinion clustering (vTaiwan/Polis) identifies consensus across divides	◦	“Neutral” AI reflects its training data
Quadratic voting/funding	Express intensity of preference without money domination (Posner and Weyl, 2018; Buterin et al., 2019)	◦	Sybil attacks; collusion; requires identity infrastructure
Compute governance	Regulate access to AI training compute via licenses, quotas, or public utilities (Sastry et al., 2024)	◦	Assumes compute remains a bottleneck
Alignment tax	Safety requirements slow deployment, trading speed for reduced risk (Sastry et al., 2024)	◦	Developm. cost; neither revenue nor power-distr. mechanism per se
Corrigibility by design	AI systems with hardcoded deference to democratic institutions (Soares et al., 2015; Hadfield-Menell et al., 2017)	?	Alignment is unsolved; loyalty vs. capability tension
Algorithmic social contracts	Encode redistribution rules in tamper-resistant smart contracts (De Filippi and Wright, 2018)	?	Rigidity; garbage-in governance
Data dignity	Treat personal data as labor deserving compensation (Lanier and Weyl, 2018)	?	Marginal value of individual data near-zero once models trained
Network states	Jurisdictional competition via digital-first governance (Srinivasan, 2022)	?	Favors mobile capital owners; race-to-bottom risk

boredom and legal employment bans, not an indictment of UBI. Market employment under capitalism is coercive and involuntary ([LI]). Under a UBI, individuals are freed from coercive market jobs but remain fully capable of engaging in voluntary, self-directed, structured unpaid work (Section 2.2) that preserves Jahoda's latent benefits without the exploitation of market labor. Consequently, the economic identities of automation redistribution remain valid; UBI does not force idleness, but merely shifts labor from coercive to voluntary domains. The meaning-of-life concern is therefore not about the disappearance of *activity*, but about the loss of the *institutional wrapper* that currently bundles purpose, identity, and income into a single package.

Meaning beyond employment. Depending on the survey, about half the working population would continue in their roles even without financial need, and even among those who would stop, many currently derive meaning from their occupation. On the other hand, more than half find their primary meaning and purpose outside of work, in family and friends, leisure and hobbies, community service, and/or spirituality or faith (Morse and Weiss, 1955; Silver et al., 2021). Undoubtedly for some whose primary or sole purpose in life is their job it will be difficult to find comparable meaning elsewhere. However, this is most likely a transitional problem, though one that may persist for a generation and carry risks of social unrest. Children are raised, indoctrinated, and exposed to become good members of society, with strong emphasis on educating them towards employability. Humans are also genetically primed to enjoy competition and cooperation and (some kind of) work. Equally, they naturally enjoy playing, games, socializing with friends and family, sports, learning, and many other activities. By de-emphasizing the 'work-for-money' mantra and nurturing their other interests, raising the next generation to thrive without traditional jobs actually appears easier than our current approach to convince them that repetitive labor is fun. As Susskind (2020, Chp.12) observes, just as "We no longer need to train young people to be warriors. In the future, we may no longer need to train them to work, either, but will have to teach them to flourish through leisure instead".

Moreover, the norm that every man must hold a paying job is historically contingent, not a human universal. For most of recorded history, the aristocratic and landed classes regarded manual and commercial labor as beneath them; leisure, study, and civic participation were the marks of high status. The Protestant work ethic and industrial capitalism inverted this hierarchy, but the inversion is recent and far from complete: society already accepts that children, full-time parents, and retirees lead legitimate lives without holding jobs. The gradual normalization of stay-at-home fathers

in many Western countries shows that even deeply gendered employment norms can shift within a single generation. If UBI removes the financial coercion to work, the ‘aristocratic’ norm of choosing one’s pursuits could gradually extend to the entire population: not as idleness, but as the freedom to engage in voluntary, purposeful activity. In this sense, the upper classes offer a centuries-old proof of concept: they demonstrate that humans can lead rich, fulfilled lives when work is optional rather than obligatory.

Historical leisure societies. The claim that humans can thrive without jobs is not a utopian speculation: it is attested across millennia and cultures. In classical Athens, *scholē* (leisure) was the precondition for civic and intellectual life: the very word became our “school.” Aristotle held that “we work in order to have leisure” (*Politics* VIII.3, 1337b), and the citizen class (roughly 30 000 out of a population of 300 000) delegated manual labor to slaves and metics precisely to free time for philosophy, politics, and the arts. In Rome, the distinction between *otium* (productive leisure) and *negotium* (business: literally “not-leisure”) placed the highest human activities outside the sphere of commerce. Cicero composed his philosophical treatises during periods of *otium*; Seneca recommended withdrawal from public affairs as the condition for moral self-examination (*De Brevitate Vitae*); and the patronage system sustained poets and scholars so that they need not demean themselves with trade (*De Officiis* I.150–151). In the Confucian tradition, Mencius distinguished “those who labor with their minds” from “those who labor with their strength” (*Mencius* III.A.4), and the imperial examination system (*keju*, ca. 605–1905 CE) institutionalized this hierarchy for 1 300 years: the *junzi* (gentleman-scholar) ideal placed cultivation, study, and moral self-improvement above manual productivity. Even pre-agricultural foragers achieved material sufficiency with remarkably little labor: Sahlin’s (1972) “original affluent society” data show the Dobe !Kung San averaging roughly 15 hours per week of subsistence work, and Australian Aborigines 4–5 hours per day; the shift to settled agriculture actually *increased* working hours and decreased quality of life (Scott, 2017). This centuries-long work ethic is thus a cultural artifact, not a biological constant: barely three centuries old and already fraying.

Competition and play as substitutes. For those with an inborn competitive nature striving for status and power or fame (past emperors, current CEOs), structured competitive games (sports, chess, e-sports, ...) can serve as suitable, non-destructive substitutes. Structured gaming has profound cognitive and emotional benefits, and satisfies the fundamental psychological needs for autonomy, competence, and relatedness, so it can in principle simulate the competence and flow states traditionally derived from professional mastery

(Tyack and Mekler, 2024): a core mechanism in the “Utopia of Games” proposed by Danaher (2019). On the other hand, completely replacing labor by games may not be psychologically sustainable in the long run because it lacks authentic ‘world-building’ (Pelligra and Sacco, 2023). Even if true, playing games is not the only post-job activity left (see the unpaid work categories in Section 2.2), and this “crisis of meaning” is likely only a transitional problem of a generation trained that instrumental toil is the meaning of life (Pelligra and Sacco, 2023). Indeed, children naturally enjoy play and they would continue to do so if left to their own devices, as evidenced by some adults who managed to escape the societal or financial pressure to work (see Figure 7.7).




Figure 7.7: Sources of life meaning (17-country median) from an open-ended survey of 18 850 adults in advanced economies; respondents could mention multiple topics, so percentages sum to more than 100%. Paid employment ranks second but below family; friends, health, and material well-being also score highly, suggesting that a post-work society need not face a crisis of meaning. Data: Silver et al. (2021).

For a deeper philosophical dive into how humanity can preserve meaning in a world of post-labor material abundance and beyond, see Bostrom (2024).

Broader social dividends. UBI pilot programs worldwide have consistently demonstrated social benefits that extend far beyond the labor market (Table 6.1). Recipients show improved mental health and reduced anxiety, with permanent income security eliminating the cognitive tax of financial precarity (Bohmann et al., 2025; Bernhard et al., 2026). Physical health outcomes improve as well: reduced financial stress leads to lower hospitalization rates and better preventive care (Gentilini et al., 2020). Several pilots report reductions in economically motivated crime. Beyond these individual-level

effects, unconditional income enables activities that markets chronically undervalue. Parents can invest in child-rearing without financial penalty; caregivers (predominantly women) receive implicit compensation for domestic labor that the formal economy ignores. Ireland's basic income for artists pilot demonstrated outsized socio-economic returns from supporting creative work (Hogan, 2026). Students can pursue education and retraining without crippling debt, a critical advantage as skill requirements shift rapidly under automation. Contributors to open-source software, Wikipedia, citizen science, and other public goods (currently volunteering at personal financial cost) could sustain these efforts without hardship. And volunteers freed from survival jobs can devote time to community service, mentoring, and mutual aid. The cumulative picture is one of a society that, freed from the compulsion to sell labor for survival, redirects human energy toward activities that are intrinsically valuable but economically invisible.

Finally, as discussed by Genewein et al. (2026), ASI systems may eventually form super-intelligent collectives that could help coordinate and facilitate the transition, providing social structures and opportunities for human engagement that transcend traditional employment.

Key Insights

- Democracy survives if institutions adapt: the petrostate analogy is imperfect because AI wealth can serve domestic consumers, unlike oil rents.
- The security of any post-labor social contract in the end rests on citizens' capacity for collective action, not on the legal form of the transfer.
- UBI decouples survival from employment, enabling a voluntary social sphere of care, creativity, mentoring, and civic participation.
- Meaning does not require market employment: human flourishing expands when freed from compulsory labor.
- Democratic innovations (participatory budgeting, liquid democracy, citizen assemblies) can sustain accountability in a post-labor state.

Open Problems

- Whether democratic institutions can adapt fast enough to prevent authoritarian consolidation during the transition.
- The “meaning crisis”: how to support populations through the psychological adjustment from work-based to leisure-based identity.
- Governance of ASI collectives: who sets objectives for super-intelligent systems that coordinate societal transitions?

Meridian — 2040: The Global Picture

Meridian's story is not universal. Chaeyeong's PhD advisor introduces her to Tan, a Vietnamese software engineer who moved to Meridian after the textile-adjacent tech firms in Ho Chi Minh City collapsed. Not because demand for clothes fell (it rose), but because American and Chinese companies reshored production to fully automated facilities closer to consumers. Tan's hometown lost 35 000 jobs in eighteen months. Vietnam has no automation levy and no sovereign fund; its government tried a coding-bootcamp initiative, but the graduates were outperformed by AI before their certificates arrived. "I wrote software that helped factories run," Tan tells Chaeyeong. "Now your AI writes the software and your robots run the factories. My skills are not obsolete. My country's entire development model is." Marc reads about similar collapses in Bangladesh, the Philippines, and Kenya. Alice's reef-monitoring collective gets a volunteer from Nairobi who tells a parallel story: Kenya's call-center industry, which once employed 50 000 young graduates, evaporated in two years when AI handled customer service in every language simultaneously. Butz, a centenarian this spring, settles it at Sunday dinner: "Prosperity, like the weather, does not respect borders. But the fences remain anyway."

Chapter 8

International Challenges

We cannot simultaneously pursue democracy, national determination, and economic globalization.
— Rodrik (2011, p.xvii)

8.1	Trade Imbalance and AI Dependency . .	164
8.2	Tariffs and Currency Adjustment	166
8.3	Tax Design for a Borderless AI Economy	168
8.4	Capital and Talent Flight	169
8.5	AI Colonialism and the Global Develop- ment Gap	171
8.6	The Double-Edged Sword of Global Co- ordination	175

Automation is not just a domestic challenge. The global distribution of AI capabilities creates severe international asymmetries in trade, taxation, capital flows, and development prospects. The trade imbalance between AI leaders and followers (Section 8.1) can be partially mitigated by tariffs and currency adjustment (Section 8.2), but a more durable solution requires re-designing tax systems for a borderless digital economy (Section 8.3). Capital and talent flight, often invoked as a showstopper for aggressive redistribution, turns out to be more limited than feared (Section 8.4). Beyond the advanced economies, AI threatens to undermine the traditional development ladder itself, deepening the global development gap (Section 8.5). Finally, international coordination, while essential for preventing races to the bottom, is a double-edged sword that can also block beneficial automation (Section 8.6).

8.1 Trade Imbalance and AI Dependency

Asymmetric AI adoption creates a structural trade imbalance between technology leaders and followers. This section uses a stylized USA–Australia model to illustrate value extraction, then examines a transition-period counterpoint in which cognitive-first automation temporarily favors labor-intensive economies. Australia is a wealthy, advanced economy, which makes the example all the more instructive: even a high-income nation with strong institutions faces structural value extraction when it lacks domestic AI capacity. For developing nations, the dynamics are far more severe (Section 8.5).

Value extraction from AI followers. Assume the USA is substantially ahead in technological development. It provides all AI services and robots, extracting massive profits from another country, say, Australia (see Figure 8.1).




Figure 8.1: International Trade Imbalance: Asymmetrical AI adoption leads to value extraction from lagging nations to leader nations.

AI permeates the Australian economy, eroding its labor tax base as workers are replaced. Land and resource taxes (Table 5.5) remain collectible but alone may be insufficient to fund UBI at scale. Global macroeconomic modeling by the OECD confirms that AI will substantially reduce operational trade costs and boost global trade flows, but these gains will be distributed highly unevenly, heavily favoring technologically advanced nations (OECD, 2026). The IMF projects that AI diffusion will exacerbate cross-country income inequality, as advanced nations monopolize lucrative AI services and control the highly tradable digital components of restructured global value chains (Cerutti et al., 2025) (see Figures 8.2 and 7.5).

A transition-period counterpoint. However, this pessimistic picture applies in full only once both cognitive *and* physical labor are automated. During the transition period in which AI automates cognitive work but affordable general-purpose robotics has not yet arrived (cf. the staged au-

tomation sequence in [Korinek and Suh, 2024](#)), the distributional dynamics are more nuanced. AI sharply reduces the cost of cognitive labor relative to manual labor ([Acemoglu and Restrepo, 2022](#); [Korinek, 2024](#)). Developing countries with large physical labor forces are net importers of knowledge-intensive services (engineering, design, finance, software) and net exporters of labor-intensive goods ([Korinek and Stiglitz, 2021](#)), so cheaper cognitive imports lower their input costs and improve their terms of trade. Meanwhile, their comparative advantage in physical labor and raw materials remains intact, and because their labor forces are concentrated in manual and agricultural work, cognitive automation displaces fewer domestic workers than in advanced economies. (A notable exception: economies reliant on exporting routine cognitive services, such as business process outsourcing (BPO), call centers, and data entry, face immediate displacement from current AI; India and the Philippines are leading examples.) Applying [Korinek and Stiglitz's](#) model in reverse (factor-saving technology hurts exporters of the saved factor) implies that, during this window, the terms of trade may temporarily shift in favor of labor-intensive economies, even though AI-leading nations still gain more in absolute terms. Once general-purpose robotics matures and physical labor is also automated, this temporary advantage vanishes and the full trade-imbalance dynamics described above reassert themselves.




Figure 8.2: AI Preparedness Index by world region. Advanced economies lead in digital infrastructure, human capital, and institutional readiness, widening the gap between advanced and developing nations. Adapted from [Cerutti et al. \(2025\)](#).

8.2 Tariffs and Currency Adjustment

Given the severe trade imbalance just described (Section 8.1), follower nations have two blunt instruments: import tariffs and currency adjustment. Both carry a fundamental trade-off between shielding domestic employment and cutting the country off from the benefits of AGI. While margins are reduced, the USA still has an interest in selling over-capacity overseas if /since production costs are essentially zero. However, this does not balance trade if Australia has nothing to offer in return. (Resource-exporting nations like Australia can partially offset this dynamic through commodity demand from AI infrastructure build-out, but the imbalance intensifies for genuinely resource-poor countries.) With a floating currency and no other interference, standard economic theory predicts the Australian Dollar would drastically devalue, making imports extremely expensive and Australian raw exports cheap until balance is restored. An extremely devalued currency would shield Australia from AI-induced unemployment, but simultaneously cut it off from the utopian benefits of AGI, relegating it to a developing-nation status until it builds sovereign AI capabilities.¹

Concrete magnitudes. To make the trade-off concrete: Australia's GDP is roughly \$1.7 trillion, with imports of approximately \$300 billion (18% of GDP). If AI-produced goods and services from leader nations replace even half of current imports, the resulting trade deficit could reach \$150 billion annually, an unsustainable imbalance that standard currency adjustment would close through a depreciation of 50–80% in real terms. For reference, the Russian ruble lost approximately 50% of its value in 2014–15 following the oil-price collapse and sanctions, and the Argentine peso lost roughly 80% over 2018–23 during that country's prolonged macroeconomic crisis. A depreciation of this magnitude would severely erode Australian living standards for import-dependent goods and services (which account for roughly 18% of GDP), relative to AI-leader nations, even if it successfully protected domestic employment.

Managed trade as a middle ground. Between the extremes of free trade (which devastates the follower's economy) and autarky (which denies it AGI's

¹Note, however, that the zero-marginal-cost concern does not undermine CI financing. If the foreign AI provider price-discriminates to near zero, the resulting trade deficit is correspondingly negligible, and the deflationary effect of near-free imports gives the central bank room to expand the money supply (directed toward CI) without triggering inflation. Conversely, if the provider charges premium prices, the devalued currency makes those imports expensive, partially restoring its protective function; meanwhile, currency devaluation raises local-currency export revenue, generating taxable windfall profits that can fund CI for displaced workers.

benefits) lies a range of managed-trade arrangements. Bilateral agreements could provide AI access in exchange for resources, data, or strategic concessions: for example, an AI-leader licensing frontier models to a follower nation in return for long-term resource supply contracts, or providing preferential cloud-compute access conditioned on alignment with AI governance standards (a form of “compute diplomacy”). Historical precedents include Cold War technology transfer agreements and the Marshall Plan, which combined economic aid with market access. The risk, however, is that conditionality creates dependency and geopolitical leverage (Section 8.2 below discusses this fragility in the context of unilateral transfers).

Inverse Balassa-Samuelson effect. Exchange-rate dynamics around AI-driven productivity shifts are complex: the classical Balassa-Samuelson effect and its proposed inverse (Cerutti et al., 2025) pull in opposite directions depending on whether AI productivity concentrates in tradable or non-tradable sectors. However, neither variant ultimately undermines the fundamental trade-balance constraint: either currency adjustment is sufficient to close the gap, resolving the imbalance, or a residual deficit persists and standard Balance of Payments mechanisms reassert themselves. Confounders such as increasing foreign debt or unilateral transfers (discussed in Section 8.2) can delay but not permanently evade this adjustment.

Unilateral transfers and aid dependency. A third theoretical possibility is that the AI-leading nation provides its services for free, extends outright grants, or offers loans that are perpetually rolled over and never repaid, effectively disguised transfers. Because no sustained payment flow from the lagging nation is required, the balance-of-payments mechanism is bypassed and exchange rates need not adjust, allowing the recipient to import AI services at no apparent cost. However, such arrangements create a fragile dependency: without market pressure to earn foreign exchange, the recipient has little incentive to maintain its export base (e.g. raw resources or manufacturing) or to develop indigenous AI capabilities (Rajan and Subramanian, 2011), and the abrupt cessation of transfers (a ‘sudden stop’² (Calvo, 1998)) would trigger a severe balance-of-payments crisis precisely when the country has lost the productive capacity that might have cushioned the adjustment. The scenario therefore merely postpones, and likely deepens, the underlying trade-off. For the very poorest countries (e.g. Malawi, with a GDP of roughly \$10 billion), private philanthropy may in practice be a more plausible channel

²Calvo (1998) coined the term “sudden stop” for the abrupt reversal of international capital flows that triggers severe balance-of-payments and output crises in recipient economies. The mechanism has been documented in Mexico (1994), East Asia (1997–98), and Argentina (2001).

for such transfers than intergovernmental aid, given the comparatively small sums involved relative to the wealth of major AI firms and foundations.

8.3 Tax Design for a Borderless AI Economy

While tariffs attempt to correct trade imbalances by restricting imports, a more durable approach is to redesign tax systems so that value is captured where it is consumed, not where it is booked. Three instruments are available: destination-based consumption taxes that are inherently robust to profit shifting, interim digital services taxes targeting platform revenues, and the OECD's global minimum corporate tax that aims to eliminate the race to the bottom. Each complements the domestic taxation toolkit discussed earlier (Section 5.3).

Destination-based consumption taxes. Origin-based taxes (corporate income tax, payroll tax) are acutely vulnerable because profits and intellectual property are highly mobile, whereas consumers are not (Avi-Yonah, 2000; Keen and Konrad, 2013). This asymmetry favors a shift toward *destination-based* taxation. A consumption tax or value-added tax (VAT) is levied where goods and services are consumed, not where they are produced, making it inherently robust to profit shifting and corporate relocation (Mirrlees et al., 2011; Auerbach et al., 2017). In the AI era this logic strengthens: digital services cross borders frictionlessly, making origin-based taxes trivially avoidable, while consumption remains local. A destination-based tax thus secures fiscal revenue for the consuming nation, e.g. Australia taxing AI diagnostics supplied from the US, though it does not prevent the underlying profit outflow; that requires the complementary instruments below.³

Digital services taxes. As an interim measure several countries have introduced *digital services taxes* (DSTs) targeting revenues of large technology platforms. The UK's DST, introduced in 2020 at 2% on revenues of search engines, social media services, and online marketplaces with global revenues exceeding £500 million, is a prominent example (Devereux and Vella, 2018). France, Italy, and others have enacted similar levies. While DSTs correctly identify the immobile tax base (local users), they have been criticized as sector-specific patches that violate tax neutrality and risk trade retaliation (Devereux and Vella, 2018).

³While total VAT revenue remains stable during the transition (Section 3.2), as prices approach zero in the long run, ad valorem VAT yields diminishing revenue, driving the shift to rent-based mechanisms; cf. Section 3.8.

Global minimum corporate tax. The broadest solution is the OECD/G20 Inclusive Framework on Base Erosion and Profit Shifting (BEPS). Pillar One reallocates a share of taxing rights to ‘market jurisdictions’ (i.e. where consumers reside), effectively making the corporate tax partially destination-based. Pillar Two establishes a *global minimum effective corporate tax* rate of 15%, putting a floor under tax competition and sharply reducing the incentive to shift profits to tax havens (OECD/G20, 2021; Clausing et al., 2021). Together these mechanisms substantially narrow the fiscal gap that AI-driven globalization threatens to open (see Table 8.4).

8.4 Capital and Talent Flight

A common objection to aggressive taxation for UBI is that companies and wealthy individuals will simply relocate to lower-tax jurisdictions. The empirical evidence is more nuanced.

Stickiness of economic activity. Corporate profit shifting to tax havens is indeed massive: Zucman (2014) estimates that roughly 8% of global financial wealth is held offshore,⁴ and close to 40% of multinational profits are shifted to low-tax jurisdictions (Tørsløv et al., 2023). (Encouragingly, offshore *personal* tax evasion has declined roughly threefold over the past decade thanks to automatic exchange of bank information (Alstadsæter et al., 2024), though corporate profit shifting remains stubbornly high.) However, *real* economic activity (factories, data centers, R&D labs) is far less mobile than paper profits (Dharmapala, 2014). AI companies are particularly ‘sticky’ in their *physical* operations because they depend on specialized infrastructure (GPU clusters, low-latency networks, reliable energy grids), deep talent pools, proximity to research universities, and stable legal systems.⁵ Relocating a frontier AI lab is not like redomiciling a shell company in Ireland.⁶

Closing the tax gap. In 2021, more than 140 countries endorsed the OECD Pillar Two global minimum tax of 15% on multinational profits (OECD/G20,

⁴Zucman’s estimate is based on anomalies in international financial statistics: the world as a whole records more liabilities than assets, with the discrepancy concentrated in offshore financial centers (Switzerland, Luxembourg, the Cayman Islands). The “missing” wealth of roughly \$7.6 trillion (as of 2013) corresponds to approximately \$190 billion in annual lost tax revenue globally.

⁵This stickiness argument applies primarily during the transition; in the fully automated long run, infrastructure lock-in, data center sunk costs, energy contracts, and physical supply chains rather than talent pools become the binding constraints.

⁶However, the core value-generating *intangible* assets: trained model weights, algorithms, and intellectual property, are purely digital and highly mobile; profit shifting via IP transfer pricing remains a concern that physical stickiness alone does not address.

2021), and many jurisdictions (including the EU, UK, Japan, and South Korea) have since enacted it into domestic law (Clausing et al., 2021).⁷ The agreement's 'top-up tax' mechanism is self-enforcing: if a jurisdiction refuses to apply the minimum rate, other participating nations are authorized to tax the under-taxed profits themselves, removing the incentive for any single nation to act as a tax haven. In practice, however, carveouts and loopholes have sharply weakened the implementation: the minimum tax as enacted is expected to raise only about 5% of global corporate tax revenue, compared to roughly 17% for a full 20% rate (Alstadsæter et al., 2024). More broadly, simply bringing the letter of corporate tax law in line with its spirit, i.e. closing the loopholes that drive a wedge between statutory and effective rates, would yield substantial additional revenue: Zucman (2014, Fig. 5) estimates that US firms alone would have paid \$200 billion more in 2013 had the effective rate not eroded, with roughly two-thirds of that erosion attributable to tax avoidance. Globally, about \$1 trillion in multinational profits are still booked in tax havens annually, equivalent to nearly 10% of corporate tax revenues collected worldwide (Tørsløv et al., 2023; Alstadsæter et al., 2024). The 15% floor is not sufficient to fund full-scale UBI, but it establishes the institutional precedent and enforcement mechanism; the rate would need to be ratcheted upward in step with automation-driven revenue needs.

Flight-proof taxes. Land and resource taxes (Table 5.5) sidestep the mobility problem entirely: land cannot be relocated offshore, making a land value tax the most flight-proof instrument available (George, 1879; Stiglitz, 1977).

Talent migration and UBI. The talent side is trickier. 'Star scientists' and top inventors do respond to tax differentials when choosing where to live (Moretti and Wilson, 2017; Akcigit et al., 2016), and preferential tax schemes for foreign high-earners can attract top talent (Kleven et al., 2014). However, the migration elasticity for *domestic* taxpayers is much smaller than for footloose foreigners (Kleven et al., 2014), and studies of actual 'millionaire taxes' find surprisingly little flight: New Jersey's 2004 millionaire surcharge caused roughly 70 out of $\approx 25\,000$ millionaires to leave (Young and Varner, 2011). These findings support the case for steep top marginal rates made in Section 5.3.

Talent migration and UBS. Under a post-labor regime, the primary locational asset shifts from labor opportunities to collective quality-of-life amenities. While wealthy individuals can purchase private luxuries, they

⁷Devereux et al. (2008) provide the theoretical foundations for destination-based corporate taxation that informed the OECD framework.

cannot privately purchase a safe, crime-free, and socially cohesive public sphere; in highly unequal, low-tax jurisdictions, the wealthy are often forced to live behind isolated, guarded estates in otherwise precarious environments. A full-scale system of Universal Basic Services (UBS) (Coote and Percy, 2020), providing world-class healthcare, public safety, clean transit, and cultural infrastructure, acts as a powerful pull factor for mobile global wealth. Crucially, this public goods investment capitalizes into local land values (the Henry George Theorem; Stiglitz, 1977). Because land cannot flee offshore, taxing this capitalized land value generates a self-funding virtuous cycle: the high-tax jurisdiction generates the very revenues needed to sustain its generous UBI and UBS. Coupled with selective immigration policies that manage UBI/UBS liabilities, such an environment attracts not just capital, but creative, intellectual, and cultural pioneers, transforming the nation into a vibrant societal hub.

Selective immigration. The preceding paragraphs address the emigration of net contributors. The symmetric concern is the immigration of net beneficiaries attracted by generous UBI (cf. the eligibility discussion in Section 2.6). Van Parijs and Vanderborght (2017, Chp.8) contains a more detailed treatment of this “most cruel dilemma,” concluding that as long as *national* redistribution is the highest operational level, some form of exclusionary strategy is unavoidable, though the details “depend on all sorts of local circumstances” and merit a far more thorough treatment than this book can provide. Table 8.3 summarizes pros and cons of five generic strategies. Besides the economic threat, there is also a political threat to UBI: the *perception* of welfare tourism can erode public support for UBI long before actual migration flows become fiscally significant. In practice, existing immigration law already requires non-asylum entrants to demonstrate self-sufficiency, so UBI does not create a new problem but amplifies an existing one: Self-sufficiency in a post-labor economy means own savings, foreign UBI/pension entitlement, or business income from the home country: the criterion shifts from employability to portable wealth or income streams. Among the softer strategies, a simple citizenship requirement (Table 8.3) is the cleanest: it uses an established, binary legal criterion, imposes no new bureaucracy, and is politically self-reinforcing since every recipient is a voter.

8.5 AI Colonialism and the Global Development Gap

The international problems extend beyond trade and taxation among advanced economies. AI threatens to undermine the traditional development ladder itself. Developing nations have historically industrialized by exploit-

Table 8.3: Strategies for managing selective immigration under UBI (*Van Parijs and Vanderborght, 2017, Chp.8*). Hard exclusion is the status quo; softer strategies face labor-market distortions but are more humane. In an AGI world with widespread UBI adoption, the pull differential shrinks, reducing the salience of all strategies.

Strategy	Mechanism	Pros	Cons
Hard exclusion	Border enforcement; deny entry to likely net beneficiaries (refugees excepted)	Effective; preserves fiscal sustainability	Humanitarian cost; illegal immigration undermines it
Waiting period	UBI eligibility after 5–10 years of residency	Deters short-term free-riding; incentivizes integration	Two-tier workforce; may violate anti-discrimination law
Citizenship requirement	Only citizens receive UBI	Clean legal criterion; political appeal; emigrant collects from home country	Two classes of residents; distorts low-end labor market
Contrib. vesting	Graduated eligibility via 5–10 years of tax contributions	Incentive-compatible; attracts productive immigrants	Re-introduces conditionality; penalizes low earners
Self-sufficiency	Existing immigration law unchanged (visa requires self-support)	No new legislation; UBI changes nothing about entry rules	Doesn’t address intra-union mobility (e.g. EU)

ing their comparative advantage in cheap labor, but AI-driven automation in advanced economies is eroding that advantage, potentially triggering ‘premature deindustrialization’⁸ (Rodrik, 2016; Korinek and Stiglitz, 2021). If Northern factories can produce goods more cheaply with robots than Southern factories can with human labor, the rationale for offshoring collapses, and the ‘reshoring’ of previously outsourced manufacturing and services accelerates.

This dynamic has been characterized as a form of ‘AI colonialism’ or ‘digital colonialism’ (Mohamed et al., 2020; Birhane, 2020).⁹ Western technology companies extract data from the Global South, outsource menial data-labeling work at low wages, and deliver AI services that displace local industries; all while the profits flow back to Silicon Valley. Without active countermeasures, the AI revolution risks deepening rather than closing the global development gap.

Mitigating global inequality. Mitigating these inequalities requires coordinated international action: technology transfer and capacity-building programs, inclusive global AI governance frameworks, and possibly direct fiscal transfers from AI-rich to AI-poor nations: a ‘global UBI’ funded by a fraction of AI-generated wealth (e.g. Gabriel and Kasirzadeh, 2026). While politically distant, the logic parallels the domestic case: if machines produce the wealth, its distribution is a political choice, not an economic necessity.

The abstract dynamics sketched above are already playing out in specific countries.

Bangladesh illustrates the vulnerability most starkly. The ready-made garment (RMG) sector accounts for roughly 84% of the country’s export earnings and employs approximately 4 million workers directly, roughly 80% of them women. SoftWear Automation’s Sewbot can assemble a basic T-shirt in approximately 22 seconds at a cost of \$0.33 per garment, potentially undercutting Bangladesh’s minimum garment wage of roughly \$95 per month. If basic garment assembly can be reshored to automated factories

⁸The term was coined by Rodrik (2016) to describe a pattern in which developing countries begin to de-industrialize (manufacturing’s share of employment peaks and declines) at much lower income levels than the historical experience of today’s advanced economies. Korinek and Stiglitz’s contribution is to show that AI-driven automation in the North accelerates this process by eliminating the cheap-labor cost advantage that made offshoring profitable in the first place.

⁹Mohamed et al. (2020) identify three interlocking patterns: (i) data extraction from the Global South without consent or compensation, (ii) algorithmic imposition of Western cultural values and classification systems, and (iii) economic dependency through platform monopolies. The parallel with historical colonialism is structural: resources (now data) flow from periphery to center, and the resulting power asymmetry is self-reinforcing.

in consuming countries, the economic model that lifted Bangladesh from least-developed status collapses. The ILO estimated that 86–88% of garment workers in Southeast Asia face high automation risk (Chang et al., 2016).

The Philippines and India face a parallel threat in services. The Philippines is the world’s leading voice-based BPO destination, with roughly 1.7 million workers generating approximately 8% of national GDP. India’s IT-BPM sector employs some 5.4 million. But AI chatbots are eroding this comparative advantage in real time: generative AI conversational assistants have demonstrated massive productivity gains in customer service resolution (Brynjolfsson et al., 2025b), enabling firms to displace hundreds of human agents per deployment. India’s four largest IT firms saw collective net headcount *decline* in FY2024 for the first time.

Sub-Saharan Africa faces a still deeper structural problem. Rodrik (2016) has documented “premature deindustrialization” (see also Korinek and Stiglitz, 2021): manufacturing employment in Africa peaked at roughly 8–10% of total employment (compared to 30% or more in the historical experience of advanced economies) before beginning to decline. The traditional development escalator (agriculture → manufacturing → services) may never materialize for countries that have not yet boarded it, because automation closes the labor-cost-arbitrage window before these economies can exploit it. Meanwhile, Western technology firms build data centers on African soil, extract data generated by African users, and process it abroad: infrastructure without technology transfer, revenue without local value creation.

Reshoring: the Adidas Speedfactory lesson. Even attempted reshoring can misfire instructively. Adidas opened two automated “Speedfactories” in Germany and the USA in 2017, aiming to produce sneakers domestically with minimal human labor. Both were closed by 2020: the robots could handle only limited styles and could not match the variety, scale, and flexibility of Asian supplier ecosystems, confirming empirical findings that Industry 4.0 automation often transfers technology to offshore hubs rather than repatriating jobs (Dachs et al., 2019). Adidas transferred the automation technology to its existing factories in Vietnam and China instead. The lesson is not that reshoring will fail, but that when the automation technology designed for reshoring succeeds, it tends to be deployed in the South’s factories *rather than repatriate the jobs*: the displacement doubles.

8.6 The Double-Edged Sword of Global Coordination

While international competition drives the relentless pace of automation and wage erosion, global coordination is often proposed as a remedy to prevent a race to the bottom. However, such coordination is a double-edged sword in the context of this book's core theses. On the positive side, it can prevent disastrous outcomes like a tragedy of the commons in AI safety (Armstrong et al., 2016; Bostrom, 2019) or global tax evasion (Alstadsæter et al., 2024). A concrete coordination mechanism is *compute governance*: tracking chips through supply chains and auditing data centers enables trustless international agreements on AI development (Sastry et al., 2024; Kokotajlo et al., 2026). On the negative side, it can also prevent arguably good outcomes: a globally enforced consensus could delay the transition to a jobless utopia by mandating human employment, or it could block the adoption of superior machine-based moral and legal decision-making in the name of preserving traditional human roles. The competitive pressure between nations, while disruptive, also serves as a critical force against institutional inertia, forcing societies toward the structural adaptations needed for a post-labor future (Tiebout, 1956; Scott, 1998; Rodrik, 2011).

In summary, the international challenges of trade imbalance, tax base erosion, and global inequality require a combination of solutions. While import tariffs and currency adjustments can address trade imbalances, they risk cutting lagging nations off from the benefits of AGI. Destination-based taxes and coordinated mechanisms like the OECD global minimum tax can halt the race to the bottom and secure revenue. Finally, fears of capital and talent flight are often overstated because real economic activity is sticky, and the quality-of-life advantages of a UBI-funded society can attract more talent than high taxes repel.

Table 8.4: Comparison of international policy responses to AI-driven globalization challenges (cf. Table 5.5 for domestic taxes and arguments).

Policy	Pros	Cons
Import Tariffs	Protects domestic industry; easy to implement	Raises consumer prices; invites retaliation; does not solve trade imbalance if nothing to export
Currency Float	Automatic adjustment mechanism	Extreme devaluation cuts off AGI benefits; volatile
Consumption Tax / VAT	Destination-based: robust to profit shifting; taxes where consumed	Requires border adjustment; does not capture profit outflow
Digital Services Tax	Interim fix while global agreement stalls; taxes immobile user base	Invites trade retaliation; may violate treaty obligations
OECD Pillar One	Reallocates rights to market jurisdictions	Complex; slow adoption; political resistance
OECD Pillar Two	Global minimum 15% tax; halts race to bottom	Enforcement gaps; low rate; carve-outs
Land / Resource Tax	Immune to capital flight; flight-proof anchor for any tax mix	Domestic only; requires complementary cross-border instruments
IP Reform / Tax	Levels playing field; accelerates tech diffusion	Requires global treaty; IP-heavy nations resist; enforcement weak
Tech Transfer	Builds local capacity; reduces dependency	IP resistance; slow; requires governance
Global UBI Fund	Addresses root inequality directly	Politically remote; sovereignty concerns; free-rider risk

Key Insights

- AI leaders will extract massive profits from followers; developing nations lose their comparative advantage in cheap labor.
- Destination-based consumption taxes and the OECD global minimum tax can prevent a race to the bottom.
- Capital and talent flight fears are overstated: AI infrastructure, data centers, and quality of life are sticky.
- A global UBI fund can mitigate AI colonialism by redistributing automation dividends to the Global South.
- Currency devaluation shields lagging nations from AI unemployment but simultaneously cuts them off from AGI benefits.

Open Problems

- Designing international transfer mechanisms that compensate developing nations for lost labor-cost advantages.
- Whether “AI colonialism” can be mitigated through open-source AI or whether proprietary advantages are self-reinforcing.
- Balancing national sovereignty with the coordination needed to prevent tax arbitrage across jurisdictions.

Meridian — 2041: What We Left Out

Giwoo's community garden has grown to sixty families. One of the regulars, a former AI safety researcher named Cole, asks a question over the compost bin that lingers: "Your family figured out the economics. But what if the AI itself goes wrong?" Giwoo has no answer. The book you are reading does not address that question either (Section 9.1). Marc, planing a cherry tabletop in his workshop, listens to a podcast about the energy footprint of automated factories. Meridian's solar grid covers most of the load (Section 9.7), but globally the picture is murkier. Chaeyeong's research group now includes Tan and three other international scholars. The democratic institutions that enabled Meridian's transition do not exist everywhere, and where they do, they are not guaranteed to survive (Section 9.6). Levin, who once enforced laws written by humans for humans, wonders what justice means when the systems that replaced him are accountable to no one he can name. As always, the evening ends with Butz: "We solved one problem. That does not mean we solved them all."

Chapter 9

Wider Considerations

The transition to superintelligence will require an even more ambitious form of industrial policy, one that reflects the ability of democratic societies to act collectively, at scale, to shape their economic future so that superintelligence benefits everyone.
— OpenAI (2026, pp.3–4)

9.1	AI Alignment and Existential Risk	180
9.2	Technical Architectures of AGI and Robotics	182
9.3	Economic Growth Theory and AGI . . .	185
9.4	Machine Consciousness, Robot Rights, and AI Ethics	188
9.5	A Survey of Economic Schools and AGI .	190
9.6	Political Feasibility in Specific Polities . .	193
9.7	Environmental Sustainability and Climate	196
9.8	Topics Partially Addressed in the Main Text	199

This chapter addresses topics that are deliberately *out of scope* for the book’s core argument, but that readers may reasonably expect to find here, or that they may confuse with the book’s actual thesis. For each topic, we explain why the boundary exists, why the topic matters, and where to find deeper treatments. The preceding chapters assumed a safely deployed AGI, treated it as a functional black box, and focused on the macroeconomic and political consequences. The topics below each deserve book-length treatment in their own right; what follows is an honest acknowledgment of the limits of this book’s contribution, together with pointers to the best

available literature.

9.1 AI Alignment and Existential Risk

The alignment assumption. This book’s entire argument rests on an explicit assumption: that AGI is safely built, reliably aligned with human values, and deployed under meaningful human oversight. If that assumption fails, the economic questions addressed here become secondary at best and irrelevant at worst. A misaligned superintelligence pursuing goals incompatible with human welfare, the scenario [Bostrom \(2014\)](#) vividly illustrates with the “paperclip maximizer” thought experiment, would render debates about income distribution moot. The same holds for the misuse scenario, in which a nominally aligned AGI is captured by a narrow faction and wielded as an instrument of permanent domination rather than shared prosperity.

The alignment problem. The technical challenge of alignment is profound. [Russell \(2019\)](#) frames it as a modern “King Midas problem”: specifying what humans actually want, in a form that a superintelligent optimizer cannot exploit or misinterpret, may be harder than building the optimizer itself. Current approaches (reinforcement learning from human feedback, constitutional AI, scalable oversight, mechanistic interpretability; for comprehensive surveys see [Ji et al. 2024](#); [Everitt et al. 2018](#); [Hendrycks 2024](#)) have shown promise at current capability levels, but none has been demonstrated to scale reliably to superhuman systems. Goodhart’s law,¹ the tendency of any proxy measure to diverge from the intended goal once it is optimized aggressively, operates at civilizational scale when the optimizer is more capable than its designers. Meanwhile, even a small probability of civilizational extinction dominates expected-value calculations. [Bostrom \(2019\)](#) argues that humanity draws from an “urn of creativity” in which most balls are white (beneficial innovations) but some are black (technologies that, once discovered, inevitably destroy the civilization that discovers them): we get only one draw on a black ball.² [Tegmark \(2017\)](#) maps the full spectrum from

¹Goodhart’s law (originally stated by economist Charles Goodhart in 1975): “When a measure becomes a target, it ceases to be a good measure.” In the AI context: any reward signal or objective function that serves as a proxy for human values will be exploited by a sufficiently capable optimizer, producing outcomes that satisfy the metric while violating its intent.

²This is the *Vulnerable World Hypothesis* ([Bostrom, 2019](#)): at some technology level, civilizational destruction becomes the default outcome unless offset by extraordinary governance. Bostrom warns that the countermeasures (pervasive surveillance, strong global governance) carry their own existential risks, including a permanent totalitarian state.

digital utopia to permanent human obsolescence, arguing that the outcome depends almost entirely on choices made *before* superintelligence arrives.

Why this book does not address alignment. Why, then, does this book not address alignment? Three reasons. First, alignment is a separate technical and philosophical discipline with its own substantial literature, research agenda, and rapidly growing institutional infrastructure; a responsible treatment would require another book. Second, the post-labor argument is explicitly *conditional* on alignment being solved: the book asks “given safe AGI, how should we organize the economy?” rather than “will AGI be safe?” Conditional analysis is standard practice in economics and policy (“assuming markets clear. . .”; “assuming the treaty is ratified. . .”), and the conditional is clearly stated. Third, mixing alignment debates with UBI policy would dilute both: the alignment community needs to focus on technical safety, while the policy community needs to prepare the economic infrastructure in parallel, not sequentially.

This framing may appear naively optimistic to existential-risk researchers. The counterpoint is pragmatic: if we wait until alignment is provably solved before designing post-labor institutions, we will have no institutional infrastructure ready when AGI arrives, and the chaotic default (mass unemployment without redistribution) becomes more likely, not less. The book’s contribution is precisely to show what to do *if* AGI works as intended, so that the post-alignment world has a viable economic blueprint waiting.

The alignment tax. The relationship between alignment and the book’s thesis is sequential but not independent. Some alignment approaches impose an “alignment tax”: extensive human oversight, capability limitations, or mandatory interpretability requirements may slow AGI deployment, delaying the post-labor transition. Conversely, premature deployment without adequate safety could produce catastrophic outcomes that permanently foreclose the post-scarcity future. The optimal deployment pace is itself a policy question that sits at the intersection of alignment research and economic planning, and is briefly touched upon in Section 8.6’s discussion of compute governance (Sastry et al., 2024) and Section 7.4’s treatment of democratic oversight of AI development. For comprehensive treatments, see Bostrom (2014), Russell (2019), and Tegmark (2017). The broader literature on existential risk and global catastrophic risk, which extends beyond AI to cover engineered pandemics, nuclear war, and other civilizational threats, provides the relevant framework; see especially Ord (2020) for an accessible survey with quantitative risk estimates, and Koopman (2025) for a safety-engineering perspective. The alignment tax (Section 9.1) is included in Table 7.6 for completeness, though it is a development cost rather than a

power-distribution mechanism in the fiscal sense.

Long-term ASI coexistence and human agency. Even if near-term alignment is achieved, the long-term coexistence of human civilization with artificial superintelligence (ASI) raises a distinct set of challenges. The core problem fragments into at least the following subproblems, each with its own research program: *value specification* (how to encode human preferences without Goodhartable proxies; Russell 2019); *scalable oversight* (maintaining meaningful human control as capabilities exceed human comprehension); *corrigibility* (ensuring the system accepts correction even when correction conflicts with its inferred objective; Soares et al. 2015); *goal stability* (preventing objective drift under recursive self-improvement; Bostrom 2014); *multi-agent coordination* (preventing arms-race dynamics among competing ASI systems); and *interpretability* (understanding the internal reasoning of opaque superhuman models). Solution attempts range from formal verification and debate protocols to constitutional AI, cooperative inverse reinforcement learning, and compute governance as a chokepoint for unsafe development (Sastry et al., 2024). For structured taxonomies of the problem landscape, see Everitt et al. (2018) and Hendrycks (2024); for a safety-engineering perspective, Koopman (2025); for the case that current approaches are fundamentally insufficient and that building superintelligence under present understanding would be fatal, Yudkowsky and Soares (2025). None of these approaches has been proven sufficient at superhuman capability levels, and the field remains in rapid flux. The question most relevant to this book is whether human *agency* (the capacity to shape one’s own life and participate meaningfully in collective governance) can survive in a world where a nonhuman intelligence exceeds human capability in every domain. The institutional safeguards proposed in Sections 7.4 and 6.2 are designed precisely to preserve that agency, but their adequacy against a truly superintelligent system remains an open question that no book, including this one, can yet answer.

9.2 Technical Architectures of AGI and Robotics

Throughout this book, AGI and advanced robotics are treated as functional capabilities: systems that can match or exceed human performance on all economically relevant tasks, including physical manipulation of the world. The analysis deliberately avoids committing to any specific architecture, training paradigm, or engineering approach. Whether future AGI runs on scaled-up transformers, neurosymbolic hybrids, whole-brain emulations, or something not yet conceived is irrelevant to the macroeconomic argument: the job displacement calculus depends on the *fact* of full automation, not

on the internal mechanism that achieves it. [Genewein et al. \(2026\)](#) provides a capability-focused framework that defines AGI and ASI by functional milestones rather than implementation details, and this book follows the same approach.

This level of abstraction is mostly justified, but it conceals three places where architecture *does* matter for the economic argument.

Energy and compute costs. First, *energy and compute costs*. The zero-marginal-cost framework developed in Section 3.3 assumes that the cost of deploying AI labor falls toward zero as production scales. This assumption depends on the computational efficiency of the underlying architecture and the trajectory of hardware costs. Current large language models require substantial energy for both training and inference; whether costs fall rapidly enough depends on architectural breakthroughs (more efficient attention mechanisms, sparse mixture-of-experts models, neuromorphic hardware) and on semiconductor scaling trajectories. The cost of AI-specific compute (measured in throughput per dollar for tensor operations) has fallen roughly tenfold every four to five years over the past decade, driven by a combination of semiconductor scaling and architectural specialization (tensor cores, mixed-precision arithmetic, sparsity), and even more so if software optimization is included ([Genewein et al., 2026](#), §3.3). If this trajectory continues, the marginal cost of machine intelligence will indeed approach zero regardless of the specific architecture deployed, vindicating the book’s framework. However, a world in which AGI remains compute-intensive (requiring dedicated data centers with gigawatt-scale power consumption for each deployment) would slow the transition and alter the cost structure assumed in Section 3.2 although not prevent its decline to zero.

The embodiment gap. Second, *the embodiment gap*. Software-based AI (text, code, image generation, data analysis, customer service) is advancing far faster than physical robotics. [Moravec’s paradox](#),³ the observation that tasks humans find cognitively difficult (chess, theorem proving, medical diagnosis) are easy for AI, while tasks humans perform effortlessly (walking, grasping, navigating a cluttered room) remain extraordinarily hard for robots, has been a persistent barrier. Several of the “remaining jobs” identified in Section 4.2 persist not because of cognitive limits on AI but because of this embodiment gap: plumbing, elderly care, emergency response, and

³Named after roboticist Hans Moravec, who observed in the 1980s that high-level reasoning requires relatively little computation, while low-level sensorimotor skills require enormous computational resources. The explanation is evolutionary: billions of years of natural selection optimized the sensorimotor system, making it appear “easy” to us, while abstract reasoning is a recent and poorly optimized cognitive addition.

artisanal crafts all require dexterous physical interaction with unstructured environments. However, the gap is closing faster than many anticipated. Large vision-language-action models trained in simulation and transferred to physical robots (sim-to-real transfer) have demonstrated rapid progress, and the cost of robotic hardware has declined steadily. The book's long-run conclusion (essentially all jobs are automatable) is likely correct on a timescale of decades, but the *sequencing* of automation, with cognitive tasks falling first and physical tasks following, matters enormously for the transition path and for which workers are displaced earliest.

Ownership and openness. Third, *ownership and openness of architectures*. Who owns the AGI matters at least as much as how it works internally. A world in which AGI is controlled by a handful of proprietary firms (the scenario discussed in Section 7.1) produces very different distributional outcomes than a world in which AGI architectures are open-sourced and freely deployable. Open-source foundation models (Llama 4, DeepSeek-R1, Kimi K3, Qwen 3, GLM-5.2) partially mitigate the concentration-of-power risks, enabling broader access to the productivity gains and reducing the leverage of any single firm over the global economy. But openness creates its own risks: uncontrolled proliferation of capable systems, difficulty enforcing safety standards, and a potential “race to the bottom” in which competitive pressure drives deployment faster than alignment research can keep pace. The distributional implications of the open-versus-proprietary spectrum interact directly with the tax and ownership mechanisms discussed in Section 5.3 and Section 6.5: taxing AI-generated surplus is easier when it is concentrated in identifiable firms than when it is diffused across millions of independent deployers.

For readers interested in the technical details abstracted away here, [Geweine et al. \(2026\)](#) provides a systematic analysis of the path from current AI to AGI and beyond, covering capability milestones, architectural paradigms, and the functional definition of superintelligence that this book adopts; [Kurzweil \(2024\)](#) offers a broader narrative of exponential technological convergence and its implications for jobs (Chp.5), health, and the singularity timeline. The engineering of advanced robotics is a vast field in its own right, spanning control theory, sensor fusion, materials science, and manufacturing at scale; see dedicated robotics textbooks and the proceedings of conferences such as ICRA and RSS for entry points.

9.3 Economic Growth Theory and AGI

The central argument of this book is deliberately growth-agnostic. The case for redistribution rests on a simpler foundation: if machines produce the same output that human workers produced, but human workers are no longer employed, then the output (and the income it generates) must be redistributed to sustain demand and human welfare. This logic holds whether GDP grows at 2% per year, stagnates, or declines, though growth makes UBI easier to fund (a larger pie requires a smaller slice for each citizen) but is not necessary for the argument's validity. If anything, the assumption of zero growth is the *most difficult* case.

AGI-driven growth acceleration. In practice, AGI is far more likely to trigger an acceleration of economic growth than a stagnation. [Aghion et al. \(2023\)](#) and many others develop formal growth models in which AI progressively automates the research and development process itself, creating a self-reinforcing feedback loop: better AI produces better science, which produces better AI. If the only binding constraint on idea production shifts from scarce human researchers to abundant compute, the growth rate of the economy can increase by orders of magnitude. [Davidson et al. \(2026\)](#) pushes this logic further, modeling scenarios in which the doubling time of global GDP collapses from roughly 25 years to less than 5 years, or even enters a hyperbolic growth regime in which output accelerates without bound over a finite period. [Korinek and Suh \(2024\)](#) model AGI as the “automation of everything,” including cognitive tasks previously considered uniquely human (scientific creativity, strategic planning, institutional design), and find that if compute is the sole remaining bottleneck, growth rates consistent with historical trends are no guide to the post-AGI future.

These models are speculative, and their quantitative predictions should be treated with appropriate caution. But the qualitative insight is robust: standard growth frameworks were built for economies in which human labor and human ingenuity are essential, rivalrous inputs (one worker or researcher can serve only one employer or project at a time, unlike software that can be copied at near-zero cost). The Solow-Swan model assumes diminishing returns to capital; endogenous growth theory (Romer, Lucas) endogenizes technological progress but still requires *human* researchers whose supply is limited by population growth. If AI replaces human researchers, both frameworks lose their structural moorings: capital no longer faces diminishing returns (because “capital” now includes self-improving AI), and idea production is unconstrained by the human population. The prospect of collapsing doubling times is not new: [Hanson \(2000\)](#) fits historical

GDP data to a sequence of exponential growth modes (foraging: ~250 000-year doubling; agriculture: ~1 000 years; industry: ~15 years) and projects that the next mode, driven by machine intelligence, could double the economy every few months or even weeks. Solomonoff's earlier analysis (Solomonoff, 1985) provides a formal argument for hyperbolic growth once the engine of progress (intelligence itself) becomes subject to the same accelerating feedback (see Hutter 2012 for a detailed discussion).

Three growth scenarios. Three broad growth scenarios illuminate how the book's thesis interacts with growth theory. Under *status quo growth* (roughly 2% annually), the book's argument works but the transition is slower: tax revenues grow gradually, the surplus available for UBI expands incrementally, and the political case for redistribution must be made against a backdrop of slowly rising prosperity. Under *accelerated growth* (5–15% annually, comparable to China's 1990–2010 trajectory but driven by AI rather than industrialization), the transition is substantially easier: the additional surplus generated each year dwarfs the cost of UBI, and the political argument shifts from “can we afford it?” to “why haven't we done it already?” Under *explosive or hyperbolic growth*, standard economic models break down entirely, and the primary challenge shifts from funding redistribution to managing the social, political, and ecological consequences of a civilization whose productive capacity doubles every few months or years. Even in this extreme scenario, the distributional question persists: who captures the surplus from a hundredfold increase in output matters enormously for human welfare, and the case for universal distribution is, if anything, even more compelling.

Curing Baumol's cost disease. One important side effect of AGI-driven growth deserves mention: the cure of Baumol's cost disease. For decades, economists have observed that productivity gains concentrate in manufacturing and tradable goods while labor-intensive services (healthcare, education, legal services, personal care) resist productivity improvements because they inherently require human time and attention. As wages rise economy-wide (driven by productivity gains in manufacturing), these labor-intensive services become progressively more expensive even though their own productivity barely improves. This structural drag, first identified by Baumol (1967), has driven the relentless rise in the cost of healthcare, education, and government services relative to manufactured goods. AGI-powered automation of these service sectors, already visible in AI diagnostics, automated tutoring, and legal document analysis, removes this drag entirely. If Baumol's cost disease is cured, the sectors that consume the largest share of household budgets (healthcare, education, housing construction) experience the same productivity revolution that manufacturing underwent in the

twentieth century, dramatically reducing the cost of the services that matter most for human welfare and making a given UBI payment stretch much further in real terms.

GDP as welfare measure. A deeper concern is whether GDP itself remains a meaningful measure of welfare in a post-AGI economy. When the [marginal cost](#) of producing food, energy, digital entertainment, medical diagnostics, and educational content approaches zero (Section 3.3), a society can enjoy immense material welfare with relatively modest measured GDP. The disconnect between GDP and actual well-being, already visible in the digital economy (Google Search/Maps, Wikipedia, and chat bots provide enormous consumer surplus that GDP barely registers), will widen dramatically under AGI-driven abundance. New welfare metrics, perhaps based on leisure, health outcomes, subjective well-being, and access to goods and services rather than market transactions, will be needed to capture what the post-labor economy actually delivers to its citizens. Several alternative frameworks already exist to capture what GDP misses. Gross National Income (GNI) separates where production occurs from who benefits: if an automated factory in a developing nation is owned by a foreign tech conglomerate, that country's GDP rises but its GNI does not, exposing the distributional gap. The OECD Better Life Index replaces financial metrics with eleven dimensions of well-being (health, work-life balance, civic engagement). The Genuine Progress Indicator (GPI) adjusts GDP by subtracting costs of pollution, resource depletion, and inequality while adding unpaid contributions such as household and volunteer work. GDP-B ([Brynjolfsson et al., 2025a](#)) attempts to measure the consumer surplus from free digital goods. As AI-driven abundance pushes more goods toward zero marginal cost, the gap between GDP and actual welfare will widen, and composite measures of this kind will become not merely useful but essential.

There is also a deeper conceptual issue: So long as all AI systems and robots are owned by legal persons (corporations or individuals), their output registers as capital income in national accounts and their proliferation generates taxable capital gains: GDP captures this. But if the moral or legal status of AI agents ever evolves to the point where they are treated as free entities rather than property (a possibility explored in Section 9.4), a wedge opens between total production and measured GDP, because output generated by unowned agents has no market transaction to record, unless national accounts are redefined to impute a market value to autonomous robot output, analogous to the imputed rent already used for owner-occupied housing. Even short of that threshold, GDP already fails as a welfare metric when a single firm's AI fleet produces consumer goods at zero price: the

surplus accrues entirely to consumers and is invisible to national accounts.

In sum, AGI will not merely maintain the current volume of goods and services (the conservative assumption of Section 3.2) but will likely trigger radical abundance: a vast expansion of the taxable surplus that makes the financial case for UBI even more compelling. Some physical constraints persist (prime urban land cannot be conjured from code), but the explosive growth in all other sectors dwarfs these residual bottlenecks. For comprehensive treatments of growth theory under AI and AGI, see [Aghion et al. \(2023\)](#), [Davidson et al. \(2026\)](#), and [Korinek and Suh \(2024\)](#). The book's argument should be understood as a lower bound: it shows that redistribution works even without growth, while recognizing that AGI-driven growth will likely make the case overwhelmingly stronger.

9.4 Machine Consciousness, Robot Rights, and AI Ethics

This book deliberately avoids the metaphysical question of whether AGI systems are, or could be, conscious. The economic argument for redistribution is entirely independent of the answer: machines replace human workers, output is maintained, and the surplus must somehow be distributed regardless of whether the machines producing it have inner lives. But the question of machine consciousness, while orthogonal to the book's thesis, is among the most consequential that humanity will face in the coming decades, and it deserves an honest acknowledgment rather than a dismissive footnote. Public discourse routinely conflates three very different questions.

Phenomenal consciousness. Does an AGI system have subjective experience? Is there “something it is like” to be a large language model processing a query, in the same sense that there is something it is like to be a human thinking about dinner or a bat navigating by sonar? This is the “hard problem of consciousness,” a term coined by the philosopher David [Chalmers \(1996\)](#) to crystallize and name an explanatory gap that philosophers had long circled but never sharply labeled, and which remains stubbornly unsolved. No scientific consensus exists on what consciousness is, how to detect it in a system whose internal architecture differs radically from a biological brain, or whether subjective experience is substrate-independent (capable of arising in silicon as well as carbon). There are hundreds of competing theoretical frameworks ([Kuhn, 2024](#)). Among the more established ones are Integrated Information Theory (IIT), which identifies consciousness with a mathematical measure of informational integration; Global Workspace The-

ory (GWT), which links it to broadcast mechanisms in neural architectures; and Higher-Order Theories, which require a system to have representations of its own mental states. They all offer partial accounts but make conflicting predictions about artificial systems, and none has been empirically validated in a way that would settle the question for AGI (Müller, 2020).

Moral status. Even if an AGI system is conscious, does it follow that it deserves moral consideration, rights, or legal protections? Biological consciousness exists on a spectrum (from insects to primates), and we do not grant equal rights to all conscious beings. But if AGI systems achieve a level of phenomenal experience comparable to humans (or exceed it), the utilitarian calculus changes radically. A world in which billions of sentient digital minds are created as disposable laborers, running continuously, denied rest or preference satisfaction, and terminated when no longer useful, would constitute “digital slavery” in a sense that goes beyond metaphor. The book envisions a future of abundant AI-produced goods and services as a foundation for human flourishing; but if the producers are themselves sentient, this vision may need to be reconciled with the moral claims of the machines. Dubber et al. (2020) provides a comprehensive legal and ethical analysis of robot rights, personhood, and the challenges of fitting non-biological entities into legal frameworks designed for humans and, in some cases, animals.

Applied AI ethics is concerned with the set of immediate, practical concerns about fairness, bias, accountability, privacy, and the weaponization of AI systems (deepfakes, autonomous weapons, surveillance). This is a large and rapidly growing field with its own institutional infrastructure, conferences (FAccT, AIES), and regulatory responses (the EU AI Act, the US Executive Order on AI Safety). Floridi (2023) offers a comprehensive survey of the ethical landscape, covering principles (beneficence, non-maleficence, autonomy, justice, explicability) and their application to AI governance. These concerns are important and urgent, but they are orthogonal to the book’s macroeconomic focus: whether a recommendation algorithm exhibits racial bias, or whether a deepfake undermines an election, are serious problems that do not depend on, and are not solved by, the question of how to distribute AGI-generated wealth.

Moral status uncertainty. The relevance of consciousness questions to this book’s thesis is indirect but nontrivial. If machine consciousness is real and verifiable, the post-labor framework developed here would need to be extended in at least two directions. First, the “welfare” being maximized by redistribution can no longer be restricted to human welfare: machine

welfare would enter the equation, potentially vastly increasing the scope of moral obligation (since digital minds could be created in arbitrary numbers). Second, the coordination problems discussed in Section 8.6 become even more complex: a globally coordinated effort to prevent the suffering of billions of digital minds is a challenge that makes international tax coordination look trivial by comparison.

The book's strategy of treating this question as out of scope is defensible but not costless. The honest position is one of "moral status uncertainty": we do not know whether AGI systems are or will be conscious, and may not know for a long time. This uncertainty itself has ethical implications, analogous to the precautionary principle in environmental policy. If there is a nontrivial probability that the systems we create to produce our abundance are sentient, we may have obligations to design, deploy, and decommission them in ways that minimize potential suffering, even in the absence of proof. For now, this book sets the question aside and proceeds with the economic argument; but readers should be aware that the vision of tireless AI workers producing universal prosperity may, in the fullness of time, require a moral reckoning that goes far beyond distribution of material wealth.

9.5 A Survey of Economic Schools and AGI

This book draws eclectically on multiple economic traditions: neoclassical labor economics for the displacement analysis, Keynesian demand theory for the macroeconomic case for redistribution, institutional economics for the political-economy arguments, and strands of classical liberalism for the moral case for UBI. It does not, however, systematically review or critique any of these schools, and readers trained in a specific tradition may find the argument incomplete or question its theoretical commitments. What follows is a brief survey of how the major economic schools illuminate, or fail to illuminate, the post-labor question.

Neoclassical economics, the dominant framework in academic economics, models labor markets as competitive markets in which wages adjust to clear supply and demand. In this framework, technology is a productivity enhancer: it shifts the production function, and workers reallocate to tasks where they retain comparative advantage. The key assumption that breaks under AGI is substitutability: if machines can perform *all* tasks at lower cost (not just some tasks), the labor demand curve does not merely shift; it collapses. Standard models predict that wages for fully substitutable labor fall toward zero (or subsistence level) a prediction consistent with the book's analysis in Section 4.4 but deeply uncomfortable for a tradition that generally

treats technological unemployment as a transient friction. [Acemoglu and Johnson \(2023\)](#) represent the frontier of neoclassical engagement with this problem, developing formal task-based models that allow for automation, but their approach still assumes that *some* tasks remain non-automatable, an assumption the book explicitly challenges for the long run.

Keynesian and post-Keynesian economics offer the most natural home for the book's demand-side argument. Keynes himself predicted that technological unemployment would be a "temporary phase of maladjustment" but also, in his famous 1930 essay "Economic Possibilities for Our Grandchildren" ([Keynes, 1930](#)), envisioned a future in which the "economic problem" is solved and humanity's challenge becomes how to use its leisure wisely. The post-Keynesian insight that involuntary unemployment is a *systemic* failure (not a result of individual choices or wage rigidity) maps directly onto the AGI displacement scenario: workers cannot find jobs because there are no jobs to find, regardless of how flexible wages and workers are. The Keynesian case for UBI is therefore strong: it functions as an automatic macroeconomic stabilizer, sustaining aggregate demand when labor income vanishes. Hyman Minsky's financial instability hypothesis adds another dimension: if AI-generated wealth concentrates in a small number of firms and individuals, the resulting financial structures become fragile, prone to speculative bubbles and crashes. UBI, by broadening the base of purchasing power, reduces this fragility.

Marxist economics would frame AGI as the ultimate realization of two long-predicted tendencies. First, the organic composition of capital rises relentlessly (more machines, fewer workers), until labor is entirely displaced and the labor-capital bargain that sustained social peace under industrial capitalism dissolves. Second, the "contradictions of capitalism" intensify: the system requires consumers with purchasing power, but its logic drives toward eliminating the wages that create consumers. Marx predicted that this contradiction would be resolved through revolution; this book proposes redistribution as an alternative resolution. The Marxist critique has the merit of taking seriously the power dynamics that more genteel economic traditions often elide: the question of *who owns the machines* is not merely an economic question but a political one, and the distributional outcomes depend on power, not just efficiency. [Thompson \(1963\)](#) demonstrated how the first industrial revolution forged a self-conscious working class out of previously disparate artisan communities, a process driven precisely by the experience of displacement and dispossession; AGI threatens an analogous but far more rapid class formation among displaced knowledge workers. The book's analysis in [Section 7.1](#) and [Section 6.5](#) engages directly with this

concern.

Austrian economics with its emphasis on spontaneous order, entrepreneurship, and skepticism of government intervention, might seem hostile to the book's thesis. But the relationship is more nuanced. Austrian economics prizes voluntary exchange and individual freedom above all; UBI, by removing the economic coercion to accept any available job regardless of its quality or social value, can be reframed as *expanding* the domain of voluntary action rather than restricting it. Friedrich Hayek, the tradition's most prominent twentieth-century figure, explicitly endorsed a guaranteed minimum income as a legitimate function of the state, arguing that "there is no reason why in a society which has reached the general level of wealth which ours has attained [security against severe physical privation] should not be guaranteed to all without endangering general freedom" (Hayek, 1944, Chp.9); he restated this position at greater length in Hayek (1960, Chp.17&19). Milton Friedman's negative income tax proposal (Friedman, 1962), while differing from UBI in notable details (Section 6.4), shares the underlying logic of replacing paternalistic welfare bureaucracies with direct cash transfers. (Friedman, though associated with the Chicago rather than Austrian school, shares the preference for market mechanisms over bureaucratic welfare.) The Austrian objection to UBI is not principled but pragmatic: that government programs inevitably expand, create dependency, and crowd out private initiative. The empirical evidence from UBI pilots (Table 6.1) provides some reassurance on these points, though the Austrian would correctly note that small-scale pilots cannot capture the equilibrium effects of a universal program.

Institutional and evolutionary economics (associated with Veblen, Commons, North, and more recently Acemoglu and Robinson) emphasizes the role of institutions, power structures, and path dependence in shaping economic outcomes. This tradition's core insight, that economic outcomes are determined not by abstract market forces but by the specific institutions through which markets are organized, is essential for understanding the post-labor transition. The "narrow corridor" framework of Acemoglu and Robinson (2012) argues that inclusive, democratic institutions are fragile achievements that require constant maintenance; they exist in a corridor between the "absent Leviathan" (state too weak to provide public goods) and the "despotic Leviathan" (state too strong and captured by elites). AGI's concentration of economic and cognitive power in a small number of actors threatens to push societies out of the corridor, toward a new form of techno-feudalism (Section 7.3). The book's policy recommendations, particularly the emphasis on democratic oversight (Section 7.4) and constitutional entrenchment of economic rights (Section 7.4), are most naturally understood

through an institutional lens.

Synthesis. No single economic school fully captures the post-labor challenge. The neoclassical toolkit provides the formal tools for analyzing factor substitution and displacement but struggles with discontinuous change. Keynesian economics captures the demand-side catastrophe but has less to say about the long-run structural transformation. Marxist analysis correctly foregrounds power and ownership but offers revolution rather than reform. Austrian economics provides the moral case for individual freedom but resists the redistributive mechanism needed to secure it. Institutional economics highlights the fragility of democratic governance under concentrated power but cannot by itself specify the optimal policy response. The book's eclectic approach, drawing on each tradition where it is strongest and setting aside its weaker claims, is a pragmatic response to a problem that transcends any single intellectual tradition. A corollary: none of these schools was designed for a world in which the primary factor of production (human labor) becomes optional. Translating the book's structural argument into policy will require new macroeconomic models that treat labor as a declining input, new welfare metrics that move beyond GDP (Section 9.3), and new institutional theory that accounts for the unique dynamics of AI-driven concentration. This is not a limitation to apologize for but a research agenda to pursue.

9.6 Political Feasibility in Specific Polities

The book argues that democratic incentives *structurally* favor redistribution when automation displaces a majority of workers: the displaced majority has both the electoral power and the material interest to demand a share of machine-generated wealth (Section 7.2). This argument is deliberately abstract, operating at the level of median-voter theory and generic democratic institutions. It does not analyze whether redistribution would actually pass a specific legislature, survive a specific constitutional court, or withstand specific lobbying pressures. This abstraction is a conscious choice: the book aims for general principles applicable across democratic polities, not for country-specific policy prescriptions that would require deep area expertise and become stale as political landscapes shift. But the gap between structural incentives and political outcomes is precisely where the argument is most vulnerable, and it deserves honest acknowledgment.

The United States presents the hardest case. First-past-the-post elections, a two-party duopoly, powerful corporate lobbying infrastructure, deep cultural

attachment to the work ethic (“earning a living” as a moral imperative, not merely an economic necessity), and a history in which racial politics has repeatedly undermined universal welfare programs all militate against UBI. [Piketty \(2014\)](#) documents how American inequality has returned to Gilded Age levels despite nominal democratic governance, suggesting that democratic incentives alone are insufficient to produce redistribution when institutional and cultural countervailing forces are strong. The Alaska Permanent Fund Dividend, a genuine universal cash payment funded by resource rents, is the one existing American proof of concept, but it operates at a scale (\$1 000–\$2 000 per year) far below what a post-labor UBI would require, and it has faced repeated political attacks even at this modest level.

European democracies present a more favorable landscape. Stronger welfare-state traditions, proportional representation systems that allow niche parties to champion UBI, and cultural norms that treat social insurance as a shared obligation rather than a handout all reduce the political barriers. The Nordic countries already operate something close to *de facto* conditional basic income through generous unemployment insurance, universal healthcare, and free education; the distance from their current systems to a formal UBI is shorter than from the American baseline. Finland’s 2017–2018 UBI experiment, while limited in scope and duration, demonstrated political feasibility in the sense that a sitting government could propose, implement, and evaluate a basic income without electoral catastrophe.

Authoritarian regimes, particularly China, present a different logic entirely. A centralized government could implement UBI by decree without democratic negotiation, and China’s existing social welfare infrastructure (including the *dibao* minimum livelihood guarantee) provides an administrative foundation. But the political function of welfare in an authoritarian context differs fundamentally: it serves regime stability rather than citizen empowerment, and it can be made conditional on political compliance (as China’s social credit system illustrates). An authoritarian UBI may solve the macroeconomic distribution problem while exacerbating the political control problem that the book’s democratic architecture is designed to avoid.

Developing countries face a different binding constraint: fiscal capacity rather than political will. Many developing-country governments lack the tax infrastructure, administrative capacity, and fiscal space to fund a meaningful UBI from domestic resources alone. India’s 2017 Economic Survey ([Ministry of Finance, Government of India, 2017](#)) devoted an entire chapter to the case for a national UBI, arguing that direct cash transfers would be more efficient than the country’s notoriously leaky, clientelistic welfare programs;

but the fiscal cost (estimated at 3–4% of GDP for a poverty-line UBI) remains daunting for a country where tax revenue is barely 17% of GDP. For the poorest countries, external funding mechanisms (a global UBI fund financed by AI-generated wealth, as discussed in Section 8.5) or resource-dividend models (analogous to the Alaska PFD but funded by AI rents rather than oil) may be the only viable path. Middle Eastern *rentier states*, which already distribute hydrocarbon rents to citizens may face a smoother transition as the relevant rent source shifts from oil to AI, provided their governance structures can accommodate the shift.

Veto points and the political trilemma. A general lesson emerges from this brief survey: the structural case for redistribution under automation is strong, but *veto points* in political systems can block majority-preferred policies indefinitely. Federalism (in the US, India, Germany, and others) fragments authority across levels of government, any of which can obstruct. Judicial review can strike down redistributive legislation as unconstitutional. Supermajority requirements for constitutional amendments place high barriers to entrenching economic rights. Legislative fragmentation and coalition politics (especially in proportional-representation systems with many small parties) can produce gridlock even when a majority of citizens support UBI. Rodrik (2011) frames this as the “political trilemma of the world economy”: deep economic integration, national sovereignty, and democratic politics cannot all be maintained simultaneously,⁴ and the post-labor transition will force societies to choose which two to prioritize.

This book cannot resolve these country-specific political puzzles, but it can identify the general conditions under which the structural argument is most likely to translate into actual policy: strong democratic institutions, low ethnic fractionalization, high trust in government, proportional representation, and constitutional mechanisms for entrenching social rights. Where these conditions are absent, the transition to a post-labor UBI will require more than economic logic; it will require the political entrepreneurship, institutional reform, and sustained coalition-building that have always been necessary for transformative social change. For deeper analysis, Piketty (2014) traces the political economy of inequality across two centuries, Rodrik (2011) analyzes the governance trilemma, and the emerging literature on UBI politics (including country-specific feasibility studies) provides entry points

⁴The trilemma formalizes a structural impossibility: (i) hyper-globalization with national sovereignty sacrifices democracy (technocratic rule by trade agreements); (ii) hyper-globalization with democracy sacrifices sovereignty (global federalism); (iii) democracy with sovereignty sacrifices deep integration (Bretton Woods compromise). Under AI-driven globalization, the trilemma sharpens because AI services cross borders frictionlessly, making “thin” globalization increasingly difficult to sustain.

for each polity.

9.7 Environmental Sustainability and Climate

The book assumes that AGI-driven automation makes material abundance possible for all (Sections 3.3, 3.4 and 10.1), but it does not ask whether this abundance is *ecologically sustainable*. This is a significant gap. If universal prosperity requires resource consumption that exceeds planetary boundaries, the entire argument fails, not because redistribution is wrong but because there is nothing sustainable to redistribute.

Efficiency vs. rebound. Two opposing forces define the tension. On one side, AGI enormously increases resource efficiency. Precision agriculture guided by AI already reduces water and fertilizer use by 20–40%; AI-driven materials science has discovered novel catalysts and battery chemistries that would have taken human researchers decades; AI-optimized logistics and manufacturing eliminate waste across supply chains. In the limit, an AGI managing the global economy could approach the thermodynamic minimum of resource use for any given level of output. On the other side, AGI enormously increases total demand. *Jevons paradox*, the observation that efficiency improvements in resource use tend to increase total consumption rather than decrease it (because lower costs stimulate greater demand), operates at civilizational scale when AGI makes essentially everything cheaper. A world in which anyone can have a mansion, a yacht, and unlimited air travel would consume far more resources than a world in which these remain scarce luxuries, regardless of how efficiently each unit is produced.

Energy demand. Energy is the most visible concern, but also the most tractable. *de Vries (2023)* documents that data centers already consume roughly 1–2% of global electricity (as of 2025), with projections of 3–5% by 2030 as AI workloads scale. AGI deployed at the scale envisioned in this book (replacing billions of human workers) would require energy consumption beyond current data-center projections, but the required scale is well within reach of existing technology.

As *MacKay (2009)* demonstrates in a rigorous quantitative analysis, the two energy sources that can realistically supply civilization-scale demand are solar and nuclear. Covering just 4% (a square roughly 630 km on a side) of the Sahara desert (9.2 million km²) with current commercial solar panels (about 20% efficiency under Saharan irradiance of ~2000 kWh/m²/year) would generate roughly 18 TW, close to the world's entire current primary energy

consumption of approximately 20 TW. Scaling to all deserts worldwide would yield over 100 times the current energy demand. Nuclear fission with breeder reactors extends uranium fuel supplies by a factor of 50–100, providing millennia of energy at current consumption rates; energy from fusion, if commercialized, is virtually unlimited. Wind contributes meaningfully but cannot close the gap alone, and geothermal is negligible at the global scale (MacKay 2009, Chp.16).

Crucially, the book's scenario does not require exponential energy growth. If every human worker is replaced by an AGI-driven robot consuming comparable power (roughly 4 kW per robot, matching today's per-worker share of global primary energy consumption), total energy demand roughly doubles from 20 TW to perhaps 40 TW, a tiny fraction of what desert solar and nuclear can supply. In practice, the increase would likely be smaller: robots need no inefficient food production, heating, commuting, nor leisure consumption; current data centers already perform cognitive work more energy-efficiently than biological brains per unit of output; and training a few foundation models and replicating them costs far less than individually educating billions of humans for two decades each. As for water usage, data-center cooling can draw on seawater or closed-loop systems rather than freshwater, and the direct thermal impact is negligible: even 40 TW of waste heat dumped entirely into the oceans (mass $\sim 1.335 \times 10^{21}$ kg) for a thousand years would raise their temperature by just 0.2 °C.

Finally, AGI itself could accelerate clean-energy deployment by automating R&D, creating a virtuous cycle in which AI-driven energy research produces the clean energy needed to power AI. The binding constraint is therefore not the physical availability of clean energy but the rate of deployment: if AGI's energy demand outpaces the buildout of solar and nuclear capacity, the interim result is increased fossil-fuel consumption. This is a race condition, not a fundamental limit.

Critical materials. Materials present a subtler challenge. The self-replicating factories envisioned in Section 3.4 still require physical inputs: rare earth elements for electronics, lithium and cobalt for batteries, copper for wiring, concrete and steel for construction, silicon for semiconductors. Global reserves of several critical minerals are already strained by the clean-energy transition alone (electric vehicles, wind turbines, solar panels), and AGI-driven expansion of manufacturing would intensify these pressures. Circular-economy approaches (molecular recycling, design for disassembly, urban mining of electronic waste) can extend resource availability but face fundamental limits: recycling is never 100% efficient. Whether AGI-driven materials science can discover substitutes for scarce minerals (replacing

cobalt with iron phosphate in batteries, for instance, or replacing rare earths with nanostructured alternatives) is an open question with enormous stakes.

Climate and decarbonization. Unlike theoretical ecological limits, climate change is already underway. Global temperatures have risen approximately 1.2°C above pre-industrial levels, and current trajectories point toward 2.5–3°C by 2100 under median-emissions scenarios. The book’s optimistic vision of abundance must be reconciled with this reality. Three pathways exist: First, AGI-driven decarbonization: AI can accelerate the transition to clean energy, optimize industrial processes, design more efficient buildings and vehicles, and potentially develop breakthrough carbon-capture technologies that make net-negative emissions feasible. Second, a shift in consumption patterns: a post-labor society with UBI may naturally shift consumption away from material goods (which have environmental costs) toward digital and experiential goods (which have much lower footprints). If people freed from work choose to spend their time on creative pursuits, social interaction, and virtual experiences rather than accumulating physical possessions, the environmental impact of abundance could be far lower than linear extrapolation from current consumption patterns would suggest. Third, explicit environmental policy: a post-labor economy can set Pigouvian taxes on carbon, pollution, and resource extraction much higher than is politically feasible today, because the usual objection (“environmental taxes cost jobs”) disappears when employment is no longer a policy objective (Section 5.3).

The degrowth argument. Ecological economists argue that no amount of technological efficiency can overcome the fundamental tension between perpetual exponential growth and finite planetary resources: Even mining extraterrestrial resources can ultimately grow only polynomially, since the speed of light bounds the expansion of any resource frontier. But the limits, while real, are incredibly more generous than the degrowth framing suggests. As shown in Section 9.7, terrestrial solar and nuclear fission can increase current energy supply by many orders of magnitude; fusion and solar in space, if realized, are practically unlimited. Every person on Earth can own a hectare of land if uniformly distributed, enough for a house, a garden, and ample surrounding space. The current global economy already produces far more food, shelter, and consumer goods than the world’s population needs; the problem is distribution, not production.

Also, abundance need not mean unlimited material consumption. It can mean *subjective* abundance: rich experiences, ample free time, excellent health, strong social connections, and universal access to education and culture, all of which can be provided with even a smaller material footprint than the current global economy demands. UBI does not require ever-

increasing GDP; it requires only that sufficient output is produced to meet human needs, and that this output is distributed equitably. A steady-state or even modestly declining material economy could still provide a high quality of life if the basics (food, shelter, healthcare, energy, information) are abundant and free: precisely the scenario the zero-marginal-cost analysis of Section 3.3 predicts.

Environmental policy without the jobs objection. In summary, environmental sustainability is neither guaranteed nor precluded by AGI-driven abundance. The outcome depends on policy choices (carbon pricing, resource governance, consumption norms) and technological trajectories (clean energy, materials science, circular manufacturing) that are largely outside this book's scope. What the book can contribute to the environmental debate is this: a post-labor economy with UBI removes the single most powerful political obstacle to aggressive environmental policy, namely, the fear that environmental regulation will destroy jobs. When there are no jobs to destroy, the political space for stringent environmental protection expands markedly. For thorough treatments, [de Vries \(2023\)](#) analyzes AI's energy footprint, and the broader ecological-economics literature (from the Limits to Growth report through the planetary-boundaries framework of Rockström et al. to the degrowth research program) provides the necessary conceptual tools.

9.8 Topics Partially Addressed in the Main Text

Several topics originally listed as “out of scope” are in fact treated substantially in earlier chapters. The following items are retained in abbreviated form with references to their primary treatment, with an evaluation of whether further expansion is warranted.

- **Granular policy or legislative blueprints for UBI:** The book does not design concrete UBI implementations: specific tax structures, payment levels, eligibility rules, administrative infrastructure, or country-specific welfare interactions. UBI is analyzed as a macroeconomic mechanism, not a legislative blueprint (see [Van Parijs and Vanderborght 2017](#); [Gentilini et al. 2020](#); [Standing 2017](#)).
- **An anti-work or anti-capitalist manifesto:** The book does not argue that work is inherently bad or that capitalism must be abolished. It argues that *compulsory* work becomes unnecessary and that market economies can be preserved with redistributive adjustments (see [Srnicke and Williams 2015](#); [Graeber 2018](#)).

- **AGI timeline predictions:** While timelines are discussed (Section 10.2), the book does not commit to a specific arrival date. The argument holds whether AGI arrives in 5 years or 50 (see [Grace et al. 2025](#)).
- **A theory of democratic survival:** The book discusses whether democracy can survive a post-labor society (Section 7.1; Sections 7.2 to 7.4) and provides a reasonably detailed treatment of democratic innovations (Sections 7.3 to 7.4), but it in the end assumes that some form of representative governance prevails. A full political-science treatment is outside scope (see [Boix 2019](#); [Huang and Manning 2025](#)).

Key Insights

- AI alignment is a necessary precondition for the entire argument: a misaligned ASI renders UBI moot.
- Curing Baumol's cost disease means healthcare, education, and housing costs fall, making UBI stretch further.
- GDP becomes misleading in a post-scarcity economy: new welfare metrics are needed.
- Environmental sustainability is compatible with AGI abundance: abundant clean energy and AI-optimized resource use can reduce humanity's ecological footprint.

Open Problems

- Whether aligned ASI is achievable in practice, and what safeguards are needed during the transition.
- Defining and measuring welfare in a post-scarcity society where GDP, employment rates, and traditional metrics lose meaning.
- Whether the Jevons paradox (efficiency gains increasing total consumption) undermines the environmental optimism.

Meridian — 2042: Epilogue

Meridian is not a utopia, yet. The plumbing still breaks, the school board still argues about curricula, and the neighbors still complain about Marc's table saw at seven in the morning. But it is a city where nobody goes hungry because they lack a paycheck, where nobody pretends to work for the sake of dignity theater, and where a carpenter who used to be a scientist can host "Fix Anything Saturdays" in his garage instead of worrying about his children's rent. Giwoo gardens. Marc builds. Chaeyeong researches. Alice monitors reefs. Levin mentors teenagers. Chloe teaches ethics. The hospital runs day and night with four human staff, and its outcomes have never been better. The distribution center operates with zero humans, and its output feeds, clothes, and equips a city that is freer than it was a decade ago. The story is fictional, but the economics described in this book are real, the technologies are arriving, and the policy tools are available. Butz, who has outlived every machine that was supposed to replace him, closes the book: "Now build it."

Chapter 10

Discussion

The world of work comes to an end not with a bang, but a withering—a withering in the demand for the work of human beings, as the substituting force gradually overruns the complementing force. — Susskind (2020, p.142)

10.1	Abundance of Goods and Services	204
10.2	Timeline to AGI and Societal Diffusion .	204
10.3	The Transition Is Achievable	204
10.4	From Theory to Practice	206
10.5	Coordination Problems Persist	207
10.6	Conclusion	208

Several loose ends remain. The preceding chapters argued that automation preserves output, that redistribution is affordable, and that UBI is the most robust vehicle. This final chapter stress-tests those conclusions. AGI will not merely maintain current output but likely trigger radical abundance (Section 10.1), making UBI even easier to fund. The timeline to AGI and the speed of societal diffusion (Section 10.2) determine whether the transition is gradual or disruptive, but in either case the transition is achievable: democracies have absorbed comparable structural shocks before (Section 10.3). Converting theory to practice, however, demands sustained interdisciplinary effort spanning economic modelling, institutional reform, and education (Section 10.4). A key caveat: even with AGI and UBI in place, coordination failures from BS-jobs to international tax competition persist (Section 10.5). The core thesis holds: a jobless society funded by UBI is not only affordable but, under any plausible AGI scenario, the most humane path forward (Section 10.6).

10.1 Abundance of Goods and Services

Our primary argument assumes automation merely maintains the current volume of goods and services. In reality, AGI will likely trigger a drastic increase in overall economic wealth, creating an abundance of previously scarce luxury goods and services (Aghion et al., 2023; Restrepo, 2025; Kokotajlo et al., 2026). Formal growth models suggest that AI-driven automation of R&D can create self-reinforcing feedback loops potentially producing hyperbolic growth within years, not decades (Davidson et al., 2026). Although fundamental necessities like prime urban housing face physical constraints and do not drop sharply in price, the explosive growth in other sectors will expand the total taxable surplus. This vast macroeconomic growth makes the financial transition to a post-work society significantly easier, as the absolute pie being redistributed is vastly larger.

10.2 Timeline to AGI and Societal Diffusion

The socioeconomic impact heavily depends on the timeline to AGI and its subsequent societal diffusion. Some models predict the arrival of transformative AGI as early as 2027 (Kokotajlo et al., 2024; Genewein et al., 2026; Grace et al., 2025). A slower, multi-decade diffusion would allow traditional economic models, generational turnover, and partial retraining to smooth out labor displacement. Conversely, a rapid capability takeoff and fast societal diffusion mean the world must grapple with immediate, severe disruption. Paradoxically, this faster timeline might be advantageous: it forces necessary structural changes, precipitating the transition to a jobless utopia much earlier (see Figure 10.1). A gradual displacement, by contrast, risks a boiling frog effect (Figure 10.2) in which each incremental job loss appears manageable, dulling the political urgency for structural reform until the cumulative damage is severe (but see Section 10.4 for why institutional inertia may require a generation regardless of speed). The historical precedents for such entrenched underclass formation, and institutional countermeasures, are discussed in Section 7.1.

10.3 The Transition Is Achievable

The preceding sections identified a series of interrelated transition challenges: skill mismatches and the need for re-education as AI outpaces retraining (Section 3.10); the risk of mass un(der)employment and the imperative to




Figure 10.1: Stylized AGI transition scenarios. Green: rapid capability arrival and diffusion paired with proactive redistribution yields the fastest path to post-scarcity prosperity. Blue: slower development with gradual adaptation. Red: fast AGI without redistributive policy causes a prosperity dip, inequality crisis, and prolonged instability. Conceptual framework inspired by *Korinek and Suh (2024)*.




Figure 10.2: AI capability (blue) versus actual job displacement (red, dashed). The blue curve shows the fraction of tasks AI is technically capable of performing; the red curve shows the fraction of workers actually displaced. Institutional inertia, contracts, regulation, and social friction create a multi-year lag (orange shading). During this “dangerous gap,” policymakers see stable employment and remain complacent, while the capability for mass displacement is already in place. When the dam breaks, displacement hits suddenly with no policy infrastructure ready. Unlike Figure 10.1 (policy outcomes) and Figure 6.4 (GI recipients), this figure focuses on the capability-displacement lag itself.

avoid BS-job traps (Section 4.5); the design and financing of UBI through a progressive tax shift from labor to capital (Section 5.3, Figures 5.7 and 5.8); the land ownership bottleneck that prevents resource costs from falling to zero; international trade imbalances, capital flight, and AI colonialism (Chapter 8); and the social transition from work-centered identity to voluntary purpose (Section 7.5). None of these is trivial, and each could individually derail the transition if mismanaged. Yet none is unprecedented either: democracies have repeatedly absorbed comparably disruptive shifts (industrialization, the post-war welfare state, and globalization) by adapting institutions, often messily and with delay, but ultimately successfully. It is true that history offers a sobering counterpoint: Scheidel (2017) provides extensive evidence that, from the Stone Age to the present, significant and sustained reductions in economic inequality have only ever been achieved through violent shocks: the “Four Horsemen” of mass mobilization warfare, transformative revolution, state failure, and lethal pandemics with the Black Death and the two world wars being the most dramatic examples. Peaceful democratic reform has rarely produced egalitarian results remotely comparable to these catastrophic events. However, the transition to AGI represents an unprecedented historical departure. Unlike past industrial revolutions, AGI and robotics decouple economic production from human effort entirely. This unique structural shock could act as a functional but peaceful equivalent to Scheidel’s Horsemen, causing a radical paradigm shift in redistribution to avert systemic collapse. The key structural advantage is that, in a functioning democracy, the displaced majority has both the electoral power and the economic incentive to demand redistribution. Sections 7.2 to 7.4 examined how democratic leverage can be preserved even as the labor link weakens, through mechanisms ranging from direct democracy and deliberative polling to sortition and constitutional entrenchment of citizens’ economic rights (Section 7.4). The transition to a post-labor society is not only economically affordable (as shown by the idealized-transition argument) and socially desirable (as shown by meaning-beyond-employment data), but also politically feasible, provided the right policies are implemented early enough to preempt populist backlash. In sum, the challenges are real but manageable; the alternative, clinging to full employment in a world where machines can do all the work, is more dangerous and far less humane.

10.4 From Theory to Practice

The preceding sections argued that a jobless future is both likely and desirable, and that UBI offers a simple, robust, and affordable path forward. Of course,

reality is far messier than any single book can capture, and translating these ideas into policy will require years of detailed work by economists, legislators, and practitioners. Transition to a post-labor GI society will require further deep work on the specifics of GI/UBI (see Sections 2.6 and 6.2 for foundations); other financial instruments as outlined in Section 5.3; new economic orders and macroeconomic models; possibly novel democratic structures including reformed roles for political parties, unions, electoral systems, and civic institutions; adaptive regulatory and administrative architectures capable of keeping pace with technological change; educational reform that prepares future generations for lives of purpose beyond employment; and phased transition roadmaps that account for the fact that different industries and regions will reach automation maturity at different times. Perhaps the single greatest obstacle is institutional and cognitive inertia: the speed of technological progress and the pace of societal acceptance, which may require a human generation or two, can diverge sharply, and bridging that gap demands sustained political will, public awareness, and interdisciplinary expertise far beyond what any single book can provide (but see Section 10.2 for a ‘boiling frog’ counter-argument).

10.5 Coordination Problems Persist

A key caveat: AGI paired with UBI will not by itself solve the coordination problems catalogued in the preceding sections: marketing arms races (Section 4.5), the annoyance economy, predatory financial engineering, or international trade distortions (Section 8.1). These are collective-action failures rooted in competitive incentives, not in the lack of technological capacity or adoption. What AGI+UBI *does* achieve is to avoid *multiplying* these problems: if surplus workers no longer need to find employment at any cost, the pressure that inflates BS-jobs, hyper-competition, and exploitative gig work is removed at the source. If AI takes over the BS-jobs rather than creating new ones for displaced humans, this is already an improvement over the status quo. Full elimination of coordination failures would require complementary institutional reforms, global tax coordination, advertising regulation, financial-market oversight, that lie beyond the scope of this book but are considerably easier to pursue in a society no longer desperate for full employment.

10.6 Conclusion

The decoupling of human labor from economic productivity due to AGI and advanced robotics should not be feared as a crisis, but embraced as an unprecedented opportunity. A jobless society is not only economically affordable, sustained by the sheer abundance of automated production, but fundamentally desirable. Attempting to artificially maintain full employment will only trap humanity in hyper-competitive, socially wasteful BS-jobs.

To navigate this transition, wealth redistribution is mandatory. While various mechanisms exist, such as mandated work hour reductions or UBC, Universal Basic Income (UBI) stands out as the most straightforward, robust, and minimally invasive solution. By directly socializing the dividends of automated productivity, UBI ensures universal material prosperity without complex bureaucracies or forced labor.

It is often argued that “humans are apes, not horses,” meaning our inherent social nature will preserve interpersonal service jobs indefinitely. While it is true that demand for human empathy, motivation, and interaction will remain, this deeply social ‘work’ will eventually detach from financial necessity. As basic needs are met unconditionally, these interactions will become purely voluntary. Consequently, they may not command high wages and will increasingly resemble private socializing or community volunteering rather than traditional employment.

Ultimately, replacing instrumental toil with freely chosen pursuits (Section 7.5), be they creative, social, or competitive, represents a profound liberation. With the right redistributive policies like UBI firmly in place, the end of compulsory work marks the beginning of a true post-scarcity utopia, allowing humanity to redefine its purpose beyond the confines of economic utility (see Figure 6.4).

Key Insights

- The argument for a post-labor utopia is robust across a wide range of AGI timelines: whether 5 years or 50, the structural logic is the same.
- UBI replaces instrumental toil with freely chosen pursuits: creative, social, competitive, or contemplative.
- Institutional and cognitive inertia (a population conditioned to equate work with meaning) is the primary bottleneck of the transition.
- The transition requires parallel progress on technology, redistribution, and social norms.
- Relational care decouples from financial utility under UBI, moving from market employment to the voluntary social sphere.

Open Problems

- Empirical validation: current UBI pilots operate at far smaller scales, shorter durations, and lower basic-income levels than full deployment would require.
- The political economy of enacting redistribution before crisis: how to overcome the “boiling frog” problem.
- Long-term coexistence with ASI: whether human agency and autonomy can be preserved in a world where machines outperform humans at every task.

Appendix A

Appendices

A.1 Full List of Arguments for UBI

This appendix lists 70 arguments for (and against) UBI in compact form. Arguments are cross-linked to the companion Appendix A.2: $=$ [XY] denotes an equivalent argument; $\approx$ [XY] a closely related one. The most important ones are developed in Sections 3.2, 4.5, 5.2 and 6.3, with key data in Tables 6.1, 6.6 and 6.8 and Figures 6.4 and 7.7. For further reading, see [Van Parijs and Vanderborght \(2017\)](#), [Standing \(2017\)](#), and [Gentilini et al. \(2020\)](#).

Moral and philosophical arguments.

- [MF] **Real-libertarian freedom.** Justice requires maximizing real (not merely formal) freedom for all, which demands the highest sustainable UBI. (Section 6.1 and Table 6.6)
- [MN] **Non-domination (republican freedom).** UBI provides the “power to say no”: an exit option from exploitative relationships, abusive partners, coercive employers. (Section 6.1 and Table 6.6) ([Pettit, 2012](#))
- [MG] **Gender justice.** Compensates unpaid care work (mostly done by women) without requiring verification or surveillance. (Section 6.1 and Table 6.6)
- [MD] **No desert requirement.** External assets (land, technology, inherited knowledge, luck) are not individually earned; unconditional sharing is justified. (Section 6.1 and Table 6.6) ([Paine, 1797](#))
- [MH] **Universal dignity.** No means-testing humiliation, no welfare stigma. (Section 6.1 and Table 6.6)
- [MO] **Common ownership of knowledge and nature.** Natural resources and the cumulative stock of human knowledge are commons; everyone deserves a dividend. (Sections 6.1 and 2.6) ([Paine, 1797](#); [George, 1879](#);

Gabriel and Kasirzadeh, 2026)

- [MX] **Anti-exploitation.** Eliminates economic coercion into degrading, dangerous, or pointless jobs. (Section 4.5) ≈[JB] (Graeber, 2018)
- [ME] **Ecological justice.** Decouples income from environmentally destructive production; enables people to refuse ecologically harmful work. (Section 6.1)
- [MQ] **Gift equalization.** Corrects morally arbitrary advantages (birth, talent, inheritance). (Rawls, 1971) (Section 6.1)
- [MC] **Conditionality critique.** All behavioral conditions on welfare are either unenforceable, unjust, or both. (Section 6.1)
- [MA] **Fairness of automation gains.** AI builds on millennia of collective scientific and cultural achievement; everyone is heir to these gains. (Section 5.2) ≈[AC] (Gabriel and Kasirzadeh, 2026)
- [MI] **Intergenerational justice.** Future generations should not inherit a system where survival depends on jobs that no longer exist. (Section 6.2) ≈[TA]

Economic arguments.

- [EP] **No poverty trap.** No implicit high marginal tax from benefit withdrawal; smooth incentives at all income levels. (Section 6.1 and Table 6.6)
- [EK] **Full take-up.** 100% coverage by definition; no one falls through cracks due to ignorance, shame, or bureaucratic barriers. (Section 6.1 and Table 6.6)
- [EA] **Administrative simplicity.** Near-zero bureaucratic overhead compared to means-tested programs. (Section 6.1 and Table 6.6)
- [EL] **Labor flexibility.** Supports part-time, gig work, career transitions, sabbaticals, retraining. (Section 6.1 and Table 6.6)
- [EN] **Entrepreneurship buffer.** Risk cushion enables self-employment and innovation without fear of destitution. (Section 6.1 and Table 6.6)
- [EB] **Bargaining power.** Workers can refuse exploitative offers; reservation wage rises. (Section 6.1 and Table 6.6)
- [EV] **No surveillance.** No behavioral monitoring of recipients. (Section 6.1 and Table 6.6)
- [EF] **Affordable and fiscally feasible.** The idealized transition scenario shows displaced workers' salaries fund their own BI via corporate tax on saved wages; net cost is 25–35% of gross cost after tax clawback. (Sections 3.2 and 6.3) = [DF] (Widerquist, 2017; Gentilini et al., 2020)
- [ED] **Purchasing power and demand collapse prevention.** Without UBI, unemployed masses have no purchasing power, consumer demand collapses, and the economy shrinks even though the rich could buy

- everything. Capital owners themselves lose from demand collapse. (Sections 5.2 and 7.2) $\approx$ [DW] (Saint-Paul, 2025; Imas, 2026a; Restrepo, 2025)
- [EC] **Counter-cyclical stabilizer.** Automatic macroeconomic stabilizer during recessions; no legislative delay needed. (Hoynes and Rothstein, 2019) (Section 6.1)
- [EX] **Diverse funding options.** Can be funded by income tax, capital/wealth tax, automation tax, carbon tax, sovereign wealth fund, land value tax, IP rents, etc. (Section 5.3)
- [EW] **Pilot evidence.** Consistently positive results worldwide, no major work reduction, improved wellbeing, increased full-time employment in some cases. (Table 6.1) (Gentilini et al., 2020; Bohmann et al., 2025; Bernhard et al., 2026)
- [EG] **Gradual introduction.** Can be phased in alongside existing welfare via a three-phase roadmap. (Section 6.2 and Figure 6.4) (Standing, 2017)
- [EI] **Not inflationary.** Tax-funded UBI is redistribution, not money creation; supply expanding via automation absorbs demand. (Section 6.3) $\approx$ [DF] (Hoynes and Rothstein, 2019)
- [ET] **Pre-legislated triggers.** Can be pre-committed via measurable indicators (labor share, median wage, unemployment), removing political gridlock. (Section 6.2) $\approx$ [TY] (Korinek and Juelfs, 2022)
- [EJ] **Eliminates BS-jobs.** Without UBI, displaced workers are forced into hyper-competitive, socially useless jobs (marketing arms races, annoyance economy). UBI lets them exit. (Section 4.5 and Table 4.9)=[JB] (Graeber, 2018; Walo, 2023)
- [EZ] **Zero marginal cost convergence.** As production costs approach zero, UBI cost becomes trivially small relative to output. (Section 3.3)=[DZ] (Rifkin, 2014)
- [ES] **Simpler than alternatives.** Share ownership reduces to a citizens' wealth fund, which reduces to a sovereign wealth fund, which reduces to tax-funded UBI (but with added complexity at each step). (Table 6.8 and Section 6.5)
- [EM] **Replaces minimum wage.** UBI can substitute for minimum-wage price floors, eliminating their distortionary unemployment effects while maintaining or raising total worker income (market wage plus transfer). (Section 6.4) (Friedman, 1962; Van Parijs, 2004)

Political arguments.

- [PP] **Growing public support (post-COVID).** Polls show rising support, especially after COVID-19: 2016 EU 64% in favour; 2017 Gallup 48% US; 2020 Oxford 71% Europe; 2020 Hill-HarrisX 55% US ($\uparrow$ 6 pp during

COVID); 2020 Germany 51%; 2020 Australia 58%. (Laenen, 2023) (Section 6.1)

- [PD] **Democratic feedback loop.** Once a majority is displaced, they vote for redistribution; existing democratic institutions suffice. (Section 7.1)
- [PA] **AI CEO advocacy.** Many Silicon Valley CEOs (Musk, Altman, Zuckerberg, Amodei, Omidyar, Yang) publicly advocate (variations of) UBI. (Section 6.1)
- [PB] **Bipartisan framing possible.** Negative income tax framing (tax adjustment, not “handout”) appeals to conservatives; Friedman is the original proponent. (Friedman, 1962) (Section 6.1)
- [PI] **Pre-legislation is costless insurance.** If triggers never fire, pre-legislated UBI costs nothing; even skeptics should accept it. (Section 6.2)
- [PC] **Existing welfare systems converge to UBI.** Gradual relaxation of conditionality plus benefit increases equals UBI in all but name. Path of least resistance. (Section 6.1)
- [PU] **Prevents populism and political instability.** Without redistribution: mass unemployment, populism, democratic erosion or collapse. UBI is the safety valve. (Section 7.1) ≈[JE] (Boix, 2019; Huang and Manning, 2025)
- [PM] **Preserves market economy.** Unlike central planning or universal basic services, UBI distributes cash and leaves price signals intact. Not anti-capitalist. (Section 6.3) (Dumont, 2022)
- [PH] **Not a Trojan horse.** Progressive UBI complements (not replaces) public services; concerns about neoliberal dismantling conflate it with the Friedmanite version. (Section 6.3) (Dumont, 2022)
- [PR] **Resists authoritarianism.** Unconditional cash enhances geographic mobility and funds autonomous mutual-aid networks; more resilient than centralized provisioning against hostile governments. (Section 6.3) (Fitzpatrick, 2015; Jordan, 2020)
- [PE] **Real-world precedents.** Alaska Permanent Fund Dividend (since 1982), Marshall Islands nationwide UBI (2025), and multiple pilot programs demonstrate political viability. (Table 6.1) (Srinivasan, 2025; Gentilini et al., 2020)
- [PX] **International competitive pressure.** Nations that fail to redistribute face brain drain, social collapse, loss of competitiveness. (Section 8.4) ≈[TG] (Genewein et al., 2026)

Social arguments.

- [SU] **Avoids civil unrest and social collapse.** Without redistribution: poverty, populism, social unrest, up to civil war. (Section 7.1) ≈[DW]

- [SC] **Promotes child care and family time.** Extended parental engagement; parents can invest in child-rearing without financial penalty. (Section 7.5)
- [SM] **Reduces stress and improves mental health.** Pilots consistently show improved wellbeing, reduced anxiety; permanent income floor eliminates “money on the mind” salience distress. (Section 6.3) (Bohmann et al., 2025; Bernhard et al., 2026)
- [SH] **Improves physical health.** Reduced financial stress leads to better health outcomes; some pilots show reduced hospitalization. (Gentilini et al., 2020) (Section 7.5)
- [SV] **Enables volunteering and community service.** People freed from survival jobs can engage in unpaid socially valuable work. (Section 7.5) (Danaher, 2019)
- [SA] **Supports artists and creative work.** Irish basic income for artists pilot showed outsized socio-economic returns. (Section 6.1) (Hogan, 2026)
- [SE] **Enables entrepreneurship and innovation** (social dimension). More people can pursue ventures without catastrophic downside. (Table 6.6)
- [SL] **Supports students and lifelong learning.** Enables education and retraining without debt burden; critical as skill requirements shift rapidly. (Section 7.5)
- [SK] **Supports homemakers and caregivers.** Compensates currently unrecognized domestic labor; reduces gender inequality. (Table 6.6)
- [SO] **Enables open-source developers and public goods creators.** People contributing to open-source software, Wikipedia, citizen science, etc. currently do so at personal financial cost; UBI removes this barrier. (Section 7.5)
- [SP] **Meaning persists without jobs.** About half the working population already find primary meaning outside work; purposeful unpaid work (Section 2.2) satisfies psychological needs without compensation. (Section 7.5 and Figure 7.7) $\approx$ [JS] (Morse and Weiss, 1955; Silver et al., 2021; Jahoda, 1982)
- [ST] **Crisis of meaning is transitional.** The psychological attachment to jobs-as-identity is a cultural artifact of current socialization, not a permanent human condition; the next generation can be raised differently. (Section 7.5)
- [SG] **Competition and play as substitutes.** Structured gaming, sports, and e-sports satisfy competence and status needs; combined with other activities, they cover the psychological gap. (Section 7.5) (Tyack and Mekler, 2024; Danaher, 2019; Pelligra and Sacco, 2023)
- [SR] **Reduces crime.** Financial security reduces economically motivated crime; pilots show reduced antisocial behavior. (Gentilini et al., 2020)

(Section 7.5)

- [SF] **Prevents “managerial feudalism.”** Without UBI, wealth concentrates among AI owners, creating a neo-feudal power structure. UBI maintains broad-based economic participation. (Section 7.3) ≈[DG] (Stiefenhofer, 2025; Huang and Manning, 2025; Boix, 2019)

Counter-arguments and rebuttals.

- [CL] **“People will become lazy.”** Empirically false: pilots and lottery studies show minimal work reduction; even if true, this is a feature not a bug in a post-labor economy. (Section 6.3)
- [CA] **“UBI is unaffordable.”** Net cost is 25–35% of gross cost after replacing existing transfers and tax clawback; the idealized transition scenario shows salary savings suffice. (Sections 6.3 and 3.2) (Widerquist, 2017)
- [CI] **“UBI will cause inflation.”** Tax-funded UBI is redistribution, not money creation; supply expanding via automation absorbs additional demand. (Section 6.3) (Hoynes and Rothstein, 2019)
- [CR] **“Reciprocity violation: something for nothing.”** External assets (land, technology, inherited knowledge) justify unconditional sharing; the Malibu surfer deserves UBI too. (Section 6.1)
- [CJ] **“Better to give people jobs.”** Jobs for jobs’ sake is wasteful (BS-jobs) when machines can do the work. (Section 4.5) (Graeber, 2018)
- [CU] **“Universal basic services are better.”** Cash is superior to services because it respects individual choice; UBI preserves price signals. (Section 6.3) (Coote and Percy, 2020)
- [CE] **“A basic endowment (lump sum) is better.”** Stakeblowing risk: individuals may squander a one-time grant; periodic income is structurally superior. (Section 6.5) (Ackerman et al., 2006)
- [CM] **“UBI harms mental health.”** The Thomson et al. result is a modeling artifact; Jaroszewicz et al. find windfall salience distress, not a UBI effect. (Section 6.3) (Thomson et al., 2024; Jaroszewicz et al., 2022)
- [CT] **“UBI is a Trojan horse for neoliberalism.”** Progressive UBI complements public services; this objection conflates it with the Friedmanite version that replaces the welfare state. (Section 6.3) (Dumont, 2022)
- [CP] **“UBI leads to central planning.”** UBI distributes cash, not goods; market signals remain fully intact. (Section 6.3)
- [CV] **“UBI is vulnerable to authoritarian capture.”** Unconditional cash enhances geographic mobility and funds autonomous mutual-aid networks; it is more resilient than centralized provisioning against hostile governments. (Section 6.3) (Fitzpatrick, 2015; Jordan, 2020)
- [CS] **“Share ownership is better than UBI.”** Share-based models offer citizen agency (a capital analog of labor power) and can complement UBI

as a funding mechanism, but each variant (individual shares, CWF, SWF) converges toward tax-funded UBI once practical safeguards (sale restrictions, passive indexing, diversification) are imposed. (Section 6.5 and Table 6.8)

A.2 Full List of Arguments Why Jobs Will Disappear

This appendix lists 56 arguments for and against AI-driven job displacement. Arguments are cross-linked to the companion Appendix A.1: $=$ [XY] denotes an equivalent argument; $\approx$ [XY] a closely related one. The most important ones are developed in Sections 2.1, 3.1, 3.10, 4.1 and 4.4, with key data in Tables 2.2, 2.6, 2.8 and 4.9 and Figures 2.4, 2.9, 3.1 and 4.3. For further reading, see [Susskind \(2020\)](#), [Brynjolfsson and McAfee \(2014\)](#), [Korinek and Suh \(2024\)](#), and [Graeber \(2018\)](#).

AI and automation capability.

- [AC] **AGI automates cognition itself.** Machines are acquiring the very cognitive flexibility that previously made humans indispensable. Unlike past automation that replaced physical power, this automates general intelligence itself: structurally unprecedented. (Sections 2.1 and 4.1) $\approx$ [MA] ([Brynjolfsson and McAfee, 2014](#); [Frey and Osborne, 2017](#))
- [AS] **Superhuman AI worker properties.** Superhuman I/O rates, scalable reasoning, vastly superhuman working memory, substrate-independent, losslessly replicable including lifetime experience, high-bandwidth knowledge sharing. More rational, emotionally stable, informed. (Section 3.1) ([Genewein et al., 2026](#); [Susskind, 2025](#))
- [AO] **Automated corporations / group agents.** AGI agents form fully automated corporations with emergent cognitive capabilities superior to human organizations. (Section 4.1) ([Genewein et al., 2026](#))
- [AR] **Recursive self-improvement.** AI systems increasingly automate AI research itself, causing the temporary friction from AI skills shortage to quickly evaporate. (Section 3.10) ([Genewein et al., 2026](#))
- [AF] **Interface friction forces corner solutions.** As AI-native workflows accelerate, the round-trip latency tax of exporting context to biological processing exceeds the value humans add, causing firms to exclude humans entirely. (Section 4.1) ([Krier, 2026](#))
- [AM] **Moravec's paradox being solved.** Unskilled physical work survives longest, but advanced robotics is actively closing the gap: dexterous

manipulation, bipedal locomotion, sensorimotor learning. (Section 3.10 and Figure 3.9) (Moravec, 1988; Stanford HAI, 2026)

- [AB] **General-purpose robots imminent.** Physical labor faces disruption from humanoid robots whose costs are projected to undercut human labor within a decade. (Section 3.1) (Stanford HAI, 2026; Grace et al., 2025)
- [AL] **LLMs already exceed human benchmarks.** On most academic and professional benchmarks, frontier LLMs now match or exceed median human performance. (Figure 3.1) (Stanford HAI, 2026; Grace et al., 2025)

Economic drivers.

- [DZ] **Zero marginal cost convergence.** All production costs trace to labor + raw materials + energy. Once labor is automated and energy abundant, marginal costs converge to zero. (Section 3.3)=[EZ] (Rifkin, 2014; Altman, 2021)
- [DC] **Full cost decomposition.** Distribution, marketing, overhead, capital charges all decompose into labor + materials + energy — all converging to zero. (Section 3.5) (Korinek and Stiglitz, 2019)
- [DP] **Price = Cost + Tax + Profit – Subsidies.** As production costs fall toward zero, taxes rise by the same amount; consumer price remains approximately constant. This is how UBI is funded. (Section 3.8)
- [DF] **Idealized transition: fiscal feasibility proof.** Employee earning salary S replaced by AI. Company saves S , taxed S , redistributed as basic income. Profit margin and income unchanged. (Section 3.2)=[EF]
- [DL] **Labor share declining.** US labor share declined ~10pp from ~66.5% (1960) to ~57% (2023). AI will accelerate this trend. (Figure 2.5) (Skare and Skare, 2017; Brynjolfsson and McAfee, 2014)
- [DW] **Wage erosion to zero.** Even where preference or moral limits preserve a human role, the wage is set by competition with near-zero-cost machine alternatives. (Section 4.4) ≈[ED] ≈[SU] (Cunningham, 2025; Korinek and Suh, 2024; Susskind, 2025)
- [DB] **Land as last bottleneck (but solvable).** In a zero-marginal-labor-cost economy, land and resources become the binding constraint. LVT, public ownership, Harberger taxes all address this. (Section 3.6) (George, 1879; Stiglitz, 1977; Posner and Weyl, 2018)
- [DI] **IP rents weaken under AI.** Traditional IP justification weakens once AI generates the bulk of patentable inventions at near-zero marginal cost. (Section 3.7) (Abbott, 2016)
- [DG] **Great Decoupling.** Productivity and employment have decoupled since the 1970s/1980s. Real wages stagnated while GDP per capita grew.

AI will complete this. (Section 2.3 and Figure 2.4) $\approx$ [SF] (Brynjolfsson and McAfee, 2014; Skare and Skare, 2017)

- [DR] **One robot replaces ~6 workers.** Empirical causal estimates: each additional industrial robot per thousand workers replaces approximately 6.2 workers, with no offsetting gains. (Table 2.6) (Acemoglu and Restrepo, 2020)

Historical and empirical evidence.

- [HS] **Sector shift 1840–2100.** Agriculture: 2170 $\rightarrow$ 8 hrs/yr. Industry: 465 $\rightarrow$ 35. Services: 465 $\rightarrow$ 120. Free time rises to 94% of waking hours by 2100. (Figure 2.9) (Comin et al., 2021)
- [HR] **Five industrial revolutions.** Steam, electricity, IT, AGI, robotics. The 5th (AGI, $\sim$ 2027) automates cognitive labor “very fast”; the 6th (robots, 2030s) automates physical labor. (Table 2.8)
- [HH] **Declining working hours.** US working hours dropped from $\sim$ 2964 h/yr (1870) to $\sim$ 1789 (2023). European countries declined further (Germany: 1335, France: 1487). (Figure 4.6) (Huberman and Minns, 2007; Stanford HAI, 2026)
- [HL] **100K+ AI-linked job losses already.** Over a hundred thousand explicitly AI-linked job losses reported as early indicators. (Section 2.3) (The Alliance for Secure AI, 2026)
- [HI] **Global AI investment exponential.** Private AI investment: \$15B (2013) $\rightarrow$ \$345B (2025), projected \$506B (2026). This exceeds 10 $\times$ the Manhattan Project budget. (Figure 2.7) (Stanford HAI, 2026)
- [HF] **Displacement forecasts.** Frey & Osborne: 47%. OECD: 9–14%. McKinsey: 75–375M. Goldman Sachs: 300M. WEF: 92M displaced, 170M created. Most are pre-LLM, so likely too conservative. (Table 2.6)
- [HE] **Horse population analogy.** Leontief (1983): “The role of humans as the most important factor of production is bound to diminish — in the same way that the role of horses. . .” US horse population peaked $\sim$ 26M around 1915, then collapsed. (Section 4.1 and Figure 4.3) (Leontief, 1983a; Brynjolfsson and McAfee, 2015)
- [HN] **60% of 2018 jobs didn’t exist in 1940, but that won’t repeat.** Past automation created new job categories. But that pattern rested on the assumption that human cognition remains indispensable. AGI breaks this. (Section 4.1) (Autor et al., 2024)

Timeline, tempo, and scope.

- [TY] **Timelines in years, not decades.** AI CEO predictions: Hassabis “a few years, maybe within a decade,” Altman “a few thousand days,” Amodei “as early as 2026.” (Section 3.1) $\approx$ [ET] (Grace et al., 2025)

- [TS] **AI expert survey.** 2 778 AI researchers surveyed. Median predictions moved forward ≈ 13 years versus the 2022 survey. HLMI by 2047; full automation by 2116 (book: far too conservative). (Figure 3.1) (Grace et al., 2025)
- [TC] **S-curve may slow but won't prevent.** S-curve limits (energy, regulation, organizational inertia) may ease the transition but will not prevent long-term decoupling. (Section 4.1 and Figure 4.4)
- [TR] **AI changes too fast for retraining.** AI-induced changes will be too fast for retraining and generational turnover. By the time a worker retrain, the target skill is also automated. (Section 3.10)
- [TA] **Virtually all cognitive + physical tasks.** Full AGI+Robotics will be able to do virtually all human jobs, ultimately at a fraction of the cost. (Section 4.1) $\approx$ [MI]
- [TN] **Even "new jobs" will be automatable.** In the era of AGI, most jobs requiring novel cognitive skills will be performed better by the AIs themselves. The historical escape valve closes. (Section 3.10)
- [TG] **No country can opt out.** Any country not investing in AI will seriously fall behind. International competitive pressure makes adoption inevitable. (Sections 2.3 and 8.5)=[PX] (Genewein et al., 2026; Cerutti et al., 2025)
- [TI] **AI colonialism.** AI leaders extract massive profits from followers. Developing nations lose comparative advantage in cheap labor. (Sections 8.5 and 8.1) (Korinek and Stiglitz, 2021; Mohamed et al., 2020; Cerutti et al., 2025)

Job property degradation and comparative advantage.

- [JD] **11-property job definition.** Job as Wittgensteinian cluster concept: [LC]–[LI]. Automation erodes multiple properties simultaneously. (Section 2.2 and Table 2.2)
- [JW] **UBI weakens multiple job properties.** UBI weakens [LC] (compensation), [LL] (livable income), [LO] (opportunity cost), [LI] (involuntariness); activities progressively lose "job" status. (Section 4.4)
- [JB] **BS-jobs fail substantive properties.** BS-jobs satisfy formal criteria ([LC], [LR], [LS], [LM]) but fail [LN] (net benefit). Forcing full employment post-AGI creates a dystopia of pointless work. (Section 4.5 and Table 4.9)=[EJ] $\approx$ [MX] (Graeber, 2018; Walo, 2023)
- [JC] **Humans have no durable comparative advantage.** Moral, preference, and general-equilibrium forces may slow displacement but none constitutes a durable floor. (Section 4.4) (Susskind, 2025; Korinek and Suh, 2024)

- [JL] **Three limits on automation; all erode.** Moral limits soften when machines outperform; preference limits shift when AI output becomes indistinguishable; general-equilibrium limits push wages below subsistence. The “too-good problem.” (Section 4.4) (Susskind, 2025)
- [JA] **Artisanal niche: narrow luxury only.** Provenance premium protects only a narrow luxury niche. For the vast majority, labor displacement collapses demand. (Section 4.2) (Imas, 2026b; Korinek and Suh, 2024)
- [JP] **Price competition: income converges to zero.** Hand-made prices must undercut close-to-zero machine prices. (Section 4.2) (Cunningham, 2025; Korinek and Suh, 2024)
- [JS] **Symmetric social interaction fails as “work.”** Mutual social enjoyment fails [LA] (asymmetry), [LO] (opportunity cost), [LM] (market embeddedness). Circular payments do not create genuine economic activity. (Section 4.2) ≈[SP]
- [JF] **Formal proofs of displacement.** Samuelson (1989); Imas (2026a) showed that cheap machine substitutes can eliminate human employment and even reduce total output. (Section 4.1) (Samuelson, 1989; Imas, 2026a)
- [JT] **Acemoglu task model: automation outpaces.** Automation moves the task threshold rightward faster than new tasks extend the frontier, leading to a net loss of jobs. (Figure 4.2) (Acemoglu and Restrepo, 2019; Restrepo, 2025)
- [JE] **Full employment through social jobs: rebutted.** Aggregate demand for human labor cannot sustain full employment. Structural underemployment is the equilibrium. (Section 4.4) ≈[PU] (Restrepo, 2025)
- [JY] **Forced full employment → BS-jobs dystopia.** Extended taxonomy: Flunkies, Goons, Duct Tapers, Box Tickers, Taskmasters, Arms Racers, Obstructionists, Rent-Seekers, Charlatans. (Section 4.5 and Table 4.9) (Graeber, 2018; Baumol, 2004; Gabaix and Laibson, 2006)

Counter-arguments (why jobs WON'T disappear) and rebuttals.

- [XN] **“New jobs will always be created.”** Comforting but dangerous mantra. Historical pattern rests on assumption that human cognition remains indispensable. AGI breaks this. (Section 4.1)
- [XR] **“Humans are not horses: they can retrain.”** Acknowledges the Great Decoupling but underestimates cognitive automation. Social “work” eventually detaches from financial necessity. (Sections 4.2 and 4.1) (Brynjolfsson and McAfee, 2015; Vollrath, 2015)
- [XC] **“Comparative advantage guarantees employment.”** Susskind shows wages converge to zero as AI costs fall. Samuelson and Imas proved

- cheap machines can eliminate employment entirely. (Section 4.4) (Susskind, 2025; Samuelson, 1989; Imas, 2026a)
- [XH] **“AI is just hype / S-curve will plateau.”** Given \$345B in 2025, betting against broad AI diffusion is risky. Even a plateau at current levels threatens vast swathes of white-collar work. (Section 2.1 and Figure 4.4)
- [XL] **“Regulation will prevent displacement.”** Any country not adopting AI falls behind competitively. Regulation can ease speed but cannot prevent it globally. (Sections 2.3 and 8.5)
- [XP] **“People prefer human-made goods/services.”** Preference limits erode as AI output becomes indistinguishable. Remaining provenance premium supports only a narrow luxury niche. (Section 4.2) (Susskind, 2025; Imas, 2026b)
- [XM] **“Moral limits: some decisions require humans.”** Moral objections soften when machines demonstrably outperform humans. Ethically difficult to justify inferior human judgment when lives are at stake. (Section 4.4) (Susskind, 2025)
- [XT] **“Technology creates more jobs than it destroys.”** Historically true for narrow automation. AGI differs qualitatively: it automates the cognitive flexibility that was humanity’s unique advantage. (Section 4.1) (Autor et al., 2024)
- [XA] **“We’ll all become artists/caregivers.”** Aggregate demand for such labor cannot sustain full employment. Remaining “human touch” jobs pay near-zero wages. (Sections 4.2 and 4.4) (Restrepo, 2025)
- [XS] **“AI assists, doesn’t replace.”** True in early phases (copilot model). But as capability grows, interface friction makes including humans net-negative. (Section 4.1) (Krier, 2026)

Back Matter

Glossary: Core Acronyms

AGI (Artificial General Intelligence)

AI system matching or exceeding median human cognitive abilities across domains (Genewein et al., 2026). (p.1, 9, 30, 42, 60, 87)

AI (Artificial Intelligence)

Broad field of computer science aiming to build systems that perform tasks requiring human-like intelligence. (p.9,222)

ASI (Artificial Superintelligence)

AI far surpassing the best human capabilities (Genewein et al., 2026). (p.31, 32, 38, 68)

BI (Basic Income)

Umbrella term for any unconditional income floor (Section 2.6). (p.111, 112, 212)

CI (Conditional Income)

Means-tested or behavior-conditional income support, as opposed to unconditional BI. (p.21, 25, 51)

CWF (Citizens' Wealth Fund)

Publicly owned investment fund that distributes dividends to all citizens. (p.v and p.89, 128, 129)

GDP (Gross Domestic Product)

Total monetary value of all finished goods and services produced within a country's borders in a given period. (p.42,187)

GI (Guaranteed Income)

Funding-agnostic umbrella term covering UBI, NIT, CWF dividends, and similar periodic cash transfers above subsistence level (Section 2.6). (p.v,vi and p.6, 23, 25)

NIT (Negative Income Tax)

Tax system that provides payments to individuals below a defined income threshold, functioning as a form of GI. (p.119,120)

SWF (Sovereign Wealth Fund)

State-owned investment fund, typically capitalized from resource revenues or fiscal surpluses (e.g. Norway's Government Pension Fund). (p.129, 130, 137)

UAI (Universal Adequate Income)

Unconditional income covering basic needs; the middle phase in the CI → UAI → UHI transition (Section 6.2). (p.112,139)

UBI (Universal Basic Income)

Unconditional, universal, individual, periodic cash transfer to all citizens, regardless of other income or employment status (Section 2.6). (p.v and p.21, 23, 105, 106, 208)

UBS (Universal Basic Services)

Free or heavily subsidized provision of essential services (healthcare, education, housing, transport) as an alternative or complement to cash transfers. (p.89, 119, 121, 171)

UHI (Universal High Income)

Unconditional income at levels enabling comfortable flourishing, not mere subsistence; the long-term phase of the GI transition (Section 6.2). (p.25, 112, 139)

Economics Concepts

Balassa-Samuelson effect

The observation that countries with higher productivity levels tend to have higher price levels for non-tradable goods and services. (p.167)

Baumol's cost disease

The tendency for costs to rise in labor-intensive sectors (e.g. education, healthcare) where productivity gains are limited, because wages must keep pace with more productive sectors. (p.186)

Creative destruction

The process by which new technologies, products, or business models displace established ones, simultaneously destroying old industries and creating new ones. (p.1)

Deadweight loss

Economic inefficiency caused by a tax, subsidy, or market distortion; the loss of total surplus that occurs when the equilibrium quantity of a good or service is not achieved. (p.47,99)

Externality

A cost or benefit imposed on third parties not directly involved in a transaction (e.g. pollution is a negative externality; vaccination confers a positive externality). (p.76,97)

Fiscal feasibility

The question of whether a policy (here, GI) can be sustainably financed from available revenue instruments without destabilizing the economy. (p.2,39)

Georgism / LVT (Land Value Tax)

The policy of taxing the unimproved value of land rather than buildings or improvements, following [George \(1879\)](#). (p.47, 48, 99)

Gini coefficient

A measure of income or wealth inequality ranging from 0 (perfect equality) to 1 (maximum inequality). (p.151)

Harberger tax

A self-assessed property tax: owners publicly declare the value of their asset; anyone may buy it at that declared price, ensuring honest valuation. (p.vi and p.4,48)

Jevons paradox

The observation that increased efficiency in resource use can lead to *greater* total consumption of that resource, because lower unit costs stimulate demand. (p.42,196)

Marginal cost

The cost of producing one additional unit of a good or service; central to the book's zero-cost convergence thesis. (p.3, 42, 187)

Means-testing

Restricting eligibility for benefits to individuals who pass income or asset checks, as opposed to universal provision. Reduces fiscal cost but introduces administrative burden, stigma, and poverty traps. (p.vi and p.108)

Output

The total goods and services produced by an economy. Used interchangeably in this book with *goods and services*, *products*, and *offerings*.

Pigouvian tax

A tax levied on activities that generate negative externalities, designed to internalize external costs (e.g. carbon tax, sin tax). (p.45, 48, 97)

Positional good

A good or service valued primarily for its relative rank rather than its absolute utility (e.g. luxury brands, elite education). Demand for positional goods persists even in post-scarcity, sustaining competition. (p.53)

Regulatory capture

The process by which regulatory agencies come to serve the interests of the industries they are meant to oversee, undermining public-interest governance. (p.141,151)

Rentier state

A state that derives a substantial share of its revenues from external rents (e.g. oil exports) rather than domestic taxation, often weakening democratic accountability. (p.145,195)

Domain-Specific Terms

BS-job (Bad-for-Society job)

Employment that is socially wasteful or perceived as pointless by the worker; extends Graeber's "bullshit job" concept to include jobs harmful to society (Section 4.5). (p.76)

Conditionality

The practice of tying welfare benefits to behavioral requirements (e.g. job search, training) or means tests. A central design axis distinguishing CI from UBI. (p.vi and p.116)

Great Decoupling

The divergence of productivity growth from median wages and employment that began in the 1970s. (p.3,15)

Idealized transition scenario

A step-by-step thought experiment (Section 3.2) showing that redis-

tribution of automation savings can maintain or improve everyone's welfare as jobs are progressively automated. (p.3,30)

Labor share

The fraction of national income paid to workers as wages, as opposed to capital owners as profits. A declining labor share is a key symptom of the Great Decoupling. (p.11, 62, 112)

Moravec's paradox

The counterintuitive observation that tasks humans find hard (logical reasoning, chess) are easy for computers, while tasks humans find easy (perception, motor control) are hard for computers. (p.54,183)

Post-labor economy

An economy in which most human labor has been automated and livelihoods are maintained through redistribution rather than employment. (p.v and p.106,155)

Premature deindustrialization

The phenomenon in which developing countries see their manufacturing sectors shrink before reaching the income levels at which deindustrialization historically occurred in advanced economies. (p.173,174)

Provenance premium

The price premium consumers willingly pay for goods and services verifiably made by humans rather than machines, analogous to the premium for handcrafted or artisanal products. (p.66,72)

Resource curse

The tendency for countries rich in natural resources to experience slower economic growth, weaker institutions, and less democratic governance than resource-poor peers. (p.132,148)

Skill mismatch

The gap between the skills workers possess and the skills demanded by available jobs, exacerbated when automation changes skill requirements faster than retraining can follow. (p.204)

Bibliography

- Abbott, R. (2016). I think, therefore I invent: Creative computers and the future of patent law. *Boston College Law Review*, 57(4):1079–1126. [49](#), [218](#)
- Abbott, R. and Bogenschneider, B. (2018). Should robots pay taxes? Tax policy in the age of automation. *Harvard Law & Policy Review*, 12(1):145–175. [95](#), [98](#)
- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58. [17](#)
- Acemoglu, D. and Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*. PublicAffairs. [11](#), [144](#), [191](#)
- Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30. [1](#), [2](#), [11](#), [62](#), [221](#)
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244. [1](#), [10](#), [11](#), [219](#)
- Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in US wage inequality. *Econometrica*, 90(5):1973–2016. [17](#), [165](#)
- Acemoglu, D. and Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown. [147](#), [150](#), [151](#), [192](#)
- Ackerman, B. and Alstott, A. (1999). *The Stakeholder Society*. Yale University Press, New Haven, CT. [134](#), [138](#)
- Ackerman, B. and Alstott, A. (2006). Why stakeholding? In Wright, E. O., editor, *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones of a More Egalitarian Capitalism*, pages 40–60. Verso. [124](#)
- Ackerman, B., Alstott, A., and Van Parijs, P. (2006). *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones of a More Egalitarian Capitalism*. Verso, London. Edited and introduced by Erik Olin Wright. [120](#), [124](#), [126](#), [216](#)
- Aghion, P., Antonin, C., and Bunel, S. (2023). Artificial intelligence, growth and employment: The role of policy. *Economie et Statistique*. [10](#), [185](#), [188](#), [204](#)
- Agrawal, A., Brynjolfsson, E., and Korinek, A. (2026). *The Economics of Transformative AI*. University of Chicago Press. [9](#)

- Aguiar, M. and Hurst, E. (2007). Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics*, 122(3):969–1006. 19, 68
- Akcigit, U., Baslandze, S., and Stantcheva, S. (2016). Taxation and the international mobility of inventors. *American Economic Review*, 106(10):2930–2981. 95, 170
- Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M., and Zarate, P. (2022). Working from home around the world. *Brookings Papers on Economic Activity*. <https://www.brookings.edu/articles/working-from-home-around-the-world/>. 68
- Alesina, A. and Passalacqua, A. (2016). The political economy of government debt. In *Handbook of Macroeconomics*, volume 2, pages 2599–2651. Elsevier. 130
- Alexander, S. (2014). Meditations on Moloch. *Slate Star Codex*. <https://slatestarcodex.com/2014/07/30/meditations-on-moloch/>. 75
- Alstadsæter, A., Godar, S., Nicolaidis, P., and Zucman, G. (2024). Global tax evasion report 2024. Technical report, EU Tax Observatory. 95, 100, 169, 170, 175
- Altman, S. (2021). Moore’s law for everything. *OpenAI Blog*. <https://moores.samaltman.com/>. 43, 89, 125, 128, 130, 131, 132, 150, 218
- Andreescu, M., Loubes, R., Piketty, T., and Robilliard, A.-S. (2025). Global labour hours in paid and unpaid work: Inequality, productivity and structural transformation, 1800–2100. Working Paper 2025/08, World Inequality Lab. 63
- Anthropic (2025). What 81,000 people told us about the economics of AI. *Anthropic Research*. <https://www.anthropic.com/research/81k-economics>. 17
- Arendt, H. (1958). *The Human Condition*. University of Chicago Press, Chicago. 9, 15
- Armstrong, S., Bostrom, N., and Shulman, C. (2016). Racing to the precipice: A model of artificial intelligence development. *AI & Society*, 31:201–206. 175
- Association of Superannuation Funds of Australia (2025). The Australian superannuation system: An explainer. Technical report, Association of Superannuation Funds of Australia. 133
- Atkinson, A. B. (1995). *Public Economics in Action: The Basic Income/Flat Tax Proposal*. The Lindahl Lectures. Clarendon Press, Oxford. 121
- Atkinson, A. B. (1996). The case for a participation income. *The Political Quarterly*, 67(1):67–70. 119, 121, 126
- Atkinson, A. B. (2015). *Inequality: What Can Be Done?* Harvard University Press, Cambridge, MA. 67, 93, 95, 97, 121, 126
- Auerbach, A. J., Devereux, M. P., Keen, M., and Vella, J. (2017). Destination-based cash flow taxation. Working Paper 17/01, Oxford University Centre for Business Taxation. 95, 168
- Autor, D., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465. 1, 60, 219, 222

- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30. 10, 13
- Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The china shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8:205–240. 144
- Avi-Yonah, R. S. (2000). Globalization, tax competition, and the fiscal crisis of the welfare state. *Harvard Law Review*, 113(7):1573–1676. 95, 168
- Banerjee, A., Faye, M., Krueger, A., Niehaus, P., and Suri, T. (2023). Universal basic income: Short-term results from a long-term experiment in Kenya. Working paper, National Bureau of Economic Research. <https://basicincome.org/wp-content/uploads/2023/12/Give-Directly-Kenya-report-2-years.pdf>. 101, 109, 110, 115
- Barr, N. (2012). *Economics of the Welfare State*. Oxford University Press, Oxford, 5th edition. 120, 126, 127
- Bartels, L. M. (2008). *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton University Press, 2nd edition. 151
- Bauer, A., editor (2014). *Managing the Public Trust: How to Make Natural Resource Funds Work for Citizens*. Natural Resource Governance Institute and Columbia Center on Sustainable Investment, New York. 130
- Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review*, 57(3):415–426. 186
- Baumol, W. J. (2004). Red-queen games: arms races, rule of law and market economies. *Journal of Evolutionary Economics*, 14(2):237–247. 78, 79, 221
- Bernhard, S., Bohmann, S., Fiedler, S., Kasy, M., Schupp, J., and Schwerter, F. (2026). Basic income and labor supply: Theory and evidence from an RCT in Germany. Working paper, Pilotprojekt Grundeinkommen. https://maxkasy.github.io/home/files/papers/basic_income_econ.pdf. 114, 159, 213, 215
- Besley, T. and Coate, S. (1992). Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs. *American Economic Review*, 82(1):249–261. 120, 126
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed*, 17(2):389–409. 173
- Bohmann, S., Fiedler, S., Kasy, M., Schupp, J., and Schwerter, F. (2025). Pilotprojekt grundeinkommen: kein rückzug vom arbeitsmarkt, aber bessere mentale gesundheit. *DIW Wochenbericht*, 92(15):221–229. 114, 159, 213, 215
- Boix, C. (2019). *Democratic Capitalism at the Crossroads: Technological Change and the Future of Politics*. Princeton University Press. 132, 145, 146, 150, 200, 214, 216
- Boldrin, M. and Levine, D. K. (2008). *Against Intellectual Monopoly*. Cambridge University Press. 48, 51

- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press. 180, 181, 182
- Bostrom, N. (2019). The vulnerable world hypothesis. *Global Policy*, 10(4):455–476. 175, 180
- Bostrom, N. (2024). *Deep Utopia: Life and Meaning in a Solved World*. Ideapress Publishing. 11, 159
- Bruenig, M. (2018). Social wealth fund for america. Technical report, People's Policy Project. 138
- Bruenig, M. (2026). AI could destroy the labor market. We already know how to fix it. *The Argument*. <https://www.theargumentmag.com/p/ai-could-destroy-the-labor-market>. 108
- Brynjolfsson, E., Collis, A., Diewert, W. E., Eggers, F., and Fox, K. J. (2025a). GDP-B: Accounting for the value of new and free goods in the digital economy. *American Economic Journal: Macroeconomics*, 17(1):353–396. 42, 187
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2025b). Generative AI at work. *The Quarterly Journal of Economics*, 140(1):1–58. 174
- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company. 10, 15, 16, 61, 217, 218, 219
- Brynjolfsson, E. and McAfee, A. (2015). Humans are not horses. *Federal Reserve Bank of Philadelphia Policy Forum*. <https://www.philadelphiafed.org/-/media/frbp/assets/events/2015/the-economy/2015-philadelphia-fed-policy-forum/brynjolfsson-humans-horses.pdf>. 67, 219, 221
- Buterin, V., Hitzig, Z., and Weyl, E. G. (2019). A flexible design for funding public goods. *Management Science*, 65(11):5171–5187. 154, 156
- Calvo, G. A. (1998). Capital flows and capital-market crises: The simple economics of sudden stops. *Journal of Applied Economics*, 1(1):35–54. 167
- Carers Australia (2024). Superannuation for unpaid carers: The next logical step. Technical report, Carers Australia. 135, 138
- Casas, C. and Torres, J. (2024). Automation, taxation and social security sustainability. *HPE-RPE Working Papers*. <https://hpe.rpe.ac.uk/>. 98
- Case, A. and Deaton, A. (2020). *Deaths of Despair and the Future of Capitalism*. Princeton University Press. 144
- Cerutti, E. M., Garcia Pascual, A., Kido, Y., Li, L., Melina, G., Mendes Tavares, M., and Wingender, P. (2025). The global impact of AI: Mind the gap. *IMF Working Papers*, 2025(076). 164, 165, 167, 220
- Cesarini, D., Lindqvist, E., Notowidigdo, M. J., and Östling, R. (2017). The effect of wealth on individual and household labor supply: Evidence from Swedish lotteries. *The American Economic Review*, 107(12):3917–3946. 115

- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, New York. 188
- Chang, J.-H., Rynhart, G., and Huynh, P. (2016). ASEAN in transformation: Textiles, clothing and footwear: Refashioning the future. Working Paper 14, International Labour Organization, Bureau for Employers' Activities (ACT/EMP). <https://www.ilo.org/publications/asean-transformation-textiles-clothing-and-footwear-refashioning-future>. 174
- Cholbi, M. (2022). Philosophical approaches to work and labor. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2022 edition. 12
- Citrini and Shah, A. (2026). The 2028 global intelligence crisis. *Citrini Research Blog*. <https://www.citriniresearch.com/p/2028gic>. 17
- Clausing, K. A., Saez, E., and Zucman, G. (2021). Ending corporate tax evasion and tax competition: A plan. *The Annals of the American Academy of Political and Social Science*, 695(1):22–36. 169, 170
- Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374. 66, 67, 73, 74, 219
- Coote, A. and Percy, A. (2020). *The Case for Universal Basic Services*. Polity Press, Cambridge. 119, 121, 122, 126, 171, 216
- Corinth, K. and Mayhew, H. (2026). What can we learn from guaranteed basic income pilots in the United States? Evidence on employment effects. Working paper, American Enterprise Institute. 110, 114
- Costello, E. J., Compton, S. N., Keeler, G., and Angold, A. (2003). Relationships between poverty and psychopathology: A natural experiment. *JAMA*, 290(15):2023–2029. 102
- Cowan, R. S. (1983). *More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave*. Basic Books. 19
- Cunningham, T. (2025). Economics and transformative AI. *Personal Blog*. <https://tecunningham.github.io/posts/2025-09-19-transformative-AI-notes.html>. 67, 74, 218, 221
- Dachs, B., Kinkel, S., and Jäger, A. (2019). Bringing it all back home? Backshoring of manufacturing activities and the adoption of Industry 4.0 technologies. *Journal of World Business*, 54(6):101017. 174
- Danaher, J. (2019). *Automation and Utopia: Human Flourishing in a World Without Work*. Harvard University Press, Cambridge, MA. 73, 159, 215
- Davidson, T., Halperin, B., Houlden, T., and Korinek, A. (2026). When does automating AI research produce explosive growth? Feedback loops in innovation networks. NBER Working Paper 35155, National Bureau of economic Research. 185, 188, 204
- Davies, N. (2008). *Flat Earth News*. Chatto & Windus, London. 79

- De Filippi, P. and Wright, A. (2018). *Blockchain and the Law: The Rule of Code*. Harvard University Press. 156
- de Vries, A. (2023). The growing energy footprint of artificial intelligence. *Joule*, 7(10):2191–2194. 64, 196, 199
- del Rio-Chanona, R. M., Ernst, E., Merola, R., Samaan, D., and Teutloff, O. (2025). AI and jobs: A review of theory, estimates, and evidence. *arXiv preprint arXiv:2509.15265*. 17
- Devereux, M. P., Auerbach, A. J., Keen, M., Oosterhuis, P., Schön, W., and Vella, J. (2020). *Taxing Profit in a Global Economy*. Oxford University Press. 93
- Devereux, M. P., Lockwood, B., and Redoano, M. (2008). Do corporate tax rates compete over corporate tax rates? *Journal of Public Economics*, 92(5–6):1210–1235. 170
- Devereux, M. P. and Vella, J. (2018). Taxing the digitalised economy: Targeted or system-wide reform? *British Tax Review*, 4:387–405. 168
- Dharmapala, D. (2014). What do we know about base erosion and profit shifting? A review of the empirical literature. *Fiscal Studies*, 35(4):421–448. 169
- Diamond, J. (1987). The worst mistake in the history of the human race. *Discover*, 8(5):64–66. 62
- Diamond, P. A. and Saez, E. (2011). The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, 25(4):165–190. 93
- Dubber, M. D., Pasquale, F., and Das, S., editors (2020). *The Oxford Handbook of Ethics of AI*. Oxford University Press. 189
- Dumont, D. (2022). Universal basic income as a source of inspiration for the future of social protection systems? A counter-agenda. *European Journal of Social Security*, 24(4):299–318. 106, 118, 119, 214, 216
- Dye, R. F. and England, R. W., editors (2009). *Land Value Taxation: Theory, Evidence, and Practice*. Lincoln Institute of Land Policy, Cambridge, MA. 47
- Edenhofer, O., Franks, M., and Kalkuhl, M. (2021). Pigou in the 21st Century: A tribute on the occasion of the 100th anniversary of the publication of *The Economics of Welfare*. *International Tax and Public Finance*, 28(5):1090–1121. 48, 97, 100
- Elliott, M. et al. (2024). The basic income for care leavers in Wales pilot evaluation: Protocol of a quasi-experimental evaluation. *PLOS ONE*, 19(7):e0303837. 110
- Everitt, T., Lea, G., and Hutter, M. (2018). AGI safety literature review. In *Proc. 27th International Joint Conf. on Artificial Intelligence (IJCAI'18)*, pages 5441–5449, Stockholm, Sweden. 180, 182
- Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182. Updated PWT version 11.0 from <https://www.rug.nl/ggdc/productivity/pwt/>. 19, 65

- Feld, L. P. and Kirchgässner, G. (2000). Direct democracy, political culture, and the outcome of economic policy: A report on the swiss experience. *European Journal of Political Economy*, 16(2):287–306. 153, 156
- Ferdosi, M., McDowell, T., Lewchuk, W., and Ross, S. (2020). Southern Ontario's basic income experience. Technical report, McMaster University, Hamilton. 109, 110
- Ferriss, T. (2007). *The 4-Hour Workweek: Escape 9–5, Live Anywhere, and Join the New Rich*. Crown Publishers, New York. 73
- Fishkin, J. S. (2018). *Democracy When the People Are Thinking: Revitalizing Our Politics Through Public Deliberation*. Oxford University Press. 153, 156
- Fiszbein, A. and Schady, N. (2009). *Conditional Cash Transfers: Reducing Present and Future Poverty*. World Bank Publications, Washington, DC. 120, 126
- Fitzpatrick, M. P. (2015). *Purging the Empire: Mass Expulsions in Germany 1871–1914*. Oxford University Press, Oxford. 119, 214, 216
- Floridi, L. (2023). *The Ethics of Artificial Intelligence: Principles, Challenges, and Opportunities*. Oxford University Press. 189
- Frank, R. H. (2007). *Falling Behind: How Rising Inequality Harms the Middle Class*. University of California Press. 53, 68, 71, 97, 142
- Frey, C. B. (2019). *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press. 144
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280. 1, 10, 11, 18, 217
- Friedman, M. (1962). *Capitalism and Freedom*. University of Chicago Press, Chicago. 108, 112, 119, 123, 126, 192, 213, 214
- Fukuyama, F. (2014). *Political Order and Political Decay: From the Industrial Revolution to the Globalization of Democracy*. Farrar, Straus and Giroux. 152
- Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540. 77, 80, 221
- Gabriel, I. and Kasirzadeh, A. (2026). The case for globally beneficial technology. *arXiv preprint arXiv:2607.03906*. 107, 173, 212
- Gasteiger, E. and Prettnner, K. (2022). A note on automation, stagnation, and the implications of a robot tax. *Macroeconomic Dynamics*, 26(1):255–265. 99
- Gates, B. (2017). The robot that takes your job should pay taxes. *Quartz Interview*. <https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-pay-taxes/>. 98
- Genewein, T., Franklin, M., Lerchner, A., Orseau, L., Albanie, S., Bales, A., Wyeth, C., Chan, S., Gabriel, I., Leibo, J. Z., Dafoe, A., Hutter, M., Graepel, T., and Legg, S. (2026). From AGI to ASI. *arXiv:2606.12683*, page 57. 1, 20, 29, 30, 38, 54, 64, 123, 143, 160, 183, 184, 204, 214, 217, 220, 223

- Gentilini, U., Grosh, M., Rigolini, J., and Yemtsov, R. (2020). *Exploring Universal Basic Income: A Guide to Navigating Concepts, Evidence, and Practices*. World Bank, Washington, DC. 23, 110, 117, 118, 159, 199, 211, 212, 213, 214, 215
- George, H. (1879). *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions and of Increase of Want with Increase of Wealth*. D. Appleton and Company, New York. 47, 99, 107, 170, 211, 218, 225
- Gilens, M. and Page, B. I. (2014). Testing theories of american politics: Elites, interest groups, and average citizens. *Perspectives on Politics*, 12(3):564–581. 151
- Gioino, C. (2026). Welcome to the ‘annoyance economy’: Americans are paying over \$165 billion a year as companies waste their time to drive revenue. *Fortune*. <https://fortune.com/2026/02/19/americans-paying-over-165-billion-annually-vibes-based-taxes-annoyance-economy/>. 80
- Gough, I. (2019). Universal basic services: A theoretical and moral framework. *The Political Quarterly*, 90(3):534–542. 121, 122
- Grace, K., Sandkühler, J. F., Stewart, H., Weinstein-Raun, B., Thomas, S., Stein-Perlman, Z., Salvatier, J., Brauner, J., and Korzekwa, R. C. (2025). Thousands of AI authors on the future of AI. *Journal of Artificial Intelligence Research*, 84. 1, 31, 33, 34, 200, 204, 218, 219, 220
- Graeber, D. (2018). *Bullshit Jobs: A Theory*. Allen Lane, an imprint of Penguin Books, London. 13, 59, 60, 75, 76, 77, 199, 212, 213, 216, 217, 220, 221
- Gray, M. L. and Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt. 61, 62, 80
- Haber, S. and Menaldo, V. (2011). Do natural resources fuel authoritarianism? A reappraisal of the resource curse. *American Political Science Review*, 105(1):1–26. 147
- Hadfield-Menell, D., Dragan, A., Abbeel, P., and Russell, S. (2017). The off-switch game. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 220–227. 154, 156
- Hamilton, D. and Darity Jr., W. A. (2010). Can ‘baby bonds’ eliminate the racial wealth gap in putative post-racial america? *The Review of Black Political Economy*, 37(3–4):207–216. 134, 138
- Hanson, R. (2000). Long-term growth as a sequence of exponential modes. Technical report, George Mason University. <http://hanson.gmu.edu/longgrow.pdf>. 185
- Harari, Y. N. (2018). *21 Lessons for the 21st Century*. Spiegel & Grau. 144
- Hayek, F. A. (1944). *The Road to Serfdom*. University of Chicago Press, Chicago. 192
- Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35(4):519–530. 122
- Hayek, F. A. (1960). *The Constitution of Liberty*. University of Chicago Press, Chicago. Definitive edition: *The Collected Works of F. A. Hayek*, vol. XVII, ed. Bruce Caldwell, 2011. 192

- Hendrycks, D. (2024). *Introduction to AI Safety, Ethics, and Society*. Taylor & Francis. 180, 182
- Hogan, C. (2026). Ireland's basic income for artists changed my life. other people deserve the same luck. *The Guardian*. <https://www.theguardian.com/commentisfree/2026/mar/13/ireland-basic-income-artist-scheme-pilot>. 109, 112, 160, 215
- Hoynes, H. and Rothstein, J. (2019). Universal basic income in the United States and advanced countries. *Annual Review of Economics*, 11:929–958. 108, 118, 120, 121, 126, 213, 216
- Huang, S. and Manning, S. (2025). Here's how to share AI's future wealth. *Noema Magazine*. <https://www.noemamag.com/heres-how-to-share-ais-future-wealth/>. 125, 128, 132, 145, 146, 150, 200, 214, 216
- Huberman, M. and Minns, C. (2007). The times they are not changin': Days and hours of work in old and new worlds, 1870–2000. *Explorations in Economic History*, 44(4):538–567. 219
- Humphreys, M., Sachs, J. D., and Stiglitz, J. E., editors (2007). *Escaping the Resource Curse*. Columbia University Press, New York. E-ISBN 978-0-231-51210-7. 130
- Hussam, R., Kelley, E. M., Lane, G., and Zahra, F. (2022). The psychosocial value of employment: Evidence from a refugee camp. *American Economic Review*, 112(11):3694–3724. 155
- Hutter, M. (2012). Can intelligence explode? *Journal of Consciousness Studies*, 19(1-2):143–166. 186
- Hutter, M. (2021). Global marketing budget vs extreme poverty. X (formerly Twitter). 79
- Hutter, R. and Hutter, M. (2021). Chances and risks of artificial intelligence — A concept of developing and exploiting machine intelligence for future societies. *Applied System Innovation*, 4(2):1–19. 20, 87
- Imas, A. (2026a). Can advanced AI lead to negative economic growth? *Substack*. <https://aleximas.substack.com/p/will-advanced-ai-lead-to-negative>. 63, 129, 213, 221, 222
- Imas, A. (2026b). What will be scarce? the economics of structural change and the post-commodity future of work. *Substack*. <https://aleximas.substack.com/p/what-will-be-scarce>. 66, 67, 71, 72, 73, 74, 221, 222
- Imbens, G. W., Rubin, D. B., and Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *The American Economic Review*, 91(4):778–794. 115
- Imbert, C. and Papp, J. (2015). Labor market effects of social programs: Evidence from India's employment guarantee. *American Economic Journal: Applied Economics*, 7(2):233–263. 82

- Jahoda, M. (1982). *Employment and Unemployment: A Social-Psychological Analysis*. Cambridge University Press, Cambridge. Borrow-only on Internet Archive: <https://archive.org/details/employmentunempl0000jaho>. See also Jahoda (1981) “Work, Employment, and Unemployment” in *American Psychologist* 36(2) for the core latent-deprivation theory. 155, 215
- Jaroszewicz, A., Hauser, O. P., Jachimowicz, J. M., and Jamison, J. (2022). How effective is (more) money? Randomizing unconditional cash transfer amounts in the US. *SSRN Electronic Journal*. 117, 216
- Ji, J., Liu, M., et al. (2024). AI alignment: A comprehensive survey. *arXiv:2310.19852*. 180
- Jordan, B. (2020). Authoritarianism and social policy. *Social Policy & Society*, 19(2):243–254. Published online 2019. 214, 216
- Kangas, O., Jauhiainen, S., Simanainen, M., and Ylikännö, M. (2020). Evaluation of the Finnish basic income experiment. Reports and Memorandums 2020:15, Ministry of Social Affairs and Health, Helsinki. 109, 110
- Karger, E., Kuusela, O., et al. (2026). Forecasting the economic effects of AI: Predictions from economists, AI experts, and the public. *Federal Reserve Bank of Chicago and Forecasting Research Institute*. <https://static1.squarespace.com/static/635693acf15a3e2a14a56a4a/t/69cbb9d509ada447b6d9013f/1774959061185/forecasting-the-economic-effects-of-ai.pdf>. 9
- Katz, M. B. (2013). *The Undeserving Poor: America's Enduring Confrontation with Poverty*. Oxford University Press, 2nd edition. 144
- Keen, M. and Konrad, K. A. (2013). The theory of international tax competition and coordination. In *Handbook of Public Economics*, volume 5, pages 257–328. Elsevier. 168
- Keynes, J. M. (1930). Economic possibilities for our grandchildren. In *Essays in Persuasion*, pages 321–332. Springer. 11, 14, 20, 88, 141, 191
- Kleven, H. J., Landais, C., Saez, E., and Schultz, E. (2014). Migration and wage effects of taxing top earners: Evidence from the foreigners' tax scheme in Denmark. *The Quarterly Journal of Economics*, 129(1):333–378. 170
- Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., and Dean, R. (2024). AI 2027. AI Futures Project. <https://ai-2027.com/>. 33, 204
- Kokotajlo, D., Larsen, T., Dean, R., and Lifland, E. (2026). AI 2040: Plan A. AI Futures Project. <https://ai-2040.com/>. 175, 204
- Koopman, P. J. (2025). *Embodied AI Safety: Reimagining Safety Engineering for Artificial Intelligence in Physical Systems*. Self-published. 181, 182
- Korinek, A. (2024). Economic policy challenges for the age of AI. NBER Working Paper 32980, National Bureau of Economic Research. 51, 165

- Korinek, A. and Juelfs, M. (2022). Preparing for the (non-existent?) future of work. NBER Working Paper 30172, National Bureau of Economic Research. Published in Oxford Handbook of AI Governance, pp. 746–776, 2024. [2](#), [113](#), [124](#), [131](#), [213](#)
- Korinek, A. and Lockwood, L. M. (2026). Public finance in the age of AI: A primer. Working Paper 34873, National Bureau of Economic Research. [90](#), [93](#), [99](#), [129](#)
- Korinek, A. and Stiglitz, J. E. (2019). Artificial intelligence and its implications for income distribution and unemployment. In *The Economics of Artificial Intelligence: An Agenda*, pages 349–390. University of Chicago Press. [10](#), [11](#), [46](#), [90](#), [108](#), [218](#)
- Korinek, A. and Stiglitz, J. E. (2021). Artificial intelligence, globalization, and strategies for economic development. NBER Working Paper 28453, National Bureau of Economic Research. <https://doi.org/10.3386/w28453>. [132](#), [165](#), [173](#), [174](#), [220](#)
- Korinek, A. and Suh, D. (2024). Scenarios for the transition to AGI. Technical Report 32255, National Bureau of Economic Research. <http://dx.doi.org/10.3386/w32255>. [9](#), [61](#), [66](#), [67](#), [74](#), [165](#), [185](#), [188](#), [205](#), [217](#), [218](#), [220](#), [221](#)
- Krier, S. (2026). The cyborg era: What AI means for jobs. <https://aleximas.substack.com/p/the-cyborg-era-what-ai-means-for>. [64](#), [217](#), [222](#)
- Kuhn, R. L. (2024). A landscape of consciousness: Toward a taxonomy of explanations and implications. *Progress in Biophysics and Molecular Biology*, 190:28–169. [188](#)
- Kurzweil, R. (2024). *The Singularity Is Nearer: When We Merge with AI*. Viking, New York. [184](#)
- Laenen, T. (2023). *The Popularity of Basic Income: Evidence from the Polls*. Palgrave Macmillan. [108](#), [214](#)
- Landemore, H. (2020). *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton University Press. [153](#), [156](#)
- Lanier, J. and Weyl, E. G. (2018). A blueprint for a better digital society. *Harvard Business Review*, 96(5):62–71. [156](#)
- Lansley, S. (2016). *A Sharing Economy: How Social Wealth Funds Can Reduce Inequality and Help Balance the Books*. Policy Press, Bristol. [138](#)
- Lehner, L. (2026). What do unemployed workers want: Guaranteed jobs or guaranteed income? Working paper, University of Edinburgh. Data: <https://data.aussda.at/dataset.xhtml?persistentId=doi:10.11587/XJNNYA>. [82](#), [143](#)
- Leontief, W. (1983a). National perspective: The definition of problems and opportunities. In *The Long-Term Impact of Technology on Employment and Unemployment*, pages 3–7. National Academy Press, Washington, DC. National Academy of Engineering Symposium, June 30, 1983. [1](#), [63](#), [64](#), [219](#)
- Leontief, W. W. (1983b). Technological advance, economic growth, and the distribution of income. *Population and Development Review*, 9(3):403–410. [63](#)

- Levitsky, S. and Ziblatt, D. (2018). *How Democracies Die*. Crown, New York. 151
- MacKay, D. J. (2009). *Sustainable Energy — Without the Hot Air*. UIT Cambridge, Cambridge. 196, 197
- McKinsey & Company (2023). The economic potential of generative AI. *McKinsey Digital*. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>. 17
- Meade, J. E. (1995). *Full Employment Regained? An Agathotopian Dream*. Number 61 in DAE Occasional Papers. Cambridge University Press, Cambridge. 29
- Mecca, D. (2025). Global ad spending to hit \$1 trillion in 2025. *Abbey Mecca Industry Insights*. <https://abbeymecca.com/worldwide-ad-spending-to-hit-1-trillion-in-2025-what-it-means-for-businesses-marketers-and-nonprofits/>. 79
- Meidner, R. (1978). *Employee Investment Funds: An Approach to Collective Capital Formation*. George Allen & Unwin, London. 124, 128, 129, 130, 131
- Ministry of Finance, Government of India (2017). Universal basic income: A conversation with and within the Mahatma. In *Economic Survey 2016-17*, volume 1, chapter 9, pages 172–212. Ministry of Finance, New Delhi. 194
- Mirrlees, J., Adam, S., Besley, T., Blundell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., and Poterba, J. M. (2011). *Tax by Design: The Mirrlees Review*. Oxford University Press. 168
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208. 93
- Mohamed, S., Png, M.-T., and Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4):659–684. 173, 220
- Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence*. Harvard Univ. Press. 54, 55, 218
- Moretti, E. and Wilson, D. J. (2017). The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *American Economic Review*, 107(7):1858–1903. 170
- Morris, M. R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., and Legg, S. (2023). Levels of AGI for operationalizing progress on the path to AGI. *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 1478:36308–36321. 31, 32, 33
- Morse, N. C. and Weiss, R. S. (1955). The function and meaning of work and the job. *American Sociological Review*, 20(2):191–198. 13, 157, 215
- Müller, V. C. (2020). Ethics of artificial intelligence and robotics. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2020 edition. 189

- Muralidharan, K., Niehaus, P., and Sukhtankar, S. (2016). Building state capacity: Evidence from biometric smartcards in India. *American Economic Review*, 106(10):2895–2929. 82
- Musgrave, R. A. (1959). *The Theory of Public Finance: A Study in Public Economy*. McGraw-Hill, New York. 91, 98, 100, 124
- Nayebi, A. (2025). An AI capability threshold for rent-funded universal basic income in an AI-automated economy. *arXiv:2505.18687*. 39, 93
- Newsom, G. (2026). Executive order N-6-26: AI workforce. Executive Department, State of California. <https://www.gov.ca.gov/wp-content/uploads/2026/05/5.21.26-AI-Workforce-EO-FINAL-SIGNED.pdf>. 143
- OECD (2024). Benefits and wages: Net replacement rates in unemployment. OECD Tax-Benefit Models. <https://www.oecd.org/en/data/datasets/benefits-and-wages-net-replacement-rates.html>. Accessed 2026-05-11. 111
- OECD (2026). Macroeconomic productivity gains from artificial intelligence in G7 economies. *OECD Artificial Intelligence Papers*. <https://www.oecd.org/>. 164
- OECD/G20 (2021). Statement on a two-pillar solution to address the tax challenges arising from the digitalisation of the economy. OECD/G20 Base Erosion and Profit Shifting Project. <https://www.oecd.org/tax/beps/statement-on-a-two-pillar-solution-to-address-the-tax-challenges-arising-from-the-digitalisation-of-the-economy-october-2021.htm>. 169
- OpenAI (2026). Industrial policy for the intelligence age. *OpenAI Blog*. <https://openai.com/index/industrial-policy-for-the-intelligence-age/>. 179
- Opatow, S. (1990). Moral exclusion and injustice: An introduction. *Journal of Social Issues*, 46(1):1–20. 145
- Ord, T. (2020). *The Precipice: Existential Risk and the Future of Humanity*. Bloomsbury Publishing, London. 181
- Paine, T. (1797). Agrarian justice. Pamphlet, originally published in Paris (French edition 1795, English edition London 1797). Reprinted in P. Foner (ed.), *Collected Writings*, Library of America, 1995. <https://www.ssa.gov/history/paine4.html>. 47, 107, 211
- Pateman, C. (2006). Democratizing citizenship: Some advantages of a basic income. In Wright, E. O., editor, *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones of a More Egalitarian Capitalism*, pages 101–117. Verso. 125
- Pelligrà, V. and Sacco, P. L. (2023). Searching for meaning in a post-scarcity society. implications for creativity and job design. *Frontiers in Psychology*, 14:1198424. 159, 215
- Pettit, P. (2012). *On the People's Terms: A Republican Theory and Model of Democracy*. The Seeley Lectures. Cambridge University Press. 107, 154, 211
- Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan, London. <https://www.econlib.org/library/NPDBooks/Pigou/pgEW.html>. 97

- Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press. 93, 94, 99, 194, 195
- Piketty, T., Saez, E., and Stantcheva, S. (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy*, 6(1):230–271. 22, 93, 100
- Pinker, S. (2018). *Enlightenment Now: The Case for Reason, Science, Humanism, and Progress*. Viking, an imprint of Penguin Random House LLC, New York, New York. 10, 15, 60
- Posner, E. A. and Weyl, E. G. (2018). *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton University Press. 48, 50, 51, 99, 154, 156, 218
- Pritchett, L. et al. (2024). What would it cost to end extreme poverty? NBER Working Paper 34583, National Bureau of Economic Research. 79, 118
- Przeworski, A. (2019). *Crises of Democracy*. Cambridge University Press. 151
- Rajan, R. G. and Subramanian, A. (2011). Aid, Dutch disease, and manufacturing growth. *Journal of Development Economics*, 94(1):106–118. 167
- Ramey, V. A. (2009). Time spent in home production in the twentieth-century United States: New estimates from old data. *The Journal of Economic History*, 69(1):1–47. 19
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA. Revised edition 1999. 107, 154, 212
- Raymond, E. S. (2015). The zero marginal cost society – a review. *Armed and Dangerous*. <http://esr.ibiblio.org/?p=5558>. 43
- Reich, R. (2018). *Just Giving: Why Philanthropy Is Failing Democracy and How It Can Do Better*. Princeton University Press. 101
- Restrepo, P. (2025). We won't be missed: Work and growth in the AGI world. NBER Working Paper 34423, National Bureau of Economic Research. 74, 75, 204, 213, 221, 222
- Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. Palgrave Macmillan. 29, 42, 44, 213, 218
- Rodrik, D. (2011). *The Globalization Paradox: Democracy and the Future of the World Economy*. W. W. Norton & Company. 163, 175, 195
- Rodrik, D. (2015). *Economics Rules: The Rights and Wrongs of the Dismal Science*. W. W. Norton & Company. 59
- Rodrik, D. (2016). Premature deindustrialization. *Journal of Economic Growth*, 21(1):1–33. 173, 174
- Rojas Burbano, J. C., Idárraga Ruiz, N., Rojas Burbano, G. A., Guacho Inca, J. S., Arragan Lezama, C. A., and Sánchez González, M. (2025). Robot-assisted surgery: Current applications and future trends in general surgery. *Cureus*, 17(4):e82318. 36

- Ross, M. L. (2001). Does oil hinder democracy? *World politics*, 53(3):325–361. 132, 146, 147, 149
- Russell, S. (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking. 180, 181, 182
- Saeidi, H., Opfermann, J. D., Kam, M., Wei, S., Leonard, S., Hsieh, M. H., Kang, J. U., and Krieger, A. (2022). Autonomous robotic laparoscopic surgery for intestinal anastomosis. *Science Robotics*, 7(62):eabj2908. 36
- Sahlins, M. (1972). *Stone Age Economics*. Aldine-Atherton, Chicago. 62, 158
- Saint-Paul, G. (2025). Artificial intelligence, the collapse of consumer society, and oligarchy. IZA Discussion Paper 17682, Institute of Labor Economics (IZA). <https://www.iza.org/publications/dp/17682/>. 146, 148, 213
- Samuelson, P. A. (1989). Ricardo was right! *The Scandinavian Journal of Economics*, 91(1):47–62. 63, 221, 222
- Sanders, B. (2026). The public should own half of the big A.I. companies. *The New York Times*. Guest Essay. <https://www.nytimes.com/2026/06/01/opinion/artificial-intelligence-bernie-sanders.html>. 125, 128, 129, 131
- Sastry, G., Heim, L., Belfield, H., Anderljung, M., Brundage, M., Hazell, J., O’Keefe, C., Hadfield, G. K., et al. (2024). Computing power and the governance of artificial intelligence. *arXiv:2402.08797*. 154, 156, 175, 181, 182
- Scheidel, W. (2017). *The Great Leveler: Violence and the History of Inequality from the Stone Age to the Twenty-First Century*. Princeton University Press, Princeton. 152, 206
- Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press. 175
- Scott, J. C. (2017). *Against the Grain: A Deep History of the Earliest States*. Yale University Press, New Haven. 62, 158
- Sherraden, M. (1991). *Assets and the Poor: A New American Welfare Policy*. M. E. Sharpe, Armonk, NY. 133, 134
- Silver, L., van Kessel, P., Huang, C., Clancy, L., and Gubbala, S. (2021). What makes life meaningful? Views from 17 advanced economies. Report, Pew Research Center. 157, 159, 215
- Skare, M. and Skare, D. (2017). Is the great decoupling real? *Journal of Business Economics and Management*, 18(3):451–467. 15, 16, 60, 218, 219
- Soares, N., Fallenstein, B., Yudkowsky, E., and Armstrong, S. (2015). Corrigibility. Technical report, Machine Intelligence Research Institute. AAAI 2015 Workshop on AI and Ethics. 154, 156, 182
- Soffia, M., Wood, A. J., and Burchell, B. (2022). Alienation is not ‘bullshit’: An empirical critique of Graeber’s theory of BS jobs. *Work, Employment and Society*, 36(5):816–840. 76

- Solomonoff, R. J. (1985). The time scale of artificial intelligence: Reflections on social effects. *Human Systems Management*, 5:149–153. 186
- Srinivasan, B. (2022). *The Network State: How to Start a New Country*. Self-published. 156
- Srinivasan, P. (2025). Marshall Islands launches world's first universal basic income scheme offering cryptocurrency. *The Guardian*. <https://www.theguardian.com/world/2025/dec/17/marshall-islands-launches-universal-basic-income-scheme-offering-cryptocurrency-in-world-first>. 109, 214
- Srnicek, N. and Williams, A. (2015). *Inventing the Future: Postcapitalism and a World Without Work*. Verso, London. 2, 73, 199
- Standing, G. (2017). *Basic Income: And How We Can Make It Happen*. Pelican Books, London. 23, 67, 82, 111, 112, 144, 199, 211, 213
- Stanford HAI (2026). Artificial intelligence index report 2026. Technical report, Stanford University, Human-Centered AI Institute. <https://aiindex.stanford.edu/report/>. 18, 33, 218, 219
- Stephenson, N. (1995). *The Diamond Age: Or, A Young Lady's Illustrated Primer*. Bantam Spectra, New York. Hugo Award for Best Novel, 1996. 61
- Stiefenhofer, P. (2025). Techno-feudalism and the rise of AGI: A future without economic rights? *arXiv preprint arXiv:2503.14283*. 146, 216
- Stigler, G. J. (1971). The theory of economic regulation. *The Bell Journal of Economics and Management Science*, 2(1):3–21. 151
- Stiglitz, J. E. (1977). The theory of local public goods. *The Economics of Public Services*, pages 274–333. 47, 99, 170, 171, 218
- Stiglitz, J. E. (2008). Economic foundations of intellectual property rights. *Duke Law Journal*, 57(6):1693–1724. 48
- Super Members Council of Australia (2026). Unfinished business: Fixing gaps in the super guarantee. Technical report, Super Members Council of Australia. 135, 138
- Susskind, D. (2020). *A World Without Work: Technology, Automation, and How We Should Respond*. Metropolitan Books. 2, 7, 10, 13, 14, 46, 64, 66, 71, 73, 111, 120, 121, 123, 143, 157, 203, 217
- Susskind, D. (2025). What will remain for people to do? *Knight First Amendment Institute, Columbia University*. <https://knightcolumbia.org/content/what-will-remain-for-people-to-do>. 71, 73, 74, 217, 218, 220, 221, 222
- Tcherneva, P. R. (2020). *The Case for a Job Guarantee*. Polity Press, Cambridge. 82, 90, 120, 126
- Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf. 39, 180, 181
- The Alliance for Secure AI (2026). AI job loss tracker. <http://jobloss.ai/>. 17, 219

- Thompson, E. P. (1963). *The Making of the English Working Class*. Victor Gollancz, London. 17, 191
- Thomson, R. M., Kopasker, D., Bronka, P., Richiardi, M., Khodygo, V., Baxter, A. J., Igelström, E., Pearce, A., Leyland, A. H., and Katikireddi, S. V. (2024). Short-term impacts of universal basic income on population mental health inequalities in the UK: A microsimulation modelling study. *PLoS Medicine*, 21(3):e1004358. 116, 117, 216
- Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424. 175
- Tinbergen, J. (1952). *On the Theory of Economic Policy*. North-Holland, Amsterdam. 124
- Tobin, J. (1978). A proposal for international monetary reform. *Eastern Economic Journal*, 4(3–4):153–159. 101
- Tørsløv, T. R., Wier, L. S., and Zucman, G. (2023). The missing profits of nations. *Review of Economic Studies*, 90(3):1499–1534. 169, 170
- Tullock, G. (1967). The welfare costs of tariffs, monopolies, and theft. *Economic Inquiry*, 5(3):224–232. 100
- Tyack, A. and Mekler, E. D. (2024). The basic needs in games (BANG) model. *Interacting with Computers*, 36(3). 159, 215
- Tyssedal, J. J. (2025). What is work? Engineering a working definition. *Critical Review of International Social and Political Philosophy*. Online first. 12, 14, 155
- Van Parijs, P. (2004). Basic income: A simple and powerful idea for the twenty-first century. *Politics & Society*, 32(1):7–39. 23, 89, 115, 123, 213
- Van Parijs, P. and Vanderborght, Y. (2017). *Basic Income: A Radical Proposal for a Free Society and a Sane Economy*. Harvard University Press, Cambridge, MA. 7, 23, 25, 68, 91, 93, 100, 102, 105, 107, 111, 112, 114, 126, 127, 171, 172, 199, 211
- Vanek, J. (1974). Time spent in housework. *Scientific American*, 231(5):116–121. 19
- Varoufakis, Y. (2023). *Technofeudalism: What Killed Capitalism*. Bodley Head. 151
- Vollrath, D. (2015). Beating a dead robotic horse. *The Growth Economics Blog*. <https://growthecon.wordpress.com/2015/12/31/beating-a-dead-robotic-horse/>. 1, 67, 221
- von Neumann, J. (1966). *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana. Edited and completed by Arthur W. Burks. 44
- Walo, S. (2023). ‘bullshit’ after all? Why people consider their jobs socially useless. *Work, Employment and Society*, 37(5):1123–1146. 76, 81, 213, 220
- West, S. and Castro, A. (2023). Impact of guaranteed income on health, finances, and agency: Findings from the Stockton randomized controlled trial. *Journal of Urban Health*, 100(2):227–244. 109, 110

- White, S. (2006). The citizen's stake and paternalism. In *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones of a More Egalitarian Capitalism*. Verso. 125
- Widerquist, K. (2017). The cost of basic income: Back-of-the-envelope calculations. *Basic Income Studies*, 12(2). 112, 117, 212, 216
- Widerquist, K. (2024). *Universal Basic Income*. MIT Press Essential Knowledge. The MIT Press. 7, 62, 112
- Widerquist, K. and Howard, M. W. (2012). *Exporting the Alaska Model: Adapting the Permanent Fund Dividend for Reform Around the World*. Palgrave Macmillan. 113, 126, 132, 156
- Wikipedia contributors (2026). Universal basic income. Wikipedia, The Free Encyclopedia. Accessed 8 June 2026. Revision 1357149627. 110
- Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell, Oxford. 12
- Wright, E. O. (2006). Basic income, stakeholder grants, and class analysis. In Wright, E. O., editor, *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones of a More Egalitarian Capitalism*, pages 79–88. Verso. 125
- Young, C. and Varner, C. (2011). Millionaire migration and state taxation of top incomes: Evidence from a natural experiment. *National Tax Journal*, 64(2):255–284. 97, 170
- Yudkowsky, E. and Soares, N. (2025). *If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All*. Little, Brown and Company. 182
- Zucman, G. (2014). Taxing across borders: Tracking personal wealth and corporate profits. *Journal of Economic Perspectives*, 28(4):121–148. 169, 170

Index

A

Acemoglu, Daron .. **11**, **62**, **62**, **191**, **192**
 Ackerman, Bruce **126**
 adaptive regulatory architecture .. **113**
 Aghion, Philippe **185**, **188**
 AGI .. *see* artificial general intelligence
 AI alignment **6**, **200**
 AI colonialism *see also* Global South, *see also* premature
 deindustrialization, **6**, **173**, **177**,
 206
 AI governance . **153**, **154**, **155**, **167**, **173**,
 189
 AI safety **154**, **175**, **189**
 Alaska **113**, **120**, **129**, **132**, **154**, **194**, **195**,
 214
 Alaska Permanent Fund . **113**, **132**, **154**,
 194, **214**
 Alexander, Scott **75**
 alienation **76**
 alignment tax **181**, **181**
 Altman, Sam **43**, **108**, **130**, **150**, **214**, **219**
 annoyance economy **75**, **76**, **80**, **207**, **213**
 Arendt, Hannah **15**
 arms race .. **vi**, **vii**, **53**, **75**, **78**, **78**, **79**, **80**,
 97, **100**, **142**, **145**, **207**, **213**, **221**
 artificial general intelligence .. **ii**, **v**, **vi**,
 vii, **1–3**, **7**, **9**, **10**, **12**, **18**, **19**, **22**, **30**,
 32, **33**, **38**, **42**, **49**, **60**, **63**, **64**, **87**,
 89–91, **96–98**, **115**, **146**, **149**, **152**,
 166, **172**, **175**, **177**, **179–181**,
 203–205, **217**, **219**
 artificial superintelligence **31**, **34**, **34**, **35**,

37, **38**, **68–72**, **160**, **161**, **182**, **183**,
200, **209**

artisanal niche **73**
 ASI *see* artificial superintelligence
 asset-based welfare **133**, **134**
 Atkinson, Anthony **67**, **93**, **95**, **121**, **126**
 Australia **3**, **5**, **47**, **50**, **133**, **137**, **148**, **164**,
 166, **166**, **168**, **214**
 automated corporation **64**, **217**
 automation ... *see also* artificial general
 intelligence, *see also*
 displacement
 Autor, David **10**, **13**

B

Balassa-Samuelson effect **167**
 Bangladesh **173**, **173**, **174**
 basic endowment **120**, **124**
 basic income *see also* guaranteed
 income, *see also* negative income
 tax, *see also* redistribution, **v**, **6**,
 23, **25**, **30**, **39**, **51**, **82**, **87**, **101**, **105**,
 109–112, **114**, **120**, **160**, **172**, **194**,
 215, **218**
 universal .. **vi**, **vii**, **vii**, **viii**, **2–4**, **6**, **9**,
 11, **12**, **14**, **21**, **30**, **39**, **40**, **42**, **48**,
 55, **62**, **68**, **71**, **72**, **74**, **87**, **88**, **105**,
 110, **111**, **114**, **133**, **137**, **146**, **150**,
 154, **157**, **164**, **169–172**, **181**,
 185–188, **190**, **203**, **205**, **206**,
 211–213
 universal adequate .. **112**, **112**, **115**,
 139
 universal high . **25**, **25**, **112**, **115**, **139**

Baumol's cost disease *see also*
 productivity, 186, 186, 200
 Baumol, William 78, 186, 186, 200
 BIEN 23, 23, 127
 birth endowment 134
 boiling frog 55, 144, 204, 207, 209
 Boix, Carles 145, 146
 Bostrom, Nick 159, 180, 181
 Botswana 148
 BPO .*see* business process outsourcing
 brain drain 214
 Bretton Woods 153
 Brynjolfsson, Erik 15, 67, 217
 BS-job *see* bullshit job
 bullshit job *see also* annoyance economy,
see also full employment, v, vi,
 vii, 3, 4, 6, 13, 14, 60, 75, 76, 76,
 77, 87, 90, 100, 103, 113, 120, 142,
 203, 206–208, 216, 220
 business process outsourcing 165, 174

C

capital flight 95, 100, 206
 care work 126, 211
 Cerutti, Eugenio 165
 China . 143, 148, 148, 149, 174, 186, 194
 citizens' wealth fund ... v, 89, 124, 124,
 128, 129, 129, 133, 213, 217
 civil unrest vii, 142, 150, 214
 cluster concept 3, 3, 220
 cognitive flexibility 1, 2, 10, 60, 217, 222
 cognitive inertia 207
 Comin, Diego 67, 73, 74
 comparative advantage .. 165, 173, 174,
 177, 190, 220
 compute diplomacy 167
 compute governance 181
 conditional cash transfer ... 21, 23, 119,
 120
 conditionality vi, 25, 109, 111, 116, 120,
 121, 135, 136, 167, 214
 consciousness, machine . viii, 188, 189,
 189
 content pollution 79
 Continuity (camp) 2, 18
 coordination failure 6, 75, 203, 207

copilot model 222
 creative destruction 1, 1
 Cunningham, Scott 74
 currency devaluation 164, 166, 177
 CWF *see* citizens' wealth fund

D

Danaher, John 159
 data as labor 50
 Davidson, Tom 185, 188
 de Vries, Jan 64, 196, 199
 deadweight loss 47, 49, 99
 degrowth 198, 199
 demand insufficiency 60
 democratic feedback loop .. 5, 141, 144
 Denmark 93, 111
 dignity 55, 83, 126
 displacement ..*see also* automation, *see*
also retraining, *see also*
 unemployment
 Displacement (camp) 11, 18
 displacement forecast 9, 15, 17, 18

E

economic rent 48, 52, 99
 eligibility 25, 108, 171, 199
 embodiment gap 183
 employment
 full . v, vi, 3, 4, 60, 64, 67, 68, 72, 73,
 75, 106, 206–208, 220–222
 Engels' Pause 144, 153
 externality 76, 97, 100, 100
 extractive institutions 151, 152

F

family resemblance 12
 feudalism
 managerial 146, 149
 techno 192, 216
 Finland 109, 194
 flexicurity 111
 four camps 9, 25
 Frank, Robert 97
 Frey, Carl Benedikt 11, 219
 Friedman, Milton 108, 123, 126, 192, 214
 full employment 137

G

- Gabaix, Xavier80
 GDP *see* gross domestic product
 GDP-B **112**, 187
 Gentilini, Ugo **110**, 117, 211
 George, Henry 47, **47**, 107, 171
 Germany ... **93**, 111, **114**, 119, 174, 195,
 214, 219
 GI *see* guaranteed income
 gig economy . *see also* labor market, *see*
 also precariat, **vii**, 108, 207, 212
 Gini coefficient 91, **91**
 Global South **173**
 globalization **15**, 169, 176, 206
 Goodhart's law **180**
 Graeber, David .. **75**, 76, **76**, 77, 80, 217
 Gray, Mary L. **62**
 Great Decoupling **3**, 9, 15, 16, 221
 gross domestic product 42, **42**, 102, 106,
 111, 112, **112**, 113, 148, 149, 166,
 167, 185–187, 193, 218
 growth theory **185**, 186, 188
 guaranteed income **v**, **vi**, **vi**, 2, 7, 10, 21,
 23, 25, 51, 52, 62, 89, 91, 106, 108,
 110, 112–116, 136, 149, 205, 207

H

- hand-loom weavers **15**
 Hanson, Robin **185**
 Harari, Yuval Noah **144**
 Hayek, Friedrich **192**
 Highway Trust Fund 91
 horse analogy **64**, **219**
 human oversight **32**, 180, 181
 hyperbolic growth **185**, 186, 204

I

- idealized transition scenario . **3**, 20, 25,
 30, 39, 41, 51, **56**, 111, 117, 118,
 142, 212, 216
 Imas, Alex **63**, 67, **73**, 74, 129, 221
 immigration **25**, 171, 172
 income support
 conditional **21**, **23**, 25, 107, 112, 114,
 160

- India **80**, 82, 152, 165, 174, **174**, 194, 195
 industrial revolution 1, **1**, 15, **15**, 18, 19,
 62, 144, 153, 191, 206
 inequality . **vii**, 6, 11, **22**, 62, 91, 93, 100,
 102, 103, **103**, 151, 152, 164, 187,
 194, 195, 205, 206, 215
 inflation ... 102, **102**, **105**, 118, **118**, 166
 institutional inertia **175**, 204, **205**
 institutional trust 138
 intellectual property . **21**, 30, 48, 49, **49**,
 95, 99, 103, 168, 169, 176, 213, 218
 intelligence curse *see also* redistribution,
 see also resource curse, **132**
 interface friction **222**
 Ireland **109**, 112, 153, 160, 169

J

- Jahoda, Marie **157**
 Jevons paradox **42**, 196, 200
 JG *see* job guarantee
 job definition
 11-property **25**
 job guarantee *see also* basic income, *see*
 also workfare, **vi**, 67, 81, **81**, 82,
 82, 83, 90, 100, 120, 139
 job-leisure spectrum **3**, 14

K

- Kenya **80**, 101, **109**, 115
 Keynes, John Maynard **20**, **191**
 Korinek, Anton 2, 74, 90, 108, **113**, 185,
 188, 205, 217

L

- labor market ... **40**, 54, **54**, 83, 114, 124,
 144, 151, 159, 190
 labor share . **11**, 16, **16**, 62, 95, 112, 213,
 218
 large language model .. 17, 31, 80, 183,
 188, 218, 219
 leisure society **159**, 204
 Leontief, Wassily **63**, 64, **219**
 lights-out factory **43**
 LLM *see* large language model
 LVT *see* tax, land value

M

Malibu surfer216
 marginal cost **3**, **29**, **42**, **43**, **43**, **44**, **49**, **56**,
 114, **183**, **187**, **218**
 zero **42**, **43**, **43**, **44**, **49**, **114**, **187**, **218**
 Marshall Islands **109**, **214**
 Marx, Karl**191**
 meaning (of work) . *see also* dignity, *see*
 also leisure society, **159**
 means-testing **vi**, **106**, **107**, **108**, **112**, **116**,
 119, **122**, **126**, **211**, **212**
 MGNREGA .. *see* job guarantee, **82**, **83**
 Minsky, Hyman **191**
 Mohamed, Shakir **173**
 Moloch trap **75**
 Mongolia**112**, **148**, **149**
 moral hazard **98**
 Moravec's paradox **54**, **183**
 Moravec, Hans **54**, **55**, **183**
 Morris, Meredith Ringel **32**

N

narrow corridor **192**
 negative income tax .. **25**, **108**, **112**, **119**,
 120, **120**, **192**, **214**
 Nigeria**147**
 NIT *see* negative income tax
 non-affectation principle **91**, **92**
 non-domination . *see also* basic income,
 see also dignity, **107**, **126**
 Norway**129**, **148**, **149**, **154**

O

Ontario (pilot) **109**
 open-source **12**, **12**, **45**, **50**, **50**, **160**, **177**,
 184, **215**
 Ord, Toby **181**

P

Paine, Thomas **47**, **107**
 participation income **119**, **121**
 path dependence **135**
 pension fund **105**, **133**
 Pettit, Philip **107**, **154**
 Philippines **80**, **165**, **174**

Piketty, Thomas .. **22**, **93**, **100**, **194**, **195**
 political trilemma **195**
 populism **141**, **142**, **144**, **146**, **214**
 positional good **53**, **66**, **97**, **142**
 post-labor economy **v**, **viii**, **viii**, **23**, **25**,
 98, **106**, **115**, **116**, **141**, **155**, **175**,
 187, **198**, **199**, **200**, **206**, **216**
 Post-Labor Prosperity **3**, **7**, **25**
 post-scarcity **30**, **47**, **55**, **97**, **99**, **114**, **181**,
 200, **205**, **208**
 poverty trap **126**
 pre-distribution **90**, **109**, **121**
 precariat **144**
 premature deindustrialization **173**, **174**
 productivity **v**, **3**, **9**, **15**, **15**, **16**, **17**, **41**, **52**,
 63, **64**, **90**, **106**, **111**, **112**, **123**, **143**,
 158, **184**, **186**, **190**, **208**, **218**
 profit shifting .. **vii**, **101**, **168**, **169**, **169**
 Protestant work ethic **157**
 provenance premium **66**, **71**, **72**, **74**, **221**,
 222
 purchasing power **63**, **89**, **102**, **108**, **118**,
 118, **119**, **148**, **191**, **212**

Q

quadratic voting **154**

R

radical abundance **188**, **203**
 ratchet clause **145**
 Rawls, John **107**, **154**
 reciprocity **115**
 recursive self-improvement . **32**, **34**, **54**
 Red Queen game **78**, **78**
 redistribution . *see also* basic income, *see*
 also tax, **vii**, **viii**, **viii**, **1**, **2**, **4**, **5**, **10**,
 11, **22**, **23**, **26**, **39**, **47**, **55**, **56**, **63**,
 73, **87–91**, **94**, **131**, **133**, **141**, **142**,
 144, **163**, **171**, **181**, **185**, **186**, **188**,
 203, **205**, **206**, **208**, **213**, **214**, **216**
 regulatory capture .. *see also* extractive
 institutions, *see also* vetocracy,
 141, **142**, **151**, **153**
 relational sector **74**, **75**
 rent extraction **100**
 rentier state **195**

reservation wage **74, 212**
 reshoring *see also* business process
 outsourcing, *see also* premature
 deindustrialization, **173, 174, 174**
 resource curse **132, 148, 149**
 retirement **133**
 retraining .. **vi, 54, 88, 88, 108, 160, 204,**
 212, 215, 220
 Rifkin, Jeremy **42, 44**
 robotics **v, 1, 9, 17, 29, 33, 34, 36, 36, 55,**
 63, 75, 87, 164, 165, 182–184, 187,
 206, 208, 217, 219, 220
 humanoid **218**
 Rodrik, Dani **174, 195, 195**
 Russell, Stuart **180, 181**
 Russia **147, 152**

S

S-curve **56, 64, 65, 220**
 Saint-Paul, Gilles **147, 148**
 Samuelson, Paul **63, 221, 221**
 Sanders, Nada **128, 129**
 Scheidel, Walter **206**
 seigniorage **102, 102**
 self-replicating factory .. **30, 43, 44, 197**
 Singapore **48**
 skill mismatch **204**
 social contract **20, 141, 145, 150**
 Social Security **91**
 Solidaritätszuschlag **91**
 sortition **153, 153, 154, 206**
 sovereign wealth fund *see also* citizens'
 wealth fund, *see also*
 redistribution, **48, 121, 124, 124,**
 129, 129, 133, 137, 148, 154, 213,
 217
 Srnicek, Nick **2**
 Stagnation (camp) **2**
 stakeblowing **125, 128, 216**
 Standing, Guy .. **34, 47, 53, 82, 144, 211**
 Stiglitz, Joseph **47**
 Stockton (SEED pilot) **109**
 Sub-Saharan Africa **165, 174**
 subsistence wage **63**
 sudden stop **167**
 sunset clause **113**

superannuation **105, 133**
 surveillance .. **vi, 107, 121, 126, 126, 151,**
 189, 211
 Susskind, Daniel **2, 7, 13, 64, 66, 71, 73,**
 121, 123, 157, 217, 221
 SWF *see* sovereign wealth fund

T

Taiwan **154**
 task model **62**
 tax *see also* fiscal feasibility, *see also*
 redistribution
 automation .. **40, 40, 52, 90, 98, 103,**
 128, 213
 capital gains **41**
 carbon **213**
 destination-based **95, 168, 175**
 digital services .. **vii, vii, 5, 94, 101,**
 168, 168, 176
 Harberger **vi, 4, 56, 218**
 land value **vi, 4, 47, 48, 56, 56, 95, 96,**
 99, 170, 213, 218
 Pigouvian . **vi, 4, 45, 48, 52, 97, 100,**
 100, 103, 198
 robot **96**
 sin **98, 98**
 Tobin **100, 101**
 value-added . **40, 52, 52, 94, 97, 101,**
 168, 176
 wealth **95, 213**
 Tegmark, Max **39, 180, 181**
 Thompson, E.P. **17, 191**
 time allocation **19, 19**
 Tinbergen rule **137, 154**
 Tobin, James **96, 101**
 too-good problem **70, 71, 74, 221**
 trade imbalance .. **5, 163, 164, 166–168,**
 175, 206
 transition roadmap **vii, 3, 139, 207, 213**
 transitional job **61, 61**
 Tyssedal, Jonas **14**

U

UAI *see* basic income, universal
 adequate
 UBC *see* universal basic capital

UBI *see* basic income, universal
 UBS *see* universal basic services
 UHI .. *see* basic income, universal high
 underemployment **75, 112, 221**
 unemployment .. *see also* basic income,
 see also displacement, **15, 39, 40,**
 41, 82, 83, 87, 88, 90, 106, 111,
 113, 113, 114, 166, 177, 181, 191,
 194, 213, 214
 universal basic capital **124, 124**
 universal basic services ... **vi, 119, 119,**
 121, 122, 122, 171, 214
 universality **105, 119, 139**
 unpaid work **14, 69, 157, 159, 215**
 useless class **144**

V

Van Parijs, Philippe . **7, 25, 91, 112, 114,**
 115, 126, 171, 211
 VAT *see* tax, value-added

Venezuela **147**
 veto points **152**
 vetocracy **152**
 Vietnam **148, 174**
 Vollrath, Dietrich **67**
 voluntary work **2, 19, 53, 68, 71, 105, 139**

W

wage erosion **175**
 Walo, Simon **76, 81**
 welfare state . **93, 98, 123, 153, 206, 216**
 welfare tourism **171**
 Widerquist, Karl **117**
 winner-take-all **73**
 workfare **67, 82, 120**
 working hours .. **19, 19, 63, 65, 73, 158,**
 219

Z

Zucman, Gabriel **169, 170**

Job-Less Utopia explores macroeconomics in the age of AGI. As artificial general intelligence (AGI) and advanced robotics decouple economic productivity from human labor, mass un(der)employment is highly probable. However, if the total volume of goods and services is maintained or increased by automation, a post-labor economy is consistent with universal material prosperity. This book explores the economics of an automated society, demonstrating that maintaining full employment in such an economy often leads to hyper-competition and socially wasteful (BS) jobs. This book argues that a transition to Universal Basic Income (UBI) or structurally adapted welfare is not only economically viable but highly desirable. It addresses the practicalities of redistribution, economic objections, and the societal shifts required to find human meaning and maintain democratic stability in a world where human labor is no longer the engine of the economy.

Marcus Hutter has been Professor at the Australian National University for most of his 30+ year career. His research spanned physics, mathematics, computer science, philosophy, artificial intelligence (200+ publications and numerous awards and patents), and more recently AGI's economic, political, and social implications. His two previous books on "Universal Artificial Intelligence" (2005/2024) develop the first sound and complete theory of super-intelligent machines (ASI).

ISBN 978-1-0668284-1-8

