A Constitutional Model for the Persistent Taxation of Concentrated Asset Holdings

Stress-Testing Pass: Does the Architecture Survive Contact with Reality?

Drew Griffiths

July 2026

This is the companion stress-testing pass to the architecture paper. Its purpose is not to defend the model but to try to break it, and to report honestly what breaks. Each test is classified on a four-way scale: ***fatal*** *(the model cannot work as designed),* ***serious but addressable*** *(a real problem with a known fix that must be adopted),* ***manageable*** *(a cost or limit to be accepted and disclosed), or* ***unresolved*** *(genuinely open, needing work this pass cannot complete). The two tests most likely to be fatal — revenue sufficiency and human-rights durability — are taken first, because if either failed, the rest of the pass would be beside the point. Neither returned fatal. The headline finding is that, on the evidence available to this review, no fatal defect has been identified. The model is not presently disqualified on revenue or on law, but its viability remains conditional upon specific design safeguards identified below and upon external work — a government costing and counsel's opinion — that this pass cannot complete. That no fatal defect has been found on the available evidence is a more modest statement than that none exists, and the document is careful to keep to it.*

Anyone asked to support this proposal will ask two questions before any other: will it raise the money, and will it survive legal challenge. If the answer to either is no, nothing else matters. This pass therefore leads with those two, states its findings in the harshest defensible terms, and only then proceeds to the secondary tests. Throughout, the burden of proof is placed on the model: a test is not passed because the model can be defended, but because a serious critic would concede the point.

A word on method. This pass uses the best publicly available UK data — HM Revenue and Customs receipts, the Wealth Tax Commission's base estimates, and the European Court of Human Rights case-law on taxation and property. It does not have access to the confidential taxpayer-level data on which a government costing would rest. Where a finding depends on data only government holds, that is stated, and the finding is marked accordingly. No number in this pass should be treated as a costing; the numbers are order-of-magnitude tests of whether the model is in the right region, not forecasts.

> **Confidence and limits of this pass, stated once and plainly.** Three things a reader relying on this document must know up front. *First*, the revenue figures rest on soft data: the count of people above £10m is reasonably robust, but the aggregate wealth they hold is under-captured at the very top by the survey sources it derives from, and the base figure used here is a consistency check against advocacy estimates, not an independent measurement — a circularity disclosed in Section 2. *Second*, the entity-layer contribution to revenue is directionally positive but *unquantified* in this pass; the public data does not permit sizing it, and no figure for it should be relied upon. *Third*, the legal findings are conditional on safeguards the architecture paper does not yet contain: the accurate statement is that reciprocal purchase is *capable of being made* compliant, not that it is compliant as currently drafted. Addressable-in-principle is not addressed-in-fact. This pass identifies conditions; it does not certify that they have been met.

1. Test 1 — Revenue: does it replace what it abolishes?

Classification: serious but addressable. The model can clear the replacement hurdle, but only at a rate at or above roughly 1.5 per cent, and only if the base estimate

holds after behavioural response. At 1 per cent on the personal base alone it does not clear the hurdle.

1.1 The hurdle

The charge is proposed as a replacement for Capital Gains Tax (CGT) and Inheritance Tax (IHT). The replacement hurdle is therefore the combined receipts of those two taxes, which must be read as a rising multi-year band rather than a single figure:

- CGT: approximately £14.5bn in 2023–24, falling to around £13bn in 2024–25. The fall despite rate rises is itself evidence of the realisation-based lock-in the model is designed to remove.
- IHT: £7.5bn in 2023–24, £8.2bn in 2024–25, and a record £8.5bn in 2025–26, forecast by the Office for Budget Responsibility to keep rising as thresholds stay frozen.

The combined hurdle is therefore roughly **£21–23bn and trending upward**. A replacement that merely matched the low end in year one would fall behind within a few years as IHT climbs. The model must aim to clear the upper end and grow with it.

1.2 The base

The base rests on two anchors of very different strength, and honesty requires separating them.

The *first* anchor is the **population count**: independent estimates converge on roughly **20,000 to 32,000 individuals** with personal net wealth above £10m. This figure is reasonably robust — several sources reach the same region by different routes — and can be leaned on.

The *second* anchor is the **aggregate wealth** that population holds, and it is much softer. The Wealth Tax Commission and subsequent analyses imply that the excess of this group's wealth above the £10m threshold aggregates to approximately £1.0–1.2 trillion. But this figure is derived by back-solving from advocacy estimates — that a charge of around 2 per cent on the excess would raise about £24bn implies an excess base of about £1.2 trillion — and those estimates rest partly on survey data known to under-capture wealth at the very top. **Using a figure in that region to then test this model is therefore a consistency check against the advocacy numbers, not an independent measurement.** The circularity is disclosed rather than hidden: the true base could be materially larger (top-end under-capture would mean more wealth than the surveys show, helping revenue) or the effective base materially smaller (behavioural response and valuation disputes could shrink it, hurting revenue). No weight should be placed on the precise trillion figure; the honest claim is only that the base is of an order that makes the model worth testing, not that it is known.

The proposed model is a **cliff**, not an allowance: once the threshold is crossed, the full averaged value is charged, not merely the excess. This adds back the first £10m for each liable person — roughly £200–320bn across the population, this increment resting on the more robust population count — giving a *personal* cliff base of the order of £1.2–1.5 trillion, subject to the same disclosed softness in the underlying wealth figure, and before the entity layer is counted at all.

1.3 The first-cut arithmetic

On a mid-estimate personal cliff base of about £1.3 trillion, before behavioural response:

- at 1.0 per cent: gross £13bn
- at 1.5 per cent: gross £19.5bn

- at 2.0 per cent: gross £26bn

Applying honest behavioural haircuts — 10 per cent (optimistic), 20 per cent (central, the Wealth Tax Commission's working assumption), and 35 per cent (pessimistic, reflecting a first-of-its-kind annual charge with no exemptions) — gives net figures of:

- at 1.0 per cent: £8.5–11.7bn net — **does not clear the hurdle**
- at 1.5 per cent: £12.7–17.6bn net — clears it only under the optimistic haircut, and only on the personal base
- at 2.0 per cent: £16.9–23.4bn net — clears the hurdle under optimistic-to-central haircuts on the personal base alone

1.4 The entity layer is the model's revenue advantage — and its accounting hazard

The figures above count only personally-held wealth. The model's distinctive feature is that it also charges companies, trusts, and other liable persons independently. That entity layer is additional base not captured by conventional personal wealth taxes, and it is the reason the model can, in principle, out-raise them. This is a genuine structural advantage.

But it cannot simply be added to the personal figure, and honesty requires saying so plainly. Where a shareholder's shares and the company's underlying assets are both charged, the same economic value appears at both levels. How much *net* revenue the entity layer adds therefore depends entirely on the interaction rule — the very question the architecture paper left, correctly, to Parliament:

- under **full stacking**, the entity layer adds a substantial increment, because the same value is charged at both levels — but at the cost of taxing that value twice in full, which is where the human-rights proportionality question (Test 2) bites hardest.
- under a **credit or offset**, related value is charged once at the higher level, the entity layer adds little, and revenue rests mainly on the personal base — keeping the required rate toward 1.5–2 per cent.

A claim in an earlier draft of this pass — that full stacking "comfortably clears the hurdle at 1 per cent" — is withdrawn as unestablished. Sizing the entity layer requires knowing how much of the £10m-plus personal wealth is held *through* UK-resident companies and trusts rather than directly, at a granularity the public data does not provide. The honest statement is therefore that the entity layer's net contribution is **directionally positive but unquantified in this pass**, and no revenue figure should rest on it. What can be said is the qualitative point that follows.

This is the single most important interaction in the whole model between the revenue question and the legal question: *the rate the model needs and the legal risk it runs move in opposite directions as the interaction rule changes.* Full stacking raises more and raises proportionality risk; crediting lowers legal risk and lowers the entity layer's revenue contribution. This trade-off should be stated openly to anyone considering the model, because it cannot be designed away — and, as Section 3 notes, it is not the only place where a revenue dial is also a legal-risk dial.

1.5 Honest verdict on revenue

The model is in the right region. On the personal base alone, at a rate of 1.5–2 per cent, it plausibly matches or exceeds the £21–23bn hurdle under central behavioural assumptions, and the entity layer provides genuine additional headroom. It is *not* a model that obviously loses money, unlike rate-aligned CGT reform. But three cautions are binding:

1. at 1 per cent it does not clear the hurdle on the personal base, so the "1 per cent" figure common in wealth-tax advocacy is too low for a *replacement* tax.
2. the base estimate is the Wealth Tax Commission's, built partly on survey data that is known to under-capture the very top; the true base could be larger (helping) or the behavioural response larger (hurting).
3. only government, with taxpayer-level data, can convert this order-of-magnitude test into a costing. This pass can say the model is not disqualified on revenue; it cannot say it will raise a specific sum.

2. Test 2 — Human-rights durability: does the charge, and reciprocal purchase, survive A1P1 and Article 6?

Classification: serious but addressable, with one component (reciprocal purchase) that is conditionally survivable only if specific safeguards from the case-law are built in. Without those safeguards it is fatal to that mechanism; with them it is defensible.

2.1 The favourable baseline

The starting position is more favourable than instinct suggests. Under Article 1 of Protocol 1 (A1P1, protection of property), the European Court of Human Rights affords the legislature a *wide margin of appreciation* in tax matters. Few cases succeed under Article 1 in the tax field; where they do, there is an element of clear arbitrariness, and they have almost all arisen in the *enforcement* field rather than the imposition of a charge. There is no general principle against even retrospective taxation. An annual wealth charge, as such, is well within the margin the Court allows. So the *existence* of the charge is not the problem.

2.2 The operative constraint: individual and excessive burden

The binding A1P1 constraint is the requirement of a *fair balance* between the general interest and the individual's rights, breached where a measure imposes an "individual and excessive burden." This is where two features of the model — neither of them the charge itself — attract risk:

- **The cliff.** Charging the full value once the threshold is crossed, rather than the excess, produces a discontinuity that could, at the margin, be argued to impose an excessive individual burden on someone just over the line. This is manageable (see Test 3) but should be understood as the model's most exposed ordinary feature.
- **The illiquid holder.** A person whose averaged wealth exceeds the threshold but who lacks the cash to pay — the private-company or farm owner — is the classic "excessive burden" claimant if forced to sell to pay an annual charge. The architecture's answer (liability determined normally, but payment deferrable with interest) is the correct one and is close to essential for A1P1 compliance; it should be treated not as an optional kindness but as a durability requirement.

2.3 Reciprocal purchase and the Hentrich problem

The sharpest legal finding of this pass concerns reciprocal purchase, and it comes from *Hentrich v France* (1994). France operated a revenue right of pre-emption: where the authorities considered a declared sale price too low, they could buy the property at the declared price. The Court found a violation of A1P1 — and a linked violation of Article 6 (fair trial) — but the *reasons* are what matter, because they map precisely onto the model's design choices.

The violation was *not* that state pre-emption at a declared value is inherently incompatible with property rights. It was that:

1. the applicant was a **"selected victim"** of a *discretionary* power applied to her individually and arbitrarily; and

2. she had **no effective opportunity to challenge** the measure — the procedure did not let her contest the authorities' view or present her case on equal terms.

The Court expressly said the burden "could have been rendered legitimate only if she had had the possibility — which was refused her — of effectively challenging the measure."

This is the decisive conditional for the model. Reciprocal purchase is *Hentrich-compliant only if* it is built to avoid exactly the two defects the Court identified:

- **It must not be arbitrary or selective.** The power must operate under clear, general, published rules applying uniformly to all who make a substitute-value election — not as a discretionary power the state aims at chosen individuals. The architecture's design helps here: the power is triggered by the *taxpayer's own election* of a low value, not by the state selecting a victim. That is a materially stronger position than Hentrich, where the state initiated. But the rules governing when and how the state may exercise the resulting option must be general and transparent, or the Hentrich defect reappears.

- **It must afford an effective challenge and equality of arms.** The taxpayer must have a genuine, timely, adequately-resourced route to contest valuation and process, on terms that do not place them at a substantial disadvantage against the state. This is why the architecture's separation of the ordinary challenge (no purchase exposure) from the substitute-value election (purchase exposure) matters: the ordinary challenge route is what supplies the "effective possibility of challenging" that Hentrich required and France lacked.

There is also an Article 6 length-of-proceedings point in Hentrich (the French proceedings ran over seven years). The architecture's "pay throughout, correct forward with interest" rule must be paired with a genuinely expeditious challenge process, or a long dispute during which the taxpayer pays on a contested value could itself found an Article 6 complaint.

2.4 The second, independent condition: Article 6 and the timing of a contested charge

There is a further requirement in Hentrich that is distinct from non-selectivity and deserves to be treated as a separate condition rather than a sub-clause of it. The Hentrich proceedings ran over seven years, and the Court found a violation of Article 6 on the length alone. This matters acutely for the present model *because its design deliberately runs the charge during a dispute*: under the "pay throughout, correct forward with interest" rule, a taxpayer contesting a valuation pays on the state's figure for as long as the dispute lasts. A perfectly non-selective, rule-based reciprocal purchase could still generate an Article 6 violation purely through delay, if that delay is long and the taxpayer is paying on a contested value throughout.

So there are *two* independent conditions on the mechanism, not one:

1. **Non-selectivity and effective challenge** (the A1P1 condition): rule-based, uniform, with a genuine route to contest.

2. **Expeditious resolution** (the Article 6 condition): the challenge and any purchase process must be genuinely fast, precisely because the charge accrues on a contested value while the dispute runs. A slow process is not merely inconvenient; on the authority of Hentrich it is independently unlawful.

A supporter should see both, because a design that satisfied the first and failed the second would still be struck down.

2.5 Honest verdict on durability

The charge survives. Reciprocal purchase is **capable of being made compliant** — but the accurate statement is precisely that, not that it *is* compliant as the architecture currently stands. The architecture paper does not yet contain the anti-selectivity rules or the expeditious-resolution guarantee; it describes the mechanism but not the safeguards that make it lawful. So the correct finding is: the mechanism can be drafted to survive A1P1 and Article 6, and because its trigger is the taxpayer's own election rather than state selection it starts from a stronger position than Hentrich — but it survives *only if* both conditions above are written into the charging legislation, and a version lacking either would fail. These are not refinements; they are the difference between a compliant mechanism and a struck-down one, and they are exactly the points a competent tax silk will raise first. This must be disclosed to anyone supporting the model.

3. Test 3 — The cliff and the excessive-burden margin

Classification: manageable, provided quarterly averaging and payment deferral are retained and disclosed.

The entry-threshold (cliff) design is a deliberate choice, defended in the architecture paper as a trade-off for a broader base. Under stress it holds, but only with its accompanying safeguards treated as load-bearing:

- quarterly averaging prevents a single day's valuation from determining liability on a whole holding, which both reduces arbitrariness (helping A1P1) and defeats one-day boundary management;
- payment deferral with interest addresses the liquidity case that would otherwise be the strongest "excessive burden" claim;
- the discontinuity itself remains, and should be disclosed as an accepted cost — a marginal person just over the threshold bears a real jump. This is defensible as a matter of legislative line-drawing (all thresholds draw lines) but it is the model's most criticised ordinary feature and should never be hidden.

A residual honest point: if the excessive-burden risk at the margin proves sharper than expected, the fallback is a taper band just above the threshold, which softens the discontinuity without abandoning the broad base. Adopting a taper would be a legislative choice, not an architectural change, and it is worth naming as the available mitigation.

There is also a connection to Test 1 worth surfacing, because it completes a pattern. The revenue analysis found the model needs a rate of roughly 1.5–2 per cent on the personal base. But the rate is not only a revenue dial — it is simultaneously a *legal-risk* dial, because a higher annual charge strengthens any "individual and excessive burden" argument, most sharply at the cliff edge and for the illiquid holder. So two of the model's key parameters both trade revenue against proportionality risk in the same direction: the **interaction rule** (Test 1) and the **rate** (here). Pushing either toward more revenue pushes the model toward more legal exposure. A decision-maker should see this as a single coherent constraint rather than two separate ones: the model's revenue and its human-rights durability are coupled, and cannot both be maximised at once.

4. Test 4 — Avoidance migration

Classification: serious but addressable; partly unresolved pending data.

The architecture paper claims resistance to *realisation-based* avoidance only, and that claim holds: a charge that does not wait for a disposal cannot be escaped by not disposing. But avoidance migrates, and the pass must say where:

- **Beneficial-ownership concealment** becomes the primary avoidance channel. The model's Principle 3 (economic ownership over legal form) is the intended answer, but where ownership is genuinely unattributable — offshore discretionary trusts, layered foreign entities — it has less to bite on. This is the model's real frontier, shared with every wealth tax, and it is only partly addressable by a single jurisdiction. *Unresolved* to the extent it depends on international beneficial-ownership transparency.
- **Valuation gaming** shifts from "declare low" (disciplined by reciprocal purchase) to disputes over what is included, minority discounts, and control premiums. Manageable through valuation rules, but real.
- **Debt manufacture** is addressed by the bounded-net rule (related-party and artificial liabilities disregarded), which is the correct design and must be robustly drafted.
- **Threshold planning** — fragmenting holdings across connected persons — is addressed by connected-person aggregation, but that aggregation rests on a normative rule (that genuinely independent persons are not aggregated merely by connection) which the architecture paper honestly flagged as needing further specification. *Serious but addressable*, requiring careful anti-fragmentation drafting.

The honest overall finding: the model removes the largest avoidance class at its root and forces the remainder into channels that are visible and, mostly, addressable — but the offshore-opacity channel is a genuine residual limit that no domestic design closes, and the model should claim no more.

5. Test 5 — Capital flight and situs

Classification: manageable for immovable and UK-situated wealth; unresolved for portable financial wealth.

Because liability attaches to assets by situs, UK land, property, and domestic business assets cannot leave, and the charge continues under any owner — so the flight response that undermines conventional wealth taxes is genuinely blunted for the immovable majority of the base. This is a real structural advantage and survives stress.

The residual is portable financial wealth — listed securities and cash — which a departing owner can move. A charge on realised value leaving the jurisdiction is a partial answer; full capture is beyond any single state. This is a genuine limit, not a fatal one, because the immovable base is large and anchored; but the model should not claim to solve flight for portable wealth, and this pass does not.

6. Test 6 — Administrative feasibility

Classification: serious but addressable; a multi-year build with real cost.

A continuously-maintained register of beneficial ownership and standing values across 20,000–32,000 liable persons, plus entities, plus those near the threshold who must file, is a substantial institutional undertaking. The Wealth Tax Commission estimated setup costs of the order of £580m for a conventional annual wealth tax; this model, with its register, valuation function, reciprocal-purchase capability, and disposal machinery, is at least as demanding and probably more so.

Two specific requirements are load-bearing rather than incidental:

- **A credible, funded state purchase-and-disposal capacity.** Reciprocal purchase disciplines valuation only if the state can and would exercise it. Without a standing, funded acquisition-and-disposal function, the discipline is decorative — and, as the architecture paper noted, acquisition of operating businesses can itself destroy value, so the capability must be asset-class-differentiated. This is a real institutional cost that must be budgeted, not assumed.
- **Valuation capacity at scale.** Setting and defending standing values for illiquid assets — private companies, land, art — is the function that has historically overwhelmed wealth-tax administrations. The burden-shift and reciprocal-purchase design is the model's answer, and it is a good one, but it does not eliminate the need for a serious valuation capability; it changes its shape.

Feasible, then, but only as a multi-year build with a real budget, and no honest presentation should imply otherwise.

7. Test 7 — Cumulative taxation and fairness

Classification: manageable; a designed feature to be defended, not denied.

The model deliberately charges the same underlying value at more than one level (company and shareholder). The correct defence, confirmed by this pass, is not to deny cumulative taxation but to locate it: the charges fall on different liable persons for a distinct legislative purpose, exactly as corporation tax and dividend tax already do on the same corporate profits. That precedent is directly on point and is a strong answer. The fairness question — whether the resulting aggregate burden is justified — is a political one the model does not pretend to settle, and the interaction-rule choice (Test 1) is where it is actually decided.

8. What would make a supporter safe — and what would not

Given the stakes in putting one's name to this, the pass ends with a direct statement of what a supporter can and cannot safely claim.

A supporter can safely say:

- that the model is a genuine, original answer to realisation-based avoidance and to the legal-form escape, not a re-packaged wealth tax;
- that it is not disqualified on revenue — at 1.5–2 per cent it is in the region of replacing CGT and IHT on the personal base, with the entity layer as directionally positive additional headroom;
- that it is not disqualified on human-rights grounds — the charge sits within the wide margin the ECtHR allows, and reciprocal purchase is capable of being made compliant if built with the Hentrich safeguards;
- that its limitations are disclosed rather than hidden, which is itself a mark of seriousness.

A supporter must not claim:

- that it raises a specific sum — only government can cost it, the base rests on soft top-end data, and the honest statement is "in the right region, rate-dependent, subject to behavioural response";

- that the entity layer raises a quantified amount — its contribution is directionally positive but unsized in this pass, and the "clears at 1 per cent under stacking" claim is withdrawn as unestablished;
- that it is airtight against avoidance — offshore opacity remains a real residual limit;
- that reciprocal purchase is safe as currently drafted — it is capable of being made safe, but only once the non-selectivity and expeditious-resolution safeguards are written into legislation the architecture does not yet contain, and a version without them would fail on Hentrich;
- that the cliff has no cost — it has a real discontinuity, defensible but not costless;
- that revenue and legal safety can both be maximised — the interaction rule and the rate each trade one against the other;
- that it can be stood up quickly or cheaply — it is a multi-year institutional build.

9. Questions now answered — and questions still open

It is worth separating, at the end, what this pass has settled in principle from what still requires evidence it could not obtain. The distinction matters because a supporter needs to know which objections they can already meet and which remain genuinely live.

Answered in principle by this pass:

- the model's constitutional identity — a persistent charge on the holding of concentrated assets by any liable person, distinct from event-driven taxes;
- that it is not obviously disqualified on revenue: the personal base at 1.5–2 per cent is in the region of the replacement hurdle;
- that the charge itself falls within the ECtHR's wide margin of appreciation in tax;
- that reciprocal purchase has a route to compliance, via the taxpayer-election trigger and the Hentrich safeguards;
- that the principal avoidance channel it removes (realisation-based) is genuinely closed, and where the residue migrates.

Still requiring external evidence this pass could not supply:

- a government revenue costing on taxpayer-level data, including the size of the entity layer, which the public data cannot size;
- counsel's opinion on reciprocal purchase and on the cliff's excessive-burden exposure, which a case-law reading can frame but not settle;
- an administrative costing of the register, valuation function, and state purchase-and-disposal capacity;
- behavioural-response modelling specific to this design rather than borrowed from lower-threshold historic wealth taxes;
- the transition path from CGT and IHT, which this pass has not examined at all.

The honest shape of the project is therefore that the architecture is settled, this pass has cleared the two questions most likely to be fatal, and what remains is specialist external work — not further authorship.

10. Overall finding

The architecture survives the stress-testing pass *conditionally*, and the word conditionally is doing real work. No test returned fatal. On the evidence available to this review, no fatal defect has been identified — which is a different and more modest statement than that none exists, because two of the findings rest on data and law this pass could frame but not finally settle. Two tests returned “serious but addressable” on the load-bearing questions of revenue and durability, each with a specific, nameable condition: a rate at or above roughly 1.5 per cent on a base whose top end is softly evidenced, and reciprocal purchase built with both the non-selectivity and the expeditious-resolution safeguards. The remainder are manageable costs or genuine but bounded limits, with offshore opacity the one materially unresolved frontier, shared with every wealth tax and not closable by domestic design alone.

The honest one-sentence verdict, suitable for someone deciding whether to put their name to it: *on the evidence available to this review, this is a serious, original, and defensible model that is not disqualified on any of the grounds that usually disqualify wealth taxes — provided it is presented with its conditions and limits intact, never oversold beyond them, and understood to await the government costing and counsel’s opinion that only they can supply.*

This pass used public data: HMRC receipts for the replacement hurdle; the Wealth Tax Commission and subsequent analyses for the base; and the European Court of Human Rights case-law, principally Hentrich v France (1994), for the durability tests. It is an order-of-magnitude and legal-principle assessment, not a costing or a legal opinion. A government costing requires taxpayer-level data this pass does not have; a definitive human-rights assessment requires counsel’s opinion. What this pass establishes is that the model is not disqualified on the two grounds most likely to disqualify it, and that its remaining risks are identified, classified, and — with the stated conditions — addressable.