

T H E L O N G E R L O O K

thelongerlook.com

What the Inheritance Tax Reform Means for UK Tech

*A plain English explainer for founders, angels, VCs,
and the people who fund and build UK tech companies*

Doug Scott

30 April 2026

About this work

Doug Scott is not a lawyer or an accountant. He is a founder. A friend shared a policy document about the April 2026 inheritance tax reform with him, and he decided to see what AI tools could do with it. He prompted four AI tools — Claude, ChatGPT, Grok, Gemini — across multiple parallel sessions with simple continuation-style cues, and answered when the tools prompted back. The AI tools produced the writing, the analysis, the citations, and the cross-critique. Doug scanned the output and decided to ship. No human expert reviewed any of this work before publication. The instructions he gave were simple ones, repeated across the work: *be factual, be truth-seeking, do not flinch from where the evidence leads*. The goal he set was that all of the information should be in the public domain and every argument tested, so that a government — and the citizens it serves — can make the decision in the long-term benefit of the country. This piece is one of several that resulted.

What this is and is not. It is the product of a non-specialist, working with AI tools, on a question that affects him directly. It is not a legal opinion. It is not financial advice. It is not an HMRC, HMT, or Treasury document. The author owns shares in unlisted UK companies and would be affected by parts of what is discussed. Readers should weigh the analysis with that knowledge.

Scope. This piece is about UK tech specifically — the cohort the country has said it wants to grow more of. The reform also affects agriculture and other family businesses. Those are real questions handled in other places. This piece is about UK tech.

The change, in one paragraph

From 6 April 2026 the UK caps Business Property Relief at £2.5 million per person (£5 million per couple, since the allowance now transfers to a spouse). Above that, you get only 50 per cent relief — meaning an effective 20 per cent inheritance tax rate on the excess. Tax can be paid in ten annual instalments, interest-free. AIM-listed shares now get only 50 per cent relief from the first pound, no allowance. The cap was originally going to be £1 million when announced in October 2024; the government raised it to £2.5 million on 23 December 2025 after lobbying. The £2.5 million cap refreshes every seven years for individuals (ten for trusts). That is the reform. The argument is about whether the way the tax is collected — the death-event valuation of shares in private trading companies — is the right mechanism for the kind of equity tech founders, their early staff, and their early investors actually hold.

Why this matters for tech specifically

Founder equity in a growth-stage UK tech company is not like other illiquid assets. Three things make it unusual.

The value at any given moment is a paper number against an outcome that has not yet happened. A founder of a Series B SaaS company, valued at £80 million in the last round, holds shares that might eventually be worth £400 million if the company succeeds, £80 million if it muddles along, or zero if it fails. The death-event mechanism forces a single number — the death-date valuation — on top of that distribution. There is no observable market price; the value is whatever HMRC and the estate eventually agree on, after a process that often takes years. This is structurally different from a piece of farmland, where the death-date value is a reasonable proxy for what the heir will eventually realise.

The cohort is mobile. A UK tech founder can build the same company from Dubai, Singapore, Lisbon, Austin, Zurich, or Dublin. The same is not true of UK farmers, whose land is in the UK by definition. It is also less true of family-business owners in mature sectors whose customers, suppliers, and operations are physically anchored. The founder cohort that the reform most affects is the cohort with the most international optionality, and the cohort the country is most explicitly trying to retain.

The capital stack interacts with the tax in non-obvious ways. By the time a UK tech company has reached a valuation where the founder's stake breaches the £2.5 million allowance, the cap table is rarely simple. Preference shares held by VCs sit ahead of common stock held by the founder. Liquidation preferences mean the founder's economic interest in a downside outcome is much smaller than the headline valuation suggests. Anti-dilution provisions, tag-along and drag-along rights, vesting cliffs, and ratchets all affect what the founder's equity is actually worth in different scenarios. The HMRC valuation team will, in principle, take all of this into account through "discount for lack of marketability" and minority-discount adjustments. In practice, the negotiation is contested and slow, and the outcome varies materially with the quality of advice the estate can afford. None of this is a problem in cash-generative private companies where the dividend stream is observable. It is a meaningful problem in venture-backed equity.

How the reform hits each part of the UK tech funding stack

Founders

The most directly affected group, but with the widest range of outcomes depending on company stage and exit trajectory. A pre-seed or seed-stage founder with paper equity below the £2.5 million allowance is unaffected in current terms; the question is whether the company will grow into the cap. A Series B or later founder with paper equity above

the cap is in scope. The death-event mechanism is operationally awkward for this cohort because the asset is illiquid, the valuation is contested, and the funding routes (instalments out of dividends, share buy-back from the company, life insurance) work differently for venture-backed companies than for cash-generative private businesses. The instalment option is meaningful — paying a £2 million IHT bill at £200,000 a year over ten years, interest-free, is more manageable than the same bill due immediately — but only if there is income to fund it. For many growth-stage founders, the company is reinvesting all its cash and dividends are not an option until exit.

The behavioural response from the founder cohort is the question every other part of this analysis depends on. The Office for Budget Responsibility's January 2025 costing of the related non-dom reform assumed 25 per cent of the most affected non-doms would leave the UK as a result of that reform. That figure is not directly transferable to the tech-founder cohort — different population, different reform, different mobility characteristics — but it is the closest published official assumption on UK high-net-worth behavioural response, and it is meaningfully large. Companies House data, reported by the Financial Times, shows 3,790 UK directors changed their primary residence to abroad between October 2024 and July 2025, a 40 per cent increase on the same period a year earlier. Most of those directors are not tech founders, and the period coincides with several other tax changes, but the named cases that have been individually reported — Nik Storonsky of Revolut (UAE residency registered); Herman Narula of Improbable (reported as preparing to emigrate, citing a mooted exit-tax rather than the BPR reform specifically); and others — are concentrated at the top of the UK tech founder cohort. Adviser surveys reported in Sifted in May 2025 cite four named UK tax-advisory firms confirming a marked uptick in tech-founder enquiries about UAE relocation, with one firm reporting that UAE relocation features in 15 to 20 per cent of all new business enquiries.

The widely-cited figure of 16,500 millionaires leaving the UK has been forensically debunked by Tax Policy Associates and Tax Justice Network and should not be cited as evidence by anyone serious about this question. The narrower findings above — the OBR's 25 per cent assumption (with the caveat that it is for a different population), the Companies House director-departure data (with the caveat about contamination), and the named tech-founder cases — are not subject to those objections. Something is happening; the scale is uncertain; the IHT reform is one of several pressures rather than the sole driver.

Angel investors

UK angel investors, including those investing through the Enterprise Investment Scheme (EIS), face a more complex picture than founders. EIS shares qualify for BPR if held for at least two years, so an EIS portfolio is potentially relievable on death — but only up to the new £2.5 million combined APR/BPR allowance. An angel with a portfolio of twenty EIS positions across ten years of investing, totalling £4 million in cost basis and perhaps £8 million in current value, is materially in scope. The reform changes EIS economics for the most active angels in two ways: it caps the inheritance tax shelter at £2.5 million, and it interacts with the existing EIS rules (income tax relief, CGT

exemption on disposal, loss relief on failures) in ways that the practitioner literature is still working through.

For active angels, the practical effect is that the EIS scheme remains attractive on the front end (income tax relief, CGT exemption, loss relief) but the back-end inheritance tax shelter is less generous than it was. Most angel portfolios are not concentrated enough in a single position to produce death-date valuation disputes of the kind founders face. The EIS structure also means investments are spread across many companies, smoothing the outcome distribution at the portfolio level. For the most active and most capitalised angels — the cohort that funds a meaningful share of UK seed-stage tech — the reform is meaningful but not existential.

The harder question is downstream. If founders relocate, the deal flow that angels see thins out. If the most experienced angels relocate, the seed-stage funding ecosystem loses both capital and the operator-investor expertise that makes UK angel money distinctively useful. The reform's first-order effect on angels is moderate. Its second-order effect — through the founder cohort it interacts with — is the question that matters more.

Venture capital and growth equity (LP interests)

Limited partner interests in venture and growth equity funds qualify for BPR if the fund is investing in qualifying trading companies, subject to the two-year holding period. For UK-resident individuals investing as LPs in UK venture funds, the interest is BPR-eligible up to the £2.5 million allowance. A high-net-worth UK-resident LP with £10 million committed across several funds is materially in scope above the cap.

The mechanism issue here is even sharper than for founders. An LP interest in a closed-ended venture fund is illiquid by structural design — there is no realisation event until the fund itself realises the underlying portfolio, which can be eight to twelve years after commitment. Death during the fund's life produces a contested valuation against an asset class with no public market price, no liquidity option, and no buy-back mechanism. The instalment option helps but does not solve the underlying problem: the heir owes tax against a paper valuation of an asset they cannot sell and from which they may not see distributions for years.

The behavioural risk in this cohort is that high-net-worth UK individuals reduce their LP commitments to UK venture funds, or relocate to jurisdictions where their LP interests are not exposed to UK IHT in the same way. The funds themselves are mobile — most large UK-resident LPs already invest across UK, US, and European funds — and the marginal commitment that goes elsewhere is the commitment most affected by the reform. The aggregate scale of this effect is hard to estimate without HMRC microdata, but the direction is clear.

Private equity in growth tech

Private equity investing in growth-stage UK tech companies — buyout-style PE, but also growth equity at later stages — interacts with the reform mainly through the personal

holdings of partners and the management equity in portfolio companies. The fund-level economics are not directly affected; the personal IHT exposure of UK-resident partners is. For the largest UK-resident PE partners, personal carry interests and direct co-investments above the £2.5 million allowance are in scope. The cohort is small but well-advised, and structural responses (trusts, holding companies, residence planning) are typically already in place. The reform's first-order effect is to accelerate planning that this cohort was largely already doing.

The effect on portfolio companies is indirect. Where founders of PE-backed UK tech companies have material personal stakes, those stakes are affected on the same terms as direct founder equity. The PE fund's interest in retaining UK founders for the duration of the hold period interacts with the founder's interest in not facing a punitive death-event tax during the hold. This is a private negotiation between PE sponsors and founders, and the reform changes its dynamics at the margins.

Early employees with vested equity

An often-overlooked group. UK tech employees who joined growth-stage companies early and now hold material vested equity above the £2.5 million allowance are in scope on the same terms as founders. This cohort includes several thousand people across the UK tech ecosystem — early engineers and operators at companies that have grown to unicorn or near-unicorn valuations. The death-event valuation problem is the same for them as for founders; the structural responses available to them (trusts, life insurance, residence planning) are similar but typically less sophisticated and less well-advised. For this cohort, the practical effect of the reform is that early-employee equity has become a more complex tax-planning question than it was, with downstream effects on how easy it is for UK tech companies to recruit and retain senior operators.

The three positions, in plain English

Three serious positions are being argued by people who have looked at this carefully.

Hold the system, fix the rough edges. Keep tax-at-death. Adopt practical fixes — closing avoidance routes, publishing valuation methodology, tightening rules on temporary non-residence, providing clearer guidance on company buy-backs. Resource HMRC's Shares and Assets Valuation team adequately. The cohort directly affected is small; mechanism change for one asset class would be a special carve-out for the wealthiest holders. The strongest version of this position is harder than the technocratic one: that the operational-mismatch argument is the same argument the wealthy have always made about every tax that affects them, that the historical record is that these arguments are usually overstated, and that the present argument should be treated the way comparable arguments have always been treated — with measured scepticism, while still adopting genuinely sensible operational improvements.

Change how the tax falls on this kind of asset. For unlisted UK trading-company shares above the allowance, replace tax-at-death with capital gains tax when the heir

actually sells. No base-cost uplift. Long-stop deemed disposal at year ten or fifteen. The principle is the same — tax on generational wealth transfer — but the mechanism collects against optimised exit values rather than contested death-date numbers, with much less behavioural pressure on the original founder to leave the UK before they die. Australia has run a comparable regime for forty years. There is a structural feature of this option that its supporters have not engaged with: if the long-stop is deferrable until actual liquidity, the state ends up holding contingent equity exposure across hundreds and over time thousands of UK private companies. Whether that is a good thing or a bad thing is a question the public debate has not asked.

Adopt the practical fixes, decide on the bigger question later. Do the four practical fixes everyone agrees on. Commission proper modelling. Decide on the mechanism question in two to three years when there is real evidence rather than projection. The strongest version commits the government, in advance, to switching to the second option if defined evidence thresholds are breached at a published review point. The weakness is that conditional frameworks of this kind tend to fail in practice — the trigger never quite fires, the data is always contested, and the regime drifts in its current form regardless of what the underlying empirics actually show.

What is already happening on the ground

The behavioural response started when the reform was announced in October 2024, eighteen months before it took effect. By April 2026 the most mobile founders had already made decisions — some left, some restructured, some put plans on hold to see what the final regime looked like. The £1 million-to-£2.5 million concession in December 2025 changed the calculation for some of them; the concession history matters because it tells us the regime is responsive to evidence, but also that responsiveness has limits.

The UK venture-funded tech ecosystem in early 2026 is in a particular state. Capital deployment is steady but cautious. Founders considering where to incorporate the next company are weighing UK alongside Delaware, Singapore, and the UAE more openly than they were two years ago. Senior operators considering where to take their next role are factoring inheritance tax exposure into compensation conversations in a way they did not need to before. Angels and VCs are watching their portfolio founders to see who relocates and who stays. None of this is catastrophic. All of it is real, and all of it is downstream of policy choices.

What we do not yet know — what no one yet knows — is the scale. The OBR's 25 per cent figure is the central case for one cohort; the realised behavioural response in the tech founder cohort specifically may be lower (because UK-anchoring matters for operating companies) or higher (because tech founders have particularly strong international optionality). HMRC modelling on the question has either not been done or has not been published. The next two to three years will produce the data. Until then, anyone arguing the regime is broken or arguing the regime is fine is arguing from inference rather than evidence.

Why this matters beyond the cohort directly affected

The UK has been explicit, across multiple administrations and industrial-strategy documents, that growing more large UK tech companies is a national economic priority. The reasoning is straightforward: large successful tech companies generate disproportionate productivity gains, anchor regional ecosystems, attract international capital, and produce the kind of skilled employment growth that traditional industrial policy struggles to deliver. The country wants more Revolut, more Improbable, more Wise, more DeepMinds. The country wants the next generation of those companies to stay UK-headquartered, UK-listed, UK-employing.

The IHT mechanism applied to founder equity is one of several policy choices that interact with that objective. It is not the most important one — the cost of capital, the listing environment, the talent pipeline, the regulatory regime, and the cost of energy all matter as much or more. But it is one of the choices, and it is one that the country can change without legislation that touches anyone outside the affected cohort.

The honest summary of the question is this: the reform is defensible in principle and the operational mechanism is contested for one specific asset class — the asset class held by the cohort the country has said it wants to grow more of. There is a defensible case for keeping the regime as it is. There is a defensible case for changing the mechanism. There is a defensible case for waiting for evidence. Reasonable people, working in good faith, will reach different views depending on which empirical outcomes they consider most likely, which normative framings they treat as primary, and how much weight they put on the country's stated industrial-strategy objectives versus the country's stated fairness-across-asset-classes objectives. The answer depends partly on what you think the regime is for, and that is a political choice rather than an empirical finding.

If you want to go deeper

The article in its tax-policy register and the policy options paper in HMT format provide the same analysis at higher technical depth. The longer plain English version walks through the same content in more analytical detail, including the international comparators (Australia, Canada, the United States, Germany, France, Japan, South Korea) and what each of their regimes can and cannot tell us about the UK choice. The methodology piece explains how this analysis was made, what AI tools did well and badly, and the disclosures the author thinks readers should weigh.

None of those pieces will tell you what to think. The author has a personal interest in the question and is honest about that. The work tries to give you the materials to think it through for yourself.

Source and corrections

Canonical version. This document is downloaded from thelongerlook.com. The article version of the same analysis sits at thelongerlook.com/articles/2026-04-30-uk-tech-iht-readable.html.

Corrections log. The publication maintains a public corrections page at thelongerlook.com/corrections.html, listing every correction the publication has made since launch with the date, what changed, and why. If you are reading a copy of this document that you downloaded some time ago, the corrections page will tell you whether anything in this document has been corrected since.

If you find an error. The publication invites corrections from specialist readers. AI cross-critique catches some errors and misses others; the errors AI cross-critique misses are exactly the ones a tax practitioner, IFS economist, CIOT or STEP member would catch on a careful read. If you find an error in this document, please send the correction; the publication will post it on the corrections page with attribution and update the canonical version.