A Constitutional Model for the Persistent Taxation of Concentrated Asset Holdings

A recurring charge on the holding of assets, independent of income, gains, and transfers

Drew Griffiths

July 2026

The proposed charge is imposed on qualifying liable persons by reference to the aggregate net standing value of the taxable assets attributable to them under the charging legislation. Liability does not depend upon income, profits, gains, transfers or death, and it is not avoided by the adoption of a different legal form for holding those assets. This paper sets out the constitutional architecture of such a charge on its own terms. It settles what follows from the model's principles, marks what remains a matter for Parliament, and reserves what only evidence can determine for a later stress-testing pass. It does not argue that wealth ought to be taxed; its contribution is the mechanism, not the politics.

Everything in this paper elaborates a single proposition, stated above. It is not a replacement for Capital Gains Tax and Inheritance Tax dressed up as something grander (though it could replace them), nor a conventional wealth tax organised around a valuation register. It is a distinct model: a recurring charge on the continued holding of concentrated assets, in which legal form is irrelevant to liability.

A note on terminology

Three terms are used precisely and kept distinct throughout.

- **Liable person** — the taxpayer under the charge: the person or arrangement the legislation designates as bearing liability. This includes legal persons (natural persons, companies) and other asset-holding arrangements that are not separate legal persons in the same sense — notably trusts, where the trustees are the legal owners and the charge attaches to the trust arrangement as the legislation provides. "Liable person" is used wherever the taxpayer is meant, precisely because not every taxpayer is a "legal person" in the strict company-law sense.
- **Beneficial holder** — used where what is valued and charged is a beneficial interest (a shareholding, a partnership interest, a beneficial interest under a trust) rather than legal title to the underlying asset.
- **True holder** — used only in the anti-avoidance discussion (Principle 3), to denote the person whose economic reality a sham legal form is designed to disguise. It is deliberately confined to that context and does not stand in for "liable person" elsewhere.

How to read this paper: three kinds of question

Tax design mixes three different kinds of question, and much confusion comes from answering one in the language of another. This paper keeps them separate throughout, and labels which is which:

1. **What follows from the principles.** Given the constitutional principles in Section 2, certain things are entailed — the taxable base, who is liable, the reach of the anti-avoidance principle. These are architectural.

2. **What Parliament must choose.** The architecture permits a range of policy settings — the rate, the threshold figure, how multiple liabilities interact, which classes of person are exempt. These are legislative choices; the architecture is compatible with several, and the paper does not pretend otherwise.

3. **What only evidence can settle.** Whether the charge raises enough, how behaviour responds, whether it survives legal challenge — these need modelling and data, and are reserved for a separate stress-testing pass.

The strength of a constitutional paper lies partly in refusing to answer questions of the second and third kinds as though they were of the first. Where this paper defers, it says so, and says into which category the question falls.

1. The problem this model solves

The United Kingdom's taxes on wealth and its transfer are *event-driven*: Capital Gains Tax fires on disposal, Inheritance Tax on death or a gift within seven years of death, Stamp Duty on purchase. Because the triggering event is generally within the taxpayer's control, the tax can be deferred, timed, disguised, or escaped by not triggering it:

- **Lock-in**: CGT is paid only on sale, so disposals are deferred and gains extinguished at death. The rational strategy is never to sell.
- **Timing**: realisations are rushed or delayed around announced changes.
- **Disguise**: one form of event is dressed as another.
- **Non-triggering**: "buy, borrow, die" realises no taxable event across a lifetime.

Each manipulates the *event*. A charge that accrues through time, on the holding of assets rather than on any transaction, has no realisation event to manipulate. That is the core of the model, and it makes the design **structurally resistant to realisation-based avoidance** — the largest single class. It does not end avoidance as such; Section 7 sets out honestly where avoidance migrates instead.

A second, connected weakness of conventional design is that it can often be escaped by **changing the legal form** in which assets are held — moving wealth into a company, trust, or foundation so that no natural person appears to hold it. This model closes that route not by hunting for the "ultimate" owner, but by making **every liable person that holds concentrated assets independently liable**, so that there is no form into which assets can be moved to escape the charge.

2. The constitutional principles

Three principles define the model. The first two establish the charge; the third prevents avoidance by legal form. The third is a distinct normative choice, not a deduction from the first two.

> **Principle 1. Every taxable asset has one legally operative standing value.**
>
> **Principle 2. Every liable person whose aggregate net standing value of taxable assets exceeds the statutory threshold bears the annual asset charge.**
>
> **Principle 3. Legal form does not prevent the charge applying to the true holder of taxable assets.**

On Principle 1. Each taxable asset carries a single standing value that is *legally in force* at any time — recorded in the register, not re-derived at every moment. It is this legally operative value, not a fresh valuation, that the system uses (and averages across the year — Section 4.2).

Interests in assets are themselves valued: a shareholding, a partnership interest, or a beneficial interest in property is an asset with its own standing value, distinct from the underlying thing, carrying its own premiums, discounts, and restrictions.

On Principle 2. Liability attaches to a *liable person* — a natural person, a company, a partnership, or a trust arrangement (through its trustees, as the legislation provides) — by reference to the *aggregate* net standing value of the taxable assets that person holds, once the aggregate exceeds the threshold. The term "liable person" is used rather than "legal person" precisely because not every taxpayer is a legal person in the strict sense: a trust is generally not a separate legal person in UK law, yet the model must be able to charge assets held in trust, so the legislation designates the liable person (typically the trustees or the trust arrangement itself) accordingly. Two features matter:

- It is the *aggregate* that counts, never individual assets crossing the threshold. A person holding £4m of property, £3m of shares, £2m of land, and £2m of art holds no single asset near the threshold but an £11m aggregate, and is liable. A tax on individually large assets would be defeated by simply holding many smaller ones.
- The charge falls on the person *as a holder of assets*. A company holds its assets; a shareholder holds a different asset (the shares); each is a liable person holding taxable assets, and each above the threshold is independently liable. This is developed in Section 3.

On Principle 3. Legal form neither exempts assets (a genuine company's assets are still held by a liable person — the company) nor shields them (a sham arrangement does not hide the true holder). As a general anti-avoidance principle it reaches, among others:

- nominees and bare trustees,
- retained-benefit arrangements,
- artificial fragmentation of holdings, and
- similar sham structures whose purpose is to place the legal form at odds with the economic reality.

The principle is stated generally so that novel devices are caught without amendment; the list is illustrative, not exhaustive, and is deliberately kept in the commentary rather than written into the principle itself.

3. Cumulative liability is a designed feature, not an interaction to be explained away

Because liability attaches to *every* liable person holding concentrated assets, the same underlying economic value may support liability in more than one person. A company worth £100m holds assets above the threshold and is liable; a shareholder whose shares are worth £60m holds an asset above the threshold and is also liable. These are different liable persons holding different assets — the company's factory and cash, the shareholder's shares — so both fall within the charge.

This is not an awkward interaction requiring defence. It is a direct consequence of the base Parliament would adopt:

> The proposal accepts that different liable persons holding different taxable assets may each incur liability under the charge. This is a consequence of the tax base adopted by Parliament, not an unintended interaction.

The model does not track a unit of economic value through the economy to ensure it is taxed only once. It charges liable persons who hold concentrated assets. That two such persons exist in relation to related value is the point of closing the legal-form escape, not a flaw in it.

The interaction rule is a legislative choice, not an architectural one.

> Where different liable persons are independently liable in respect of different taxable assets whose economic value overlaps, Parliament must determine whether liabilities accumulate in full, interact through credits or offsets, or follow another statutory interaction rule. The constitutional architecture proposed here is compatible with each approach. The choice concerns the intended economic burden of layered holdings rather than the definition of the tax base itself.

Full accumulation would treat holding wealth through a vehicle as bearing a heavier charge; a credit or offset would ensure related value is charged once at the higher level; other rules are possible. Each is a policy decision about the intended burden of layered holdings. The architecture defines the base and the liable persons; it does not dictate how their liabilities interact, and this paper does not pretend it does.

4. Valuation, timing, and the discipline against understatement

4.1 The state sets the legally operative value

The register records each asset's standing value, set by the state, on which the charge is due. That the state sets the value is a **policy choice**, not a requirement of the principles: the value could be taxpayer-declared, set by an independent office, computed by formula, or established by transaction. It is assigned to the state here because that keeps a single operative value continuously in force and pairs with the discipline in 4.3.

4.2 Quarterly averaging

Liability is determined by the **arithmetic mean of the person's aggregate net standing value at four statutory quarterly reference dates** during the tax year. If that average exceeds the threshold, the charge applies (on the basis in Section 5).

This does not require four full revaluations a year. The average is taken of the standing values *legally in force* at each date: quoted securities update from observable prices; illiquid assets carry their standing value forward until a recognised revision event (a scheduled reassessment, a market transaction supplying evidence, a material event, or a revision initiated by taxpayer or authority). Averaging concerns the person's recorded aggregate across time; it does not churn the register quarterly.

Averaging is chosen over a single annual date because, under the entry-threshold model (Section 5), a single date would let one day's movement determine liability on the whole holding, and would invite one-off boundary management. Four dates make boundary management something that must be sustained across the year, while keeping the calculation tractable. Entry and exit occur annually on this same averaged basis; no multi-year entry or exit tests are used, as they would introduce a second discontinuity harder to justify than averaging.

4.3 The reciprocal-purchase election

A standing value is credible only if it cannot simply be declared low. The discipline is a **taxpayer election** distinct from an ordinary challenge (Section 4.4):

Where a taxpayer does not merely dispute the state's value but **affirmatively elects a lower substitute value**, that election opens a time-limited power for the **state** to purchase the asset (or interest) at the declared figure, on the same legal and economic terms that informed the valuation.

Honest qualifications, each stated rather than glossed:

- It is **not perfect symmetry**. The state, seeing the declaration, chooses whether to buy, and acquires only what is useful; the taxpayer bears the downside of understatement regardless. It is a credible putative-sale condition attached to a taxpayer-initiated downward valuation, with procedural protections — not symmetry.
- It **disciplines understatement only**. It does not resolve state over-valuation, disputes over what rights are included, minority discounts, control premiums, contingent liabilities, or transfer restrictions. Declaring a *tax* value is not volunteering to *sell*; that is why the election route is separate from ordinary challenge.
- It **transacts in the asset as valued, net of attached liabilities**. If a £15m property carries a genuine £10m secured debt, the equity interest's standing value is £5m; a £4m substitute election exposes the equity — subject to the £10m debt — to purchase at £4m (total state exposure £14m), not a debt-free property for £4m.
- It is **available only where economic equivalence can be preserved**. Where an interest cannot be transferred on equivalent terms — personal guarantees, bespoke loans, non-transferable licences, change-of-control triggers that would destroy value — reciprocal purchase cannot operate simply, and an alternative valuation discipline applies. It works cleanly for land, securities, and transferable property; for operating companies it targets the ownership interest, not the operating assets, with continuity protections; where transfer is legally impossible it does not apply.
- It is **state-only, not assignable**. Assignability would create a market in purchase rights — speculation, hostile acquisition, collusion — destroying the legitimacy the mechanism draws from being the state relying on the taxpayer's own figure.

The power is an **integrity mechanism, not an acquisition programme**: it exists to keep declared values honest, and where declarations are realistic it will seldom be used. Its credibility depends on the state genuinely having funded capacity to buy and to manage disposal (Section 7); without that, the discipline is decorative.

4.4 Two distinct routes: challenge and election

- **Ordinary challenge**: the taxpayer asserts the state's value is *wrong*. This leads to independent adjudication, with **retrospective correction from the date of the challenged assessment** and repayment of overpayments **with interest**. It does **not** expose the taxpayer to purchase — a person must be able to challenge a state error without risking forced sale at the state's own figure.
- **Substitute-value election**: the taxpayer *affirmatively proposes* a lower value and elects to have the register altered to it, opening the purchase window (4.3).

Only the election exposes the taxpayer to acquisition. The charge remains payable while any dispute proceeds (the register cannot pause); a successful challenge corrects the value from the assessment date, with interest on overpayments and penalties or enhanced interest where either party behaved unreasonably; hardship deferral may prevent irreparable harm without altering the standing value.

4.5 Debt: bounded net valuation

The standing value of an asset or interest is its market value **after deducting genuine, legally enforceable liabilities directly attached to the acquisition, retention, or value of that asset**:

- **Direct third-party liabilities** (genuine secured borrowing to acquire, improve, or preserve the asset, owed to an independent creditor) are normally deductible.
- **General personal liabilities** do not automatically reduce each asset's standing value.
- **Related-party or artificial liabilities** are disregarded unless commercial substance, actual transfer of funds, ordinary repayment terms, and genuine enforcement risk are established.

This makes the charge a **bounded net tax** — neither a gross-assets tax nor an unrestricted net-worth tax — closing the incentive to manufacture leverage to erode the base.

5. The threshold and the entry-threshold model

5.1 The threshold

The charge applies to a liable person whose averaged aggregate net standing value exceeds a high threshold — for illustration, **£10 million per person, £20 million for a married couple** — with connected persons and interposed vehicles aggregated so that holdings cannot be fragmented across connected parties to keep each below the line. The precise figure, and its justification against the distribution of UK wealth, is an **empirical question** for the stress-testing pass.

The threshold is an **entry threshold, not an allowance**: once the averaged aggregate crosses it, the charge applies to the **full** averaged value, not merely the excess.

5.2 The entry-threshold model is a deliberate trade-off

This produces a genuine **discontinuity**: at the assessment measure, a person just below pays nothing and a person just above pays on their whole holding. Quarterly averaging ensures this turns on a representative annual figure rather than a single day, and makes boundary planning harder — but it does **not** abolish the discontinuity, and nothing in the design does.

It is not claimed, as an earlier draft did, that the cliff is "self-correcting." The valid observation is narrower: **the threshold is not necessarily a stable shelter for continuing accumulation**, because retained returns and appreciation tend to move an accumulating holder across it over time. But this is a behavioural tendency, not an architectural property; paying the charge later does not retrospectively reduce the value that established it, holders with strong returns may remain above the threshold despite it, and illiquid assets can generate liability from appreciation without the means to pay.

The honest defence is that the entry-threshold model is a **deliberate policy trade-off**: it excludes ordinary and moderately-substantial wealth entirely while bringing the full holdings of those above the boundary into the base; the discontinuity is accepted in return for a broader base and a clearer exclusion of those below. Whether to prefer this to an allowance model is itself a **legislative choice**; this paper adopts the entry threshold and states its cost plainly.

5.3 Liquidity is separate from liability

Valuation and liability are determined normally; inability to pay is addressed separately — through deferral, instalments, security over the asset, or payment on a later liquidity event, with

interest or indexation on deferred amounts so that hardship treatment is not an interest-free avoidance route. A holder whose value lies in an illiquid company or farm is liable; how the liability is discharged flexes, whether it exists does not.

6. Scope: which persons and arrangements are within the charge

The model's default is that **every liable person holding aggregate taxable assets above the threshold is within the charge** — natural persons, companies, partnerships, and trust arrangements. Which classes should nonetheless be exempt is a matter of legislative scope, not constitutional architecture:

> Whether particular classes of person or arrangement, including pension schemes, registered charities and other public-benefit entities, should be wholly or partly exempt is a matter of legislative scope rather than constitutional architecture.

The architecture establishes that such bodies and arrangements are capable of holding assets and so, by default, within the charge. Whether Parliament wishes to exempt them — to protect retirement saving, or public-benefit activity — is a policy decision to be made and justified on its own terms, as an exemption *from* a clear rule rather than an ambiguity *in* it. Connected-person aggregation (5.1) remains the mechanism against fragmentation; it is distinct from the question of which classes of person are exempt, and rests on Principle 3 only so far as fragmentation is artificial — genuinely independent persons are not aggregated merely through connection, which would require a further normative rule this paper does not assume.

7. Honest limitations

Stated at architecture level; the stress-testing pass will quantify them.

Avoidance shifts, it does not disappear. Removing realisation-based avoidance pushes effort toward: concealment of beneficial ownership; manipulation of debt; fragmentation among connected persons; disputes over situs and asset classification; threshold planning; movement of portable financial wealth; and non-compliance with the register. The model is resistant to the realisation-based class specifically, and claims only that.

Offshore opacity. Where beneficial ownership is genuinely unattributable — discretionary trusts, layered foreign entities — Principle 3 has less to attach to, and entity-level charge on the holding vehicle becomes the practical backstop. Piercing genuine opacity needs beneficial-ownership transparency and international cooperation; some wealth still escapes, so the system may capture visible domestic holders more reliably than the deliberately obscured.

Portable financial wealth and situs. Because liability attaches to assets, it can be anchored by where an asset is situated: UK-situated assets — land, property, domestic businesses — cannot leave, so the charge continues regardless of the owner's residence, and a sale is merely a change of holder in the register while the asset stays in the base. This makes the base resistant to capital flight for immovable wealth. The residual risk is genuinely portable financial wealth, addressable by a charge on realised value leaving the jurisdiction but not fully capturable by any single jurisdiction.

Credible state purchase capacity. The reciprocal-purchase election disciplines valuation only if the state can and would exercise it, requiring funded acquisition-and-disposal capability. Acquisition can itself destroy value in operating businesses (change-of-control clauses, covenants, licences, key-person departures), which is why the election must be asset-class-differentiated and cannot be assumed uniform across the base.

The entry-threshold discontinuity. As in 5.2, the model creates a genuine discontinuity with boundary-planning incentives and liability volatility near the threshold. Averaging limits but does not remove these; it is an accepted cost of the chosen model.

Cumulative taxation. Income tax falls on receipts as a flow; this charge falls on the holding of a large stock of assets. A critic may object that the same resources are taxed as income and again as held assets. The answer is not to deny cumulative taxation but to locate it: the charges fall on different persons or bases for different purposes, and — per Section 3 — cumulative liability across liable persons is a designed feature of the base. Whether the resulting burden is justified is part of the contested political question the paper does not claim to settle.

Revenue sufficiency. If the charge replaces CGT and IHT, it forgoes roughly £22–23bn of current annual revenue; whether it raises at least as much depends on rate, captured base, and behaviour — an empirical question, genuinely open, for the stress-testing pass.

Legal durability. The state-set-value, pay-throughout structure and the purchase election must satisfy procedural-fairness, proportionality, and property-rights standards, including under the European Convention on Human Rights. The challenge/election separation, retrospective correction with interest, and hardship deferral are designed with this in mind, but the analysis is for the dedicated pass.

8. What the model claims

The claim is bounded, and each boundary is a deliberate act of the discipline set out at the start.

It does **not** claim wealth ought to be taxed more heavily — that is a contested political question left open.

It does **not** claim the mechanisms follow deductively from the principles — the state setting values, quarterly averaging, the interaction rule, and scope exemptions are institutional requirements or legislative choices, labelled as such.

It does **not** claim the system is airtight — avoidance shifts to other channels, catalogued in Section 7.

It does **not** claim record-keeping generates the tax — the charge is an independent choice a register makes *possible*, not one it *entails*.

What it **does** claim is this: if Parliament chooses to charge concentrated wealth, that charge can be imposed on every liable person holding aggregate taxable assets above a threshold, by reference to a continuously maintained standing value rather than to realisation events — removing realisation-based avoidance at its root, closing the legal-form escape by making every qualifying holder independently liable, and anchoring the base against flight for immovable wealth — at the accepted cost of an entry-threshold discontinuity, cumulative liability across liable persons, and the residual international and opacity limits no single jurisdiction escapes.

The originality is the shift from event-driven, single-owner wealth taxation to a persistent charge on the holding of concentrated assets by any liable person, in which legal form is irrelevant to liability. The specific mechanisms — bounded net valuation, quarterly averaging, the reciprocal-purchase election, and the challenge/election split — make such a charge administrable and reasonably fair without the realisation events that conventional taxes wait for and that sophisticated holders manage away.

This is the architecture pass, presenting the model on its own terms. It settles what follows from the principles, labels what remains a legislative choice (the interaction rule, scope exemptions including pension schemes and charities, the allowance-versus-entry-threshold question, the rate), and reserves what only evidence can settle (revenue and rate modelling, threshold justification against UK wealth distribution, behavioural and avoidance response, treaty interaction, transition, *and constitutional and human-rights durability) for a separate stress-testing pass. Keeping those three categories distinct is the paper's method, not an evasion of its harder questions.*