

AI: pain vs gain - the evidence

Evidence base for the nomination. Every claim carries a source and an evidence grade (A robust → D contested).

The discipline this issue most needs: separate the empirical from the value question

AI is the kind of subject where the cardinal error the Method exists to prevent - presenting a value choice as a technical necessity - is most tempting, and most dangerous. So the two questions are kept apart throughout. Every factual claim below carries a source and an evidence grade; the value questions are surfaced rather than buried; and the candidate mechanisms are presented as a neutral menu for the public to weigh, not as positions Pragma advocates.

- **The empirical question** - *Will AI materially erode the labour-based tax base and concentrate the returns to a transformative technology in a narrow set of hands?* This is contestable with evidence, and the evidence is mixed: robust on the structure of the fiscal base, genuinely uncertain on the scale and speed of labour displacement.
- **The value question** - *If the gains are large, how should they be shared, and who should bear the cost of the transition?* No amount of evidence settles this. It is a question about fairness and the social contract, and it belongs to the public and their representatives, not to economists or to Pragma.

A respectable case can hold that the empirical risk is real *and* that the right response is to do very little. Keeping the two questions separate is what lets that argument be had honestly.

1. What is happening

General-purpose AI has moved, within roughly three years, from a research curiosity to a technology being adopted across white-collar work. The economic claim that follows is not that "the robots are coming for every job" - that framing is both unprovable and a distraction. It is narrower and more structural: a technology built substantially on publicly-created foundations (public research, publicly-funded compute and skills, and the collective written and visual output of the population) may (a) shift the returns to economic activity from labour toward capital and a small number of firms, and (b) do so faster than the institutions that fund public services and protect workers can adapt. Whether that shift is large or small is the empirical question below. That it *could* be large is enough to ask whether the UK is prepared for it.

2. What the UK public actually thinks - and why it reframes the case

It is tempting to assume the public's chief AI fear is for their jobs. **The evidence says otherwise**, and getting this right matters for how any proposal is framed.

In the Ada Lovelace Institute / Alan Turing Institute 2025 survey (*How do people feel about AI?*, n=3,513), "job cuts" was named as a concern about large language models by **42%** of respondents - **sixth of eight** listed concerns. Ranked above it: erosion of problem-solving and critical thinking (66%), bias (50%), offensive or harmful content (47%), unclear responsibility when things go wrong (46%), and copyright (45%). A separate question found **94%** very or somewhat concerned about the spread of harmful AI-

generated content online. Demand is for *regulation*, not enthusiasm: 72% said laws and regulation would increase their comfort with AI, up from 62% in 2022/23. [Grade B - robust survey, large sample, but attitudes are early and shifting; the percentages are multi-select, not exclusive rankings.]

Two consequences follow:

1. **The case for this nomination does not rest on public sentiment** - it rests on the structural fiscal fact in section 3, which stands regardless of what the public currently fears.
2. **The public-facing framing must not lead with "AI will take your job."** It should lead with the fiscal and fairness structure: who pays for public services if the tax base narrows, and how the gains from a publicly-underwritten technology are shared. The displacement question is real but contested, and over-claiming it would forfeit credibility.

There is also a useful inversion here: low public salience on the *economic* dimension, combined with the policy gap in section 5, is precisely what makes this a *neglected* problem - which strengthens, rather than weakens, the case under the Method's persistence test.

3. The structural fiscal vulnerability (the spine of the case)

The UK raises the great majority of its revenue from taxing labour income, which makes the public finances structurally exposed to any sustained shift of returns from labour to capital.

- Income tax is forecast to raise **£329bn in 2025-26 - 26.7% of all receipts** and the single largest revenue source. National Insurance contributions raise **£205.4bn (16.7% of receipts)** and are levied *only* on labour income. [OBR; Grade A.]
- Together, income tax and NICs are expected to raise **~£535bn - roughly 43% of total public-sector receipts** (£1,235bn). Adding capital gains tax, income tax + CGT + NICs are **59% of all HMRC receipts** in 2025-26, and have averaged **56% over the past decade**. [OBR; HMRC annual bulletin; Grade A.]
- The tax burden is already at a post-war high and rising: National Accounts taxes climb from 35% of GDP (2024-25) to **an all-time high of just over 38% from 2029-30**, with **roughly two-thirds of the increase since 2019-20 coming from rising personal (labour-income) taxes** - chiefly threshold freezes and employer-NIC rises. [OBR EFO, November 2025; Grade A.]

The interpretation - that a tax base this concentrated on labour is exposed if AI shifts the factor share toward capital - is Pragma's synthesis, not an OBR forecast, and is offered as such. But the underlying structure is not in doubt: when the same activity is performed by capital (compute, models, datacentres) rather than by salaried workers, much of it falls outside the bases that raise 43% of government revenue. Corporation tax and capital taxes do not automatically fill the gap, and a substantial share of AI value accrues to firms domiciled outside the UK.

4. The labour-displacement evidence - real, but genuinely uncertain

This is the part of the case where honesty about evidence strength matters most. The range of credible forecasts is enormous, and anyone citing a single number is overselling.

- IPPR's scenario analysis (of ~22,000 tasks) spans **zero net job losses** (full augmentation, +13% GDP) to **7.9 million** (worst case, no GDP gain), with a central scenario of **4.4 million** jobs displaced alongside +6.3% GDP. Task exposure rises from 11% of current worker-tasks to a potential 59% in a deep-integration "second wave." [IPPR, March 2024; **Grade C** - scenarios, not predictions; the source pre-dates the most recent capability jumps.]
- Official analysis (DSIT / AI Security Institute, January 2026) relays an estimate that **~70% of UK workers** are in occupations containing tasks AI could perform or enhance - higher than the ~60% for comparable economies, reflecting the UK's service-heavy structure - split roughly half "high complementarity" (AI assists the worker) and half "low complementarity" (AI may substitute). [DSIT/AISI relaying an IMF estimate; **Grade C**.]

The verification of this evidence base actively *killed* two over-strong claims that circulate in commentary (a "only 2% think AI will take their job" line, and specific job-advert-decline figures), which is itself a caution: the displacement discourse is full of figures that do not survive scrutiny. The honest position is a **range, not a point estimate**, and an explicit statement that the magnitude and timing are not yet known.

5. What the UK government is - and is not - doing

The UK has an active *industrial* AI strategy and effectively no *distributional* or *fiscal-adjustment* AI strategy. That asymmetry is the policy gap.

What exists (AI Opportunities Action Plan, January 2025, and one-year-on update): five AI Growth Zones reporting £28.2bn of investment; a plan to expand public compute roughly twentyfold by 2030 (Isambard-AI and others); a **£500m Sovereign AI Unit** taking *market-terms* equity in UK AI firms and offering funded compute; ~1m AI upskilling courses delivered against a 10m-by-2030 target; and the AI Security Institute scaled past 100 technical staff. [gov.uk; Grade A for what the plan contains.]

What is absent. On the documentary evidence, there is **no** measure to (a) capture the economic returns of AI for the public, (b) tax AI profits, windfalls or compute, or (c) fund a workforce transition beyond voluntary upskilling. The plan names a "Future of Work Unit" - but for *analysis and evidence*, not delivery. The Sovereign AI fund subsidises private labs without a public upside-sharing mechanism. [Grade B - an inference of absence drawn from the primary documents.] The government's own commissioned commentary (IPPR, April 2026) warns of a public backlash "unless it can show who benefits from AI."

In short: the state is investing public money to *accelerate* AI, while the questions of who captures the returns and who bears the transition cost are unaddressed.

6. The cross-category map - four candidate sub-problems

AI's policy implications do not sit in one Register category. Pragma will likely need separate nominations over time; this is the first, with the rest flagged.

1. **The fiscal base and the distribution of returns** (*the primary nomination*). How the UK funds public services if the labour-tax base narrows, and how the gains from a publicly-underwritten technology are shared.
2. **Displacement and worker protection** (section 8). Whether UK employment law adequately protects workers reorganised, managed or replaced by AI - sharpened by the contrast with the recent Chinese rulings.

3. **Safety, liability and frontier governance.** The cluster the public most worries about (misinformation, harmful content, accountability) and where the UK is most active (AIS). Least neglected, so lowest priority for Pragma.
4. **Compute, energy and infrastructure.** The grid, water and planning demands of datacentres - overlapping the Register's existing energy and planning entries.

7. The mechanism menu - what a UK government could actually do

Presented as a neutral menu of *who-captures-the-returns* / *who-funds-the-transition* options, with precedent, feasibility and the main objection to each. Pragma's role is to cost and present these, not to choose among them.

- **Windfall / excess-profits levy.** A one-off or time-limited tax on abnormal profits. *Precedent (UK):* 1997 utilities windfall tax (~£5.2bn); 2022 Energy Profits Levy (25%, 65% combined headline rate, ~£5bn/yr). [Grade A.] *Main objections:* defining "excess" profit is hard; investment disincentive and competitiveness; retroactivity sits uneasily with rule-of-law expectations of advance notice.
- **Compute levy.** A tax on large-scale training runs or datacentre compute - the physical bottleneck of the AI economy, and harder to relocate than reported profit. *Precedent:* none directly; analogous to other input/extraction levies. *Main objections:* could deter the very investment the Action Plan courts; measurement and avoidance via offshoring of training.
- **Conditionality.** Public warrants, royalties or revenue-share attached *whenever* a firm benefits from public data, public compute, grid priority or public procurement - capturing upside in exchange for a public input already being given. *Precedent:* procurement clawback and equity-for-grant arrangements exist in other sectors. *Main objections:* administrative complexity; risk of deterring uptake of public compute offers.
- **Sovereign wealth fund / National Wealth Fund route.** A standing, arms-length fund holding (ideally non-voting) equity, capturing long-run upside for the public. *Precedent:* Norway's Government Pension Fund Global is the archetype; the UK's **National Wealth Fund (£27.8bn)** is a *policy bank rather than a sovereign wealth fund*, but it *can* and does hold equity stakes. [IfG; Grade A/B.] *Main objections:* the UK has no funding surplus to endow a Norway-style fund; conflicts of interest if the state both owns and regulates (see below).
- **AI / citizens' dividend.** Distributing captured returns directly to households as an annual payment, rather than into general spending. *Precedent:* Alaska Permanent Fund dividend. *Main objections:* depends entirely on there being returns to distribute; macroeconomic and inflation questions at scale.
- **"Robot tax" / automation tax.** Taxing automation that displaces labour. *Main objections:* widely criticised by economists as taxing productivity itself; hard to define a "robot"; likely to slow beneficial adoption.
- **Worker-protection / liability route** (section 8). Not a return-capture mechanism but a transition mechanism: strengthening rights and consultation where AI reorganises or replaces work.

The US comparison, in passing

US Senator Bernie Sanders proposed (June 2026) that the public take a **~50% equity stake** in major AI labs (OpenAI, Anthropic, xAI) via a sovereign wealth fund, justified on the ground that the models are built on "the collective output of humanity." [Grade B.] The commentator David Shapiro, among others, accepted that underlying case but argued the equity-stake *mechanism* is the weakest available - it targets only the visible model labs (missing chips, cloud and energy), creates a government that is "both referee

and investor," and is administratively near-impossible to value and define - preferring instead a value-chain windfall levy, a compute levy, conditionality, an arms-length *non-voting* fund, and household dividends. [Grade C - opinion.] The point for Pragma is narrow: **the equity-stake mechanism does not transfer to the UK at all** (the UK cannot take equity in US firms), but the alternatives Shapiro and others converge on are exactly the UK-actionable levers above. The US debate is context; it is not the subject.

8. Displacement, worker protection and the liability route

A development worth Pragma's attention: courts in **Hangzhou, China** ruled it **unlawful** for a firm to dismiss an employee (a quality-assurance supervisor, surname Zhou, on ~300,000 yuan/≈£35,000) *solely* in order to replace him with AI - the employer was found to be improperly "shifting the cost of the technological transformation to the employee," and the dismissal did not meet the legal conditions for redundancy. The ruling built on a December 2025 precedent. [NPR, Fortune, Caixin, Bloomberg; **Grade B** - multiple reputable outlets, not the primary judgment.]

The UK contrast is instructive. UK workers are covered by general law (UK GDPR, the Information and Consultation Regulations, health-and-safety duties, the ECHR) but there is **no AI-specific protection**: the TUC identifies gaps in transparency, protection against discriminatory algorithms, and worker voice, and its draft *Artificial Intelligence (Regulation and Employment Rights) Bill* - proposing a Workplace AI Risk Assessment, consultation before high-risk decisions, a right to explanation and a right to human reconsideration - is **not** part of the government's Employment Rights Bill. [TUC; Grade B.] Notably, **86% of UK working adults oppose AI being used to make firing decisions**. [TUC; Grade B.] And under UK redundancy law, technological reorganisation can constitute a *genuine* redundancy - so, on the face of it, a UK worker replaced by AI may have **weaker** protection than the Chinese precedent now provides. [Grade C - Pragma's reading; not legal advice.]

9. The value questions for the public

These are the questions to put to the public rather than to answer for them:

- Should the returns from a technology built on publicly-created foundations be shared with the public - and if so, by how much and through what vehicle?
- Where the gains accrue but the displacement costs fall on particular workers and places, who should bear the transition cost - the firms, the general taxpayer, or those displaced?
- How should any captured return be used - to fund public services, to cut other taxes, to endow a fund, or to pay a direct dividend?
- How much weight should be given to *not deterring investment and adoption* relative to *capturing returns and protecting workers*? (This is where the fiscal-conservative, social-democratic and libertarian perspectives most sharply diverge, and the disagreement should be shown, not resolved.)

10. Evidence summary

Claim	Grade
Labour-income taxes are ~43% of receipts; IT+CGT+NICs 59%	A
Tax burden rising to ~38% of GDP, two-thirds from personal taxes	A
Job loss is not the UK public's top AI concern (42%, 6th of 8)	B
No current UK policy on AI return-capture, taxation or transition funding	B
Windfall-tax precedents and objections	A
National Wealth Fund can hold equity but is not a sovereign wealth fund	A/B
Scale/speed of AI labour displacement	C (contested)
~70% of UK workers in AI-exposed occupations	C
Chinese courts' AI-dismissal rulings	B
UK worker-protection gap vs the Chinese precedent	C
Sanders proposal / Shapiro critique (US context)	B / C

Sources

- Ada Lovelace Institute / Alan Turing Institute, *How do people feel about AI? (2025)* - <https://attitudestoai.uk/findings-2025> ; <https://www.adalovelaceinstitute.org/press-release/ada-turing-survey-2025/>
- OBR, tax-by-tax (income tax; NICs) and public-finances brief guide - <https://obr.uk/forecasts-in-depth/tax-by-tax-spend-by-spend/income-tax/> ; <https://obr.uk/forecasts-in-depth/tax-by-tax-spend-by-spend/national-insurance-contributions-nics/> ; <https://obr.uk/forecasts-in-depth/brief-guides-and-explainers/public-finances/>
- OBR, *Economic and Fiscal Outlook*, November 2025 - <https://obr.uk/efo/economic-and-fiscal-outlook-november-2025/>
- HMRC, tax and NICs receipts annual bulletin - <https://www.gov.uk/government/statistics/hmrc-tax-and-nics-receipts-for-the-uk/hmrc-tax-receipts-and-national-insurance-contributions-for-the-uk-new-annual-bulletin>
- IPPR, *Transformed by AI* (March 2024) - <https://www.ippr.org/media-office/up-to-8-million-uk-jobs-at-risk-from-ai-unless-government-acts-finds-ippr>
- DSIT / AI Security Institute, *Assessment of AI capabilities and the impact on the UK labour market* (January 2026) - <https://www.gov.uk/government/publications/assessment-of-ai-capabilities-and-the-impact-on-the-uk-labour-market/assessment-of-ai-capabilities-and-the-impact-on-the-uk-labour-market>

- UK Government, *AI Opportunities Action Plan* and one-year-on - <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan> ; <https://www.gov.uk/government/publications/ai-opportunities-action-plan-one-year-on/ai-opportunities-action-plan-one-year-on> ; <https://www.sovereignai.gov.uk/>
- Institute for Government, windfall taxes; National Wealth Fund explainers - <https://www.instituteforgovernment.org.uk/article/explainer/windfall-taxes> ; <https://www.instituteforgovernment.org.uk/explainer/national-wealth-fund>
- TUC, *Artificial Intelligence (Regulation and Employment Rights) Bill* and AI Bill Project - <https://www.tuc.org.uk/research-analysis/reports/artificial-intelligence-regulation-and-employment-rights-bill> ; <https://www.tuc.org.uk/research-analysis/reports/ai-bill-project>
- China AI-dismissal rulings - NPR (<https://www.npr.org/2026/05/01/nx-s1-5807131/tech-worker-china-ai>), Fortune (<https://fortune.com/2026/05/03/chinese-court-layoffs-workers-ai-replacement-labor-market/>), Caixin (<https://www.caixinglobal.com/2026-04-30/chinese-courts-rule-companies-cannot-fire-workers-simply-to-replace-them-with-ai-102439602.html>)
- Bernie Sanders, *The Public Should Own Half of the Big A.I. Companies* - <https://www.sanders.senate.gov/op-eds/the-public-should-own-half-of-the-big-a-i-companies/> ; David Shapiro, *Bernie Wants Half of Big AI* - <https://daveshap.substack.com/p/bernie-wants-half-of-big-ai>