

T H E L O N G E R L O O K
· A R T I C L E O N E ·

Inheritance Tax and the Companies the Country Says It Wants More Of

*What is being argued, what the disagreement turns on,
and what different evidence would mean.*

By Doug Scott

30 April 2026

thelongerlook.com

EDITORIAL · FROM THE PUBLISHER

What this piece is, and what it is not

The Longer Look exists for questions that public debate treats too quickly. It is not a newsletter, a feed, or a blog. Pieces appear when they are ready, and not before.

This is Article One. It concerns the reformed inheritance tax regime that took effect in April 2026, and specifically how that regime applies to one particular asset class — shares in unlisted UK trading companies. The regime is contested. Three serious positions have been advanced, and the public debate has been conducted at high volume on incomplete evidence.

The piece does not argue that any of the three positions is correct. It sets out what each is, what they share, what their differences turn on, what comparable jurisdictions actually do, and what different empirical outcomes would imply for which position the evidence ultimately supports. The intent is to help readers think clearly about a question they will form their own view on.

An earlier version of this article presented one framework — Position C with hard triggers toward Position B — as a recommendation. A reader noted that this was advocacy with the architecture of analysis. The reader was right. This version is rewritten in a different register: the framework is presented as one possible response to the current uncertainty, and a new section walks through what each plausible empirical outcome would mean.

Two notes before the argument

ON HOW THIS PIECE WAS MADE

This piece was written by Doug Scott working with four AI tools as builders.

The author is the architect: he defined the question, judged every draft, ran the substantive disagreements until they resolved, and rejected directions that did not survive scrutiny. Across the production of this piece, four large-language-model AI systems — Claude (Anthropic), ChatGPT (OpenAI), Grok (xAI) and Gemini (Google) — acted as builders, used across multiple parallel sessions, with each tool drafting text, modelling numbers, structuring arguments, finding evidence, and stress-testing the case from positions the author asked it to take.

The work is the author's in the sense that the substantive judgments are his. The writing is collaborative in the sense that it was produced through iterative work with AI tools rather than alone. The Longer Look will be transparent about authorship and AI involvement on every piece it publishes.

ON THE AUTHOR'S POSITION

The author is a UK-resident participant in the cohort whose tax position this paper analyses. Several of the directions canvassed in this paper, if adopted, would benefit him directly. He has written it because the policy debate is currently being conducted on incomplete evidence. The piece sets out the question, the open evidence, and what different empirical outcomes would imply — without arguing for one direction over the others. Readers should weigh this disclosure accordingly.

Summary

The reformed inheritance tax regime that took effect in April 2026 brings unlisted UK trading-company shares — founder equity, venture LP interests, EIS portfolios — into the inheritance tax base above a £2.5 million per-person allowance and £5 million combined allowance for couples. The reform is contested in a way most reform debates are not: not over whether the change should have happened, but over whether the mechanism the reform uses is operationally fit for the asset class it now applies to.

Three serious positions have been advanced by knowledgeable people in good faith. **Position A** holds that the existing mechanism is correct in principle and that operational issues should be resolved through four practical measures within the existing regime. **Position B** holds that the death-event valuation mechanism is mismatched to illiquid private-company shares and should be replaced for that asset class with a Capital Gains Tax charge on actual realisation by the heir. **Position C** holds that the four practical measures should be adopted now and the mechanism question deferred until the evidence on relocation, dispute caseload, and receipt outcomes can be assessed.

This piece does not argue that any of these positions is correct. It sets out what each is, what they share, what their differences turn on, what comparable jurisdictions actually do, and — most importantly — what different empirical outcomes from HMRC modelling and observed receipts data would imply for which position the evidence ultimately supports.

One framework — Position C with hard triggers toward Position B — is described in detail in Section 3 as a possible response to the current uncertainty. It is not the recommendation of this article. It is one workable design for acting under uncertainty, included so readers can see how a conditional commitment could be operationalised. Other designs are equally legitimate.

1. What is Actually in Dispute

The public debate over the IHT reform reads as a clash between irreconcilable positions: on one side, founder representatives arguing the reform will drive talent abroad and damage the UK's growth prospects; on the other, defenders of the reform arguing that any softening would be capitulation to the wealthy. Both framings overstate the disagreement.

Three serious positions have been advanced, and they agree on more than they disagree on. All three accept that the principle of the reform is correct. All three accept the £2.5m / £5m allowance structure. All three support adopting four practical measures: a collateral-trigger rule to close the buy-borrow-die avoidance vector; a statutory safe-harbour valuation methodology to reduce SAV dispute volume; targeted tightening of temporary non-residence rules to capture relocation; and practical guidance on structured buy-backs under existing instalment provisions for the cohort where they work. These four measures are common ground.

The disagreement is narrower. It is about whether those four measures are sufficient on their own — Position A's view — or whether they need to accompany a change to the underlying mechanism for one specific asset class — Position B's view — or whether the mechanism question should be deferred pending further evidence — Position C's view.

The three positions, on their proponents' best terms

Position A — Hold the existing mechanism. Maintain IHT-at-death as the taxable event for unlisted trading-company shares. Adopt the four practical measures. Resource SAV adequately to handle the dispute caseload. The principle of the reform — that very large illiquid private holdings should be brought into the inheritance tax base — was correct, and is delivered most cleanly by a regime that taxes wealth at the point of generational transfer. The affected cohort is small (HMRC forecasts approximately 220 BPR-only estates pay additional tax in 2026-27). Operational issues exist but are tractable. Mechanism change for one asset class would be a special carve-out for the wealthiest holders and would weaken the principle the reform is meant to embody.

Position B — Switch the mechanism to CGT on realisation. Replace IHT-at-death with Capital Gains Tax charged on actual proceeds when the heir realises the asset, with no base-cost uplift on death. Combine with a long-stop deemed-disposal rule and a collateral-trigger rule. The death-event mechanism produces problems no first-principles design would generate: a growing dispute caseload at SAV that cannot scale to industrial-strategy growth assumptions; behavioural pressure on the

affected cohort that leaks revenue through pre-death relocation; and an avoidance vector through collateralised borrowing that the existing regime does not close. The proposed mechanism collects the same principle against optimised exit values rather than contested death-date values, with lower administrative cost. Australia has operated a comparable rollover-to-realisation regime for forty years.

Position C — Hybrid: practical fixes only, mechanism question deferred.

Adopt the four practical measures. Defer any mechanism change pending evidence on the size of the operational issues over the following two to three years. The four measures are independently valuable and unobjectionable — they would be adopted under any of the three positions. Mechanism change carries political risk and design risk that the practical measures do not. The posture commits Treasury to data-driven review rather than to a specific outcome that prejudices the empirical questions.

1.5. A Note on Timing

The behavioural response to the announced reforms began before the reforms took effect. This is observable in the published evidence, and any analysis or framework that treats April 2026 as the start of the response — including the trigger criteria proposed in Section 3 — is calibrated against the wrong baseline.

The official UK government assumption

The Office for Budget Responsibility's January 2025 supplementary forecast on the non-dom reform sets out the official behavioural assumption underpinning the £33.9 billion the reforms are projected to raise: the OBR expects approximately 25 per cent of non-doms with excluded property trusts (the wealthiest cohort, most directly comparable to the population the IHT reform also affects) to leave the UK as a result of the reform, alongside 12 per cent of other non-doms. These are not lobby figures. They are the central behavioural assumptions in the official UK government costing.

Companies House evidence on directors leaving

Financial Times analysis of Companies House records, published in 2025, recorded 3,790 UK company directors moving their primary residence abroad between October 2024 and July 2025 — a 40 per cent increase on the 2,712 who moved in the same period a year earlier. April 2025 alone saw 691 director departures, 79 per cent above April 2024 and 104 per cent above April 2023. Approximately 150 UK-based directors moved specifically to the UAE between April and June 2025.

Sector and size profile of the named cases

Sector and size data are partial. The departures that are individually reportable and have been verified by name include several at the top of the UK technology cohort: Nik Storonsky, founder of Revolut; Herman Narula, co-founder of Improbable (a company most recently valued at approximately £700 million); and a number of named non-tech billionaires. Trade-press reporting from Sifted (May 2025) cites four UK tax-advisory firms — Wilson Partners, Evelyn Partners, Founders Law, and Capital Partners — confirming a marked uptick in UK-based tech-founder enquiries about Dubai relocation. Founders Law specifically reported that UAE relocation now features in 15–20 per cent of all new business enquiries.

On the contested figures

The widely-cited Henley & Partners projection of 16,500 UK millionaire departures in 2025 should be cited with caution. Tax Policy Associates published a detailed

forensic critique in July 2025 — examining methodology changes, statistical anomalies in the data, and contradictions with official UK figures — and concluded that the Henley reports should be treated as marketing material rather than evidence. Tax Justice Network reached similar conclusions independently. Position A's defenders are right to be sceptical of the headline 16,500 figure. The narrower, harder-data findings — the OBR's own 25 per cent assumption, the Companies House director-departure trend, and the named tech-founder cases — are not similarly contested.

What this means for the rest of the article

The framework in Section 3 commissions HMRC modelling on the relocation channel; the four scenarios in Section 5 describe what that modelling could show. Both treat the data as forward-looking. The honest reading is that significant pre-positioning has already occurred. The trigger thresholds in Section 3 — even taken as illustrative — are calibrated against a population that has already partly responded to the announced reform. So is any modelling HMRC will subsequently produce. The early-mover cohort has already, at least in part, exited the dataset. This means measured behavioural response will understate the true response, possibly by a wide margin. Any framework that does not engage with this — including the one in Section 3 — is incomplete in a way that matters for the policy choice.

2. What the Disagreement Actually Turns On

The three positions cannot all be right. They rest on differing answers to a small number of questions. If those questions could be answered with rigour, the disagreement would resolve. The questions are partly empirical and partly normative, and they have not been answered with rigour.

QUESTION 1

How large is the relocation channel under the existing mechanism?

If the proportion of the affected cohort that relocates pre-death is small, Position A's revenue forecast is robust and the mechanism issue is overstated. If the proportion is large, the existing mechanism leaks materially and the case for change strengthens. The question is empirical and answerable through dynamic behavioural modelling using HMRC data. It has not been answered with rigour.

QUESTION 2

How quickly does SAV dispute caseload grow as the cohort grows?

If the caseload scales linearly with affected estates and SAV can be resourced to keep pace, the operational problem is solvable through capacity expansion. If the caseload scales superlinearly, resourcing alone cannot keep pace. The question depends partly on the rate of cohort growth — which Government industrial strategy is actively trying to accelerate — and partly on whether disputes resolve faster or slower as the safe-harbour methodology beds in.

QUESTION 3

How effective is the collateral-trigger rule against sophisticated structuring?

If the rule cleanly captures economic realisation through borrowing, the buy-borrow-die vector is closed under any of the three positions. If sophisticated advisers can route around the rule through structures the legislation does not capture, the rule's effectiveness is overstated. The question is technical and depends on drafting quality.

QUESTION 4

Can a 10- or 15-year long-stop deemed disposal be administered without recreating the original problem?

Position B's long-stop assumes a deemed disposal at year 10 or 15 produces an administrable computation against an observable value. If the asset remains illiquid at the long-stop date, the long-stop may recreate the same valuation and liquidity problem the original mechanism produced. The honest reading: the long-stop is

structurally weaker than its proponents acknowledge, but closer to acceptable when paired with a payment-on-realisation provision.

QUESTION 5

What is the regime for?

This is the question that cannot be settled by modelling. If the primary purpose of the reform is revenue, the position that maximises receipts net of leakage and administrative cost wins. If the primary purpose is fairness across asset classes — that all wealth above the allowance is taxed at the point of generational transfer regardless of liquidity — Position A is uniquely consistent. If the primary purpose is to support industrial strategy by retaining the cohort that builds growth-stage UK companies, Position B is most consistent. The question is normative, the answer depends on which objective is treated as primary when they conflict, and ministers — not analysts — should answer it.

Reasonable, informed, good-faith readers reach different conclusions on these five questions. The disagreement between the three positions is largely a disagreement about how the questions should be answered, weighted, and prioritised. A reader who reaches a clear view on the questions will, by implication, reach a clear view on the policy.

3. One Possible Framework for Acting Under Uncertainty

NOT A RECOMMENDATION

This section describes one framework — Position C with hard triggers toward Position B — that some commentators have proposed as a way of acting under the current uncertainty. It is set out here so readers can see how a conditional commitment could be operationalised, not as a recommendation. Other designs are equally legitimate; some are noted at the end of the section.

The framework's core move is to commit to the four practical measures immediately, commission HMRC modelling that would answer the open empirical questions, and pre-commit to a mechanism switch if defined evidence thresholds are breached at a published review point. The thresholds make the trigger operational rather than rhetorical.

Step 1. Adopt the four practical measures within ninety days

- **Collateral-trigger rule.** Use of inherited unlisted trading-company shares as collateral for borrowing above 25 per cent of the asset's assessed value triggers a deemed disposal proportionate to the encumbered share.
- **Statutory SAV safe-harbour methodology.** HMRC publishes a statutory valuation methodology covering preference-stack treatment, minority discount ranges, DLOM ranges by company stage, and aggregation under IHTA s.161. Estates that adopt the safe-harbour formula receive fast-tracked SAV approval.
- **Targeted temporary-non-residence tightening.** Extends the qualifying TNR period from five to ten years for individuals whose UK-acquired assets exceed a defined threshold.
- **Structured buy-back guidance.** HMRC and HMT publish guidance on the use of company purchase of own shares under CTA 2010 s.1033 to fund IHT instalments out of operating cash flow. Works for cash-generative private companies; does not work for early-stage venture-backed holdings or LP interests in closed-ended funds.

Step 2. Commission HMRC dynamic modelling within thirty days

HMRC's behavioural team is asked to produce dynamic forecasts on three specific questions: the relocation elasticity of the affected cohort, segmented by company

stage and holder age; the projected SAV caseload over a five- to ten-year horizon; and the realistic distribution of realisation events that would crystallise CGT under a Position B mechanism.

Step 3. Set a twelve-month formal review with calibrated triggers

If any two trigger criteria breach defined thresholds, the mechanism switch (Position B legislation drafted under Step 4) takes effect at the next available Finance Bill.

Trigger criterion	Threshold
1. Relocation evidence	Documented departures by primary residence change among the affected cohort exceed 60 in the 12-month review window.
2. SAV dispute caseload	Pending SAV caseload exceeds 350 active files, with average resolution time above 24 months.
3. Receipts underperformance	Audited BPR-only estate IHT receipts fall more than 20 per cent below Budget forecast in the first full operational year.
4. Adviser-survey evidence	More than 35 per cent of advised affected estates report active relocation planning underway in a structured anonymous survey.
5. Valuation discounts	Settled SAV cases show average final-assessed value at less than 65 per cent of opening assessment.

These thresholds are illustrative. They are calibrated against comparator data — Canadian deemed-disposal regime caseload, observable UK private-company exit timings, OECD data on high-net-worth mobility — but Treasury modelling could justify adjustments.

Step 4. Draft the alternative legislation in parallel

Treasury legal drafts the realisation-based legislation alongside the Step 1 measures, so that if the trigger fires the legislation is ready for the next Finance Bill rather than commissioned at that point. Without this step the trigger is effectively two years rather than twelve months from review to implementation.

Other designs are equally legitimate

Position A's defenders would argue, with force, that no trigger is needed: the four measures should be adopted, the regime should run for two to three years, and the question should be revisited only if observed data subsequently demands it. Position B's defenders would argue, with equal force, that the mechanism switch should happen now, on the available evidence, without waiting for confirmation that may arrive too late to retain the cohort. Each design is internally coherent. Each has consequences explored in Section 5.

4. International Comparators

The UK's chosen combination — a £2.5m / £5m allowance, 50 per cent relief above, death-event valuation, no business-specific conditional relief above the threshold, no realisation route — is unusual in international terms. This section sets out what comparable jurisdictions actually do. It is not an argument that the UK should copy any of them. It is a check on the implicit claim that the UK's approach is the natural or default treatment.

Australia · No inheritance tax; CGT rollover at death

Australia does not levy an inheritance tax. When an Australian resident dies, capital gains tax is not triggered at the point of death. Instead, the heir inherits the deceased's cost base, and CGT crystallises only when the heir later disposes of the asset. This is the closest international precedent for what Position B proposes. The Australian regime has been operating for forty years. What the comparison does not show: Australia raises substantially less from intergenerational wealth transfer than the UK does.

Canada · No inheritance tax, but deemed disposition at death

Canada has no inheritance tax in the formal sense, but the Canada Revenue Agency treats the deceased as having sold all capital property at fair market value immediately before death — the "deemed disposition" rule. CGT crystallises on the deceased's final tax return. The heir then inherits the asset at the stepped-up cost base. This is structurally the opposite of Position B. The capital gains inclusion rate is 50 per cent for the first \$250,000 of annual gains and 66.67 per cent above; combined with provincial top marginal rates, the effective tax on large unrealised gains at death can reach 35 per cent or more.

United States · High exemption with step-up in basis

The US position changed materially under the One Big Beautiful Bill Act of July 2025. Effective 1 January 2026, the federal estate tax exemption is \$15 million per person and \$30 million per couple, made permanent and indexed to inflation. The top federal estate tax rate above the exemption remains 40 per cent. The step-up in basis on death survives unchanged. In practical terms, US federal estate tax exposure now begins at a level (approximately £11–12m equivalent per person) far above the UK's new £2.5m allowance.

Germany · Inheritance tax with conditional business relief

Germany operates a recipient-taxed inheritance tax with personal allowances per beneficiary (€500,000 for a spouse, €400,000 for a child, €200,000 for a grandchild). The mechanism that matters for the UK comparison is the conditional business-asset relief under §§13a and 13b ErbStG.

- **Standard relief: 85 per cent exemption.** Conditional on a 5-year hold and cumulative payroll over the period of at least 400 per cent of the original annual payroll.
- **Optional full relief: 100 per cent exemption.** Conditional on a 7-year hold, cumulative payroll of 700 per cent, and a maximum 10 per cent share of administrative assets.

Germany is the clearest example of a regime that combines a death-event inheritance tax with structural carve-outs for genuinely-operating family businesses.

France · Pacte Dutreil — 75 per cent exemption with retention conditions

France's Pacte Dutreil regime provides a 75 per cent exemption from transfer taxes on qualifying business shares. Conditions include a collective undertaking to retain the shares for at least 2 years before transfer, an individual undertaking by each heir to retain the transferred shares for a further 6 years (extended from 4 by the 2026 Finance Act), and a management role exercised by at least one heir for 3 years. Combined with €100,000 allowances per parent per child every 15 years, effective transfer tax on a French family-business succession can fall to around 11-12 per cent.

Switzerland and the UAE · No inheritance tax

Switzerland has no federal inheritance tax; cantonal taxes vary, but most cantons exempt direct descendants entirely. The UAE has no inheritance tax. Both are frequently cited as destinations for UK relocators in the affected cohort. They are relevant to Question 1 (the relocation channel) rather than to the mechanism question.

What this comparison actually shows

The UK's post-April-2026 position is unusual. Among major developed economies, no peer jurisdiction combines all of: a relatively low exemption, full death-event valuation, no realisation route, and no substantial conditional relief for genuinely-operating businesses above the threshold.

- The **US** takes the high-exemption-with-step-up route.
- **Germany and France** take the conditional-relief route.
- **Australia** takes the realisation-based route.
- **Canada** takes the death-event-with-stepped-up-basis route.

The UK after April 2026 takes a route none of these jurisdictions takes. That does not make the UK regime wrong — there is room to argue the combination is novel rather than mistaken — but it does mean the regime cannot rely on international precedent for either the principle or the mechanism. Position A's claim that the reform is straightforwardly correct in design has to do its own work; the international evidence does not do it for them. Equally, Position B's claim that "this is what every other country does" is overstated; what Position B proposes is what Australia does, but Australia is one model among several.

5. What Different Evidence Would Mean

The five questions in Section 2 have answers that are partly knowable and partly contestable. This section walks through what each plausible evidence outcome would imply for the policy question, and what it would mean for the affected groups. The intent is to show readers the consequences of each branch — so they can form a view about which outcome they consider most likely, and which they would consider acceptable — rather than to predict which branch the data will land on.

The reader should hold the timing point from Section 1.5 in mind throughout this section. The four scenarios below describe what HMRC's modelling could show. They will measure a population that has already partly responded — the Companies House data and the OBR's own assumptions both indicate this — and so even the most pessimistic scenario likely understates the true behavioural response by some margin. The scenarios remain useful for thinking about the policy choice, but readers should treat the "small" scenarios with particular caution: a small measured response is consistent with either a genuinely small underlying response or a large response that has already been absorbed into the pre-reform departures.

Four scenarios are described. They are not exhaustive. They are the four that the available data and the reasonable range of expert opinion suggest are most worth considering.

Scenario one · The relocation channel is small and SAV caseload manageable

If HMRC's behavioural modelling shows annual departures from the affected cohort below roughly 30 in steady state, SAV caseload growth tracking estate-count growth linearly, and receipts within Budget forecast in year one — Position A is broadly vindicated. The four practical measures, plus adequate SAV resourcing, are sufficient. Mechanism change becomes unnecessary.

What this would mean. For Treasury, the reform delivers as designed; the £5m couple allowance is robust; the principle of fairness across asset classes is preserved. For the affected cohort, the regime is more navigable than the founder-lobby framing suggested. For industrial strategy, the regime is neutral-to-positive. The political cost of the original reform turns out to have been front-loaded — the announced response was sharper than the realised response.

The honest qualification. Even under this scenario, the venture-stage cohort (pre-revenue companies, illiquid LP interests) is not fully solved by the four practical measures. A small group continues to face genuine operational difficulty. That group

is too small to drive policy change but real enough that individual cases will surface periodically.

Scenario two · The relocation channel is meaningful but bounded

If HMRC modelling shows annual departures in the 30–80 range, SAV caseload growing somewhat faster than estate count but not unmanageably, receipts modestly below forecast (5–15 per cent) — the picture is mixed. The reform is leaking some revenue and losing some of the cohort, but not catastrophically.

What this would mean. For Treasury, the question becomes harder. Holding the regime delivers most of the revenue but loses some; switching the mechanism may collect more but at the political cost of being seen to capitulate. The choice is genuinely contested on the empirics. For the affected cohort, the regime is operationally workable for most but pressure persists for the segment most exposed to the mechanism mismatch. Some structural restructuring becomes more common as advised practice. For industrial strategy, the cost is real but bounded — the UK retains most of its founder cohort but gives up some at the margin. Whether this is acceptable depends on Question 5.

The honest qualification. This is the scenario in which design choices matter most, and in which conditional frameworks like the one in Section 3 are most likely to be useful — because the evidence does not clearly support either Position A or Position B.

Scenario three · Substantial capital flight and operational pressure

If HMRC modelling shows annual departures above 80, SAV caseload growing superlinearly with multi-year disputes the norm, receipts more than 20 per cent below forecast, and adviser-survey evidence of widespread relocation planning — the regime is in trouble. Position B's case strengthens substantially.

What this would mean. For Treasury, the choice becomes politically difficult but empirically clearer. The regime in its current form is not collecting what was forecast and is producing a measurable behavioural response. Mechanism change to a realisation basis would likely collect more but at the political price of being seen to retreat under lobbying pressure. For the affected cohort, the calculation that drove the relocation has been validated; the cohort that left does not return. For industrial strategy, the damage is significant and persistent. For the long-term tax base, mechanism change recovers some receipts going forward, but the lost cohort is lost.

The principle of fairness across asset classes is sustained nominally but undermined operationally.

The honest qualification. Even in this scenario, switching to Position B is not costless. Long-stop deemed-disposal mechanics need careful design or they recreate the original liquidity problem at the long-stop date. No mechanism is friction-free for every member of every cohort.

Scenario four · The data is contested or delayed

The likeliest real-world outcome, on the historical record of UK tax policy reviews, is that the modelling produces a mixed picture, that the trigger thresholds in any conditional framework are challenged on definition, that the political environment by month 12 differs from the one in which the framework was set, and that the formal review does not produce a clean fire-or-not answer.

What this would mean. For any conditional framework, credibility depends on the trigger firing under conditions advocates will agree have been met. If the data is genuinely contested, the trigger becomes another consultation rather than a forcing function. The political room that the framework was meant to bind closes; ministers face the choice without the cover the framework was intended to provide. For Treasury, this is the scenario most likely to result in the regime drifting in its current form for several years before any further change. For long-term policy quality, this is the worst outcome: the regime persists not because it is right but because the political conditions for changing it never quite cohere.

The honest qualification. This scenario is what makes Position A and Position B's defenders most uncomfortable with conditional frameworks, in different ways. Position A's defenders fear the trigger will fire under contested evidence and produce a hasty mechanism change. Position B's defenders fear the trigger will fail to fire under supportive evidence and entrench an inadequate regime. Both fears are reasonable. The honest answer is that conditional frameworks are not robust to data-quality failure.

6. The Limits of This Analysis

Three limits are worth naming honestly.

First, the analysis underplays strategic behaviour. Sophisticated holders of unlisted trading-company shares do not simply respond to mechanism design — they restructure to avoid the question entirely. Trust structures established years before death, holding-company restructures that move shares out of personal estate, life-insurance funding of any residual liability, and residence planning that uses the TNR window strategically rather than passively. The analysis above treats the cohort as more responsive to policy and less strategic than it actually is. The Companies House director-departure data, the Sifted adviser-survey reporting, and the OBR's own 25 per cent behavioural assumption for excluded-property-trust non-doms all point in the same direction: the most mobile and most advised members of the affected cohort are not waiting for the regime to bite before responding. The evidence on relocation, valuation discounts, and adviser-survey planning is therefore likely to understate the true behavioural response, particularly for the largest holders most able to afford bespoke planning.

Second, the analysis is asymmetric in its quantification, and the data it would rely on is partial. The friction side of the regime — SAV caseload, dispute resolution time, valuation discounts — gets specific numbers and trigger thresholds. The outcome side — relocation, restructuring volume, capital flight, long-term reinvestment — gets qualitative treatment. This asymmetry exists because the friction data is collectable and the outcome data is not, but the asymmetry is real. Worse, the outcome data that does exist measures a population that has already partly responded: as Section 1.5 notes, the Companies House director-departure data already shows a 40 per cent year-on-year increase between October 2024 and July 2025, before the IHT reform took effect. Any HMRC modelling commissioned now will measure the cohort that remained; the cohort that pre-positioned is already outside the dataset. The reader should adjust their weighting of any quantitative finding accordingly.

Third, the normative question is not settled by any analysis. Whether the regime's primary purpose is revenue, fairness across asset classes, or industrial-strategy alignment is a political choice, not an empirical finding. The four scenarios above describe consequences for each of these objectives separately, but they cannot tell the reader which objective should take precedence when they conflict. That is a matter for ministers, voters, and the deliberative judgment of the people most directly affected.

Closing

The reformed inheritance tax regime is the right reform in principle in the view of all three serious positions. The mechanism for one specific asset class is contested, on grounds that turn on five questions whose answers are not fully knowable today. Three positions and several conditional frameworks have been proposed; each rests on different empirical assumptions and different normative priorities; each has identifiable consequences under each plausible empirical outcome.

This piece has tried to set out the question, the open evidence, and what different outcomes would imply, without arguing for any of them. Readers will reach different conclusions depending on which empirical outcomes they consider most likely and which normative framing they treat as primary. That is the nature of policy questions answered under genuine uncertainty.

What would help the question resolve, in any direction, is the modelling that has not yet been done and the data that the regime's first operational year will produce. Until then, the honest position is that several reasonable people, working in good faith with the same evidence, will reach different conclusions about the right answer. That is not a failure of analysis; it is what the question actually looks like.

POSTSCRIPT · A NOTE ON THIS VERSION

An earlier version of this article presented one framework — Position C with hard triggers toward Position B — as a recommendation. A reader noted that this was advocacy with the architecture of analysis rather than analysis itself. The reader was right.

This version is rewritten in a different register: the framework is presented as one possible response to the current uncertainty rather than as the right answer, and a new section on consequences makes the stakes of each empirical outcome explicit rather than implicit. Where the earlier version asked the reader to accept a recommendation, this version aims to give the reader the materials to form their own view.

Two further points the original critique raised, which are not addressed in this version because they require more space than a postscript: the article assumes the principle of taxing intergenerational business wealth without defending that assumption explicitly; and it focuses on friction more than on outcomes. Both will be the subject of later pieces.

— Doug Scott. This version published in place of an earlier recommendation-format version, in response to a reader's substantive critique.

ABOUT

The author. Doug Scott is a UK-resident participant in the cohort whose tax position this paper analyses. His other work appears at the sites listed below.

The publication. The Longer Look exists for questions that public debate treats too quickly. Pieces appear when they are ready. The Longer Look will be transparent about authorship and AI involvement on every piece it publishes.

Method. This piece was written by Doug Scott working with Claude, ChatGPT, Grok and Gemini as builders, used across multiple parallel sessions. The substantive judgments are the author's. The writing is collaborative.

Licence. This work is published under a Creative Commons Attribution-NonCommercial 4.0 International licence (CC BY-NC 4.0). Share it. Translate it. Print it. Just credit Doug Scott and don't sell it for profit.

OTHER WORK BY THE AUTHOR

THE TRILOGY · APRIL 2026

If This Road · ifthisroad.com — *The wake. A quiet walk through what is shifting.*

Orphans · orphans.ai — *The diagnosis. Twenty-five years of unrecorded work.*

The Held · theheld.ai — *The disposition. The working relationship between a person and a machine.*

AND SEPARATELY

The Many Builders · themanymbuilders.com — *A place where the names of the people building AI sit alongside the names of those who held the rest.*

The Bear Was Right · thebearwasright.com — *A small picture book.*

The Bear Loved · thebearloved.com — *A small book in twenty-three days.*

By Doug Scott

thelongerlook.com · Article One · April 2026

© Doug Scott 2026 · Licensed under Creative Commons BY-NC 4.0