

The Transition Paradox

by Enzo Duit

“While we delay, life passes.” — Seneca

Musk Is Right. Sort Of.

Let's start with the model that everyone is using as the endpoint, because it is actually coherent. Let's trace what it actually implies — as a thought experiment, following the logic wherever it leads.

Elon Musk has described a future where AI drives the cost of almost everything toward zero. Robots make the goods. Agents provide the services. The economy produces abundance with minimal human labour.

In that world, a strange problem emerges: too much output, not enough money chasing it. Prices collapse into deflation. And deflation is dangerous — it causes people to defer spending, companies to stop investing, economies to contract.

To understand why printing money is the mathematically correct response, you need one equation. Economics has a fundamental identity: **Money Supply × Velocity = Price Level × Output**. In plain terms: the total amount of money moving through an economy equals the total value of goods and services produced. If AI drives output through the roof while the money supply stays fixed, prices have to crash. Dramatically. The only way to keep the equation balanced — to prevent economic collapse under the weight of its own abundance — is to expand the money supply to match.

In Musk's model, governments issue checks not to stimulate demand but to prevent deflation. You are keeping the ratio stable as the denominator explodes — not inflating the currency, but matching the money supply to the exploding denominator.

On paper, this is elegant. It is not reckless monetary policy but rather physics.

A natural objection: if costs approach zero, why would producers create abundance rather than simply protect margins? The answer is competitive dynamics. When marginal cost collapses, pricing power shifts to whoever undercuts the competition — margins compress, volume expands. This is already visible in AI inference: as token costs fall, consumption explodes, not because demand was patiently waiting, but because price was the ceiling all along. This essay accepts Musk's model on its


own terms. The abundance assumption is not added on top of the zero-marginal-cost premise — it is a consequence of it.

The problem is the gap between here and there.

The Gap Nobody Is Timing Correctly

Software displaces labour fast. Brutally fast. A team of lawyers replaced by an AI system in a single deployment cycle. A customer support department of 400 becoming two engineers and a chatbot inside a year. This is not speculation — it is documented. In July 2023, Dukaan, an Indian e-commerce platform, eliminated 90% of its customer service staff after deploying an AI chatbot. Klarna’s AI assistant, launched in February 2024, handled the equivalent workload of 700 full-time customer service agents within its first month. [Dukaan CEO Suumit Shah, CNN Business, July 2023; Klarna press release, February 2024] This is happening now, and it will accelerate.

Physical infrastructure is different. Building the energy grids, housing stock, and supply chains that a post-scarcity world requires takes decades. Nuclear plants: ten to fifteen years from approval to power. New mines for lithium and copper: similar timelines. Continental grid upgrades: generational. You cannot deploy those with a software push.



So here is what the transition actually looks like — as a thought experiment to make the gap concrete, not a calendar prediction:

Year 3: millions of workers displaced by software. Year 15: the robots and the energy abundance arrive. In the twelve years between those two dates, Musk’s model says — print money and distribute checks.

But the checks land in an economy where housing, energy, and land are still scarce. Still owned. And being bought up faster than ever by those positioned to move

capital quickly. The capturers of AI productivity gains and the owners of scarce housing are not always the same people — and through the financial system, they do not need to be. AI profits become investable capital; investable capital seeks real assets; real asset demand drives prices up before the money reaches anyone else. The check arrives. The landlord raises the rent.

This is the Transition Paradox: the tool designed to bridge the gap — cash, in any form: universal basic income, direct transfers, stimulus checks — actively widens it. Not because of bad intentions. Because of a mechanism. And that mechanism has a name — one that has been understood for nearly 300 years.

Why the Obvious Solution Fails

In 1730, an Irish-French economist named Richard Cantillon noticed something about how money actually moves through an economy. It does not arrive uniformly. It enters at specific points — whoever is closest to the source of new money — and travels outward from there.

Newly created money enters through central banks, financial institutions, large corporations, government contractors. Those closest to the source spend it before the wider market has adjusted its prices. They buy real assets — land, property, infrastructure — at yesterday's prices. By the time the money reaches ordinary people, prices have already moved to absorb it.

This is the Cantillon Effect. And it is not a historical curiosity.

In an AI transition funded by money printing, the closest nodes to the new money are not displaced factory workers or unemployed logistics drivers. They are the companies building and deploying AI, the investors funding them, and the asset owners who can move capital faster than anyone else. By the time a government check reaches a 52-year-old former warehouse worker in Leipzig, the housing market has already absorbed the money that was printed to help him.

The tool most people reach for — universal basic income, or any direct cash transfer to displaced workers — is not broken because the amounts are wrong or the timing is off. It is broken because cash, in an economy where the scarce assets are already owned, reliably flows to the owners of those assets. Every time. Through a mechanism that has not changed since Cantillon first described it.

This is the first constraint any real solution has to meet: it cannot be cash.

A reasonable objection: does this mechanism hold in highly regulated European economies? The ECB's quantitative easing programmes — and their measurable effect on Frankfurt and Munich real estate while median wages stagnated —

suggest it does. Regulation modulates the speed and entry points of new money. It does not reverse the direction of flow. The Cantillon Effect operates wherever new money enters unevenly, which is everywhere it has ever been tried.

Keep that in mind. We will return to it.

Not All Jobs Are the Same

Here is where the story becomes personal — and where the people reading this essay have more skin in the game than they probably realise.

Not all work faces the same risk from AI. The difference comes down to a simple economic concept: how does demand respond when costs fall?

If the cost of making bread drops dramatically, people do not eat ten times more bread. Demand is fixed by biology. When automation speeds up, most of these jobs just disappear; the market does not expand to replace them. Economists call this inelastic demand.

If the cost of software development drops dramatically, something completely different happens. Projects that were previously impossible — too expensive, too slow, too complex — suddenly become viable. Demand was never capped by appetite. It was capped by cost. Remove the cost barrier and demand expands without a ceiling. This is elastic demand. It follows a counterintuitive pattern first described by 19th-century economist William Stanley Jevons. Jevons observed that efficiency improvements in coal use led to more total coal consumption, not less — because cheaper production made previously unviable uses suddenly viable. Applied to software: when development costs drop, total demand for development does not stay flat. Projects that were too expensive to commission become viable. Markets too small to serve become reachable. The ceiling on demand was never appetite — it was price.

Roughly a third to close to half of developed economies sit in inelastic territory — basic food, housing, energy, primary transport. [OECD Service Economy in OECD Countries; Eurostat employment by sector, 2023] When AI displaces these workers, the market does not grow to absorb them.

IBM offers a concrete illustration of how elastic sectors behave differently. In May 2023, Arvind Krishna said he expected roughly 7,800 IBM administrative positions to be automated away within five years. His position on software developers: “I don’t intend to get rid of a single one. I’ll get more.” By June 2026, IBM’s software developers were 40% more productive than two years earlier — and the workforce had grown, because higher output per developer expands what is commercially achievable. Inelastic roles shrank. Elastic roles expanded. [Arvind Krishna,

Bloomberg, 1 May 2023; Masters of Scale, “IBM’s \$10 billion bet on what comes after AI,” 18 June 2026]

Inelastic work splits into two very different fates. Back-office and administrative work — bookkeeping, data entry, first-line customer support, claims processing — is fully automatable on short timelines. Organisations move more slowly than the technology allows, but the direction does not reverse.

The physically-bound trades are different. Electricians, plumbers, HVAC installers, construction crews. A general-purpose robot that can safely rewire a house or fit a boiler is a much harder, much later problem than a chatbot that processes claims. For years yet, this work stays stubbornly human, stubbornly local — and because the entire physical transition depends on it, these workers are both irreplaceable and routinely underpriced. Structural indispensability only translates into wages when workers have bargaining power — which requires visibility, mobility, and political representation. Trades workers have none of these in sufficient measure. A society that lets their wages stagnate while knowledge workers multiply their output and their freedom is quietly asking the least mobile people to carry the heaviest, least optional part of the transition — for less, not more, reward.

The protection from displacement is two-layered. The first is the one most often cited: general-purpose physical AI capable of replacing the full range of skilled trade work — rewiring a 1960s building, fitting a boiler in a basement not designed for automation — remains a harder problem than automating discrete digital tasks. Task-specific tools are arriving: smart excavators, AR-assisted installation, automated material handling. But these augment the skilled worker rather than replace them, and the most complex, site-specific work stays stubbornly human.

The second layer is less obvious and more durable: even when new construction becomes increasingly automated, the existing built environment will require skilled human hands for maintenance and repair for decades. The building stock of any developed economy turns over slowly — infrastructure built in the 1960s will still need an electrician in the 2040s. Automation arrives first at the frontier — new builds, standardised projects, industrial construction — and works backwards. The long tail of existing stock ensures sustained demand long after the frontier has moved.

The software parallel is instructive. New code written today is already substantially AI-assisted — the frontier has moved. But the large legacy systems that run most major organisations — the SAP installations, the COBOL codebases, the bespoke enterprise platforms — move far more slowly. The existing stock follows the new, at its own pace, measured in decades.

If the criterion is career stability over the next twenty years: the physically-bound trades remain a structurally sound option — not because physical automation is

impossible, but because the existing stock ensures sustained demand, task-specific tools expand capability rather than eliminate roles, and general-purpose replacement arrives last and slowest precisely where most of this work happens. (For the full labour market picture — vocational training infrastructure, migration patterns, and the political economy of who does this work — see the Appendix.)

The Arithmetic of the Obligation

Here is a question worth sitting with.

If the elastic two-thirds of the economy multiplied its output — even modestly — what happens to aggregate economic output when the inelastic third is fully displaced?

It almost fills the gap. The elastic two-thirds grows to absorb most of the space.

The arithmetic is instructive even before you get to heroic assumptions. If the elastic two-thirds of the economy gains 40% more productive output — roughly what IBM reports their software developers achieving with AI tools today — then two-thirds times one-point-four equals about 93 percent of the original total. Not a complete replacement of the displaced inelastic third. But close. And from a realistic, documented, already-observable productivity gain. Not speculation about ten-times transformations.

Hermann Arnold, an entrepreneur who reviewed an early draft of this essay, said something honest about where he personally stands: five to ten times his previous output. His teams, he estimates, at two to five times. He described this as commercially successful — but not yet sufficient as a societal contribution.

He is right on both counts. But here is the harder truth that the arithmetic alone does not show: multiplied output does not automatically reach displaced workers.

The United States ran strong productivity growth from 2010 to 2020. Total economic output was up substantially. Median real wages stagnated. The aggregate improved. The distribution did not follow. The gains went to shareholders, to executives, to the owners of capital — not to the people who needed the economy to grow.

There is a second asymmetry buried inside the elastic-inelastic split, and it deserves to be named directly: it maps almost exactly onto physical work versus digital knowledge work.

The person multiplying their output five times with an agent is, in practice, a knowledge worker — someone who can do the job from a laptop, from a beach, from a mountain trail with a phone signal. That freedom is precisely what makes

the multiplier available to them: relocate, restructure the day, delegate the routine parts, and the tools compound.

The electrician rewiring a house, the plumber, the crane operator, the person pouring concrete — none of that comes with the option to work from anywhere. Unlike pure back-office roles — where AI can replace the entire function — trades workers face a double burden: no output multiplier on their core work, and the physical transition literally depends on them. No agent can climb a scaffold. It requires a body in a specific location, showing up, every day, for years. No home office. No four-times output from an agent, because an agent cannot climb a scaffold or run conduit through a wall.

This is where the individual obligation becomes a societal one. The people whose labour cannot be relocated, cannot be multiplied by AI, and cannot be substituted for years yet, are the same people the entire transition physically depends on. That dependency has to be priced accordingly.

This is where the individual obligation ends and the structural question begins.

Three Paths — Two Dead Ends

Assume you accept the argument so far, even if you only skimmed it. Here it is again, in one paragraph, so nothing that follows depends on remembering the details above.

The gap between software displacing jobs and physical infrastructure catching up is real — it will take years, not months. Cash cannot bridge it — Cantillon ensures it makes things worse, because new money always reaches asset owners before it reaches the people it was meant to help. Elastic-sector workers can personally close part of the gap by multiplying their own output, but that multiplication does not automatically reach the displaced. It shows up as higher profits for the companies and investors who already own the productive assets, unless something structural forces it to be shared.

With that in view: what are the structural alternatives?

Tax the winners. Tax AI companies heavily. Redistribute. The obvious problem: the people being taxed are the same people with the most influence over tax policy. The less obvious but more important problem: taxes on corporate profits arrive slowly, irregularly, and under constant political pressure to be reduced. Workers are displaced now. The tax revenues come later, in the wrong form, and may not even be enough. There is a third problem, less often discussed: taxing success disincentivises the risk-taking that generates it. Heavy winner-taxation changes the investment calculus: capital migrates to lower-tax jurisdictions; some innovation

simply does not happen. The tax base you intend to capture is also the productive capacity you are eroding.

Regulate the speed. Force companies to slow AI adoption. Give workers time to retrain. This is the instinct of most European governments and every trade union. The problem: slowing adoption in one country does not slow it globally. Capital migrates to wherever the transition is moving fastest. The country that delays loses tax base, loses competitive position, and ends up with the same displaced workers and less capacity to help them. Unilateral slowdown is self-defeating.

The closest analogy is nuclear weapons. No nation can unilaterally disarm and expect the problem to disappear. A country that slows AI adoption does not slow the global transition — it exits as a participant and re-enters as a dependent. The race cannot be stopped. The question is only who shapes its terms.

Build sovereign capacity before the gap opens. This is the approach that has actually worked — not in an AI context, but in an analogous one. Norway discovered oil in the 1960s. Rather than spending the revenues directly on consumption — which would have caused Dutch Disease, inflating the Norwegian economy and hollowing out its other industries — Norway systematically converted finite oil revenues into permanent financial capital. The Government Pension Fund Global now holds \$1.8 trillion, equivalent to roughly \$330,000 per Norwegian citizen. The rules are strict: only three percent withdrawn per year; the principal preserved permanently; invested entirely abroad to avoid distorting the domestic economy. Norway participates in global growth without cannibalising itself.

The AI equivalent: the companies capturing AI-driven productivity gains are the new oil wells. A state that takes equity positions in those companies early — before the full value is priced in — and deploys the returns into domestic physical capacity is doing something fundamentally different from taxing and redistributing. It is owning a share of what is being created. Not extracting. Participating.

Two honest caveats. The first is timing: the window for the obvious targets — the foundational AI laboratories — has largely closed. OpenAI, Anthropic, and their peers carry valuations that already price in decades of anticipated dominance. The first wave is over. But technology transitions do not stop at the first wave.

Specialised AI companies — in robotics, in vertical market software, in physical infrastructure automation — are still in early formation. The case for sovereign participation is sharpened by the closing of the first window, not diminished. Urgency is the argument.

The second caveat is structural. Norway's oil was, by geological circumstance and Norwegian law, a national resource — the state's claim to a share of its value was uncontested. AI companies are private enterprises. Taking equity requires either willing sellers — sovereign wealth funds investing in private rounds, as Singapore's

Temasek already does — or political choices that approach compelled participation. The source of capital is different too: oil revenues were a geological windfall; the equivalent would have to be constructed from existing tax base or government efficiency gains. The Norway model is not a blueprint to copy but a demonstration of a principle: convert finite-moment gains into permanent capital, deploy as capacity not transfers. The specific mechanisms for AI will be harder to find. That difficulty is an argument for starting now, not for abandoning the approach.

There is also a version of sovereign capacity building that requires neither discovering oil nor running trade surpluses for decades. European governments currently spend between 50 and 57 percent of GDP on public services and administration. [Eurostat General Government Expenditure, 2024] Even a ten percent efficiency improvement through AI — faster processing of benefits, reduced administrative overhead, better-targeted transfers — would free the equivalent of five to six percent of GDP annually. In Germany, that is over 200 billion euros per year available for redeployment into permanent capital. The barriers are political and institutional, not economic. (A fair rejoinder: government is elastic too. Efficiency savings have a consistent history of being absorbed by expanded mandates rather than redeployed as capital. The mechanisms for ensuring gains convert into sovereign capacity — strict rules, external investment, constitutional protection of the principal — are precisely what the Norway model illustrates.)

The contrast is Great Britain after the industrial revolution. By 1900 — 30 years into the full industrial transition — Charles Booth's surveys of London found that 30% of the city's population lived below subsistence level. Seebohm Rowntree found near-identical results in York. Productive assets had been privatised during the transition itself. [Booth, Life and Labour, 1889-1903; Rowntree, Poverty: A Study of Town Life, 1901] The lesson was available. It was not applied.

Objection: Will AI companies generate enough to participate in?

A reasonable question: what if the whole AI sector commoditises? Efficient models cut inference costs. Locally-run models improve. Infrastructure overbuilding creates conditions for price wars. Technology has done this before.

But here is what the history of technology transitions actually shows: the underlying resource commoditises. The product layer does not.

AWS generates over \$100 billion in revenue with margins above 30%. Raw compute commoditised, exactly as predicted. But the APIs, the managed services, the institutional knowledge accumulated by millions of developers — those did not commoditise. The moat is not in the hardware. It is in everything built on top of the hardware.

Foundation AI models are playing out the same dynamic. The underlying model weights might eventually be replicated by open source equivalents. And in an AI era, products themselves face a new threat: if AI can build software faster and cheaper than before, does the product layer become more replicable too?

The answer is that the moat shifts — it does not disappear. In the previous era of software, the moat was the product itself: years of engineering accumulated in a complex application. In an AI era, that moat compresses. But the moats that formed around successful software — the data generated by millions of interactions, the network effects that make each additional user more valuable, the institutional embedding in workflows that makes switching costly — these do not disappear because building software gets cheaper. If anything, they become the only moats that hold. More use creates more data creates better product creates more use. The cycle is the same; the entry cost for the product is lower, but the exit cost for the user is not. > > Three plausible scenarios: **A — Race to the Bottom** (edge AI + efficient models + overbuilt infrastructure → margins collapse). **B — Product Layer Moat** (raw compute commoditises; a small number of dominant product companies maintain high margins through network effects). **C — Hybrid** (both simultaneously, with a public infrastructure layer emerging). Scenario A is historically unusual. B and C are more consistent with how previous transitions resolved — and both still produce concentration. Which is exactly why sovereign equity participation matters more, not less. Norway did not wait until oil was expensive. It started building the fund when oil was cheap.

What Sovereign Capacity Has to Look Like

There is one more constraint. And this is where Cantillon comes back.

Sovereign capacity is not one thing. The form it takes determines whether it reaches the people it is meant for.

If a sovereign fund distributes its returns as cash, Cantillon activates. New money enters the economy through citizens, travels toward asset owners, raises prices. The check arrives. The landlord raises the rent. The mechanism is unchanged — it does not matter that the money originated in a government equity stake rather than a printing press. Once it becomes cash in circulation, it moves by the same rules.

The only structural responses that escape the Cantillon trap are the ones that do not give out cash at all.


Universal Basic Services — housing, healthcare, energy, transport provided directly as public infrastructure — breaks the chain at its weakest link. There is no cash to flow through. The landlord cannot raise the rent on a publicly owned apartment the state refuses to sell. The insurance company cannot raise premiums on healthcare delivered by public institutions at cost. The asset is owned. Citizens benefit from the asset. The Cantillon distortion never gets a chance to activate.

The evidence is not theoretical.

Vienna’s Gemeindebau: more than 60% of Vienna’s residents live in city-owned subsidised housing. The cost of shelter — the single largest line item in most household budgets — is removed from private speculation. A private housing market crash does not put Viennese families on the street in the same way it does in Munich, a city of comparable economic scale, where a dysfunctional private housing market has created one of Europe’s worst affordability crises. The Cantillon chain ends where public ownership begins.

Singapore’s Housing Development Board: 80% of residents in state-built housing. [Singapore Housing Development Board, 2023] A city-state with no oil and no colonial-era capital base built one of the world’s most stable living environments by systematically removing shelter from the reach of Cantillon’s mechanism. The “resource” was disciplined policy, compounded over decades.

Norway’s Government Pension Fund does not send checks. It buys assets. When oil revenue becomes equity stakes in global companies, it compounds. When those returns fund public infrastructure, the Cantillon distortion does not apply — because the money never enters the consumer economy as cash. It goes directly into capacity.



Demis Hassabis, co-founder of DeepMind and one of the central figures of the current AI transition, arrives at the same conclusion in The Infinity Machine. He is sceptical of universal basic income as the solution to social disruption from AI. What he proposes instead he calls universal basic provision: not cash, but the direct delivery of a modern standard of living — a decent house, schooling, health care,

basic travel. “All of that costs you nothing as a citizen,” he writes. Services, not money. His framing matches the mechanism described in this essay exactly: remove the cost of living from the reach of private speculation, and the Cantillon problem solves itself.

Universal Basic Services is not a left-wing policy position. It is what Cantillon tells you to do once you understand the mechanism. The form of help is not a detail. It is the mechanism.

A fair counter-objection: the providers of universal basic services face cost increases too — in labour, in materials, in land. Those costs do not disappear; they are absorbed by the public sector rather than passed to tenants, but they still require funding. Vienna’s Gemeindebau works not because building public housing is cheap today, but because a century of accumulated public stock means the marginal unit is already built and paid for. The argument for Universal Basic Services is an argument for building the stock now, before cost curves shift further — which is, again, an argument for urgency rather than against the approach.

What This Has to Do With You

I spent 140 days running an experiment in agentic living — one person, phone only, multiple companies running, two ultramarathons, four countries. The details are in a companion piece, *Agentic Living*. The short version: no limitations found.

That experiment was the individual story. This essay is the structural one. They are connected.

The full argument, assembled:

The transition gap is real. Cash cannot bridge it — Cantillon ensures it makes things worse. Elastic sector workers can close part of the gap arithmetically — the IBM numbers demonstrate this, and they are not heroic assumptions — but that multiplication does not automatically fix the distribution. Structural responses are needed: sovereign capacity building, equity participation in AI productivity gains, and the delivery of that capacity in the only form that escapes Cantillon’s mechanism — services, not transfers.

None of that structural work happens without the individual part first. You cannot build a tax base on a hollowed-out economy. You cannot fund sovereign capacity on the revenues of a workforce that never multiplied its output.

Go agentic. Use the tools. Not to do the same work with less effort — but to do ten times more with the same effort.

What that looks like depends on where you are.


If you are an employee — especially in a large organisation with compliance constraints — the instinct is to wait for your employer to deploy AI. That instinct is expensive. Start with the administrative overhead of your own role. The person who learns to multiply their output 40% inside a constrained environment gains more negotiating power, not less.

If you are a manager, the finance team’s instinct is to read AI as a cost reduction — fewer headcount, same output. The right question: what becomes possible that was impossible before? Teams that multiply their output can serve markets that were too expensive to reach, build products that would have taken years.

If you are a citizen, vote for, fund, and demand the institutions that deploy AI gains as permanent public capacity — not as checks, but as services. Public housing. Public energy infrastructure. The Cantillon mechanism guarantees that cash transfers will not reach you intact.

Tim Ferriss did not just argue that remote work was possible. He showed, step by step, how to propose it to your employer. The individual obligation in this essay is the same kind of argument — but the field manual for how to act on it is somewhere else.

That manual is Agentic Living. Where this essay argues that elastic-sector workers have an obligation to multiply their output, Agentic Living shows what that multiplication actually looks like in practice. The argument and the implementation belong together. This essay is the structural case. Agentic Living is the execution.



Dario Amodei, co-founder and CEO of Anthropic, has made this point consistently with large enterprises. The instinct of finance teams is to reach for AI as a cost-reduction mechanism — fewer headcount, same output. The right frame is expansion: teams of ten doing what previously required teams of a thousand, markets served that were too expensive to reach, products built in weeks that would have taken years. The top-line question, not the bottom-line question.

That is the individual obligation.

And then: demand the institutions that deliver what that produces — not as checks, but as services. That is the structural one.

The individual incentive and the collective obligation point in exactly the same direction.

The gap does not close itself.

Which side of it are you building toward?

Enzo Duit Zell am See · July 2026

With thanks to Hermann Arnold, Stefan Haubner, and Felix Ambros, whose feedback across multiple rounds of review sharpened every argument in this essay. Hermann's challenges on AI company profitability, the physical AI timeline, sovereign equity valuation timing, and the limits of the Norway analogy shaped the Race to the Bottom section, the Obligation arithmetic, and the structural alternatives. Stefan's push for clarity on structure and self-explaining concepts shaped the narrative architecture. Felix tightened the argument at its foundations — his challenges on the abundance assumption, the Jevons Paradox application, the logical gap on trades workers, and the Cantillon mechanism in regulated markets all strengthened the core mechanism. The errors that remain are mine.

The elastic/inelastic demand distinction builds on standard price elasticity theory. Jevons Paradox was first described by William Stanley Jevons in *The Coal Question* (1865). The Cantillon Effect references Richard Cantillon's *Essai sur la Nature du Commerce en Général* (c. 1730). $MV=PY$ is the classical Quantity Theory of Money. The effectiveness framing references Dario Amodè's public articulations of AI's role in value creation. Sovereign wealth models reference publicly available data on Norway's Government Pension Fund Global, Singapore's Temasek Holdings and HDB, and Vienna's Gemeindebau programme. Hermann Arnold's quotes used with permission.

Appendix: The Labour Market in Detail

For readers who want the full picture on who does physically-bound trade work, what vocational training infrastructure looks like across economies, and how migration patterns shape the transition.

Vocational Training Infrastructure

Beyond wages, there is a deeper structural question: how does a society make these careers genuinely attractive? Better pay is necessary but not sufficient if the pathways into the trades have been allowed to decay for decades. Germany and Austria retain the most sophisticated vocational training infrastructure in Europe — co-designed with industry, well-funded, and genuinely prestigious. Most English-speaking economies have let the equivalent degrade. That is a policy choice, and it can be reversed. The economic case is straightforward: these roles are structurally essential to the physical transition, increasingly scarce, and the training infrastructure that produces them is cheap compared to the social cost of not having enough of them.

Who Actually Does This Work

In most developed economies, a disproportionate share of physically-bound trade work is performed by migrants and workers with migration backgrounds. In Germany, Destatis data from 2023 shows that 67% of employees in construction finishing and drywall work had a migration background — compared to 26% across the economy as a whole. In the United States, approximately 26–30% of the overall construction workforce is foreign-born, rising to 42% among construction laborers and around 18–22% for licensed trades such as electricians, plumbers, and HVAC technicians. [Sources: Destatis 2025; NAHB / Harvard Joint Center for Housing Studies, 2023–2024; U.S. Bureau of Labor Statistics, 2025]

This reflects a structural drift: in higher-income economies, native populations have increasingly avoided entering physically demanding trades — pulled toward white-collar, digital, or credentialed careers — while migrants have filled the gap, often because the entry barriers are lower and the relative income better than what they left behind.

The Political Economy

The result is a dependency that cuts across one of the most politically charged fault lines of the current decade: the same constituencies most vocal about restricting immigration are often the ones most reliant on migrant labour to wire their homes, fit their boilers, and build the infrastructure they live on. That contradiction is not a fringe tension — it is a live structural paradox that helps explain why skilled trades shortages and migration policy are more intertwined than either side of the debate typically acknowledges. Research on European populism consistently finds that

economic insecurity and wage stagnation among workers in routine and manual occupations correlate with higher support for anti-establishment parties — not because of trade shortages per se, but because of the broader sense of being structurally left behind while others capture the gains.