

AI: pain vs gain

The fiscal base and the distribution of AI's returns - who gains, and who pays.

This is not a question of whether artificial intelligence is, on balance, good or bad. It is a question of distribution: the gains from AI and the costs of the transition fall on **different people**. The returns may accrue to a small number of firms and to capital; the costs - a narrowing tax base, and displacement concentrated on particular workers and places - fall on the public. The UK has no policy that shares the first or cushions the second. That gap, not a verdict on the technology, is what this nomination is about.

- **Proposed domain:** Economy, Tax & Public Spending
- **Candidate primary market insufficiency:** Distributional failure (taxonomy category 9), with a secondary government-created-distortion element (category 10) and a missing-market element for transition/displacement risk (category 6)
- **Related problems:** pension adequacy (which rests on the same labour-contribution base), productivity and regional inequality, and employment, skills and contribution

The problem in one paragraph

The UK funds the great majority of its public services by taxing labour income, and protects most workers through institutions built for an economy in which paid human work is the dominant source of both output and revenue. Artificial intelligence threatens both foundations at once: it may shift the returns to economic activity from labour toward capital and a small number of firms - narrowing a tax base that raises 43% of government revenue - while the gains accrue disproportionately to those firms and the costs of adjustment fall on particular workers and places. The state is currently investing public money to *accelerate* AI while leaving entirely unaddressed the questions of who captures its returns and who bears the cost of the transition. The empirical scale of the shift is uncertain; the structural exposure, and the absence of any policy response to it, are not.

Gateway Test 1 - Materiality

Is the harm real, measurable and significant?

Assessment: passes on structural exposure; the displacement magnitude is honestly uncertain.

The materiality rests on two limbs, and the Method requires both to be graded honestly:

- **The fiscal exposure is robust and measurable.** Income tax (£329bn) and NICs (£205.4bn) raise ~£535bn, ~43% of all receipts; income tax + CGT + NICs are 59% of HMRC receipts; the tax burden is rising to a post-war high of ~38% of GDP, two-thirds of the increase since 2019-20 coming from personal taxes. A base this concentrated on labour is, by construction, exposed to a sustained shift of returns toward capital. [OBR/HMRC; **Grade A.**]
- **The displacement magnitude is contested.** Credible forecasts span from no net job loss to ~8 million (IPPR central: 4.4m); ~70% of UK workers are in occupations with AI-exposed tasks. These are scenarios, not predictions. [**Grade C.**]

This is the test where candour is essential. The nomination does **not** claim AI will certainly displace millions of workers. It claims that the *fiscal and distributional structure* is materially exposed, that the *potential* scale

is large enough to be serious, and that the asymmetry of who-gains versus who-bears-the-cost is itself a present harm in the making. A reviewer who believes the displacement will be modest can still accept the materiality of the fiscal-base exposure. A reviewer who believes neither limb is material has a legitimate basis to reject - and that disagreement should be recorded, not smoothed over.

Gateway Test 2 - Market insufficiency

Will markets, left to themselves, fail to resolve this - and exactly why?

Assessment: passes - primarily distributional (category 9), honestly labelled as such.

The market will, if anything, work *efficiently* and still produce an outcome the public may judge unjust: returns flowing to capital and to a concentrated set of firms, with no mechanism by which the public - whose data, research, compute and skills underwrote the technology - shares in the upside, and no market that insures workers against displacement. The relevant taxonomy categories:

- **Category 9 - Distributional failure (primary).** This is not a textbook efficiency failure; it is an openly value-laden question about how the gains from a transformative, publicly-underwritten technology should be shared, and whether displaced workers are left to bear the adjustment alone. The Method requires this to be labelled as a distributional choice, **not** disguised as technical necessity - which is exactly why a developed product would present a neutral menu of options rather than a recommendation.
- **Category 10 - Government-created distortion (secondary).** Public subsidy (compute, the Sovereign AI fund, grid priority, public data access) is being provided to private firms with no return-capture attached - a structure the state is actively creating.
- **Category 6 - Missing market (secondary).** No private market insures a worker against having their occupation automated; the risk is correlated, hard to price and long-horizon - the same family of failure that defeats private long-term-care insurance.

A private market will not self-correct any of these, because none of them is a profit opportunity for a private actor to capture.

Gateway Test 3 - Government tractability

Is there a credible lever by which government policy can materially improve the outcome?

Assessment: passes, with one honest limit.

There are multiple credible, UK-actionable levers with domestic precedent - windfall/excess-profits levies (1997 utilities; 2022 Energy Profits Levy), a compute levy, conditionality on public inputs, a standing fund holding equity (the National Wealth Fund already can), a citizens' dividend, and strengthened worker-transition rights. The government already has the fiscal and legislative machinery for most of these; the question is design and will, not capability. [Grade A on precedent.]

The honest limit: the UK cannot capture the returns of *foreign-domiciled* AI firms by taking equity in them (which is why the US equity-stake proposal does not transfer). UK levers act on activity, profit, compute and procurement *within UK reach* - meaningful, but not the whole global value chain. This is a real constraint on ambition, not a reason the problem is intractable; it is recorded so the eventual proposal does not over-claim.

Gateway Test 4 - Persistence or neglect

Has this genuinely resisted resolution - recurring, or repeatedly attempted and failed, or left unaddressed?

Assessment: passes on neglect - and this is the test that most needs argument, because the problem is emerging rather than decades-old.

The Method's persistence test is designed to keep Pragma on stuck, neglected problems and *away* from new, salient issues already being acted on. AI is plainly new and salient - so the nomination must clear this test deliberately, not by assumption.

It qualifies under the test's third limb - "**left unaddressed**" - on direct documentary evidence: the UK has an active *industrial* AI strategy and **no** *distributional or fiscal-adjustment* strategy. There is no policy to capture AI's returns for the public, to tax AI profits or compute, or to fund a workforce transition beyond voluntary upskilling (the "Future of Work Unit" is for analysis, not delivery). [Grade B.] The neglect is reinforced from an unexpected direction: because public anxiety centres on misinformation and trust rather than on jobs and fairness, the *economic* dimension attracts little political pressure - it is neglected on both the policy and the salience axes at once.

On the concern that this is "chasing a headline": the distinction the Method should draw is between a problem that is *new and already being acted on* (leave to the existing process) and one that is *new, consequential and being actively ignored on the dimension that matters here* (a proper target). This is the latter. Some sudden, structural shifts demand a methodical, timely response precisely *because* the existing process is not providing one - and the cost of waiting compounds. An absence of government action on a problem of significant national importance makes hesitation, not action, the indefensible choice.

A reviewer could legitimately argue the opposite - that it is too early, the evidence too thin, and the right course is to watch and revisit. That is a respectable position, and if the nomination is rejected it should be recorded as *considered and not admitted* under Test 1 (materiality not yet demonstrable) or Test 4 (premature), with the reasoning stated, so it can be revisited as the evidence matures.

If admitted - what a developed product would have to resolve

This nomination is for *admission to the Register*, not a finished policy. A full product would still need to:

- Cost the candidate return-capture and transition mechanisms (windfall levy, compute levy, conditionality, fund, dividend) as a **neutral menu**, with whole-life revenue, incidence and behavioural effects - never presenting one as the answer.
- Separate cleanly, in the model itself, the empirical question (the fiscal/displacement projection) from the value question (how the gains are shared and who bears the transition cost), and route the latter to the public.
- Stress-test every mechanism from the fiscal-conservative, social-democratic and libertarian perspectives, reporting agreement *and* disagreement.
- Take the displacement evidence as a **range**, with the proposal robust across the range rather than betting on a point estimate.
- Address the worker-protection limb (the UK gap relative to the recent Chinese rulings) as a distinct, possibly separately-nominated, sub-problem.
- Quantify the foreign-domicile leakage honestly, so the achievable return-capture is not overstated.

Recommendation

Admit as a Register entry under the **"left unaddressed"** limb of Test 4, on the strength of the Grade-A fiscal-base materiality and the documented policy gap - **or** record a reasoned rejection as premature, to be revisited as the displacement evidence matures. Either way, the decision and its reasoning should be on the record. Sequencing - when to develop it relative to social care, water and local-government finance - is a separate question for the roadmap, not for this nomination.

The figures, sources and the full evidence grades behind this nomination are set out in the evidence and framing that accompanies it.