Selfless Capitalism

Contents

Value Created, Not Extracted

What We Mean by Selfless

Not selfless as a virtue. Selfless as a description.

Capitalism has done an enormous amount of good, and this manifesto proposes its next evolution rather than its replacement. But as practised it is selfish in a precise and narrow sense: it optimises for one interest — the owner's return — across a bandwidth so narrow that forests, futures and communities fall outside it. Everything proposed here is a check or a balance that widens the bandwidth.

An asset-locked organisation is *literally self-less*. There is no residual claimant, no owner, no self whose interest the structure exists to serve. And the property recurs at three levels: at the apex, no owner to enrich; in management, no one on top saying *I did this*; along the value chain, no node extracting rent from the others. Every party trades. Nobody taxes.

The company has no self. The people in it very much do — they are paid properly, they own their firms, they compete hard, and they keep what they build.

The working term is **selfless exchange**, chosen because an exchange is symmetric: nobody is doing anybody a favour, and both sides gain by construction. Two lines follow. The model **co-operates in structure and competes in markets** — because competition is the discovery mechanism, and extraction, not competition, is the enemy. And its economics fit in one sentence: markets are positive-sum by design, and **extraction is precisely what drags a positive-sum economy back toward zero-sum.**

How This Manifesto Was Made

This manifesto was produced by SHALOME.AI with AI drafting support (Anthropic's Claude), working from founder-authored seed papers and SHALOME.AI's governance corpus, under a statement of work requiring adversarial review before publication. Its first red-team pass was conducted by the same AI support that drafted it. A structural limitation recorded openly here and mitigated, and then reviewed by the founder.

The manuscript has passed an SQ integrity audit — an independent-in-method review of its own language for the certainty, grandiosity and conquest-rhetoric it argues against — and the audit's findings and resolutions are retained in the project record. Every empirical claim is sourced in the chapter evidence notes and is to be re-verified at publication.

We disclose our use of AI because the covenant this manifesto describes requires it of everyone else.

The Concessions Register

A manifesto that argues *a programme with no published failures is a programme that has stopped looking for them* must begin with its own. These are the standing admissions of this text, produced by its adversarial review and published as part of it:

1. **Seed capital that will not extract is a thinner market than the conventional one.** The money is not absent; the extractive terms are what we refuse, which narrows the field to philanthropic, foundation, mission-aligned and founder capital. Some venture classes needing very large pre-revenue equity are unavailable to the form. Our long horizons, we should be clear, are not a symptom of this: civilisational work runs on its own clock.
2. **The pure form is unproven above roughly the hundred-thousand-person scale.** Beyond it, the working expression is hybrid — asset-locked control above conventional instruments — and our flagship exemplars are themselves hybrids.
3. **The asset lock removes the extraction failure mode, not the governance failure mode.** The asset-locked form's characteristic disasters — Equitable Life, the Co-operative Bank, Kids Company, the demutualisations — are governance disasters, which is why the discipline machinery is constitutive, not optional. And the extraction claim itself needs one boundary: what the lock removes is the claimant outside the organisation. Absorption of surplus by the institution itself remains available, and is priced at item 10 and Chapter Three §5.7.
4. **Our community accountability mechanisms are designed, not yet proven.** The published record will show whether we honour them, and the model's own success metric — communities needing it less over time — will be reported against.
5. **The selection claim — structure selects character — is a hypothesis with stated predictions, not an established law.** The same standing applies to its companion claim that the form *converts* the marginal actor (Chapter Four §5); see item 9.
6. **Beneficiary accountability is only as constitutional as founders choose to make it.** We set that dial high and publish the setting.
7. **This manifesto's production is self-referential in ways we mitigate but cannot remove**: AI-drafted on founder-authored sources, first-reviewed by its own drafting support, with its worked example the commissioning organisation — an organisation in formation: designed and documented, partially piloted, unproven at scale.
8. **Chapter Six's variable corporation tax proposal is the newest and least tested material in this book.** It has not been modelled, costed or piloted anywhere. We publish it as a designed instrument — with its assessment machinery, border problem and objections stated on the page — not as settled policy.
9. **The conversion argument rests partly on founder testimony, which is real evidence and partial evidence.** Chapter Four §5 argues that the extraction form converts ordinary actors and not merely that it recruits extractive ones. Its supports are a documented payoff asymmetry, an established behavioural finding about descriptive norms, and thirty-six years of one practitioner's unblinded observation. We have stated the sample objection in the text rather than around it, declined the clinical vocabulary that would have made the claim sound better evidenced than it is, and published the study that would falsify it.
10. Criteria fixed at foundation will bind successors to judgements made decades earlier, some of which will age badly — a measure may cease to be computable, a category may cease to mean what it meant. Accepted as the price of preventing pre-emptive amendment. The failure prevented is certain; the failure created is possible.

11. Organisations built in separable parts cannot cross-subsidise as freely as unified ones: each set of trustees owes duties to its own beneficiaries and must justify transfers in those terms. Covenanted funding relationships do most but not all of this work. This is the sharpest objection to the sunset design and is not fully answered.
12. The review consumes money and executive attention that would otherwise reach beneficiaries — a real transfer from present service to future integrity, not a rounding error.

Chapter One: The Failure Is in the Accounting

1. In Defence of the Principle

This manifesto begins where its critics will least expect: with a defence of capitalism.

Competitive markets are the most effective wealth-creating mechanism humanity has ever devised. The principle is brilliant in its simplicity: let market forces compete, and competition itself does the work — driving innovation, creating products and services no planner would have imagined, cutting costs, raising efficiency, rewarding those who serve others' needs better than the alternatives.

The evidence is not seriously contested. In 1900, world output per person stood at roughly $1,300 in comparable international dollars. By 2000 it had risen to around $6,000, and it has continued climbing since — a rise of an order of magnitude in real living standards over little more than a century, achieved *while the world's population roughly quadrupled* from 1.6 billion to over 6 billion. No other economic arrangement in history has come close. Life expectancy, literacy, child survival and the escape from absolute poverty tell the same story. Where markets were permitted to operate, prosperity followed; where they were suppressed, it did not.

This manifesto is therefore not an argument against capitalism. It is an argument that capitalism, as currently practised in the Anglo-American world, systematically betrays its own principle — and that the betrayal is located in a specific, identifiable, and correctable place.

The failure is not in the principle. The failure is in the accounting.

Three accounting failures, compounding one another, convert a wealth-creating machine into a wealth-extracting one. Each is stated below with its mechanism, because a diagnosis that cannot name a mechanism is an opinion, not a diagnosis.

2. The First Failure: Prices That Lie

Consider a coal-fired power station. It burns coal, generates electricity, sells it, and returns a profit to its owners. On its own books, it is a success. But its chimneys emit sulphur dioxide, which combines with rain and falls as acid on the forests downwind. The forests die. The forests' death appears on nobody's books — not the power station's, not its shareholders', not the electricity buyers'. The cost is real, but it is *external* to every account that governs decisions.

So the power station continues, rationally maximising the wealth of its owners while eroding the wealth of the world. Nobody in the chain behaves irrationally. Nobody, necessarily, behaves malevolently. The price of the electricity simply *lies*: it reports the cost of coal, labour and capital, and omits the cost of the forest.

This is not an exotic case. It is the ordinary condition of the modern economy. Carbon emitted at no charge to the emitter, warming everyone's planet. Plastics manufactured at the cost of production, disposed of at the cost of the oceans. Fisheries harvested to collapse because no fisher pays for the depletion of the stock. Soils exhausted, aquifers drained, antibiotics squandered — in every case the same signature: a private gain honestly accounted, a shared loss accounted nowhere.

Economists have understood this for a century; the textbook name is *externality*, and the textbook remedy — making the polluter's price tell the truth — goes back to Pigou in the 1920s. What the textbooks have been slower to confront is the scale. When the unpriced costs are planetary — climate, biodiversity, ocean chemistry — the lie in the prices is not a rounding error at the margin of an otherwise sound system. It is the system's largest single output.

A market economy is an information system: prices are its signals, and decisions follow the signals. Feed it false signals and it will do exactly what it is designed to do — optimise, relentlessly and ingeniously — toward the wrong destination. Capitalism is not failing when it destroys the forest. It is *succeeding at the task the false prices set it*.

3. The Second Failure: Whose Wealth?

The second failure sits one level deeper, in the definition of the thing being maximised.

When we say capitalism creates wealth, the honest question is: *whose* wealth, and wealth in *what*? The national accounts answer with production: GDP, the flow of marketed output. Firms answer with shareholder value. Both definitions share a blind spot — they count what is produced and traded, and ignore what is depleted to produce it.

A country can log its forests, drain its wetlands, exhaust its topsoil and empty its seas, and record every year of the process as *growth*. A firm can hollow out its workforce's health, its suppliers' resilience and its community's cohesion, and report every quarter of the process as *value creation*. In both cases something true is being measured and something equally true is being ignored: the balance sheet behind the income statement — the stock of assets, natural, human and social, from which all future income must flow.

The corrective has a name: **inclusive wealth** — the total value of produced capital (machines, buildings, infrastructure), human capital (health, skills, knowledge) and natural capital (ecosystems, resources, climate stability), assessed together. The landmark statement is the Dasgupta Review, commissioned by HM Treasury and published in 2021, whose central finding deserves to be quoted in substance: in recent decades produced capital per person roughly doubled, while the stock of natural capital per person declined by nearly 40%. We have been getting richer in the numerator by getting poorer in the denominator — and calling the difference growth.

Once wealth is defined inclusively, the question "is capitalism a good thing?" resolves cleanly. A capitalism that grows inclusive wealth is doing what it claims. A capitalism that grows private, produced wealth by liquidating shared, natural and human wealth is not creating wealth at all: it is *transferring* it — from the many to the few, from the commons to the account-holders, and from the future to the present.

Which brings us to time.

4. The Third Failure: The Tyranny of the Quarter

The third failure is the least visible and, as the case study closing this chapter will show, the most expensive: the mismatch between the time horizon over which consequences unfold and the time horizon over which decisions are rewarded.

Every consequential economic decision has a *time signature* — the period across which its costs and benefits play out. A marketing campaign's time signature is measured in months. A factory's in years. A reservoir's, a railway's, a power grid's, a forest's, an education system's, a climate's — in decades to centuries.

Capital, as currently structured, also has a time signature. A listed company reports quarterly and is repriced daily. Executive incentives vest in three to five years. Private equity funds return capital in five to seven. Even "patient" institutional capital is benchmarked annually against peers who are not being patient. The effective horizon of the shareholder economy — the distance at which future consequences still influence present decisions — is perhaps five years on a generous reading.

Now observe what happens when a forty-year asset is placed under five-year capital. Every pound spent on maintenance and renewal purchases benefits that will accrue mostly to *someone else's* holding period; every pound extracted as dividend or buyback accrues *now*. The structure does not merely permit under-investment and extraction — it *selects* for them, promotes the executives who deliver them, and prices the shares of the firms that practise them above those that do not. Individual virtue is no defence; a conscientious CEO who invests for year thirty is, within this structure, mispricing her own career.

The three failures now compound. False prices hide the cost (Failure One); narrow wealth definitions ensure the depletion is booked as gain (Failure Two); and short horizons guarantee that even visible, priced, admitted long-term costs are ignored, because they fall outside the window within which anyone deciding is paid to care (Failure Three). This is the full diagnosis: **a system that optimises brilliantly against signals that are false, toward a target that is wrong, over a horizon that is too short.**

None of these is a failure of markets as such. All three are failures of the *accounting architecture within which markets operate* — and architecture can be redesigned.

5. Naming the Next Evolution

What should the corrected form be called?

Begin from where Section 1 began, because the name has to carry it. Capitalism has done an enormous amount of good. It lifted the per-capita output of the world by an order of magnitude while the population quadrupled, and no rival system has come close. This manifesto proposes an evolution of that machine, not a replacement for it — and the evolution has a single defining property. **Capitalism, as practised, is selfish.** Not wicked: *selfish*, in the precise and narrow sense that it optimises for one interest — the owner's return — across a bandwidth so narrow that forests, futures and communities fall outside it. Everything this manifesto adds is a check or a balance that widens the bandwidth. The corrected form is therefore not a kinder capitalism. It is a **selfless** one.

That word was arrived at, and the route matters. The founding note of this project proposed *Compassionate Capitalism*, and an adversarial review found that the name was already taken — twice. It is the title of a 1993 book by Rich DeVos, co-founder of Amway, and of a 2004 book by Marc Benioff of Salesforce, the latter a best-practices guide for corporate philanthropy. The adjacent candidates fared no better: "inclusive capitalism" is the banner of a Vatican-affiliated council of chief executives founded by Lynn Forester de Rothschild, whose method is voluntary commitments aligned to World Economic Forum pillars; "responsible capitalism" is party-conference boilerplate; "inclusive growth" and "inclusive development" belong to the OECD, the EU and the UNDP.

We record this openly, because the pattern taught more than the collisions cost. Every warm adjective bolted to an economic noun has been claimed, and claimed largely by the extraction economy, for a reason this manifesto can now state as a law: **a word that does no structural work is free for anyone to take.** The 2019 Business Roundtable statement is the same move at pledge scale — 181 chief executives redefining corporate purpose around all stakeholders, and, as Bebchuk and Tallarita's scholarship found, changing essentially nothing. Compassion-as-adjective is available. Compassion-as-architecture is not, because it is not a claim about anyone's feelings.

There is a second reason "compassionate" was set aside, and it is the more important one. Compassion has an asymmetry built into its grammar: there is one who has it and one who receives it, and the giver remains the giver. That is the structure of charity-as-condescension — the posture that produces dependency in those it means to serve, and self-regard in those who serve. Chapter Two's fourth commitment and this project's governance corpus are both organised around removing exactly that asymmetry. A name that reimports it at the cover would be fighting the book.

Selfless carries no such asymmetry, and — this is the point on which the whole name turns — it is not offered as a moral boast. It is a structural description. An asset-locked organisation is *literally self-less*: there is no residual claimant, no owner, no self whose interest the structure exists to serve. Ordinary capitalism has a self, and every incentive traced in this chapter flows from serving it. Remove the self and the same machinery — competition, prices, enterprise, ambition — serves the purpose instead.

Selflessness then appears at three levels of the same architecture, which is why it names the system rather than a virtue:

- **At the apex**, there is no owner. The asset lock leaves no self to enrich (Chapter Three).
- **In management**, there is no one on top saying *I did this*. That is the lean insight, and it is why the "superstar" chief executive and the 300:1 pay ratio are one pathology in two costumes (Chapter Two, Chapter Four).
- **Along the value chain**, no node extracts rent from the others — not the shareholder, not the executive, not the platform, not the intermediary. Every party trades. Nobody taxes.

Hence the name on the cover: **Selfless Capitalism — Value Created, Not Extracted.** The working term inside the book is **selfless exchange**, and it is chosen for a reason that answers the objection to "compassionate": an exchange is *symmetric*. Nobody is doing anybody a favour; both sides gain, by construction. Growth, competition and enterprise are fully present — what is absent is the party taking value out of the exchange without adding to it.

Two consequences run through every chapter. The model **co-operates in structure and competes in markets** — federation, shared knowledge and cross-subsidy above; market discipline below — because competition is the discovery mechanism, and extraction, not competition, is the enemy. And the economics reduce to one line an economist cannot dodge: markets are positive-sum by design, that being the entire point of gains from trade, and **extraction is precisely the mechanism that drags a positive-sum economy back toward zero-sum**, converting mutual gains into transfers. (The nearest prior art is Marjorie Kelly's contrast between extractive and generative ownership design; this manifesto's contribution is to run that distinction to its structural conclusion and name the legal family waiting there.)

A final word on keeping "capitalism" in the name, since Section 4 has just argued that capital's rule of the core is the problem. Deirdre McCloskey — no critic of free markets — has complained for years that the word "capitalism" misleadingly implies the rule of capital rather than the liberty to exchange. She is right, and this manifesto's own position is the sharper version of her complaint: *we are not saving capitalism from itself; we are saving the market economy from capital's rule of its core.* Johan Norberg's answer to McCloskey is the one adopted here. Keep the word, because it is the word the argument actually happens in — and if those who believe in exchange do not fill it with meaning, its opponents will.

So: capitalism, kept. The self, removed.

6. What We Do Not Claim

A manifesto that has diagnosed overclaiming as a civilisational disease must not practise it. Three limits are therefore placed on the record now, before the argument builds.

We do not claim that correcting capitalism would end war. An earlier draft of this project's founding note suggested that compassionate capitalism could "eliminate all war." The defensible version is narrower. A substantial body of evidence associates economic interdependence with reduced interstate conflict, and resource scarcity with elevated conflict risk; an economy that reduces scarcity-driven desperation and resource-trap dynamics plausibly removes *one class* of war's causes. But wars are also fought over identity, security dilemmas, status and ideology — causes no accounting reform reaches. We claim less, so that what we claim can be believed.

We do not claim that architecture replaces character. Structures shape behaviour; they do not guarantee it. A later chapter argues that non-extractive structures and spiritually mature leadership are complements — each fails without the other. Anticipating that argument, we note here: nothing in this chapter implies that a correctly accounted economy runs itself. Systems that describe themselves as making wrongdoing impossible are systems that have stopped listening for it.

We do not claim novelty. Pigou saw the pricing failure in 1920. Development economists, co-operators, mutualists and the architects of steward-owned firms have built partial correctives for a century and a half. The Dasgupta Review carries the authority of HM Treasury, not of this manifesto. What this manifesto claims is narrower and, we believe, genuinely unclaimed: that these correctives converge on a single organisational form, that the form already exists in law, and that one organisation's complete governance corpus demonstrates it in working detail. That argument occupies the chapters that follow. First, the evidence that the diagnosis is not theoretical.

7. Case Study One: The True Cost of a Coca-Cola

7.1 A price with three things missing

In July 2026, in a British convenience shop, a single-serve bottle of Coca-Cola works out at about £4.60 a litre. For sugar and water.

That price is the subject of this case study, because three things are absent from it. The plastic it is sold in, which the buyer will discard and the public will pay to chase across beaches and out of rivers. The metabolic disease that follows its consumption at scale, which the buyer's health service will treat. And any relationship whatever to what it costs to make. Failures One and Two of this chapter — prices that lie, and wealth counted narrowly — are not abstractions requiring a textbook. They are on the shelf, at eye level, in every shop in the country.

7.2 The price series, and the objection to it

Coca-Cola offers economic history a rare gift: a fixed reference point. Because of an inflexible bottling contract struck in 1899, a 6.5-ounce contour bottle sold in the United States for five cents from 1886 until 1959 — seventy-three years of price stability. That makes 1959 the last clean baseline in the product's history and the only sensible place to start.

Comparing like with like means tracking what a person actually buys in order to have a drink: the *single-serve* format, whatever it happens to be made of. That format migrated — 6.5-ounce glass in the 1950s, the 12-ounce can from the 1970s, the 20-ounce PET bottle today — so prices are standardised per litre to strip packaging out of the comparison.

Year	Standard single serve	Price	Per litre
1959	6.5 oz contour bottle	$0.05	$0.26
1969	12 oz can	$0.15	$0.42
1979	12 oz can	$0.25	$0.70
1989	12 oz can	$0.45	$1.27
1999	20 oz PET bottle	$1.00	$1.69
2009	20 oz PET bottle	$1.50	$2.54
2019	20 oz PET bottle	$2.00	$3.38
2026	20 oz PET bottle	$2.48	$4.20

Against a US Consumer Price Index that rose about 1,045% from 1959 to 2026, the per-litre price of single-serve Coca-Cola rose roughly **1,515%**. Had it merely tracked general inflation, today's bottle would cost around $1.75 rather than nearly $2.50.

The objection arrives immediately and must be met, because meeting it is the whole point. *You have compared a returnable glass bottle to a premium convenience format; some of that 1,515% is a change in what is being sold, not gouging.* Correct — and the change in what is being sold is precisely the indictment. Nobody drifted accidentally from refillable glass to single-use plastic. Single-serve is where the margin is; single-serve at scale requires lightweight disposable packaging; and a single-serve portion of that size is what delivers 65 grams of sugar in one sitting. **One commercial decision, three harms**: the price premium, the plastic, and the sugar dose. They are not three coincidences. They are one strategy with three outputs, and separating them to defend the price series would be to miss what the series records.

The second half of the arithmetic is where the money went. Advances in bottling automation, logistics and supply chain over sixty years were extraordinary; in a competitive market, such gains reach the customer as lower real prices. Here they did not. Coca-Cola's gross margin runs consistently in the region of 60–63%, against roughly 30–33% for consumer staples generally. (One honest technical note: that margin is not strictly comparable to a full-chain manufacturer's, because Coca-Cola refranchised most of its capital-intensive, labour-intensive bottling operations to separate companies. Which is itself an observation about extraction architecture — the corporate structure was reorganised so that the high-margin concentrate and brand layer sits with the shareholders, and the asset-heavy, wage-paying layer sits elsewhere.)

7.3 The first externality: plastic

The transition from local refillable glass to centrally filled single-use PET moved the cost of packaging from the company's balance sheet to the world's. Break Free From Plastic's global brand audits have ranked Coca-Cola the world's leading corporate plastic polluter year after year; in recent audits volunteers recovered over 33,000 pieces of Coca-Cola-branded waste across some 40 countries, roughly three times the second-placed brand. The cost of collecting it, of the microplastics that escape collection, and of the ocean gyres where it accumulates falls on municipalities, taxpayers and charities. It appears on no beverage company's accounts.

Here a technical point matters, because it locates the decision correctly. A hostile reviewer will observe that refillable glass is heavy, and that hauling full bottles out and empties back can emit more than lightweight single-use plastic. True — and entirely dependent on distance. Refillable glass wins decisively in *local* loops and loses over long hauls. So the substantive decision was never really glass versus plastic. It was **local bottling versus centralised mega-plants** — and single-use lightweight packaging is what centralisation requires. Circularity was not abandoned because it stopped working. It was abandoned because centralising was more profitable, and circularity was the price of centralising. The oceans paid it.

That reframing carries a consequence for Chapter Two's fourth commitment: local production is not sentimentality about place. It is frequently the physical precondition of circularity.

7.4 The second externality: metabolic health

A 20-ounce bottle contains roughly 65 grams of added sugar — about sixteen teaspoons. That is close to double the American Heart Association's recommended daily maximum for an adult man and over two and a half times the limit for an adult woman, in one bottle. Liquid sugar is uniquely efficient at doing harm because it bypasses satiety: the body does not register it as food, so it does not reduce what is eaten alongside it. The epidemiological link between sugar-sweetened beverage consumption at population scale and type 2 diabetes, obesity, metabolic syndrome and cardiovascular disease is not seriously contested.

The cost of treating that disease is borne by health systems — which is to say, by the public. In the vocabulary of this chapter, it is an unfunded public subsidy to beverage margins: society pays the medical bill for the consumption that generates the profit, and the payment never appears in the price.

7.5 Founder's testimony

The manifesto now changes voice for one section. What follows is first-person testimony from SHALOME.AI's founder, given because the argument of this chapter is ordinarily made from outside the firms it describes, and can here be made from inside.

I worked on Coca-Cola's marketing strategy in the United Kingdom and in India. I also worked on the bottling, distribution and logistics strategy for Coca-Cola Schweppes Beverages in the UK. I did this as a strategy consultant, first at OC&C Strategy Consultants and later at Monitor. My job was to make Coca-Cola more money, and I was good at it.

What was I optimising for? Return on investment, as perceived by Coca-Cola's shareholders. Completely self-centred in economic terms. Nothing to do with the oceans, or with people, or with health systems. No regard for the end consumer or the environment whatsoever.

I want to be precise about this, because it is the part people get wrong. The harm was not weighed against the profit and judged acceptable. It was not concealed from a regulator or argued away by a lawyer. **It was not there.** Not in any of my spreadsheets. Not in any of my presentations. Not in anything.

I drank a lot of Coca-Cola before I did that work. Afterwards I stopped, and eventually I left the profession. It took me some years to understand that I had not been working on a beverage business. I had been working on an extraction machine that happened to sell a drink.

7.6 What the testimony establishes

The testimony is evidence for a proposition this manifesto needs and cannot obtain from outside: **the harm was absent from the model, not defeated within it.**

That distinction decides everything that follows. If externalities were weighed and dismissed, the remedy would be better values in the room — the Business Roundtable's theory, and Section 5's reason for rejecting it. But a cost that has no cell in the spreadsheet cannot be dismissed, because it was never a candidate. There is nothing to override, nothing to argue with, no moment of choice at which a better person would have chosen differently. **And that means no quantity of good character fixes it.** A more compassionate consultant produces the same deck. A chief executive with a conscience approves the same strategy, because the strategy is genuinely optimal against the numbers that exist. Everyone behaves well, by the standards available to them, and the ocean fills with bottles.

Which yields the requirement that Chapter Six exists to meet. The only intervention that reaches this failure is one that **forces the missing cell into the model** — that puts the plastic and the sugar into the return-on-investment calculation itself, where the analyst will optimise against them as diligently as she once optimised for margin. Not an appeal to conscience. A change to the arithmetic.

Chapter Five will show what a firm can build for itself today. Chapter Six sets out what only the state can do: make the prices tell the truth.

8.0 The category error

Between 1984 and 1996, the United Kingdom conducted the largest transfer of long-horizon assets into short-horizon ownership ever attempted in peacetime: telecommunications (1984), gas (1986), airports (1987), water (1989), electricity (1990), and railways (1996) were moved from state ownership into shareholder corporations.

The stated logic was sound as far as it went: the state was a poor operator; political control starved utilities of investment discipline; markets would bring efficiency and capital. And the diagnosis of state failure was substantially correct — this case study is not a nostalgia for nationalisation, and the decades of under-investment *before* privatisation are part of the record too.

But the remedy contained a category error that this chapter's third failure makes visible. A water network is a twenty-to-hundred-year asset: mains laid for a century, reservoirs for generations, treatment works for decades. A railway's embankments, tunnels and signalling are the same. These assets have the longest time signatures in the civilian economy. They were placed inside the ownership structure with one of the *shortest* — dividend-paying, quarterly-reported, equity-priced shareholder corporations — and, in the case of water, natural monopolies at that, facing no competitive discipline whatever.

The prediction from Section 4 is exact: such a structure will not invest for year forty; it will extract now and borrow against the asset, because every incentive within it points that way. What follows is what happened.

8.1 What happened: water

The ten regional water companies of England and Wales were privatised in 1989 **debt-free**, their liabilities written off and a further £1.5 billion "green dowry" of public money added. What shareholders did with a debt-free monopoly over an essential of life is now a matter of regulatory record and Financial Times analysis:

- Over the following three decades the sector paid out approximately **£78 billion in dividends** — nearly half the £190 billion it spent on infrastructure in the same period.
- Over the same period it accumulated **more than £64 billion of net debt** (rising past £70 billion by the mid-2020s) — despite having started with none. Investment was not funded by shareholders; it was funded by customers' bills and by borrowing *secured on the assets*, while dividends flowed out. By 2023, total shareholder equity across the sector stood *below* the inflation-adjusted value of what shareholders had originally injected in 1989–90: over the whole life of the experiment, the owners had taken out more than they ever put in.
- Real household bills rose by more than **40%**, a growing share of each bill servicing the debt and the payouts.
- The physical system told the same story from the other end: chronic leakage, and rivers and coastal waters receiving sewage discharges at a scale that became, by the 2020s, a national scandal.

Industry's defence must be heard, because it is partly true: investment *did* rise after privatisation — the £190 billion spent over three decades roughly doubled pre-1989 annual rates, driven substantially by European water-quality directives that the publicly owned industry, starved of capital by a Treasury in which water competed with hospitals, had failed to fund for a generation. This manifesto accepts both halves of that record. But observe what the defence concedes: the investment was funded by customers' bills and by borrowing secured on the assets — while the £78 billion left as dividends. The fair statement is not "privatisation brought investment"; it is "privatisation brought investment paid for by bill-payers and bondholders, *plus* an extraction of £78 billion that the same bills could otherwise have funded without £64 billion of debt." The relevant comparator was never the starved nationalised industry of the 1970s. It is the third form — the same regulated bills, the same borrowing capacity, and no leakage — as Wales has demonstrated since 2001.

The terminal case is Thames Water, serving sixteen million people. Privatised debt-free, it paid out £7.2 billion in dividends over its lifetime — the most aggressive extraction occurring under a private-equity ownership that loaded the company with borrowing — and by 2026 it carries close to **£20 billion of debt**, has acknowledged "material uncertainty" over its ability to continue, and is the subject of an open letter from 112 Members of Parliament demanding it be placed into special administration. Its creditors, a group of hedge funds negotiating to take control, have sought waivers from pollution fines and relaxation of performance targets as their price for rescue. Read that sentence again slowly: the owners of a water monopoly are negotiating for the *right to pollute* as a condition of continuing to own it. It is difficult to imagine a clearer photograph of the third failure.

8.2 What happened: rail

The railway ran the same experiment faster. Railtrack, the company holding the track, signalling and stations, was floated in 1996. Within four years, maintenance under-investment met its consequence at Hatfield, where a fractured rail killed four people; the emergency speed restrictions and repairs that followed exposed how little the company knew about the condition of its own network. By 2001 Railtrack was insolvent and was placed into administration — five years, in total, from flotation to collapse.

8.3 The counterfactual Britain actually ran

Here the case study turns, because Britain did not merely fail; it also, twice, quietly demonstrated the alternative this manifesto proposes — not as theory but as emergency repair.

Welsh Water. When Welsh Water's diversified parent collapsed in 2000, the business was not sold to another shareholder owner. It was acquired in 2001 — for £1, plus the assumption of £1.85 billion of debt, refinanced through what was then the largest sterling corporate bond issue — by **Glas Cymru**: a company limited by guarantee, with **no shareholders at all**. Its constitution is the asset-lock principle in utility form: there is no one to extract to. Surpluses can only be reinvested in the network, used to build reserves, or returned to customers. Its financing model is built precisely on the time-signature logic of this chapter: with no dividend stream to fund, it raises long-dated bonds from pension funds and insurers — institutions whose own liabilities run for decades and whose appetite is exactly the patient capital a water network's time signature requires. Over its first two decades, more than £440 million that would otherwise have left as dividends was retained for investment; it has run among the strongest credit ratings and balance sheets in the sector; and an Oxford Saïd Business School case study records that it reduced customer bills in real terms while outperforming most shareholder-owned peers on customer metrics.

Network Rail. When Railtrack collapsed, its replacement was not renationalisation and not another flotation. It was, again, a **company limited by guarantee without shareholders** — Network Rail, created in 2002, its surpluses recycled into the railway. The shareholder experiment on rail infrastructure lasted five years; its non-shareholder successor has now run the network for nearly a quarter of a century.

The significance of these two cases is easy to understate and must not be. The United Kingdom, having placed its longest-horizon assets under its shortest-horizon owners, twice watched the structure fail, and twice — under governments of the same broad persuasion that had privatised — reached for the *asset-locked, non-shareholder, surplus-recycling form* as the repair. The form this manifesto proposes is not an untested idea awaiting a brave government. It is the form Britain already uses when shareholder ownership of infrastructure fails, adopted in each case only *after* the failure had been paid for — once in money, once in lives.

The counterfactual must be stated with care, because one version of it is anachronistic. Glas Cymru's 2001 bond issue priced against a mature regulatory regime — a stable, transparent price-control framework — that did not exist in 1989; "ten Glas Cymrus at privatisation" assumes financing conditions the era could not yet have supplied. The defensible counterfactual is, if anything, more damning: the legal form existed (the company limited by guarantee is Victorian; the charitable incorporated organisation, its modern charitable sibling, followed in 2013), and it became demonstrably financeable *within the very regulatory regime privatisation itself created* — at lower capital cost, since the single largest expense of a capital-intensive utility is the cost of its capital, and equity demanding double-digit returns is the most expensive capital there is. From 2001 onward, the shareholder-less alternative stood proven in the same industry, on the same island. No English water company adopted it; the £78 billion continued. What was missing after 2001 was not the tool, nor the proof, but the will — and the recognition that ownership structures have time signatures, and that mismatching them is not a detail but the heart of the matter.

8.4 The honest calibration

Two concessions keep this case study honest, and both sharpen rather than weaken the thesis.

First: not all the privatisations failed, and the pattern of which succeeded is the finding. Telecommunications privatisation is broadly judged a success: prices fell, investment flowed, innovation accelerated. But telecoms differed in kind — technology change made genuine *competition* possible, and competition disciplined capital in precisely the way no water customer can discipline a regional monopoly. The failures concentrate where two conditions coincide: **natural monopoly** (no competitive discipline) and **intergenerational assets** (maximum time-signature mismatch). This yields the discriminating principle the manifesto will carry forward: *where competition can discipline capital, the shareholder form is tolerable; where it cannot — natural monopoly joined to long-horizon assets — the asset-locked form is not merely preferable but the only ownership structure whose time signature matches the asset.*

Second: the asset-locked form removes extraction; it does not guarantee excellence. Welsh Water has itself faced justified criticism of its environmental performance in recent years, documented in a former Welsh minister's review — which, notably, criticises the company's operations and its regulators while expressly crediting the ownership model with avoiding the financial excesses of its English peers. Network Rail's record carries the same double lesson: it has faced sustained criticism over cost efficiency, and its borrowing was reclassified onto the public balance sheet in 2014 — operational and fiscal imperfections this manifesto does not gloss. What neither organisation has done, in a combined half-century of operation, is pay a dividend out of a captive customer's bill. The distinction matters and this manifesto insists on it: structure determines *where the money can go*; management, regulation and accountability determine *what gets done well*. A stewardship company can be badly run. What it cannot be is systematically drained — and a badly run company that retains its surpluses can fund its own repair, where a drained one arrives at its crisis with £20 billion of someone else's debt and creditors demanding the right to pollute.

8.5 What the case study establishes

The misprivatisation record converts each of this chapter's three abstractions into a bill with a number on it. Prices that lied: bills that funded dividends and debt service while rivers took the sewage. Wealth narrowly defined: three decades booked as shareholder value creation that ended with sector equity worth less than was ever put in, and the true balance sheet — pipes, reservoirs, rivers — degraded. Horizons truncated: hundred-year assets governed by quarterly capital, with the difference extracted and the reckoning deferred to a generation that had no vote in 1989.

And it establishes the manifesto's direction of travel. The corrective is not better exhortation of shareholder companies — the Business Roundtable tested that. It is not renationalisation into the political control whose failures motivated 1989. It is the third form Britain has already twice adopted under duress: **ownership without extraction — asset-locked, surplus-recycling, accountable to those it serves, financed by capital whose patience matches the asset.** The chapters that follow generalise that form beyond utilities: first to the firm (Chapter Two), and then to its fullest expression, the charitable architecture (Chapter Three).

---

9. Sources and Evidence Notes (for WS5/WS8 review)

Verified against public sources at drafting (24 Jul 2026); full citation formatting at WS8.

1. **Global growth (E1):** Maddison Project Database — world GDP per capita ~$1,300 (1990 int$) in 1900; ~$6,000 by 2000; world population 1.6bn (1900) → 6.1bn (2000). *Note for DRI: the seed paper's figures ($1,000 → $10,000; population "just under 3 billion in 1900") have been corrected to Maddison-consistent values; the argument is unchanged and strengthened*.
2. **Dasgupta Review (E2):** *The Economics of Biodiversity* (HM Treasury, 2021) — produced capital per head ~doubled 1992–2014; natural capital per head declined ~40%. Verify exact edition figures at WS8.
3. **Water sector financials (E16):** FT analysis (2023/24): £78bn dividends over 32 years vs £190bn infrastructure spend; £64bn+ net debt from a debt-free, £1.5bn-endowed start; Hansard (Jun 2023) and Common Wealth dashboard: debt ~£72bn by 2024; sector equity (£3.4bn, Mar 2023) below inflation-adjusted original injection (£8.2bn in 2023 prices); real bills +40–44%.
4. **Thames Water (E16):** ~£20bn debt and 112-MP open letter for special administration (Jul 2026, Hansard/press); historic dividends £7.2bn (Univ. of Greenwich/Hall); Macquarie-era leverage; creditor demands incl. fine waivers to 2030 (press, Jun–Jul

2026). *Precision kept: Thames itself paid no dividends in its final years — the extraction was front-loaded; the debt remained*.

5. **Railtrack/Network Rail (E18):** flotation 1996; Hatfield 2000 (4 deaths); administration 2001; Network Rail (CLG, no shareholders) 2002. Verify death toll and dates at WS8.
6. **Glas Cymru (E17):** acquisition 2001 for £1 + £1.85bn debt assumption via ~£1.9bn bond issue; CLG, no shareholders, surpluses reinvested; >£440m dividend-equivalent retained over 21 years (Law Wales); Oxford Saïd case study: real-terms bill reductions, strong credit quality; Andrew Davies report (Dec 2024): operational/environmental criticism alongside explicit credit to the ownership model.
7. **Business Roundtable (E7):** 2019 statement; Bebchuk & Tallarita, *The Illusory Promise of Stakeholder Governance* (2020) and follow-ups.
8. **Commercial peace (E12):** interdependence–conflict literature (Oneal & Russett and successors) — supports the softened claim only.

Open items for WS3/WS4 raised by this chapter: the "time signature" concept should enter the glossary; the discriminating principle (7.5) becomes a load-bearing claim (propose as C8); Glas Cymru's bond-financing economics feed the T4 (capital constraint) answer in WS4.

Chapter Two: Five Commitments, Five Mechanisms

1. From Diagnosis to Design

Chapter One located capitalism's failure in its accounting architecture: prices that lie, wealth narrowly defined, horizons cut short. Part of that architecture is set by the state — carbon pricing, national accounts, regulation — and lies beyond any single organisation's gift. But a decisive part of it is set inside the firm: who owns it, who governs it, what it is for, how long it is built to think, and where its surpluses may go. This chapter is about that part.

The founding note of this project stated five commitments for the selfless firm: ownership shared among employees, investors and customers; products and services that genuinely help people; creation in harmony with nature; the company embedded in its community; and people management that is lean rather than hierarchical.

Stated that way, they are ends. And the record examined in Chapter One — 181 chief executives signing the Business Roundtable's stakeholder pledge and then changing nothing — establishes the rule this chapter lives by:

An end without a binding mechanism is a press release. This manifesto claims no commitment for which it cannot name the mechanism that enforces it, the exemplar that proves the mechanism works, and the residue the mechanism honestly leaves unsolved.

Before pairing each commitment with its mechanism, we set out the inheritance: the four management models the industrialised world has actually run at scale, and what each teaches.

2. The Four Inherited Models

The twentieth century ran four great experiments in organising the relationship between workers, managers and capital.

The German model (Mitbestimmung) placed workers inside governance: half the seats on the supervisory boards of large companies, and works councils with co-decision rights on the shop floor. Its lesson is a mechanism, not a mood: executive pay in Germany's largest listed companies runs at roughly 50–80 times median worker pay — a quarter of the American figure — because the people who would bear the cost of excess sit in the room where pay is decided, with votes.

The Swedish model (MBL) achieved similar compression (roughly 35–50:1) by a different route: union directors on unitary boards, a legal duty to negotiate before strategic change, and — most importantly — centrally negotiated pay scales that lift the bottom, tightening the ratio from below rather than capping it from above.

The Anglo-American model concentrated authority in management and returns in shareholders. Its virtue is real: no system reallocates capital and labour to new uses faster. Its pathology is equally real and, as Chapter One argued, mechanical: executive pay is set in equity, by committees composed of peer executives, benchmarked against other inflated packages, with no constituency in the room that pays the cost. The result — an S&P 500 average around 300:1, with FTSE 100 median CEO pay above £4 million and a lowest-paid-to-CEO ratio across the FTSE 350 of roughly 165:1 — is not an excess of the system. It is the system.

The Toyota model (lean) built the most studied innovation engine in industrial history on an economic bargain usually mistaken for a cultural trait: frontline workers volunteer thousands of productivity-improving ideas (kaizen) because an implicit promise of employment security guarantees that no improvement ever costs its author a job. Remove the security and the suggestion box empties. Japanese executive ratios (roughly 50–70:1) are held down not by worker votes but by internal promotion — leaders are lifetime employees, stewards rather than imported superstars, whose pay is historically cut before workers' in hard times.

Two observations organise everything that follows. First, the four models fail in different places — which means their strengths are largely complementary, not mutually exclusive. Second, in every model the durable behaviours are the ones with structure behind them (votes, laws, scales, security), and the fragile ones are those resting on norms alone. Norms erode; mechanisms hold. The five commitments are therefore paired below with mechanisms — and Section 4 adds the fifth column the inherited table lacks.

3. The Five Commitments and Their Mechanisms

Commitment One: Ownership Shared

The end. The firm's ownership — and therefore its returns and its ultimate control — is shared among those who create and depend on its value: employees, investors, and customers.

Why stating it fails. "Sharing value with stakeholders" is precisely what the Business Roundtable signatories pledged. Absent structure, the sharing lasts as long as the goodwill of whoever holds the actual shares. Employee share schemes bolted onto shareholder-primacy firms are routinely diluted, bought out, or reserved for executives; "stakeholder voice" without votes is consultation, revocable at will.

The mechanisms. Ownership is shared when it is *constituted*, not gifted:

- **Worker ownership with governance rights.** The co-operative form makes workers members with votes; the UK **Employee Ownership Trust** (Finance Act 2014) places a controlling stake in trust for all employees permanently — the John Lewis structure, generalised and given tax recognition. The distinction that matters is control, not certificates: an EOT-controlled firm cannot be sold over its workers' heads, and its surpluses flow to the whole workforce.

- **Broad profit participation.** Where full worker ownership is impractical, all-employee profit or equity participation — Handelsbanken's Oktogonen foundation, which for decades paid every employee, cashier to chief executive, the *same* profit share — distributes the upside to where, in a kaizen firm, the ideas actually originate. It also compresses pay ratios endogenously, without caps: when everyone holds the returns, the executive package is negotiated against co-owners rather than granted by peers.

- **Targeted co-determination.** From the German lesson, the narrow, high-leverage version: worker-elected directors holding binding votes specifically on executive compensation and on redundancy programmes — the two decisions where management's interests genuinely diverge from workers'. This captures most of the ratio-compression effect at a fraction of the governance friction of full co-determination.

The exemplar. Mondragon, the Basque federation of worker co-operatives: roughly 70,000 people, €11.2 billion revenue (2024), some 260 co-operatives, competing successfully in advanced manufacturing, retail, finance and education for seventy years — with member-agreed pay ratios ranging from 3:1 to 9:1 and averaging about 5:1. When its original co-operative, Fagor, failed in 2013, the federation redeployed 1,710 of its 1,800 workers into sister co-operatives. Ownership shared is not a constraint on competitiveness; at Mondragon it is the retention, resilience and motivation strategy.

One qualification the honest record requires, because it teaches the manifesto's own scale lesson: Mondragon's international expansion has proceeded substantially through conventionally owned overseas subsidiaries employing non-member workers — the flagship of the pure co-operative form scaled abroad partly by not being one at the edges. The reading this manifesto takes, developed in Chapter Three and priced in the Concessions Register, is that the *pure* member-owned form has an observed scale envelope, and that beyond it the working expression is the hybrid: asset-locked control at the apex, conventional instruments below.

The honest residue. Customer ownership — the third leg of the founding note — is the least solved. The forms exist and are venerable: mutuals, building societies, consumer co-operatives, and Glas Cymru's member-based governance from Chapter One. But customer-owners are structurally diffuse and passive; governance drifts toward management unless deliberately counteracted (the demutualisation wave of the 1990s, in which members voted to liquidate a century of accumulated common wealth for windfall payments, is the cautionary record). This manifesto treats customer ownership as strongest in *natural-monopoly and essential-service* settings — where the customer cannot exit and therefore most needs voice — and as an open design problem elsewhere. We flag it for the research agenda rather than claim it solved.

Commitment Two: Products That Help People

The end. The firm's products and services are constrained to genuine human benefit — health, learning, capability, wellbeing — and the firm cannot drift into harm chasing revenue.

Why stating it fails. Every corporate purpose statement claims benefit. The test is what happens when the profitable and the beneficial diverge — and under shareholder primacy that contest has a pre-announced winner. Product scope is the *first* thing commercial pressure erodes: the engagement-maximising feed, the loot box, the ultra-processed reformulation each began inside firms whose mission statements spoke of connection, play and nourishment.

The mechanisms. Purpose binds only when it is *entrenched above management's reach*:

- **Charitable objects** (the strongest form): a charity's purposes are its legal boundary; activity outside them is ultra vires, and the Charity Commission enforces the boundary. Purpose is not the strategy; it is the constitution.
- **The asset lock**: in a charity or Community Interest Company, assets and surpluses are legally dedicated to the purpose and cannot be distributed for private benefit — removing the very incentive that drives purpose-drift.
- **Steward-ownership and golden shares**: a purpose foundation or trust holds voting control (or a golden share with veto rights over purpose change and sale), so the firm cannot be redirected or sold against its mission — the structure that has held Bosch, Carlsberg and Novo Nordisk to their founders' purposes across a century in which nearly every shareholder-owned peer was bought, merged or repurposed.

- **B Corp articles** (the weakest binding form worth naming): amended articles requiring directors to weigh stakeholder impact. Better than a pledge; still revocable by the shareholders who adopted it — a fact its own history demonstrates and this manifesto does not gloss.

The exemplar. Novo Nordisk: controlled by a foundation whose charter binds it to medical purpose, it pursued diabetes and metabolic disease research across decades of patient investment that produced the GLP-1 class — arguably the most significant drug class of its generation — while remaining among Europe's most valuable enterprises. Purpose entrenchment and commercial excellence are demonstrated complements, not alternatives.

The honest residue. Entrenched purpose can entrench error: a mission fixed in 1900 may be the wrong mission for 2050, and asset-locked structures are deliberately hard to redirect. The charitable world's answer — purposes stated at the level of *human need* rather than *method*, with cy-près procedures for genuine obsolescence — is workable but slower than a shareholder pivot. This is a real cost of bindingness, and Chapter Three prices it rather than hiding it.

Commitment Three: Creation in Harmony with Nature

The end. The firm's production works within natural limits; the environmental costs of its activity appear in its own decisions, not on the commons.

Why stating it fails. This is Chapter One's first failure, restated at firm level: so long as nature is unpriced, the conscientious firm that internalises costs voluntarily is *undercut* by the competitor that does not. Exhortation asks firms to lose; they decline.

The mechanisms.

- **The circularity gate — no launch without an audited end-of-life pathway.** This is the commitment's hard mechanism, and it is a gate rather than a price. Before a product goes to market, the firm must publish how it will be recovered, reused or recycled, with the analysis stated in auditable terms: materials, collection route, realistic recovery rate, who bears the cost. No pathway, no launch. A gate has a property no levy has — **it cannot be passed through to the customer.** A firm with pricing power absorbs a charge and raises the shelf price; it cannot raise the shelf price of a design it is not permitted to sell. Chapter One's case study shows why this matters: the migration from refillable glass to single-use PET would not have survived a gate of this kind, because the pathway that existed was being abandoned rather than improved. The direction is already live in policy — Extended Producer Responsibility schemes and the EU's ecodesign regime move this way — and Chapter Six develops it as a general instrument.

- **Internal full-cost accounting**: a shadow price on carbon (and, increasingly, on water and nature) applied to every investment decision, so the firm's internal signals tell the truth even where the market's do not. This is established practice at serious firms — internal carbon pricing has been used at scale by major corporations for over a decade — and costs nothing to adopt except illusions.

- **Governance standing for nature**: environmental performance in the audited annual report and on the board agenda with the same rank as financial performance; in the charitable form, environmental purpose can be written into objects themselves, giving nature the same ultra-vires protection as beneficiaries.

- **Local production where circularity requires it**: the reuse loop is a physical system, and its economics turn on distance. Refillable containers beat single-use packaging in local loops and lose over long hauls, which means the decision to centralise production is frequently the decision to abandon circularity. Firms serious about this commitment treat the *geography* of production as an environmental variable, not merely a cost one.
- **Time-signature financing** (from Chapter One): assets financed by instruments whose maturity matches ecological payback — the Glas Cymru bond model — so that a wetland restored or a network renewed is not fighting a quarterly clock.
- **And, honestly: the state.** Firm-level mechanisms narrow the gap; only economy-wide pricing closes it. A manifesto for firms must say plainly that it is not a substitute for a carbon price, a plastics levy or a sugar levy — which is why this manifesto does not stop at the firm. Chapter Six sets out the repricing.

The exemplar. The firm-level record is genuinely mixed, and we cite it as such: internal carbon pricing changes marginal decisions but has not, alone, decarbonised any major firm. The stronger exemplars are structural — steward-owned and co-operative enterprises disproportionately populate the long-horizon land, water and energy sectors precisely because their capital can wait for ecological payback. On circularity specifically, the honest exemplar is historical rather than contemporary: the returnable-bottle systems that operated at industrial scale across the twentieth century, and still operate in parts of Latin America and Germany, demonstrate that closed loops work when production is local and deposits are real. The commitment's full redemption is joint work between firm architecture and public policy — which is Chapter Six.

Commitment Four: The Company as Part of the Community

The end. The firm is embedded in, accountable to, and mutually dependent on the community where it operates — an anchor, not an extractor.

Why stating it fails. "Corporate social responsibility" grafted onto an extractive core is philanthropy funded by the extraction: the distribution centre that hollows out a high street and sponsors its Christmas lights. Embeddedness without accountability is public relations.

The mechanisms.

- **Anchor-institution procurement and local wealth-building**: deliberately routing the firm's spending, hiring and investment through its locality — the Preston Model, in which a city's anchor institutions redirected procurement toward local and co-operative suppliers with measurable local-economy effects.
- **Revenue architecture that stays**: the 90/10 model — 90% of revenue remaining with the local practitioner or community enterprise, 10% funding shared infrastructure and cross-subsidy — makes localism a *contractual term*, not a preference.
- **Community consent with teeth**: the mechanism set developed in SHALOME.AI's Catalyst framework and adopted by this manifesto as the general standard: prior consent before deployment; a standing right of pause and withdrawal without penalty; community override where a service causes harm, regardless of operational cost; and an annual community audit, conducted by community members and *published*. Embeddedness means the community can correct the firm — and, at the limit, refuse it.

- **Accountability for harm**: the "name it, do not manage it" discipline — harms described plainly, remediation preceding expansion, the remedy defined by those harmed, the failure record published.

The exemplar. Beyond Preston, the deepest exemplar is again Mondragon, whose co-operatives are constitutionally tied to their valley towns — up to a quarter of profits allocated to community purposes by member vote — and whose region now shows Spain's lowest unemployment and lowest income inequality. Seventy years of anchored capital is a community development programme no CSR budget can imitate.

The honest residue. Community accountability mechanisms are the youngest in this chapter; the consent-and-override architecture has design maturity but not yet decades of operating record. We claim it as the standard to build to, evidenced at pilot scale, not as a proven institution.

Commitment Five: Lean, Not Traditional, People Management

The end. The firm runs on frontline intelligence and autonomy — continuous improvement from the people doing the work — rather than on supervisory hierarchy.

Why stating it fails. "Empowerment" without security is a trap workers correctly refuse: ideas that eliminate your own job, autonomy that means absorbing risk without authority. Lean adopted as a cost programme, without the bargain that makes it rational for workers to participate, curdles into surveillance and speed-up — the abundant graveyard of failed "lean transformations" attests.

The mechanisms.

- **The kaizen bargain, made explicit**: improvement never costs its author employment. This is the Toyota engine's actual fuel.
- **Flexicurity**: the Danish resolution of the bargain's tension with change — protect the *worker*, not the *job*: income continuity, funded retraining, active redeployment. The firm restructures at market speed; the person is never dropped. This single mechanism dissolves the supposed trade-off between Toyota security and Anglo agility, and its price — a genuine commitment of firm and social resources to transition — is the honest cost of frontline innovation.
- **Mission command**: leadership defines intent and constraints; teams hold real decision authority within them; consensus is reserved for decisions that cross boundaries. Buurtzorg runs ten thousand nurses in self-managing teams on roughly 1% back-office overhead; Haier reorganised into thousands of micro-enterprises holding their own P&L. Autonomy scales when authority is bounded, not when it is vague.
- **Distributed strategy tooling**: the capability that was once the moat of the executive floor and the consulting firm — analysis, modelling, planning — distributed to teams (in SHALOME.AI's architecture, this is CMind.AI's explicit function: strategy as a utility, not an elite subscription). Autonomy without analytical capability is abandonment; this closes the gap.

The exemplar. Toyota itself for the bargain; Buurtzorg for mission command; Handelsbanken — which for decades ran branch-level credit authority, no sales targets, no individual bonuses, and superior long-run returns to its bonus-driven peers — for the whole commitment operating in a conservative, regulated industry.

The honest residue. High-trust, high-security systems can breed insularity and conformity (the documented late-stage Japanese pathology). The mitigations — external hiring into the internal pay scale, rotation, porous research partnerships — are management practice, not structural guarantees. Structure enables; it does not absolve leadership, a theme Chapter Four takes up in full.

4. The Master Table: Five Models

The comparative table this project inherited had four columns. Here is the fifth.

Feature	German (Mitbestimmung)	Swedish (MBL)	Anglo-American (US/UK)	Lean / Toyota (Japan)	Asset-Locked / Charitable
Apex ownership	Dispersed shareholders + blockholders	Dispersed + institutional	Dispersed shareholders; market for control	Cross-holdings; stable insiders	**Asset-locked entity (charity/CIO, foundation, trust, CLG, co-op): no extractable ownership; no market for control**
Board structure	Dual board; 50% worker seats on supervisory board	Unitary board + union directors	Unitary board; executives + independents; zero worker seats	Unitary; internally promoted executives	**Trustee/steward board; independent chair; conflicts regime; founder one voice among equals; beneficiary/worker/community representation by design**
Workplace representation	Works councils (statutory)	Local unions	Minimal	Quality circles; enterprise unions	**Members' assemblies (co-op), all-employee trust voice (EOT), community consent-and-override mechanisms**
Managerial obligation	Co-decision on defined matters	Duty to consult and negotiate	Unilateral prerogative	Consensus (nemawashi)	**Fiduciary duty to purpose and beneficiaries — enforceable at law; stage-gated capital allocation; published accountability**

Feature	German (Mitbestimmung)	Swedish (MBL)	Anglo-American (US/UK)	Lean / Toyota (Japan)	Asset-Locked / Charitable
Job security approach	Preserve existing jobs	Retrain and redeploy (flexicurity)	At-will; low security, high agility	Implicit lifetime employment	**Kaizen bargain + flexicurity: person protected, not post; restructuring permitted, abandonment not**
Union / member role	Sectoral wage-setting	Integrated at every level	Adversarial	Cooperative enterprise unions	**Membership constitutional (co-op/mutual) or trust-mediated (EOT); community holds audit and override rights**
Where surpluses can go	Dividends + retention	Dividends + retention	Dividends + buybacks dominate	Retention-heavy; modest dividends	**Reinvestment, reserves, beneficiaries, price/bill reduction only — extraction legally impossible**
Time signature of capital	Medium (blockholders lengthen it)	Medium	Short (quarterly reporting; daily repricing)	Long (stable insiders)	**Longest: purpose is perpetual; financing by long-dated bonds and retained surpluses matches intergenerational assets**
CEO-to-worker pay ratio	~50–80:1 (DAX)	~35–50:1 (OMX)	~300:1 (S&P 500 avg); FTSE 100 median CEO £4.2m, ~113:1; FTSE 350 lowest-to-CEO ~165:1	~50–70:1 (Nikkei majors)	**~2–8:1 observed.** UK charity median CEO £59,850 (2025) vs. sector staff medians → low single digits; top-100 charity CEOs avg £192k → ~5–8:1; Mondragon capped 6:1 (range 3–9:1, avg 5:1)
Innovation engine	Engineering depth; patient blockholders	Consensus + high trust	Capital reallocation speed; venture formation	Kaizen at scale	**Purpose-driven patient R&D (Novo/Bosch pattern); frontline kaizen under the secured bargain;**

Feature	German (Mitbestimmung)	Swedish (MBL)	Anglo-American (US/UK)	Lean / Toyota (Japan)	Asset-Locked / Charitable
					distributed strategy tooling
Characteristic failure	Slowness; insider ossification	Outlier suppression; small-country scale	Extraction; short-termism; inequality	Insularity; conformity; overwork	**Capital-raising constraint; potential managerial slack absent extraction pressure; purpose ossification — priced honestly in Chapter Three**

On the fifth column's pay figure — how it was constructed, so it can be attacked fairly. The other columns report listed-company disclosures; no equivalent single statistic exists for the asset-locked sector, and a naïve comparison of medians would be a size mismatch (the median charity is far smaller than the median S&P 500 company). The construction therefore leads with the **size-matched comparison**, which is the only fair frame — and which the asset-locked form still wins decisively. At the top of each sector: chief executives of the UK's 100 largest charities average £192,000 (2025), implying mid-single-digit ratios against their own staff, while the *highest-paid charity executive in Britain* — running one of the world's largest research endowments — earns roughly one-seventh of the *median* FTSE 100 chief executive's £4.2 million package. At federation scale: Mondragon's 70,000-person system operates a constitutional cap of 6:1 (range 3:1–9:1, average 5:1). The sector-wide figures (median charity CEO £59,850 against staff medians in the high £20,000s — low single digits) are reported for completeness, not as the headline. The claim made is deliberately modest: not that every asset-locked organisation sits at 5:1, but that at *matched scale* the form's observed range (roughly 2–8:1) and the shareholder form's observed range (roughly 100–300:1) do not overlap — because the mechanism that generates the upper range (pay set among peers, funded from extractable surplus) is absent by construction.

5. The Gradient of Bindingness

Reading the five commitments together reveals the pattern that organises this manifesto's remaining argument. The mechanisms above are not equally strong. They form a gradient:

1. **Policies and pledges** — revocable by management. (Business Roundtable tier: worthless under pressure.)
2. **Contracts and schemes** — profit-sharing plans, procurement policies, consent protocols: durable until renegotiated.
3. **Entrenched governance** — amended articles, worker directors, golden shares: durable until the controlling interest changes.
4. **Controlled ownership** — EOTs, foundations, steward trusts: durable while the trust endures and its deed is honoured.

5. **Statutory purpose with an asset lock** — the charitable form: purposes enforceable at law by a regulator, assets permanently dedicated, extraction not merely discouraged but *legally impossible*, in perpetuity.

Each commitment is only as durable as the tier of its weakest mechanism. And here is the observation the next chapter builds on: **the five commitments, pursued to the top of the gradient, converge on a single organisational family.** Shared ownership at its most binding is membership or trust ownership with no external extraction; entrenched purpose at its most binding is charitable objects; harmony with nature at its most binding is purpose plus perpetual-horizon capital; community accountability at its most binding is enforceable beneficiary duty; and the lean bargain at its most binding is security guaranteed by an owner with no incentive to defect — because there is no one to extract the savings to.

That family — asset-locked, trustee-governed, beneficiary-accountable, surplus-recycling — is the charitable and asset-locked architecture. (The nearest prior art for this convergence is Marjorie Kelly's distinction between *extractive* and *generative* ownership design; this manifesto's contribution is to run that distinction to the top of the bindingness gradient and name the legal family waiting there.) Britain reached for it twice as emergency repair (Chapter One). One organisation has designed it, in complete governance detail, as a deliberate operating system. That is the argument of Chapter Three.

6. Carried Forward

- **Proposed for the convergence map:** C8 (the discriminating principle, from Chapter One) and **C9: the gradient of bindingness —** commitments are as durable as their weakest mechanism's tier.
- **Glossary additions:** time signature; gradient of bindingness; kaizen bargain (already present — cross-reference).
- **To WS4:** the fifth column's "characteristic failure" row is a promissory note — capital constraint, slack risk, and ossification must be priced there, not waved at.
- **To WS5 (red team):** the pay-figure construction in §4 is offered for attack; the customer-ownership residue (Commitment One) and the environmental exemplar concession (Commitment Three) are the chapter's acknowledged soft points.
- **To research agenda:** customer-ownership governance design for essential services (T2 remains open).

---

7. Sources and Evidence Notes (for WS5/WS8 review)

1. **Pay ratios (E3):** DAX/OMX/Nikkei ranges per inherited comparative table (S3) — refresh at WS8; S&P 500 average ~300:1 (recent figures 299–344:1 across sources/years); FTSE 100 median CEO £4.22m and ~113:1, FTSE 350 lowest-to-CEO ~165:1 (High Pay Centre, 2023–25).
2. **Mondragon (E4):** 70,085 employees, €11.213bn revenue, ~260 co-ops (2024 annual report via public summaries); ratios 3:1–9:1, average ~5:1, standard cap 6:1 (Co-operative News 2025; Columinate; CSM 2024); Fagor redeployment 1,710/1,800 (2013–14 record).

3. **Handelsbanken (E5):** Oktogonen equal profit share; decades of decentralised credit authority and superior long-run ROE — verify current status of Oktogonen contributions at WS8 (scheme was modified in 2010s/2020s; cite historical record precisely).
4. **Novo Nordisk / Bosch (E6):** foundation control and purpose charters; GLP-1 significance; Bosch patent volume — verify current figures at WS8.
5. **EOT (E10):** Finance Act 2014; controlling-interest requirement; John Lewis constitution as precursor.
6. **Charity pay (new — E19):** ACEVO Pay and Equalities Survey 2025 (median £59,850; n=703); Charity Finance 2025 top-100 survey (£192k avg; highest £580k); High Pay Centre FTSE comparators as above.
7. **Buurtzorg / Haier (E9):** ~10,000 nurses, ~1% overhead claims; Rendanheyi micro-enterprise counts — verify at WS8.
8. **Preston Model (E14):** CLES documentation of procurement redirection and local-economy effects.
9. **Demutualisation record (new — E20):** UK building society demutualisations, 1990s — supports the customer-ownership residue; add sources at WS8.
10. **Internal carbon pricing (new — E21):** prevalence among major corporations (CDP disclosures); honest-mixed effectiveness record.

Chapter Three: The Charitable Architecture

1. Where the Argument Has Arrived

Chapter One diagnosed the failure — prices that lie, wealth narrowly defined, horizons cut short — and showed, in the wreckage of the water privatisation and the quiet success of its shareholder-less repairs, that ownership structures have time signatures, and that mismatching them is ruinous. Chapter Two paired the five commitments of the selfless firm with binding mechanisms, ordered those mechanisms on a gradient of bindingness, and ended with an observation: pursued to the top of the gradient, all five commitments converge on a single organisational family — asset-locked, trustee-governed, beneficiary-accountable, surplus-recycling.

This chapter names that family, describes its anatomy precisely, argues that it — and not the reformed corporation — is the natural organisational form of selfless exchange, and then does what advocates rarely do: prices its costs, one by one, on the record.

The chapter has a second purpose. An architecture proposed in the abstract invites the reply that it would never survive contact with commercial reality. This chapter therefore works from a complete, existing blueprint: the governance corpus of SHALOME.AI, a charitable organisation designed from first principles to hold a portfolio of technology ventures inside the charitable form. The claim is not that SHALOME.AI has proven the model at scale — it is an organisation in formation, and this manifesto's own claims discipline forbids pretending otherwise. The claim is narrower and verifiable: that a complete operating system for compassionate enterprise — constitution, capital discipline, risk stratification, delivery boundaries, accountability machinery — already exists in working detail, and can be examined, criticised, copied and improved.

2. The Anatomy of the Form

The charitable architecture, in its full expression, has five components. Each solves a problem the preceding chapters identified; together they constitute the form.

2.1 The asset-locked apex. At the top sits an entity that cannot be owned: in the UK, a Charitable Incorporated Organisation (CIO) — a corporate body with limited liability, registered with and regulated by the Charity Commission, whose assets are locked in perpetuity to its charitable purposes. There are no shares, no shareholders, no market for control, and no lawful route by which surpluses become private wealth. This is Chapter Two's fifth tier of bindingness: extraction is not discouraged, capped or shamed; it is *legally impossible*. Adjacent legal wrappers achieve the same architecture with different trade-offs — the company limited by guarantee (Glas Cymru, Network Rail), the Community Interest Company with its statutory asset lock, the community benefit society, the purpose foundation — and this manifesto uses "the asset-locked form" for the family and "the charitable form" for its fullest expression, where purposes themselves are legally enforceable and the regulator holds the boundary.

2.2 Trustee governance. The apex is governed not by owners' representatives but by trustees: individuals bound by fiduciary duty *to the purpose and its beneficiaries*, personally liable for acting in the charity's best interests, subject to a statutory conflicts regime — and, under the UK's voluntary principle, generally unpaid. Pause on that last fact, because no corporate governance reform has ever proposed anything so radical: the governing body of a charity has, by default, **no financial stake in its decisions at all**. The compensation-committee capture that generates 300:1 ratios is not mitigated in this form; it has no surface to attach to.

2.3 The delivery boundary. The charity governs; it does not operate everything it governs. Where activity requires statutory, clinical or professional accountability, it is delivered by partners who hold that accountability — regulated providers, certified professionals, accountable local enterprises — under partnership agreements the trustees oversee. This single rule answers, in advance, the oldest objection to purpose-led organisations: that well-meaning generalists end up doing specialists' work badly. In this architecture they are constitutionally prevented from trying.

2.4 Trading subsidiaries. Commerce happens — vigorously — one level down. The charity owns trading companies: ordinary limited companies, competing in ordinary markets, employing at market rates, run by commercial managers. Their distinguishing feature is their destination: profits flow up to the charitable parent (in the UK, corporation-tax-efficiently under Gift Aid), where the asset lock converts them permanently into mission fuel. This is the dual structure that steward-owned giants have run for a century — foundation above, competitive enterprise below — and it dissolves the lazy dichotomy between "charity" and "business": the architecture is *both*, in the correct order.

2.5 The accountability exoskeleton. Around the whole structure runs machinery that shareholder firms lack because no one can compel them to build it: public annual reporting to a regulator with intervention powers; published accounts; and — in the full expression this manifesto advocates — the consent-and-audit architecture of Chapter Two's fourth commitment: community consent before deployment, standing rights of pause and withdrawal, community override, annual community-conducted audits, published including failures.

That is the form. Four structural consequences follow from it, and they are the argument.

3. Four Structural Advantages

3.1 Extraction is removed at the root, not restrained at the branch. Every reform of the corporation — say-on-pay votes, disclosure rules, stakeholder articles, worker directors — restrains a pressure that remains present and permanent, funded and patient, probing for the weak point. The charitable form does something categorically different: it removes the pressure's source. There is no residual claimant whose wealth grows when workers' pay is cut, when maintenance is deferred, when the community is squeezed. The 2–8:1 pay ratios observed across the asset-locked sector (Chapter Two, §4) are not the product of superior virtue; they are what pay-setting looks like when nobody in the room can extract the difference. pay ratios observed across the asset-locked sector

3.2 The time signature is perpetual. A charity's purposes do not expire, and its capital cannot demand exit. It is the only organisational form whose native horizon matches the longest-horizon assets and missions — watersheds, educational institutions, research programmes, the intergenerational commitments Chapter One showed quarterly capital destroying. It finances accordingly: retained surpluses and long-dated bonds, the Glas Cymru model, placed with pension funds and insurers whose own liabilities run for decades. Note what this does to the cost of capital: equity demanding double-digit returns is the *most expensive* funding a capital-intensive enterprise can use, and the asset-locked form is the only one that can decline it entirely.

3.3 The founder and the executive are constrained by design. Chapter Two's convergence claim C4 held that any serious form must bind its own leadership. The corporate form binds it weakly and increasingly runs the other way — dual-class shares now entrench founders *against* accountability. The charitable form binds natively: the founder sits, if at all, as one trustee among equals, subject to the conflicts regime, required to step back where personal interests arise; the SHALOME.AI constitution states the principle without decoration — *the charity must be able to function, govern, and thrive independently of the founder*. Chapter Four will argue this is also the structure most likely to *select* leaders worth having.

3.4 Accountability runs to those affected. A shareholder company is accountable, enforceably, to capital; every other accountability is discretionary. A charity is accountable, enforceably, to purpose and beneficiaries — with a public regulator, public filings, and (in this manifesto's full expression) communities holding audit and override rights. The direction of enforceable accountability is the deepest single difference between the forms, and everything else in this chapter is downstream of it.

4. The Worked Blueprint

Objection, anticipated: *charities are woolly. The form may prevent extraction, but it also prevents discipline — no owner's pressure means no rigour.* The strongest reply is not argument but exhibit. What follows walks through the SHALOME.AI governance corpus as a demonstration that the charitable form can carry a capital-allocation and risk discipline *stricter* than most corporate boards operate. Readers from the commercial world are invited to compare each element with their own governance honestly.

4.1 The constitution. The trustee framework establishes collective duties (public benefit, legal compliance, safeguarding as a governance responsibility, separation of charitable and commercial risk), an independent chair charged with constructive challenge and founder management, and named oversight roles — including one with no corporate equivalent: an **Ethics, Research and Boundary Oversight Trustee**, whose function is to police the organisation's own claims, quarantine speculative work from public representation, and "help the board distinguish between philosophy, research, and delivery." A governance role whose job is institutional epistemic honesty: the plc that creates one will be the first.

4.2 Capital discipline: the stage-gate. Every proposed venture passes an alignment gate (which charitable purpose, which measurable human outcome, who benefits first) and a set of hard blockers — safeguarding, truthful capability claims, lawful data, conservative economics, capacity, a cheap test before a big build, one named accountable individual — any one of which pauses the proposal. Scored proposals proceed only at defined thresholds: GO at ≥80%, pilot-only at 65–79%, otherwise no. Commitments are staged — exploration under 5% of investment, pilot 15–20%, scale only on evidence — with kill criteria defined *before* launch. The corpus even names the failure modes it is designed to resist, in a checklist titled the **"emotional mugging" detector**: *a donor wants it; a competitor has it; we've already spent too much; it'll be quick; it'll be great for marketing.* Boards that recognise their own last three acquisitions in that list may draw their own conclusions about which sector is woolly.

4.3 Risk stratification: the two-pillar model. Ventures are classified by risk and governed accordingly: application ventures (months to MVP, market-driven) run for adoption, revenue and operational hand-off; breakthrough ventures (long-horizon, research-driven) are quarantined — governed by scientific milestones, independent review, and explicit go/no-go gates, with their speculative status stated rather than marketed. Reputational contagion between the two is treated as a governed risk. This is portfolio governance of a sophistication most conglomerates do not practise, inside a charity.

4.4 The delivery boundary, operationalised. The corpus maps, pillar by pillar, what the charity does (frameworks, platforms, governance, partnership oversight) and what accountable partners do (clinical delivery, teaching, regulated services) — under a rule stated as non-negotiable: if a domain requires statutory, clinical or professional accountability, the charity does not deliver it directly. Trustee competence (§5.3 below) should be read against this boundary: the trustees are never asked to govern brain surgery, only the partnership under which surgeons remain accountable for it.

4.5 The humility machinery. The corpus binds itself to a claims discipline (speculative content quarantined; regulatory guidance supreme over internal framing), to community consent, override and annual community-conducted audits, and to a harm protocol of unusual candour: name harm plainly rather than in euphemism; halt expansion while any harm claim is unresolved; let those harmed define the remedy; publish the failure record, on the reasoning that *a programme with no published failures is a programme that has stopped looking for them*. And the corpus practises what it legislates: its own strategy documents carry their revision history openly, recording that earlier drafts' triumphalist language — including a claim that its systems made corruption "impossible" — was withdrawn under independent audit, with the withdrawal reasoned in public: *systems that describe themselves as making wrongdoing impossible are systems that have stopped listening for it.* This manifesto knows of no listed company that documents its own rhetorical failures with that candour, and invites correction on the point.

4.6 The distribution model. Finally, the corpus answers how the form reaches communities rather than only boardrooms: not by delivering services to passive recipients, but by certifying local practitioners — identified by capability *or by community nomination, with equal weight* — who run their own micro-enterprises on the platform, keep 90% of their revenue locally, commit to a public covenant of stewardship (including developing their own successors and disclosing their use of AI), and stand accountable to their communities through published ratings, review rights and suspension powers held by the community itself. Ownership shared, product constrained to benefit, community embeddedness and anti-dependency, operating as one design.

The blueprint's status, restated with the precision the claims discipline demands: **designed and documented, partially piloted, unproven at scale.** What it refutes is not the question "will this work everywhere?" — no document can refute that — but the assertion "no such architecture could be specified." It has been.

5. Pricing the Costs

An architecture is only as credible as its honest bill. Six costs are priced here. Five are standing objections to the charitable form, each answered to the extent the evidence allows and conceded where it does not. The sixth is not an objection outsiders make; it is one the form has to make about itself, and it is the heaviest.

5.1 The capital constraint. *Charities cannot issue equity; how do they fund growth?* Three answers and a concession.

First, the constraint is smaller than it appears, because equity is the most expensive capital there is. For the capital-intensive, long-horizon activities where this manifesto argues the form belongs, debt capital priced against stable revenues beats equity demanding double-digit exit returns: Glas Cymru financed the acquisition of an entire water utility with what was then the largest sterling corporate bond issue, has funded a quarter-century of investment from bonds and retained surpluses, and holds among the strongest credit ratings in its sector — *because* no dividend leakage competes with its bondholders. For infrastructure, the "constraint" is a discount. But this answer has an honest boundary, and the manifesto draws it rather than letting critics find it: bonds fund *asset-backed, revenue-backed* activity. They do not fund pre-revenue risk. The Glas Cymru evidence establishes the form's advantage for the economy's infrastructure core; it establishes nothing about venture-stage enterprise, which must look to the second and third answers below.

Second, the dual structure opens conventional funding one level down. Trading subsidiaries are ordinary companies: they can borrow commercially, take project finance, and — where mission-compatible and trustee-approved — admit minority equity investors at subsidiary level while the charitable parent retains control. This is the Novo Nordisk architecture (foundation control above a listed operating company) translated into UK charity law, where a charity may own trading companies outright or in part, subject to its investment duties. The manifesto flags honestly that *large-scale* versions of this hybrid remain rare in UK charity practice and deserve dedicated legal development — an item for the implementation chapter's research agenda.

Third — and here the architecture is not a handicap but an advantage rarely noticed — **the locked group moves its own capital more efficiently than a shareholder group can.** Surplus generated by one entity can be applied to another within the group without a tax event, under the ordinary Gift Aid mechanics of §4.1, and with no dividend leakage on the way. Nor is there any reason to build the machinery conventional groups assemble to hold capital "efficiently" — the offshore intermediate holdings and layered vehicles whose real function is to obscure how much is being made and by whom — because in a locked group there is nothing to obscure. No extraction means no concealment, and no concealment means none of the professional cost of concealment. Removing the self removes an entire industry of expensive complexity from the balance sheet. It does not remove every kind of complexity: the form sheds the costs of hiding what is made and acquires a different tendency to accumulate, which §5.7 addresses.

Fourth, sequenced self-funding: the corpus builds its portfolio so that earlier, cash-generative ventures fund later, capital-hungry ones. Seed capital is still required, as it is for any venture in any form, and the honest statement of the constraint is narrower than it is usually made. It is not that the form cannot be funded. It is that the founder must find seed capital *that will not extract* — and that pool is thinner and more relationship-dependent than the conventional one: philanthropic and foundation capital, mission-aligned family capital, or the founder's own. Conventional angel and venture terms are priced to extract, which is exactly why they are unavailable here. The money is not absent. The terms are the thing being refused.

That constraint then belongs to Chapter Four's argument rather than against it. **Capital selection is part of the selection loop.** A philanthropic or mission-aligned seed investor selects for alignment exactly as the absent jackpot selects for selfless leadership — the same filter, operating one layer earlier in the story. What looks like a narrower funding market is a narrower funding market by design.

One correction the manifesto owes its own evidence, because an earlier draft of this section got it wrong. Long horizons in this model are **not** symptoms of capital scarcity. The corpus's 33-year framing exists because the work is civilisational: research programmes, institutional adoption and scientific validation run on their own clock, and no quantity of capital shortens it. General relativity waited years for an eclipse to test it; the Higgs mechanism waited nearly fifty for a collider. That is the precise opposite of the three-year cycle the shareholder form rewards — raise prices, cut marketing and people, book the improved bottom line, collect the options, leave the next occupant a hollowed business. A long horizon is this architecture's purpose, not its penalty.

The concession that does stand: this form will not blitz-scale. An enterprise whose plan requires a billion pounds of risk equity before revenue will not raise it inside a locked structure, and this manifesto does not pretend otherwise. Its reply is that the domains where selfless exchange most matters — care, learning, water, food, community enterprise, long-horizon research — are precisely the domains where blitz-scaling was always the wrong clock. Where the hyperscaler model genuinely fits, the shareholder form exists, and Section 6 states the boundary honestly.

5.2 The incentive question. *If no one can get rich from surpluses, who will build?* Distinguish three groups.

Employees and executives of trading subsidiaries are paid at market: charity law permits, and commercial reality requires, competitive remuneration for commercial roles — what the architecture removes is not good pay but the *extraction multiple* on top of it, the equity jackpot funded from everyone else's forgone share. Entrepreneurial practitioners in the distribution model keep 90% of what they build — local upside, genuinely owned, at the level where Chapter Two argued the ideas originate. Breakthrough teams, Chapter Two conceded, may warrant real but capped venture-style participation at subsidiary level; the corpus's stage-gate can govern it.

What the form cannot offer is the founder's nine-figure exit. The manifesto declines to apologise for this, and offers instead a claim it is careful to state at its true strength: a **designed tendency, not an automatic law**. An architecture that cannot produce jackpots tends not to attract people primarily seeking them — Chapter Four develops this as a hypothesis, with its predictions stated so it can be wrong. But the filter has a second failure mode the honest version must name: the alternative to the greedy is not only the mission-driven — it is also the inert. A form that cannot pay tournament prizes can select low-energy administration as easily as high-purpose stewardship, and the charity sector's chronic recruitment difficulties are evidence of that risk, not of the filter working. The answer is the same as for slack (§5.4): the tendency operates only alongside *active* selection — competence-first recruitment, the Selfless Leadership Test asked of records (whose questions about repair, truth-telling and institution-building the inert cannot pass), and market-rate subsidiary pay so the form competes for builders. Two residual honest costs, then: some genuinely mission-driven builders will still choose the equity path for family-security reasons no covenant should shame; and the form must *work* for its leadership quality where the tournament firm can simply pay for applicants.

5.3 Trustee capacity. *Volunteer boards governing complex ventures — really?* The objection lands where the architecture is not followed: charities whose boards drift into operations, or govern regulated complexity directly. The architecture answers structurally — the delivery boundary keeps regulated complexity with accountable specialists; the two-pillar model matches governance mode to risk; trading subsidiaries have their own professional boards; and the stage-gate gives trustees a decision instrument rather than a general invitation to have views. What remains is real: the form demands skilled trustees, recruited for competence rather than availability, inducted seriously, and — in the specialist oversight roles — genuinely expert. And one of the form's proudest features cuts both ways here, which the manifesto concedes rather than conceals: the voluntary principle purchases incorruptibility with a recruitment constraint. The same absence of financial interest that removes capture (§2.2) also narrows the pool — unpaid, personally accountable governance of complex ventures is a proposition some of the most qualified will decline. Charity law's lever exists (payment for trustee services, properly authorised, with its own governance costs) and should be named rather than treated as unmentionable. The obligation is bounded, not unbounded — the two-pillar model and professional subsidiary boards mean the scarce commodity is a small number of genuinely skilled trustees — but it is a resourcing obligation the form must budget for, not a refutation; the corpus's induction and role specifications are the template.

5.4 Managerial slack. *Without extraction pressure, what disciplines management?* This is the strongest objection in the set, and it deserves respect: the pressure of a demanding owner, whatever its pathologies, is *pressure*, and its absence has produced sleepy institutions before. Four disciplines substitute, and the evidence is that they can exceed the one they replace. Purpose-metrics published against a regulator's gaze (beneficiaries served per pound is a harder number to massage than adjusted EBITDA). Bond-market covenants — Glas Cymru answers to credit analysts whose scrutiny of cash generation is not gentler than an equity analyst's, only longer-horizon. The internal capital-allocation machine — a stage-gate with pre-committed kill criteria disciplines more decisions than a share price does, because it operates *before* commitment rather than after. And competitive markets themselves: trading subsidiaries face customers who can leave, which no ownership structure suspends. Handelsbanken's decades of superior returns without individual bonuses stand as the long-run exhibit that internal discipline can beat extraction discipline. The concession: these disciplines must be *built* — the form permits slack in a way a raider-patrolled market does not, and an asset-locked organisation that neglects its discipline machinery will drift. The corpus exists precisely because this is true. Note that slack and capture are different failures and take different answers: slack is work not done, and the four disciplines above address it; capture is surplus absorbed, and it is the business of §5.7.

5.5 Purpose ossification. *Perpetual purpose can mean perpetual wrongness.* Conceded as a genuine cost of bindingness, and managed rather than solved: purposes drafted at the level of human need rather than method (educate, heal, sustain — not "operate schools of the 1889 pattern"); the cy-près jurisdiction for genuine obsolescence; and the periodic review of §5.7, which tests whether the purpose still requires this organisation in this form. The shareholder firm pivots faster; the charitable firm abandons less. Which failure mode is worse depends on the domain — and returns us to the boundary question.

5.6 Scale and dependency: the red team's conclusions, adopted. On scale, the adversarial review's finding is accepted and stated as the manifesto's position: the *pure* asset-locked form is unproven above roughly the hundred-thousand-person scale — and the flagship exemplars are themselves hybrids (a foundation above a listed company; a federation whose overseas plants are conventionally owned). The scalable expression of the model is therefore the hybrid this chapter has described all along: asset-locked control at the apex, conventional instruments below. The stronger claim sometimes made for decentralisation — that capability distribution makes centralised scale *unnecessary* — is held as a designed hypothesis, to be proven or indicted by the published record of the pilots built to test it, not asserted in advance. On dependency — the risk that a charitable provider hollows out state duty-of-care and local enterprise — the manifesto's accountability machinery is the designed answer, and its status is published in the Concessions Register in exactly these words: designed, not yet proven; the record will show whether we honour it.

5.7 The institution as its own claimant. *If no owner can take the surplus, where does it go?* This is the cost the form must price against itself, and the chapter states it plainly: the asset lock removes the claimant outside the organisation and leaves one inside it.

Consider a charity founded to fund research into a single disease. In year one it has three staff, a borrowed room, and one purpose; ninety-four pence in the pound reaches a laboratory. Nothing that follows is a scandal. In year four it hires a fundraiser, who brings in more than she costs. In year seven it takes a lease, because the files must go somewhere. In year nine the trustees adopt a reserves policy, correctly advised that a charity this size should have one. In year twelve the chief executive's salary is benchmarked, found low, and raised. In year fifteen a communications post is created, because a charity that cannot explain itself cannot raise money. In year twenty-two a policy function is established, on the sound reasoning that influencing national research funding is more efficient than funding research directly. In year twenty-six the reserves policy is revised to three years and the reserves are invested, and the investment income becomes a budget line the organisation comes to depend upon. By year thirty the charity has two hundred staff, a building, a brand and an annual conference, and forty-one pence in the pound reaches a laboratory. Every decision was taken by honest people who justified it in writing to a board that agreed. No trustee benefited. Fifty-three pence in the pound has nevertheless stopped reaching the purpose, and it did not go to an owner. If a shareholder had taken it, this manifesto would have a word for it.

The mechanism is the absence the form is proudest of. In the shareholder company, cost that produces no return has an enemy — an unattractive one, whose interest in efficiency is an interest in his own dividend, and whose pathologies this book has spent three chapters cataloguing. But he exists, he is informed, and he has power. Remove him and the drift above has no natural opponent. Beneficiaries rarely see the ratio and never see the accounts. Donors benchmark against organisations drifting at the same rate, so the comparator moves with the problem. The regulator confirms that nothing improper occurred, and nothing improper did. The trustees have watched each step arrive singly with a good argument attached. This is a failure with no culprit, which is why it cannot be prosecuted and must be designed against.

The answer this manifesto adopts is the same reversal it applies everywhere else. **Continuation must be earned.** An asset-locked organisation should be founded with a date on which its authority to continue expires — not a date on which it is reviewed, but one after which it exists only because people with no stake in the answer have said it should. Today an institution carries on unless somebody makes a case against it, and that case falls to people with no standing, no information and no power; which is why almost nothing ends, and why inertia, which is patient, sits on the institution's side of the table. Reverse the burden and it lands where the knowledge is: the organisation makes the case for its own continuation, in public, against criteria it did not write and cannot change, to assessors it did not choose and cannot pay. Chapter Five §7.6 gives the instrument — the cadence, the assessor design, the graduated findings and the fixed criteria — and three costs are conceded with it. Criteria fixed at foundation will bind successors to judgements that age badly. Separable organisations cannot cross-subsidise as freely as unified ones, which is the sharpest objection to the design and is answered only partly by covenanted funding. And the review itself consumes money and attention that would otherwise reach beneficiaries: a genuine transfer from present service to future integrity, which an organisation doing good work in its ninth year will feel as an imposition. All three are entered in the Concessions Register.

Most organisations will pass, and that is the expected result rather than evidence of a toothless instrument. The value is not in the closures but in what the prospect does to an institution that knows it is coming: it documents, because knowledge held in one head cannot be handed on; it keeps its parts separable, because parts that cannot stand alone cannot be released; it watches its own reserves and its own pay, because someone with no reason to be kind will be looking at them soon. And some will fail and should end. An organisation that has completed its purpose, or that has quietly become a machine for its own maintenance, does no dishonour by closing; it is doing the last thing the purpose requires, which is to get out of the way. Institutional death reads as failure only because we have confused the vessel with what it carries. The company has no self. Neither should the institution outlive the reason it was made.

6. The Boundary of the Claim

This manifesto does not propose the abolition of the shareholder corporation, and its argument is stronger for saying so plainly.

Chapter One's discriminating principle set the line: where genuine competition disciplines capital and assets are short-cycle, the shareholder form is tolerable and its reallocative speed is valuable. Where competition cannot discipline — natural monopoly, essential services, captive users — and where assets or missions are intergenerational, the shareholder form's time signature does structural harm, and the asset-locked form is not merely preferable but the only match.

To that structural line this chapter adds a human one: **the more the enterprise's users are vulnerable, and the more its product is trust, the stronger the case for the charitable form** — because vulnerable users cannot police extraction, and trust is precisely what a residual claimant is structurally incentivised to monetise. Care, health, education, children's services, community finance, essential utilities, long-horizon research: these are the charitable form's home ground. Consumer discretionary goods in competitive markets are not, and the manifesto happily leaves them to the plc — under, it adds, the full-cost accounting rules of Chapter One, which apply to everyone.

The proposal, stated in one sentence: **for the economy's trust-critical, long-horizon, monopoly-prone core, the asset-locked asset-locked form should be the default and the shareholder form the exception requiring justification — the exact reverse of today's presumption.** Thatcher's error, in Chapter One's terms, was not selling; it was selling into the wrong time signature when the right form lay unused in Victorian company law. The next such choice — and AI infrastructure, care systems and the energy transition guarantee there will be next such choices — can be made with the blueprint on the table.

7. Carried Forward

- **To WS5 (red team):** the five priced costs are offered for attack, with §5.4 (slack) flagged as the objection this chapter considers strongest; the subsidiary-equity hybrid (§5.1) needs charity-law stress-testing; T6 (scale) and the dependency risk are formally handed over.

- **To WS6:** §3.3 and §5.2 both end at the same door — structure selects and disciplines leadership but cannot replace it. Chapter Four owes the account of what kind of leader the architecture needs and how the SQ lens supplies it.
- **To WS7:** the implementation decision tree must cover: CIO vs CLG vs CIC vs community benefit society vs EOT; trading-subsidiary mechanics (Gift Aid, investment duties); the subsidiary-equity hybrid as a development item; transition pathways for existing firms and — per Chapter One — for the next infrastructure decision.
- **Glossary additions:** dual structure; delivery boundary; accountability exoskeleton; voluntary principle.
- **Proposed for the convergence map:** C14 — *the asset lock relocates the residual claimant rather than eliminating it: extraction by an external owner is made impossible, and absorption by the institution itself becomes the form's characteristic failure. The form is therefore incomplete without a periodic review in which continuation is the exception requiring justification*.
- **Proposed for the convergence map:** C10, as qualified by red-team finding R13

---

8. Sources and Evidence Notes (for WS5/WS8 review)

1. **CIO/charity law (E10/E11):** Charities Act 2011 (CIO provisions; charitable purposes; public benefit); Charity Commission guidance on trustee duties (CC3), trading (CC35), payments to trustees; Gift Aid mechanics for trading-subsidiary profit transfer. Confirm current guidance editions at WS8.
2. **Voluntary principle:** default unpaid trusteeship under UK charity law; note lawful exceptions (payment for services with authority). Precision check at WS8.
3. **Glas Cymru (E17):** bond financing, credit standing, quarter-century record — as verified at WS2.
4. **Novo Nordisk / steward-ownership dual structures (E6):** foundation-controlled listed operating company; UK translational questions flagged, not asserted resolved.
5. **Handelsbanken (E5):** long-run returns without individual bonuses — historical record, tense-checked per WS3 flag.
6. **SHALOME.AI corpus (S4–S6):** trustee roles and founder principle; stage-gate, hard blockers, emotional-mugging detector; two-pillar model; pillar mapping and delivery boundary; claims discipline; consent/override/audit; harm protocol; withdrawn-"impossible" revision record; Catalyst certification, dual-pathway identification, 90/10, Covenant of Stewardship. All quoted or paraphrased from the project documents; page-level citation at WS8.
7. **Status precision:** SHALOME.AI described as "in formation; designed and documented, partially piloted, unproven at scale" — Ethics Trustee sign-off required on this phrasing at WS9.

Chapter Four: Structure Selects Character; Character Sustains Structure

In the highest antiquity, the people did not know that there were rulers. In the next age they loved them and praised them. In the next they feared them. In the next they despised them. Thus it was that when faith in the Dao was deficient in the rulers a want of faith in them ensued... Their work was done and their undertakings were successful, while the people all said, "We are as we are, of ourselves!"

— Lao Tzu, *Tao Te Ching*, ch. 17, tr. James Legge (1891)

A note on the epigraph. The model of leadership set out in this chapter was arrived at independently, from engineering and strategy practice, and the resemblance was noticed only after the chapter was drafted. We include the passage because the convergence is itself evidence: a conclusion reached from two directions that share no vocabulary — twentieth-century management experience and fourth-century-BCE Chinese philosophy — has a stronger warrant than either route alone. Note also that Lao Tzu sets a harder standard than this manifesto does. He ranks the leader who is loved *as merely second best, on the reasoning that a people aware of their leader are to some degree oriented around him rather than around their own capability. Chapter Two's fourth commitment, and the success measure this manifesto publishes — that those served should need the organisation* less *over time — are the institutional form of that same demand.*

1. The Door the Architecture Leaves Open

Three chapters have argued for architecture: true prices, binding mechanisms, the asset-locked form. If the argument stopped there, it would repeat — at a higher level — the very error it diagnoses. For the record shows two failure patterns, not one.

The first is familiar from Chapter Two: *ends without structure*. The Business Roundtable's 181 signatories pledged stakeholder purpose and changed nothing, because nothing bound them.

The second is the one this chapter exists to face: *structure without selflessness*. This manifesto owes its readers the asset-locked form's failure roster, presented as fully as its exemplars — and here it is.

The British building societies were asset-locked mutuals, owned by their members, some for a century and a half — until, in the 1990s, members were invited to vote their own institutions out of existence for windfall payments, and management teams who stood to gain from conversion led them enthusiastically to do it. **Equitable Life**, the world's oldest mutual insurer, collapsed in 2000 under guarantees its own board had written and could not honour — a quarter of a million policyholders' trust destroyed by governance that member ownership had not disciplined. **The Co-operative Bank** came within reach of failure in 2013, its celebrated democratic structure having installed a chair with no meaningful banking qualification — the member-democratic process itself producing the governance disaster. **Kids Company** collapsed in 2015, a charity whose charismatic founder had run trustee oversight as a formality for years while the board — on paper impeccable — deferred. And the demutualisation wave stands as the summary case: sound structures dismantled *from within, by their own legitimate processes*, when the stewardship culture that was supposed to animate them had hollowed out.

Read honestly, the roster teaches one lesson four times, and it is a lesson *for* this manifesto's architecture rather than against it: **the asset lock removes the extraction failure mode; it does not remove the governance failure mode.** Every case above is a failure of precisely the discipline machinery Chapters Three and Five make constitutive of the form — skills-based trustee recruitment, institutionalised dissent, boundary oversight, founder constraint, entrenchment against windfall votes. None is a failure of the lock, and none was prevented by member democracy alone. That is why the machinery is not optional equipment, and why Chapter Three's own strongest concession (§5.4) pointed here: the asset-locked form *permits* slack, drift and capture in a way a raider-patrolled market does not — which is precisely why it demands more of the people who hold it, not less.

So the claim of this chapter, stated at the outset: **architecture and character are complements, and each is the other's failure mode.** Ends without mechanisms decay into decoration; mechanisms without selfless people are dismantled or gamed. The manifesto's own source history demonstrates both directions at once: the SHALOME.AI corpus is a set of binding mechanisms *created because* its founders assumed their own good intentions would drift — and its revision record, in which independent audit stripped triumphalist language from a document whose heart was sound, shows even high intentions requiring structural correction. Humility built the structure; the structure now guards the humility.

What, then, is the quality the architecture needs in its people — and can an organisational form do anything to find it, keep it, and survive its absence? The answer begins with a diagnostic developed, in this project's source work, far from the boardroom: in the study of why the most educated, most credentialed, most strategically fluent leaders of the modern era repeatedly produced catastrophe.

2. The Missing Quotient

Modern selection systems — for prime ministers and for chief executives alike — are good at detecting certain capabilities. Can the candidate think? (IQ: analysis, strategy, abstraction.) Can they read and move people? (EQ: empathy deployed as influence.) Can they command institutions and execute? (PQ: the power quotient.) Each is real, and each is insufficient in the same way:

Quotient	What it measures	Its failure mode, ungoverned
IQ	Cognitive intelligence; analysis; strategy	Clever cruelty; brilliant rationalisation
EQ	Emotional intelligence; reading people	Manipulation; charm without conscience
PQ	Command of institutions; execution	Domination; rule by fear or spectacle
SQ	**Moral-spiritual maturity; ego subordinated to service**	**When absent: the narcissistic-saviour pattern**

Spiritual Quotient — SQ — names the fourth dimension: the degree to which a person's exercise of power is organised around service rather than self. It is not religiosity; a person may be devout and spiritually immature, or secular and deeply mature. It is the capacity to hold power without worshipping oneself — and it is the discipline that tells the other three quotients what they are *for*. High IQ without SQ yields intelligent destruction; high EQ without it, manipulation; high PQ without it, domination; high charisma without it, a cult.

The source essay behind this chapter derived the pattern from the history of imperial and interventionist leadership — Vietnam, Iraq, Afghanistan, the repeated photographs of displacement that look identical across decades — and identified the signature of low SQ in power: the leader who believes themselves uniquely necessary; who experiences criticism as betrayal rather than possible correction; who treats people as instruments of the mission; who uses moral language while avoiding moral responsibility; whose polished realism ("hard choices must be made") sanitises the suffering their choices assign to others. And its most disturbing finding: this pattern is *not* produced by ignorance. It is produced disproportionately by the most elite institutions on earth — education without humility, brilliance without reverence, fluency without wisdom.

One boundary is inherited from the source and kept absolutely, here and everywhere this manifesto travels: **SQ is an interpretive moral lens on the exercise of power, not a clinical instrument, not a psychometric, and never a score attached to a named individual.** We read patterns in conduct; we do not measure souls. Any use of this chapter to rank, grade or diagnose persons is a misuse, and the manifesto disowns it in advance.

3. From Empire to Enterprise

Why does a manifesto about economic architecture need a framework forged on the history of war? Because the pattern translates, and the translation explains the specific leadership pathology of the extraction economy. One thing must be said before the first parallel is drawn, and said without qualification: the stakes differ by orders of magnitude. A redundancy is not a massacre; a stripped pension is not a refugee road; nothing in commerce is Saigon, and the analogy owes the dead that acknowledgement. What translates is not the magnitude of the harm but the *structure of the conduct* — promise, use, abandon; decide at a distance from the consequence; launder the choice in the language of seriousness — and it is the structure, recurring across domains and scales, that the SQ lens exists to read.

The elite pipeline translates. The same institutions that supplied the imperial mandarinate supply the executive and financial elite; the "superstar CEO" of the Anglo-American model — the indispensable outsider commanding a free-market bidding war, mythologised as the singular cause of collective outcomes — is the commercial cousin of history's indispensable-man delusion. The 300:1 ratio is not just an economic fact; it is a *theology*: a system's declaration of belief that one person is worth three hundred, which no evidence on executive marginal product supports and which the internally promoted leaders of the German, Japanese and Mondragon traditions falsify daily at 5:1 to 80:1.

And the pipeline's characteristic blindness is documented in this manifesto from the inside. Chapter One §7.5 carries the founder's testimony of strategy work for a global beverage business: high analytical capability, elite firms, work performed well by every standard the work was scored against — and the oceans, the health systems and the end consumer nowhere in any spreadsheet or presentation. That is what this chapter means by intelligence without the fourth quotient. Not malice, and not stupidity. A first-class mind, correctly optimising a model from which the human consequences had been omitted before it ever reached the analyst's desk. The elite-school disease is not that its graduates are cruel. It is that they are trained to solve, superbly, whatever problem they are handed, and never taught to ask what the problem statement left out.

The language translates. Chapter One's "prices that lie" has a rhetorical counterpart: vocabulary that launders harm into seriousness. *Leverage. Credibility. Collateral damage. Acceptable losses.* And their commercial siblings: *headcount rationalisation. Difficult decisions. Fiduciary duty* — invoked not as the modest legal standard it is, but as moral cover for choices the speaker preferred anyway. The educated extractive mind, like the educated imperial mind, rarely says "I am taking from them." It says "I am being realistic." Low-SQ leadership is hard to expose precisely because it arrives wearing the language of adulthood.

The abandonment pattern translates. The source essay's most damning exhibit was the repeated abandonment of local allies — recruited, used, and left as a visa backlog when the politics changed. Its commercial form is documented in every deindustrialised town: workforces whose loyalty was solicited for decades and liquidated in a restructuring; suppliers hollowed out by payment terms; pension schemes lighter than the promises made to them; the private-equity playbook of Chapter One's water sector, in which the asset's community wears the debt while the extraction has already left. Great powers remember their strategic interests; the abandoned remember the promises. So do towns.

And the test translates. War, the source essay argued, is the great SQ test because it reveals whether a leader can still see individual human beings inside the strategic map. The commercial equivalents are precise: the restructuring (can the leader still see the person inside the "headcount"?); the pricing decision over a captive customer; the safety-versus-schedule call; the moment the profitable and the beneficial diverge and no one outside the room would ever know. Chapter Two built mechanisms so that such moments are structurally rarer and structurally visible. This chapter's point is that no mechanism abolishes them — and what walks into that room is character.

4. Extraction Selects Its Own Psychology

Here the manifesto's two halves — architecture and character — lock together, because the relationship between them is not additive but *selective*. Organisational forms are filters: over time, each attracts, promotes and retains the psychology its incentives reward.

Consider what the shareholder-primacy form, at its Anglo-American extreme, structurally rewards. Tournament pay concentrates astronomical prizes on a single seat — and tournaments for astronomical prizes are entered disproportionately by those who most crave them. Charisma is priced: markets for "superstar" leadership select for self-presentation, overclaiming and certainty, the low-SQ signature dressed as vision. Quarterly horizons reward the leader who harvests what others planted and departs before consequences land. The extraction multiple itself — the jackpot — functions as bait whose takers are, on average, those for whom the jackpot is the point. None of this means every plc leader fits the pattern; the form contains many quiet stewards, constrained and often frustrated by it. It means the *gradient* runs the wrong way: the corporate ladder, at its extractive extreme, tends to filter *for* the narcissistic-saviour pattern and then hand it power — the inverse of what anyone designing the stewardship of consequential institutions would choose.

Now run the charitable architecture through the same lens, feature by feature, and observe that each of its structural properties doubles as a psychological filter:

- **No jackpot** (the asset lock): the prize that baits extraction psychology is absent by construction. Chapter Three called this a selection filter; this chapter completes the thought — what remains to attract a leader is *the work, the mission, fair pay and the company of peers*, which is the high-SQ motivational profile almost verbatim.
- **Founder as one among equals** (the trustee constitution): the seat offers service without a throne. The saviour needs the throne. Nobody else ever did.
- **Criticism institutionalised** (red teams, devil's advocates, boundary-oversight trustees): the low-SQ leader experiences dissent as betrayal and exits or is exposed; the high-SQ leader, for whom criticism may be correction, experiences the same machinery as support.
- **Published failure records and community override**: intolerable to the self-mythologising, ordinary weather to the servant.
- **Accountability running to beneficiaries**: the form's deepest property (Chapter Three, C10) is also its deepest filter — power held *in trust for others* is the definition of the psychology it selects.

Hence this chapter's title, and the manifesto's proposed convergence claim C11: **structure selects character, and character sustains structure.** The claim's evidentiary status is stated with the precision it deserves: it is a **hypothesis**, grounded in incentive analysis, in the organisational-behaviour tradition of attraction-selection-attrition, and in the exemplar record (the multi-decade leadership quality of the internally promoted steward traditions) — not an established law, and not self-certifying. It therefore makes predictions on which it can be wrong, and the manifesto states them: if the selection loop is real, stewardship-form organisations should show, on average, longer leadership tenures, pay-governance outcomes that stay compressed without regulatory force, and a failure typology dominated by governance lapses rather than extraction events — while shareholder-form failures at the extractive extreme should continue to cluster around extraction. The research agenda (Chapter Five §10) commissions exactly this comparison. The two failure records of Section 1 are the loop running in reverse — and the loop's virtuous direction, *actively managed rather than assumed*, is the real answer to Chapter Three's slack objection: what ultimately disciplines an asset-locked organisation is not any single mechanism but the compounding selection of people who came for the mission, governed by machinery that assumes even they will drift.

But selection is only half of what the form does to the people inside it, and it is the half with the weaker claim. The next section makes the stronger one.

5. The Conversion Machine

The selection argument, left alone, is weaker than it looks — and its weakness is worth naming before its strength is claimed, because the objection arrives immediately and is fatal if unanswered.

Selection requires unusual people. It says the ladder finds the rare type and hands it the top. Against which the defence is simple, respectable, and endlessly available: *a few bad apples*. Remove them, screen better, tighten the vetting, and the form is exonerated. Every institution facing a scandal of character reaches for this defence, and the reason it works is that it concedes the harm while protecting the architecture. If the problem is who got in, the remedy is a better door.

This manifesto makes the harder claim, which needs no rare types at all. **The extraction form does not merely recruit the extractor. It converts the marginal.**

5.1 Second testimony

The manifesto changes voice for a second time, for the same reason it did in Chapter One §7.5: the claim of this section is ordinarily made from outside the rooms it describes, and can here be made from inside them.

I have been in rooms with senior executives since I was twenty-five. I am sixty-one. That is thirty-six years of watching how people at the top of large organisations actually conduct themselves — not as an outside observer or a journalist, but as the consultant they hired, in the meetings where the decisions were taken.

What the literature calls dark-triad conduct is not, in my experience, a minority phenomenon at that level. It is the majority. And for a long time I assumed the fault must be in my sample — that I had been unlucky, or had drawn the wrong sector, or was reading ordinary decisiveness as something worse.

I want to be careful, because the objection has force. Almost none of these people would meet a clinical definition of narcissism, and I am not qualified to offer one, and it would be an abuse of the word to try. So set the diagnosis aside and ask three questions about conduct instead.

Were they extremely self-centred? Did they, in any operative sense, not care about the people their decisions landed on? Would they do whatever it took, largely irrespective of the harm it caused?

To all three, unequivocally, yes.

And here is the part I understood much later. These were not unusual men and women. They were behaving exactly as the system asked them to behave, and being paid extraordinarily well for it, and being promoted above the ones who hesitated.

5.2 What the testimony can and cannot establish

It cannot establish a prevalence rate, and this manifesto does not claim one. It is a single practitioner's observation over a long career, unblinded and unsampled, and a hostile reader is entitled to say so.

But two features of it survive that objection, and one of them cuts in an unexpected direction.

The first concerns measurement. The published prevalence estimates for these traits in executive populations should be read as a floor rather than an estimate, for reasons internal to how they are produced. The trait cluster was originally validated substantially on incarcerated and clinical populations — that is, on people who had already failed at concealment. The instruments used in corporate settings are largely self-report or colleague-rating, administered to the population that is by definition most practised at, and most rewarded for, the management of how it appears. A trait defined partly by strategic impression management, measured by asking people to describe themselves, is not a well-posed measurement problem. Where the observation of long-serving insiders and the output of the instruments diverge, it is not obvious that the instruments should win.

The second concerns the sample, and it argues for the structural reading rather than against it. A strategy consultant does not meet executives at random. He meets them at the moments they commission strategy work: the cost programme, the restructuring, the competitive push, the portfolio decision with human consequences attached. If these were stable dispositions carried between contexts, one would expect them distributed evenly across settings. If they are structurally elicited — summoned by a particular kind of room, a particular kind of decision, a particular set of rewards — one would expect them concentrated in exactly the settings this testimony describes. The observation is more consistent with the form producing the conduct than with the conduct arriving ready-made.

Which is why this manifesto states its claim in the vocabulary of the three questions and not in the vocabulary of the clinical triad. The diagnostic language is a trap for an argument like this one: it invites a dispute about thresholds, prevalence rates and construct validity that the argument does not need to win, and it concedes the ground of *conduct* — which is observable, arguable from the record, and where the case is actually strong — in exchange for the ground of *disposition*, where it is not. We are describing what people do. We leave what they are to those with the standing to say.

5.3 The asymmetry that does the work

Conversion needs two things, and the first is a payoff structure. Enormous upside; weak, delayed, or absent downside.

The documentary record is available without any recourse to psychology, and this section deliberately keeps to it — reporting outcomes rather than inner lives, in keeping with the boundary set in Section 2.

The Texas City refinery explosion of 23 March 2005 killed fifteen workers and injured a hundred and eighty. The independent panel that reported in January 2007 found that BP had not provided effective process-safety leadership, and identified an organisation measuring its safety performance overwhelmingly by occupational indicators — slips, trips, falls, vehicle accidents — while the systems standing between it and catastrophe deteriorated under sustained cost pressure. In the year of the explosion the group had asked its refineries to cut a further quarter from their budgets. The chief executive of that period left in 2007 over an unrelated matter. He had lied to a court about how he met a former partner, admitted the statements were untruthful, and resigned on the day the injunction fell; the judge declined to refer the matter for prosecution, disclosure in the judgment being punishment enough.

Going early cost him money, and the company said how much: at least £3.5 million of his retirement package, and possibly a further £12 million from the long-term share plan. He had been due to leave in July with a pay-off of £5.3 million and a pension valued at £21.7 million.

Read that ledger in the order it was written. Lying about a boyfriend cost him up to £15 million. Fifteen deaths at Texas City cost him nothing at all.

The Deepwater Horizon disaster of 20 April 2010 killed eleven men and produced the largest marine oil spill in the history of the petroleum industry. It occurred under his successor, who left the board that year with twelve months' salary in lieu of notice, a pension, and an entitlement to performance shares then worth some £1.5 million. He was denied his bonus for 2010 — the remuneration committee recorded that the scale of the spill "overrode the normal metrics for bonus outcomes" — and was given a non-executive seat on the board of the company's Russian joint venture, at a fee of around £150,000 a year.

Note what the withheld bonus tells us. The board was capable of imposing a cost, understood that it should, and said so in writing. What it imposed was the removal of a discretionary addition to a package that remained, after its removal, worth several million pounds and a further income stream. That is the ceiling of what the form permits: the sanction is real, it is recorded, and it is trivial against the harm.

In November 2012 the company pleaded guilty to eleven counts of manslaughter, two misdemeanours and a felony count of lying to Congress. No individual senior executive was convicted of any of it. The corporation confessed to the deaths; no person did.

This is the pattern that matters, and it is more instructive than any single villain would be: **the culture was built under one leader and detonated under the next, and neither bore a cost proportionate to the deaths.** No screening process detects a man who will leave before the consequences arrive. The distance between the decision and its consequence was manufactured by the form.

In banking the asymmetry is starker still, because there the punishment was administered and its scale is a matter of public record. The chief executive of the Royal Bank of Scotland presided over a loss of £24.1 billion — the largest in British corporate history — and a bank rescued at enormous public expense. His knighthood was cancelled and annulled by Letters Patent in February 2012, three years after the event. His pension was reduced.

The sequence of that reduction is worth stating precisely, because compressing it flatters him. His entitlement on leaving was some £703,000 a year, inflated by a scheme that credited him with membership from the age of twenty rather than forty, and that released at fifty the benefits he would otherwise have drawn at sixty. In February 2009 he commuted part of it, taking an estimated £2.7 million in cash and leaving £555,000 a year. In June 2009, following an internal inquiry that found no wrongdoing, he agreed to accept £342,500. Had he been dismissed rather than permitted to retire early, the figure would have been £416,000 from the age of sixty. Had the government allowed the bank to fail, the Pension Protection Fund would have paid him £28,000 a year from sixty-five.

Hold on to that last number. The rescue that cost the public tens of billions was worth, to him personally, the difference between twenty-eight thousand a year and three hundred and forty-two thousand — on top of the £2.7 million already banked. That was the punishment.

No chief executive of a British or American bank was convicted of a criminal offence arising from the crisis. One was charged — the former group chief executive of Barclays, in June 2017 — and cleared in June 2019. Elsewhere the outcome differed. Ireland sentenced the former chief executive of Anglo Irish Bank to six years' imprisonment in June 2018, for conspiracy to defraud and false accounting: he had inflated the bank's deposits by €7.2 billion in order to mislead the market about its condition. Iceland convicted the chief executive and the chairman of Kaupthing, and a number of others besides.

The obvious reply is that those were convictions for specific criminal offences rather than liability for a collapse, and the reply is correct. It is also the point. Falsifying a deposit book is a crime in London and New York exactly as it is in Dublin and Reykjavík. What differed between the jurisdictions was not the law. It was who stood close enough to a documented act to be reached by it — and the architecture of a large bank is precisely an architecture for ensuring that the people who set the incentives never stand that close.

Set the two sides next to each other. On one: enough money to keep a family for three generations. On the other: a bonus withheld, a knighthood cancelled, a pension cut — each of them arriving years after the event, and none of them leaving a mark.

Anyone can do that sum. That is the whole difficulty. You do not need a greedy man to produce this outcome. You need an ordinary one who can add up, standing where the structure has put him.

5.4 What the marginal person sees

The second requirement is a norm, and it is the more powerful of the two.

Human conduct tracks *descriptive* norms — what people observe others actually doing and getting away with — far more reliably than it tracks *injunctive* ones, which is to say what everyone agrees ought to be done. This is among the more robust findings in the study of behaviour, and it is the mechanism by which the extraction form does its work on people who arrived with ordinary consciences.

Consider the executive at the margin. Not the extractor: the one who would think three or four times before doing something that lands badly on someone else. She looks upward and observes, without ambiguity, that the conduct in question is not merely tolerated but rewarded, that its practitioners hold the seats she wants, and that the colleagues who hesitated are still where they were five years ago. Nothing in that picture tells her that hesitation is *wrong*. It tells her something more corrosive: that hesitation is *naïve* — an unworldly failure to understand how things are done, of the kind one grows out of.

And so she updates. Not because she has become a different person, and not in a single decision she would recognise as a fall. In increments, each defensible, each surrounded by people making the same increment, until a conduct she would once have found difficult is simply how the job is performed. She has not been corrupted by an unusual temptation. She has read the rules of the game accurately and begun to play them, which is what competent people do.

The system did not need to find her. It only needed her to look up.

5.5 The badge of honour

There is a final turn, and it is the one that makes the mechanism worse than mere accommodation.

The televised deception formats — the most successful of which now runs in more than twenty countries — are the closest thing to a public demonstration of this process that exists. Ordinary, pre-screened, otherwise decent volunteers adopt sustained strategic deception within days, given nothing more than a prize and social permission. That much is unsurprising. The instructive part is what happens next: the deceivers do not experience their deception as a moral cost reluctantly paid. They experience it as *skill*. They become proud of it. They seek and receive credit for how well they did it, and the audience awards that credit, and the successful deceiver leaves the format wearing the achievement as a badge.

That is not moral weakness. It is **moral inversion** — the point at which a conduct is not excused but *reclassified*, from a thing done despite its wrongness into a competence the person wishes to be known for. And once conduct is reclassified, it needs no further defending. Nobody argues for their own skill; they display it.

The commercial register is the same register. The restructuring described as courage. The hard-nosed reputation cultivated and enjoyed. The pride taken in the deal that took the counterparty apart, told as an anecdote at dinner, met with admiration. Chapter One's testimony recorded the harm being *absent from the model*; this section records something further along — the harm present, understood, and converted into a professional credential.

A system that merely tempted people would be manageable. A system that hands out esteem for the conduct has closed the loop, because the last defence an ordinary person has against doing something they think is wrong is the expectation of being thought less of. Remove that, replace it with admiration, and no conscience is required to fail. It is simply no longer being asked anything.

5.6 Why this is the more damning claim

Notice what the two arguments do to the remedy, because that is the whole reason this section exists.

If extraction required unusual pathology, the remedy would be a personnel problem — screen harder, vet more carefully, promote more wisely — and it would be soluble *inside the shareholder form*, which is precisely what two decades of governance reform have assumed. If ordinary people, most of them, reliably converge on this conduct once the jackpot and the ladder and the esteem are in place, then no quantity of screening reaches it, because there is nothing rare to screen for. **The form is not hosting the disease. The form is the disease.**

This is the same structural move Chapter One made about the missing cell in the spreadsheet, and it should be recognised as such. There, no quantity of good character fixed a cost that had never entered the model. Here, no quantity of good character survives a payoff structure that rewards its absence and an esteem economy that celebrates it. Both arguments end in the same place: the remedy has to be architectural, because the appeal to better people has already been tried by everyone who has ever run one of these institutions, including the good ones.

And now the asset lock can be seen doing something Chapter Three did not credit it with. It has been described as a filter — the jackpot removed, so the jackpot-seeker does not apply. That is true and it is the smaller half. The larger half is that the same removal takes the pressure off *the person at the margin*. Where there is no windfall to be had, hesitation stops being expensive. The executive who thinks three or four times is no longer paying for the privilege, no longer watching less scrupulous colleagues convert her caution into their promotion, no longer being taught daily that her conscience is a career-limiting eccentricity. The compression of pay ratios that Chapter Three defends on grounds of fairness turns out to have a second and possibly larger function: it removes the prize whose size was doing the converting.

The form is not a device for finding good people. It is a device for not requiring them — and, in the same motion, for not spending them. That is a claim about the ordinary majority rather than the deviant few, and it is the version of this argument the manifesto stands behind.

Evidentiary status, stated as this manifesto states such things. The conversion mechanism is a **hypothesis of the same standing as C11**, resting on a documented payoff asymmetry, on the established behavioural finding regarding descriptive norms, and on practitioner testimony that is real evidence and partial evidence. It is not a demonstrated causal chain in executive populations, and no study known to this manifesto has established it as one. Its prediction is stated so that it can fail: if conversion is real, the *same* individuals moving between extraction-form and asset-locked-form organisations should show conduct differences not attributable to selection — and the research agenda at Chapter Five §10 is amended to commission that comparison.

6. Guarding the Guardians: Humility as Architecture

That last clause matters most. The naïve reading of this chapter would be: *select high-SQ leaders and the rest follows*. The manifesto's sources forbid that reading. Power erodes the character it is given — moral licensing, deference creep, the slow conversion of "the mission needs me" into "I am the mission" — and no one is exempt, least of all founders, and least of all founders of *virtuous* projects, where the cause lends its halo to the ego wearing it. The SHALOME.AI corpus is unusually explicit that its own machinery exists on this assumption: constructive dissent is *institutionalised* rather than hoped for; listening infrastructure is *built* (skip-levels, beneficiary panels, always-on anonymous channels, a "bad news fast" protocol) rather than left to an open-door pledge; the chair's role description includes, in terms, the ability to challenge the founder; and the meta-rule of its leadership framework — *earn the right to lead by serving the mission, not your certainty* — is operationalised as **confident humility**: clarity of values and destination, genuine corrigibility about the path.

Its sharpest instrument deserves generalising: the **SQ integrity audit** — an independent review of an organisation's language and posture, hunting the specific tells of the saviour pattern (conquest metaphor, claims of certainty or impossibility of failure, speaking *about* communities rather than *with* them) and requiring their removal from documents of record. The corpus's manifesto went through four such audited revisions; the withdrawn claims are published with their reasons. This manifesto proposes the practice as a standard governance instrument for any asset-locked organisation: an audit not of the books but of the *voice* — because, as every chapter of the source history shows, the voice drifts first, and where the voice goes, conduct follows.

7. The Lens in Practice

How does a board use SQ without violating the boundary set in Section 2? Four disciplined uses, none of which scores anyone:

In trustee and executive selection, as a domain of inquiry rather than a rating: references and histories examined for the pattern-level evidence — how the candidate handled being wrong, what those with less power than them say, whether their past exits protected the vulnerable or the CV. The appendix questions are designed to be asked *of a record*, not of a soul.

In the chair's handling of the founder and the executive, as a shared vocabulary that makes the difficult conversation possible: it is easier to say "this decision needs the Section 10 questions" than "I think power is changing you" — and a leadership culture that has adopted the vocabulary in calm weather can use it in storm.

In board self-assessment, collectively and on the record: the corpus's quarterly leadership self-assessment already asks *have I created space for dissent; what am I not hearing; am I building to outlast myself* — the appendix instrument extends the set. Answered honestly in minutes that persist, these questions leave a trail that future trustees can audit.

In succession, as the tie-breaker it should always have been: between two capable candidates, the record of ego subordinated to service is the qualification, not the garnish — the Toyota-tradition insight (leaders as lifetime stewards, first among equals) elevated from culture to policy.

8. What We Do Not Claim

Four limits, on the record, in keeping with the manifesto's discipline. **SQ cannot be measured**, and any organisation claiming to quantify it has already misunderstood it; the lens reads conduct patterns, always provisionally, always open to correction — an assessment made with certainty is itself a low-SQ act. **Character does not guarantee outcomes**: high-SQ leaders inside sound architecture will still sometimes fail — markets turn, pilots disappoint, harms occur; the difference the pairing makes is not infallibility but *honest failure* — named, repaired, learned from, published. And **this chapter's typology describes patterns in the exercise of power, not verdicts on persons**: the same leader may act from high SQ in one season and low in another, which is exactly why the machinery of Section 6 never stands down. Finally, **the named cases in Section 5.3 are cited for their documented outcomes only** — findings of inquiries, sums paid, sanctions imposed — and for what those outcomes reveal about the payoff structure a form presents to everyone standing under it. They are not, and must not be read as, characterisations of the individuals named. The asymmetry is the evidence; the men are not the argument.

9. Carried Forward

- **Proposed for the convergence map:** C11 — structure selects character; character sustains structure (the selection loop as the resolution of the slack objection).
- **Proposed for the convergence map:** C13 — *the extraction form converts the marginal actor as well as selecting the extractive one; therefore the remedy cannot be screening, and the asset lock's removal of the jackpot protects the hesitant as much as it deters the acquisitive.* (§5)
- **Glossary additions:** SQ (with the T8 boundary embedded in the definition); confident humility; SQ integrity audit; narcissistic-saviour pattern (pattern-level definition only); **the conduct test** (the three questions of §5.1); **moral inversion** (§5.5); **descriptive norm** (§5.4).
- **To WS5 (red team):** invited attack on §4's selection argument (is the evidence for form-psychology sorting more than plausible inference?); on §5's conversion argument — in particular whether the second testimony's sample objection is answered or merely restated, and whether the televised-deception exhibit is evidence or illustration; on the building-society reading of §1; and on whether §3's empire-enterprise translation overreaches. The chapter's claims here are interpretive and should be tested as such.
- **To WS7:** the SQ integrity audit as a specified governance instrument (template terms of reference); selection-inquiry protocol for trustee recruitment.
- **To WS8:** this chapter, of all chapters, must pass its own test — audit its voice for the certainty and grandiosity it condemns.

10. Sources and Evidence Notes (for WS5/WS8 review)

Chapter Four's core argument is interpretive, and this note exists because §5.3 is the exception: it makes verifiable factual claims about named organisations and individuals, and those claims carry the section's weight.

1. **Texas City, 2005 (§5.3):** refinery explosion, 23 March 2005; fifteen dead, in excess of 170 injured. The independent safety review panel reported January 2007 and its findings on process-safety leadership, cost pressure, and the conflation of personal-injury metrics with process-safety performance are as characterised. **Verify: exact casualty figures, the panel's precise findings, and the wording of any quotation before publication.**

2. **BP chief executive tenure and departure terms (§5.3):** the 1995–2007 tenure, the 2007 resignation over admitted false evidence to a court, and the retained pension. **Verify: departure terms, forfeited awards, and pension quantum — reporting at the time varied and the text should state only what is documented.**

3. **Deepwater Horizon, 2010 (§5.3):** 20 April 2010; eleven dead; successor chief executive's departure later that year, exit terms, and the subsequent TNK-BP nomination. **Verify all figures.**

4. **RBS (§5.3):** the annulment of the knighthood (2012); the pension reduced from approximately £703,000 to approximately £342,000 following the 2009 political intervention, with a lump sum taken. **Verify both figures and the sequence, which the text currently compresses.** The claim that no senior UK or US bank chief executive faced criminal liability arising from the crisis should be checked against the full record and stated with any necessary qualification.

5. **Dark-triad measurement critique (§5.2):** the validation history on incarcerated and clinical populations, and the reliance on self-report and colleague-rating instruments in organisational settings, are accurately characterised at the level of generality used. **The manifesto deliberately makes no prevalence claim and cites no specific study.** If WS8 elects to add citations, note that the meta-analytic literature finds these traits weakly related to leadership *emergence* and negatively related to leadership *effectiveness* — a distinction that supports this chapter's argument and should be represented accurately rather than selectively.

6. **Descriptive versus injunctive norms (§5.4):** a standard and well-replicated distinction in social psychology. **Add citation at WS8.** The manifesto should avoid the Stanford Prison Experiment entirely, and treat the behavioural-ethics dishonesty literature with care, given documented replication and integrity problems in both.

7. **Televised deception formats (§5.5):** the international format count and the participant-behaviour observations are stated impressionistically and should either be sourced or softened to explicit illustration. **This is the weakest evidentiary element in the chapter and the red team should be invited to strike it.** The argument does not depend on it.

8. **Second founder's testimony (§5.1):** first-person, founder-authored, and to be confirmed as accurate by the founder before publication in the same manner as Chapter One §7.5.

Chapter Five: Building It — The UK Toolkit

1. From Argument to Instrument

Everything argued so far can be built with instruments that already exist in UK law — most of them for decades, one of them since Victoria. Nothing in this manifesto waits on legislation. This chapter is the workshop: the legal forms on the shelf, a decision path through them, the mechanics of the dual structure, the design of beneficiary power, the honest account of the fiscal bargain, the transition routes for organisations that already exist, and the governance instruments that keep the whole thing disciplined once built.

One orientation before the shelf. The architecture of Chapter Three has two layers — a purpose-holding apex and operating activity beneath it — and most real designs combine *two or more* instruments accordingly. The question is never "which single form is right?" but "which form holds the purpose, and which forms do the work?"

2. The Shelf: UK Legal Instruments

Instrument	What it binds	Bindingness tier (Ch2 §5)	Capital access	Built-in accountability	Best fit
Charitable Incorporated Organisation (CIO)	Purposes enforceable at law; full asset lock; public benefit requirement	5 — maximum	Donations, grants, trading via subsidiary, debt; no equity	Charity Commission regulation; public accounts; trustee fiduciary duty	Purpose-holding apex where purposes are charitable
Charitable company (company limited by guarantee + charity registration)	As CIO, with Companies House dual regulation	5	As CIO	Dual-regulator filing	Larger charities wanting company-law familiarity for lenders
Company limited by guarantee, non-charitable (CLG)	No shareholders; constitution can lock assets and purposes (entrenchment needed)	4 (5 with entrenched lock)	Bonds, debt, retained surpluses (the Glas Cymru model)	Whatever the constitution builds — membership can be designed in	Utilities, infrastructure, trade bodies; purposes wider than charity law allows
Community Interest	Statutory asset lock; community-	4–5	Limited equity (capped	CIC Regulator; annual	Social enterprises wanting lock

Instrument	What it binds	Bindingness tier (Ch2 §5)	Capital access	Built-in accountability	Best fit
Company (CIC)	interest test; dividend/interest caps (share-form)		returns), debt, grants	community-interest report	+ limited investment
Co-operative society	Member ownership and democratic control (one member one vote)	3–4 (asset lock optional — adopt it)	Member capital (withdrawable shares), community shares, debt	Member democracy; FCA registration	Worker- and multi-stakeholder enterprises
Community benefit society (CBS)	Trades for community benefit; statutory asset lock available	4–5 with lock	Community shares (a genuinely demotic capital instrument), debt	Member democracy + community purpose	Community-owned assets: pubs, energy, housing, local infrastructure
Employee Ownership Trust (EOT)	Controlling interest held in trust for all employees; sale blocked without trust consent	4	Company's own debt capacity; vendor financing on transition	Trustee duty to employees as a class; increasingly, employee voice bodies	Succession for private companies; the standard conversion route
Purpose foundation / trust + golden share	Veto or control over mission change and sale in a purpose entity	4	Unrestricted at operating-company level	Whatever the trust deed builds	Steward-ownership of commercial groups (the Bosch/Novo pattern)

Three notes the table cannot carry. First, **the CLG is the sleeper**: it is the instrument Britain reached for twice under duress (Glas Cymru, Network Rail), it can hold purposes charity law cannot (running a water network is not charitable), and its lock is only as strong as its entrenchment — draft it with the asset-lock and purpose clauses protected by supermajority or member-class consent, or a future board can undo it (the demutualisation lesson, in drafting terms). Second, **the co-operative's optional asset lock should be treated as mandatory** in any Selfless Capitalism design: an unlocked co-op is a demutualisation waiting for a windfall vote. Third, **community shares** (withdrawable share capital in a CBS, typically capped-return, one-member-one-vote regardless of holding) are the toolkit's most under-used instrument:

patient, demotic risk capital raised from the very community served — partial answer to both the capital constraint (R1) and the democratic-legitimacy finding (R13) in a single mechanism.

3. The Decision Path

For founders. Answer in order; the path composes the design.

Q1. Are the purposes charitable in law (advancement of education, health, community development, relief of need, and the other recognised heads — with public benefit)? — *Yes* → a **CIO** (or charitable company at scale) is the apex. → Q3. — *No, but purpose-led* (e.g. operating infrastructure, trading for community benefit) → **CLG with entrenched lock**, **CIC**, or **CBS**. → Q2.

Q2. Who should constitutionally hold power? — *The workforce* → co-operative, or company under **EOT**. — *The community served* → **CBS** with community shares, or CLG with a designed membership (the Glas Cymru pattern: an appointed-for-independence member panel; or an elected one — see §5). — *The mission as such, above any constituency* → CLG/CIC with purpose entrenchment, or foundation-plus-golden-share above an operating company.

Q3. Will the organisation trade substantially? — Charities may trade in direct pursuit of their purposes ("primary purpose trading") within the charity; substantial non-primary trading belongs in a **trading subsidiary** (§4). Non-charitable forms trade freely within their purpose constraints.

Q4. What capital will the work actually need? (Be honest per R1.) — *Asset-backed, revenue-backed* (infrastructure, property, established services) → long-dated **bonds and debt** at the operating level; the shareholder-free balance sheet is an advantage, not a handicap. — *Venture-stage risk* → sequenced self-funding, philanthropy/grants, community shares where the community is the beneficiary, and — the development frontier — **minority equity at subsidiary level** with the apex retaining control (§4.3). Accept the trade: this path is slower, and some venture classes are unavailable to the form (Concessions Register, item 1).

Q5. How much beneficiary power, constitutionally? (§5 — set the dial deliberately and publish the setting.)

Q6. What discipline machinery, from day one? (§7 — the lock removes extraction, not slack; the machinery is constitutive.)

4. The Dual Structure: Mechanics

4.1 The trading subsidiary. The apex (charity or locked entity) owns an ordinary limited company. The subsidiary competes, borrows, employs at market rates, and pays its people properly. Its distributable profits are donated to the charitable parent, ordinarily under **Gift Aid**, obtaining corporation-tax relief on the amount donated — the standard, Charity-Commission-recognised architecture (guidance CC35) used across the sector from retail chains to venue businesses. Points of discipline: the subsidiary must be viable on arm's-length terms (a charity may not prop up a failing trader with charitable funds); parent investment in the subsidiary must satisfy the trustees' investment duties; and governance should keep some separation between parent and subsidiary boards while preserving parental control.

4.2 Where the boundary runs. Primary-purpose trading (a school charging fees, a research charity licensing its findings) may sit inside the charity itself. The subsidiary exists for everything else — and, per Chapter Three's delivery boundary, regulated activity sits with accountable partners or appropriately regulated subsidiaries, never loosely inside the apex.

4.3 The hybrid frontier: minority equity below a controlled apex.

UK law does not prevent a charity-controlled group admitting minority investors into a trading subsidiary; trustee investment duties, conflict management and mission-protection terms (reserved matters, purpose covenants, exit provisions) govern it. Continental steward-ownership shows the pattern at the largest scale — foundation control above a listed operating company. What is missing in the UK is not legality but *practice and precedent* at scale: term sheets, valuation norms for capped-upside stakes, trustee guidance. This manifesto flags it as the single most valuable piece of legal-financial development work the model needs (Research Agenda, §10), and until that work matures, treats subsidiary minority equity as available for well-advised cases rather than as a routine tool.

5. Designing Beneficiary Power In (the R13 obligation)

The red team's sharpest structural finding: a pure charity's accountability runs enforceably to *purpose*, via regulator — but beneficiaries do not appoint trustees, and boards self-perpetuate. If Selfless Capitalism's claim is accountability to those served, that accountability must be **constitutional, not merely regulatory**. The toolkit therefore treats beneficiary power as a dial with named settings, to be set deliberately and published:

- **Setting 1 — Voice**: beneficiary panels, community consultation duties, published responses. (Floor, not destination.)
- **Setting 2 — Presence**: reserved trustee/board seats appointed from or elected by beneficiary and community constituencies; worker-elected seats per Chapter Two's targeted co-determination.
- **Setting 3 — Membership**: a designed membership (CLG/CBS/co-op) holding reserved powers — approving constitutional change, appointing a proportion of the board, receiving the annual audit. Glas Cymru's member panel is the appointed variant; a CBS's one-member-one-vote community is the elected one.
- **Setting 4 — Override**: the covenant machinery of the corpus written into binding agreements: prior community consent to deployment; rights of pause and withdrawal; suspension of harmful elements pending review regardless of cost; annual community-conducted audits, published including failures; remediation before expansion.
- **Setting 5 — Exit with the assets**: for community-anchored services, constitutional provision that on withdrawal, locally generated data, adaptations and (where designed) local infrastructure remain with the community.

Design guidance, stated as the manifesto's norm: **essential and trust-critical services should run at Settings 3–5.** Anything below Setting 2 should be regarded as the charitable form's minimum, not its expression. And one warning from the failure roster: demotic mechanisms are themselves capturable (a demutualisation is member democracy voting for windfalls) — which is why override and audit rights are paired with the asset lock rather than substituting for it. Lock *and* voice; neither alone.

6. The Fiscal Bargain (the R12 obligation)

The objection deserves its answer in full view. Yes: the architecture is tax-advantaged. Charitable reliefs apply at the apex; trading subsidiaries extinguish corporation tax on profits they Gift Aid upward; rates relief may apply. Scaled as this manifesto proposes, is that not a raid on the tax base and unfair competition with tax-paying firms?

The answer is that the reliefs are not a loophole in the social contract but a *different clause of it*. Corporation tax is society's partial claim — roughly a quarter — on surpluses that are otherwise **privately extractable**. The locked bargain replaces that partial claim with a total one: **one hundred per cent of surpluses are dedicated, in perpetuity, to public benefit**, enforceable by a regulator, with nothing left over for private extraction at all. Society does not lose 25% of the surplus; it gains the other 75% — permanently. Meanwhile the operating companies pay employment taxes, VAT, and business rates like any competitor; the relief attaches to profits precisely at the moment they cease to be private.

There is a second half to the fiscal argument, and it runs the opposite way to the objection. The locked form is not merely a fair bargain for the public purse; it is a *simpler* one. A shareholder group that wishes to move capital efficiently between entities, defer liabilities and manage the visibility of its returns builds machinery to do it: intermediate holding companies, offshore vehicles, layered structures whose stated purpose is efficiency and whose practical effect is opacity about how much is being made and by whom. That machinery costs real money in professional fees, and it costs the exchequer real revenue in ways no relief was ever legislated to grant. A locked group needs none of it. Surplus moves within the group without a tax event and without leakage, and there is nothing to conceal because there is no one to conceal it from — no owner whose take would embarrass, no bonus whose scale would provoke. **No extraction, no concealment; no concealment, no cost of concealment.** The comparison the objection ought to make is not between a relieved charity and an idealised taxpaying firm, but between a transparent structure that pays what it owes and an opaque one built to minimise it.

On the level-playing-field concern at the margin — a subsidised trader undercutting a family firm — the manifesto concedes the issue is real and is exactly why charity law polices trading boundaries; policy refinement there is legitimate and welcome. But the headline stands and should be argued with confidence: **an economy that shifts its trust-critical core into asset-locked forms is not shrinking the public's share of the surplus. It is maximising it.**

7. The Discipline Machinery: Governance Instruments

The lock removes extraction; nothing removes the need for rigour. Minimum kit, from the corpus, for any organisation adopting the model:

7.1 The stage-gate starter kit. Alignment gate (which purpose, which measurable human outcome, who benefits first); hard blockers (safeguarding; truthful capability claims; lawful data; conservative economics; capacity; a cheap test before a big build; one named accountable individual); staged commitment (exploration <5%, pilot 15–20%, scale on evidence); kill criteria written before launch; and the emotional-mugging checklist read aloud when any proposal arrives with a donor, a competitor, or sunk costs attached.

7.2 The SQ integrity audit — terms of reference (template). *Commissioned by:* the board, at least annually and before any major public document. *Conducted by:* a reviewer independent of the drafting and of the executive. *Object:* the organisation's public and internal language of record. *Hunts for:* conquest and campaign metaphor; claims of certainty, inevitability or impossibility of failure; speaking about communities rather than with them; the saviour posture (indispensability, criticism-as-betrayal, mission fused with person); euphemism where harm should be named. *Output:* findings graded critical / significant / advisory; critical findings block publication until resolved; the resolution log is retained and the existence of the audit is publishable. *Precedent:* the four audited revisions of this project's own source manifesto, including the withdrawn claim that wrongdoing had been made "impossible."

7.3 Trustee selection-inquiry protocol. Recruit for competence against a published skills matrix (finance, safeguarding, the delivery domains, and the boundary-oversight function); and inquire into the *record* using the Selfless Leadership Test's questions — how the candidate handled being wrong; what those with less power say; whether past exits protected the vulnerable or the CV. Asked of records, never scored of persons.

7.4 The listening infrastructure. Skip-level and beneficiary panels on a calendar; an always-on anonymous channel; a bad-news-fast protocol with a named route to the board; and the annual community audit of §5, Setting 4.

7.5 The published record. Annual reporting that includes: mission metrics (outcomes per pound, beneficiaries served); the pay ratio, with method; reserves held, stated against quantified liabilities rather than against a policy the organisation wrote for itself; the beneficiary-power dial setting; audit results *including failures and remediations*; and progress against the model's own success measure — whether those served need the organisation less than they did. Published in a fixed format that cannot be redesigned in a year when the figures are unflattering. Any single measure can be gamed by an organisation determined to game it; the shape they make together over a decade is harder to fake, because faking all of them consistently requires exactly the coordination that honest organisations lack.

7.6 The sunset review — terms of reference (template). *Cycle:* the tenth year after registration and every tenth year after; in no case more than twenty years between reviews. Ten years alone is too short to be the point at which everything can be undone; twenty alone is too long, because a whole leadership generation can arrive, capture and retire without meeting a review. *Presumption:* not whether cause exists to alter the organisation, but whether cause exists to continue it in its present form. *Criteria:* written at foundation and unamendable by anyone the organisation can reach — board, members, founder, or a review conducted under them. A board that can rewrite the test before sitting it has not been tested. *Conducted by:* an assessing organisation subject to the same lock and the same obligation, drawn by lot, operating in a different field, sharing no material funder, having not assessed this organisation in the preceding two cycles, and not itself assessed by it; a panel of beneficiaries constituted by a process the organisation takes no part in; and at least one member with no affiliation to any organisation of this form, whose dissent is published even when it loses. *Paid from:* a fund settled at foundation on terms the organisation cannot vary, with no other engagement permitted before or for a defined period after. *Divided by competence:* structural criteria — reserves, pay dispersion, the beneficiary-facing ratio, the direction of dependency — to the cross-field assessor, for whom independence matters more than expertise; mission quality to the beneficiary panel, whose finding is published alongside and is not subordinate to it. *Findings, graduated:* continuation as it stands; continuation for a shortened term with conditions; replacement of the governing body; separation into independent parts, each carrying this obligation forward; transfer of a purpose or capability to a body better placed to carry it; winding up under the lock. A review that can only bless or destroy will always bless.

Published: determination, reasons, evidence and dissent, in full. *If not convened by its due date:* powers to acquire assets, commit beyond twelve months, or appoint staff are suspended until it concludes.

Note on the assessor. The obvious objection is that practitioners assessing practitioners is self-regulation, and the American accounting profession's peer review programme — fifty years of it — was abandoned in 2002 for independent public oversight, on the finding that reviewers were unlikely to detect serious deficiencies and that self-regulators serve their constituents rather than the public. The mechanism there was specific: those firms competed for the same clients, shared a licensing regime, and held a common franchise in the profession's reputation, so a harsh review damaged the reviewer's own market. None of that survives translation to organisations sharing no market, no clients, no funders and no trade body. What does survive is a shared stake in the reputation of the form itself, and two provisions answer it: assignment by lot and never reciprocal, so no organisation can trade with an assessor it knows is coming; and the assessor's own record on trial at its own review, so that an organisation whose assessments have never produced an adverse finding carries that record into its own. This is the provision that makes politeness expensive. Without it, mutual accommodation is the equilibrium, and everyone arrives at it sincerely.

8. Transition Pathways

8.1 The private company → EOT. The established route, in active use across UK succession: owners sell a controlling interest to an employee trust (with capital-gains relief on qualifying disposals), typically vendor-financed from future profits. Design notes from this manifesto's standards: pair the EOT with employee voice bodies (trust ownership without voice is Setting 1), entrench against onward sale, and adopt the §7 machinery — the EOT literature's own warning is that ownership without governance change transfers the shares and not the culture.

8.2 The founder start-up → dual structure from birth. Cheaper than converting later: CIO (or locked CLG) apex, trading subsidiary from the first revenue, beneficiary dial set in the founding documents, machinery adopted before it is needed. The corpus is the worked example, and its sequencing lesson travels: fund the mission's later capital hunger from the earlier ventures' revenues, and accept the calendar that implies.

8.3 The plc → steward ownership. Stated honestly: rare, hard, and usually generational — requiring either a controlling founder-family conveying control to a foundation (the continental pattern) or an improbable shareholder vote. The manifesto does not build its programme on plc conversion; it builds on **founding, succession, and the public estate** — and lets Chapter Two's full-cost rules and the model's competitive example do their slower work on the rest.

8.4 The public estate — the next infrastructure decision. Chapter One's counterfactual becomes a live policy template the next time an asset leaves (or re-enters) public hands: water companies emerging from special administration; rail bodies; future energy, network and — the coming case — AI and data infrastructure. The template: **CLG or CIO form; entrenched asset lock and purposes; bond-financed against regulated or contracted revenues; beneficiary dial at Setting 3 or above; the §7 machinery statutory or constitutional from day one.** Neither privatisation into extraction nor renationalisation into the Treasury's shadow: the third form, chosen deliberately this time, before the failure instead of after it.

8.5 The community adoption route. For communities rather than founders, the corpus's capability-distribution pattern, generalised: begin with consent (Setting 4 machinery agreed before deployment, not after); certify local practitioners through capability *or community nomination with equal weight*; revenue architecture that stays (the 90/10 pattern); accountability sessions community-facilitated; successor development from year one; and the exit-with-assets provision so that withdrawal never strands those served. Status per the Concessions Register: designed and piloted, unproven at scale — adopt it as commitment machinery, and publish the record that will prove or indict it.

9. Beyond the UK, at Principle Level

The instruments are British; the architecture is portable wherever three functions can be assembled: **a lock** (assets and purposes beyond private extraction — foundation, non-profit, benefit-purpose entity), **a governed commerce** (operating companies under the locked apex), and **a demotic channel** (membership, community shares, or covenant machinery). In the US: a 501(c)(3) or perpetual-purpose trust above operating companies (benefit-corporation statutes assisting entrenchment); in continental Europe: the mature foundation-enterprise law this manifesto has cited throughout; in the co-operative tradition: everywhere. Jurisdiction-specific development is expressly outside this toolkit's competence and inside its research agenda.

10. Research Agenda (published, per C6)

1. **Subsidiary minority equity at scale**: model term sheets, trustee guidance and precedent for mission-protected external investment below a charitable apex — the form's highest-value missing instrument (R1/R4).
2. **Customer-ownership governance** for essential services beyond the appointed-panel pattern: designs that keep diffuse member bases engaged and capture-resistant (T2, standing).
3. **Beneficiary-power drafting library**: model clauses for the five dial settings across CIO, CLG, CIC and CBS forms (R13).
4. **The public-estate template in statute**: what an "asset-locked transfer" of infrastructure requires from government, regulators and bond markets (§8.4).
5. **Empirical programme for the selection hypothesis**: leadership tenure, pay-governance outcomes and failure typology by organisational form (R2/R10's stated predictions).
6. **The register of assessing organisations**: constitution, eligibility, the lot mechanism, cross-field matching rules and the escrow instrument — none of which exists, and which is the largest practical obstacle to §7.6 (C14).
7. **Test of the conversion hypothesis (new, per Chapter Four §5)**: a within-person design following executives who move between extraction-form and asset-locked-form organisations, measuring conduct differences not attributable to selection — the discriminating test between C11 and C13, and the one that can falsify either.

Chapter Six: The Repricing

1. The Cell That Isn't in the Spreadsheet

Everything proposed so far can be built without asking anyone's permission. Chapter Five is a workshop of instruments already sitting in UK law: found a locked company, set the beneficiary dial, run the stage-gate, build the discipline machinery. Nothing waits on legislation.

This chapter breaks that rule deliberately, and only once, because there is one failure a firm cannot fix from inside itself.

Chapter One's first case study established it from the inside of the room. The externalities of a global beverage business — the plastic, the sugar, the health systems — were not weighed against profit and judged acceptable. They were **absent from the model**. And a cost with no cell in the spreadsheet cannot be argued down, because it was never a candidate for argument. There is no moment of choice at which a better person chooses otherwise. The consultant optimises correctly. The board approves correctly. Everyone behaves well by the standards available to them, and the ocean fills with bottles.

That is why voluntary virtue fails so completely and so consistently. It is aimed at a decision that never happens.

And it is why the firm-level architecture of Chapters Two to Five, whatever its merits, is not sufficient. A locked, trustee-governed, beneficiary-accountable company still operates inside a price system that lies to it. It can put a shadow price on carbon voluntarily and be undercut by a competitor who does not. The conscientious firm is asked to lose; it declines; and the accounting boundary holds.

The remedy is not exhortation, and it is not a better class of executive. **It is to force the missing cell into the model** — to put the harm inside the return-on-investment calculation, where the same analytical talent that once optimised for margin will optimise against it with exactly the same diligence. This chapter sets out three instruments that do that, in ascending order of originality: a levy that prices the harm, a corporation tax rate that removes the gain, and a gate that forbids the design.

None of them asks anyone to care.

2. Instrument One: Levies That Price the Harm

The oldest instrument is a charge on the harmful thing itself — per tonne of virgin plastic, per gram of added sugar above a threshold, per tonne of carbon. Pigou proposed the principle in 1920. What has been missing is not theory but evidence that it changes behaviour rather than merely raising revenue.

Britain has now supplied that evidence, and the result is more instructive than its designers expected.

The **Soft Drinks Industry Levy**, introduced in 2018, was tiered rather than flat: a higher rate above 8 grams of sugar per 100 millilitres, a lower rate between 5 and 8 grams, nothing below 5. The tiering was the design insight, because it created a cliff edge the manufacturer could engineer its way off. And that is exactly what happened. Between 2015 and 2020 the sales-weighted average sugar content of levy-eligible soft drinks fell from 3.8 to 2.1 grams per 100 millilitres — a reduction of about 45% — and total sugar sold in soft drinks fell 34.3%, from roughly 135,000 tonnes to 89,000. Subsequent studies have associated the levy with reduced obesity prevalence among girls in their final year of English primary school and fewer hospital admissions for the extraction of decayed teeth in children.

Now the comparison that matters to this manifesto more than any other statistic in it. Alongside the mandatory levy, government ran a **voluntary** sugar reduction programme for other food categories, with a 20% target. It achieved 3.5%.

34.3% mandatory. 3.5% voluntary. Same government, same period, same industry, same public health objective, same good intentions on all sides. The only variable was whether the mechanism bound. Chapter Two's rule — an end without a binding mechanism is a press release — is not an inference from corporate cynicism. It has been measured.

Two design lessons carry forward. **Tier the levy, don't flatten it**, because a threshold creates an engineering problem the manufacturer can solve, whereas a flat charge creates a cost the manufacturer can pass on. And **levy the producer, not the till**: the objective is reformulation, and only the producer can reformulate.

Levies have a limit, though, and it is the reason this chapter has two further instruments. A levy makes externalising more expensive. It does not make it unprofitable. A firm with pricing power — a global brand facing no real substitute, selling at £4.60 a litre — absorbs the levy and passes it on. Notably, when British manufacturers reformulated en masse to escape the sugar levy, one product line conspicuously did not: full-sugar Coca-Cola, whose brand equity is bound to its formulation. Where brand power is strong enough, a price on harm becomes a licence fee for harm.

Which is the gap the next instrument fills.

3. Instrument Two: A Corporation Tax Rate That Varies With Harm

Here is the proposal at the centre of this chapter, and the one the authors have not seen made elsewhere in this form.

Corporation tax should not be a flat rate. It should be a function of what the company's activity costs everyone else.

At one end of the schedule: an enterprise whose products are fully circular, whose emissions are within its own accounts, whose activity imposes no measurable burden on health systems or the commons — approaching **zero per cent** on profits. At the other end: an enterprise whose profits depend on discharging plastic into the oceans and metabolic disease into public health systems — **eighty or ninety per cent.** Between them, a graduated schedule set by audited assessment, not by lobbying.

Observe what this instrument does that a levy cannot. A levy taxes the harmful *act*, and can be priced into the product. A variable profit rate taxes the **gain obtained by externalising** — and a gain that is confiscated is not a gain. It does not make pollution expensive. It makes pollution *pointless*. The 60% gross margins of Chapter One's case study exist because the plastic and the sugar are free; a rate schedule of this kind puts the plastic and the sugar back into the margin calculation, from the other end. The strategy deck changes. Not because anyone was persuaded, but because the arithmetic changed, and the same analyst who followed the arithmetic into single-use PET follows it back out.

It also converts something in this manifesto from special pleading into general principle. Chapter Five argued that the asset-locked form earns its charitable relief because 100% of its surplus is permanently dedicated to public benefit rather than a quarter being claimed in tax — a fair bargain, but one that reads as an interested party arguing for its own exemption. Under a variable schedule, that relief stops being an exemption and becomes **the bottom rung of one continuous rule that applies to everybody**: your rate is a function of what your surplus costs the rest of us. The locked circular enterprise sits at zero because it costs nothing and keeps nothing. Coca-Cola sits near ninety because it costs a great deal and keeps most of it. One rule, no exceptions, no special cases to defend.

3.1 Two axes, not one

A single harm axis leaves a loophole large enough to drive a lorry through, and it is a loophole this manifesto would create for itself: an asset-locked charity selling single-use plastic would pay nothing while polluting freely. The lock would become a pollution shelter.

So the schedule needs two dimensions:

- **The harm axis** — what the activity costs the commons and the public purse: unrecovered materials, emissions, water, health burden.
- **The extraction axis** — how much of the surplus leaves for private benefit rather than remaining in the enterprise, its workforce, its customers or its purpose.

A firm scores badly on both and pays the maximum. A polluting charity scores badly on one and pays accordingly — the lock does not excuse it. A clean shareholder company pays a moderate rate. A locked circular enterprise pays near nothing. The two axes are not a complication; they are what makes the rule impossible to game by changing legal form alone.

3.2 What this is not

It is not a windfall tax, which is retrospective, political and arbitrary. The schedule is published in advance, applies by formula, and — this is the point — a company can move down it by changing what it does. A firm that redesigns for circularity cuts its own rate. That is the incentive doing the work, and it is why the instrument must be *predictable*: an unpredictable schedule produces lobbying, while a predictable one produces engineering.

Nor is it a tax on being large, or foreign, or unpopular. Two firms of identical size in identical markets pay different rates if and only if one imposes more cost on the public than the other. That is not a penalty. It is the removal of a subsidy.

4. Instrument Three: The Circularity Gate

The third instrument is the simplest and, in some ways, the strongest, because it is not a price at all.

No product may be launched without a published, audited end-of-life pathway. Before market entry the firm states: what the product and its packaging are made of; how they will be recovered; by what route; at what realistic rate; and who bears the cost. The analysis is filed, published, and auditable. No pathway, no launch.

A gate has the property no levy has: **it cannot be passed through.** A firm with pricing power absorbs a charge and raises the shelf price. It cannot raise the shelf price of a design it is not permitted to sell.

Chapter One's case study is the argument for it. The migration from refillable glass to single-use PET would not have survived a gate of this kind — not because glass is sacred, but because the pathway that existed was being abandoned rather than improved, and a gate requires the pathway to be *stated*. And it locates the decision where Chapter One located it: a firm required to publish a real recovery route must confront the fact that circularity is a physical system whose economics depend on distance. Refillable loops work locally and fail over long hauls. The gate therefore prices in the geography of production, which is the decision that was actually taken and never examined.

The direction is not speculative. Extended Producer Responsibility schemes place recovery costs on producers in many jurisdictions, and the EU's ecodesign framework moves toward design-stage requirements. The proposal here is to complete that direction and make it a condition of entry rather than a cost of operation — and to connect it to the rate schedule, so that the audited pathway is also the evidence base for Instrument Two.

5. Assessment: How the Score Is Set

An instrument is only as credible as its measurement, and this is where a proposal of this kind either becomes policy or remains a slogan. Four requirements.

Published methodology, single authority. One statutory body, one published formula, revised on a stated cycle with consultation. Not ministerial discretion, which invites capture.

Firm-supplied, independently audited data. The firm files its own materials, recovery, emissions and product-composition data — as it already files accounts — and auditors verify it, with criminal liability for misstatement as with financial reporting. Note that this requires no new science: recovery rates, sugar content per 100 millilitres, tonnes of virgin polymer and emissions are all already measured for other purposes. The instrument mostly *joins up data that already exists*.

Sector-relative bands, absolute floors. Bands are set within sector so that a foundry is not scored against a software company, but with absolute floors so that "best in a dirty sector" does not earn a clean rate.

Appeal on the facts, not on the consequences. A firm may contest its data. It may not contest its rate on the grounds that the rate is inconvenient — which is the standing invitation that has hollowed out every discretionary environmental regime.

6. Objections

"It's protectionism, and production will relocate." The serious version of this objection concerns border treatment, and the answer is that a harm-based schedule must be levied on *consumption in the jurisdiction*, not production — with an import adjustment on the same published formula, as carbon border mechanisms already attempt. A schedule that taxes domestic manufacturing while waving imports through would be worse than nothing. That is a design requirement, not an argument against the design.

"Assessment will be gamed." Some gaming is certain; the question is whether the residual is better than the status quo, which currently scores every externality at zero. Note also that the harms in question are unusually tractable — mass of virgin material, grams of sugar, tonnes of carbon — compared with the transfer-pricing arithmetic that tax authorities already contest.

"It will raise consumer prices." For harmful products, yes, and that is partly the mechanism. But the levy evidence shows the dominant response to a well-designed threshold is reformulation rather than price rises, and the gate cannot be passed on at all. Meanwhile, the goods that become relatively cheaper are the circular ones — which is the entire point of a corrected price system.

"You are asking the state to pick winners." No: the state publishes a formula and firms pick themselves. Every company can achieve the zero band by ceasing to impose costs on the public. Nothing in the schedule names a company, a technology or a sector.

"Governments will use it to raise revenue rather than change behaviour." This is the strongest objection, because the incentives point that way: a treasury that grows fond of receipts from a harm tax has acquired an interest in the harm continuing. The structural answer is to hypothecate — receipts from the harm bands fund the remediation the harm causes, and a published projection shows revenue *falling* as compliance improves, with success defined as the schedule collecting less over time. A harm tax that grows is a harm tax that has failed.

7. Sequencing, and the Boundary of the Claim

This chapter is deliberately last, because the book's central claim must survive without it. Everything in Chapters Two to Five can be built tomorrow by anyone who chooses to, and the case for building it does not rest on the state doing anything at all.

But the honest statement of the whole argument is that the two halves complete each other. **Architecture without repricing** leaves the well-designed firm competing against subsidised harm, and asks it to lose. **Repricing without architecture** corrects the signals and leaves them running into structures built to extract, where a carbon price becomes one more cost to pass through to a captive customer. Chapter One diagnosed three failures; Chapters Two to Five address the third, the tyranny of the quarter, by changing who owns the clock. Only repricing reaches the first two.

And the sequence is not merely logical but practical, because the firms built under Chapters Two to Five are the constituency for the repricing. A circular, locked, community-accountable enterprise has nothing to fear from a harm-based schedule and everything to gain from it: its competitors' subsidy ends. Every such firm founded is one more voice for a rule the extraction economy will resist. The architecture builds the coalition; the coalition wins the repricing; the repricing rewards the architecture.

Which is where this manifesto's ambition properly ends — not with a claim about what will happen, but with a statement of what would have to be true. Prices that tell the truth. Wealth counted whole. Horizons as long as consequences. Structures with no self to serve. **Value created, not extracted.**

8. How Change Actually Happens

One last observation about method, because it runs through every instrument in this book and has not yet been named.

There is a term in the Taoist tradition, *wu wei*, usually rendered into English as "non-action" or "doing nothing." The translation is misleading. What it describes is not passivity but **action along the grain** — effort applied where the material will actually yield, rather than where the actor would prefer it to. The standing image is water. Water is the softest substance there is. It carved the Grand Canyon. Not by force, and not by confronting the rock, but by persistence along the line of least resistance, for a very long time. Nothing about that is inaction. It is simply action that does not fight what it is acting on.

Read Chapter Two's sugar levy again in that light, because it is a controlled experiment in method. The tiered threshold did not ask manufacturers to act against their own interest. It handed engineers a problem, and engineers solve problems: reformulate below five grams and the charge disappears. Total sugar sold in soft drinks fell 34.3%. Running alongside it, in the same country, in the same years, aimed at the same industry, a voluntary programme asked companies to reduce sugar because it was the right thing to do. It achieved 3.5%. One approach worked with the grain of self-interest. The other struck at it. The rock is unmarked; the water is most of the way through.

Every mechanism proposed in this book has that shape. The asset lock does not ask owners to restrain themselves; it removes the owner. The circularity gate does not ask for conscience; it changes what may lawfully be sold. The variable rate does not ask a corporation to stop optimising — it will never stop optimising — it changes what optimising produces. And the architecture of Chapters Two to Five does not ask executives to resist extraction by force of character; it removes the extraction, and then relies on nothing more than ordinary competence and ordinary ambition, of which there is no shortage.

This answers the objection that a ninety per cent tax rate is a strange thing to call harmonious. It is harmonious precisely because it makes no demand on anyone's character. To ask a beverage company to care about the ocean is the hammer, and Chapter One showed what the hammer achieves: the harm was not in the model, so there was nothing for conscience to act upon. To put the ocean into the return-on-investment calculation is the water. The same clever people, the same spreadsheets, the same ambition — and a different result, arrived at without a single conversion experience.

Which is why this manifesto proposes an evolution rather than a revolution. Not the overthrow of the market economy but its next stage: the same competitive, inventive, wealth-creating machine, running on signals that tell the truth, inside structures with no self to serve. Less waste, less conflict, fewer costs quietly transferred to people who never agreed to bear them — a system more coherent, in the plain sense that its parts stop working against each other.

The instruments exist. The evidence that binding mechanisms work and voluntary ones do not is now measured rather than argued. What remains is the decision — and decisions of this kind are not made by manifestos. They are made by people who read one, and then go and build something. Water does not argue with stone.

9. Carried Forward

- **To WS5 (second red-team pass):** the variable-rate proposal (§3) is the newest and least tested material in the manifesto and should be attacked hardest — particularly the assessment machinery (§5), the border treatment, and whether the two-axis schedule is administrable. The hypothecation answer in §6 needs fiscal-policy review.

- **Proposed for the convergence map:** C12 — *externalised harm is typically absent from the decision model rather than defeated within it; therefore no improvement in character corrects it, and only a change to the arithmetic will.*
- **Glossary additions:** the circularity gate; the harm axis and extraction axis; hypothecation.
- **To WS7:** Chapter Five's fiscal-bargain section (§6) should cross-refer forward to this chapter, since the charitable relief becomes the bottom rung of the schedule rather than a standalone bargain.
- **Research agenda additions:** design of the two-axis schedule with worked band examples; border-adjustment mechanics; comparative review of EPR and ecodesign regimes as the gate's precedent base.

---

10. Sources and Evidence Notes (for WS5/WS8 review)

1. **Soft Drinks Industry Levy design and outcomes (new — E23):** tiered structure (>8 g/100 ml higher rate; 5–8 g lower; <5 g exempt), 2018. Sales-weighted average sugar content of levy-eligible drinks 3.8 → 2.1 g/100 ml, 2015–2020 (≈45% reduction); total sugar sold in soft drinks 135,391 → 89,019 tonnes (34.3% reduction). Voluntary sugar reduction programme achieved 3.5% against a 20% target. Sources: WHO-published comparative analysis of UK sugar reduction policies (PMC11132159); systematic scoping review of SDIL impacts (2025); PLOS Medicine interrupted time series and modelling study on obesity and dental outcomes. Verify all figures and add full citations at WS8.
2. **Reformulation as dominant industry response (E23):** interrupted time series analysis of the UK soft drink marketplace 2017–2020; note the exception of full-sugar branded cola lines, which reformulated least — verify precise claim wording at WS8 before publication.
3. **Extended Producer Responsibility and EU ecodesign (new — E24):** cited as policy direction and precedent base for the circularity gate; specific scheme details to be added at WS8.
4. **Coca-Cola case-study evidence:** carried from Chapter One §7 (E25 series) — price series, CPI comparison, gross margin and bottling refranchising, Break Free From Plastic brand audits, sugar content and AHA guidance.
5. **Pigou (1920):** externality pricing principle, as cited in Chapter One.
6. **Border adjustment:** carbon border adjustment mechanisms cited as the nearest existing analogue; no claim made as to their effectiveness, which is contested.

Appendix: The Selfless Leadership Test

Adapted for enterprise governance from the SQ Leadership Test. To be asked of records and patterns of conduct — in selection, in board self-assessment, in succession — never administered as a scored instrument, never applied to a named individual as verdict. The questions are the discipline; the honesty of the answering is the result.

1. Does this leader protect the vulnerable — employees, users, communities — when doing so costs them personally or politically?
2. Do they tell the truth when the truth damages their image, their forecast, or their legacy?
3. Can they admit error without rage, denial, or the search for someone to blame — and does the record show them repairing what they broke?
4. Do they experience the suffering their decisions cause — a redundancy, a price rise on a captive customer, a harmed beneficiary — as real, or as line items?
5. Do they honour promises made to those with no power to enforce them?
6. Do they resist becoming the centre of the story — sharing credit, building successors, letting the institution outgrow them?
7. When criticised, do they treat the critic as a possible corrective or as an enemy?
8. Do they distinguish strength from domination — in negotiation, in management, in the exercise of every asymmetry the role hands them?
9. Would they make the same decision if no one would ever know, and no history would ever praise them?
10. Is the organisation more capable of functioning *without* them than it was a year ago — and are they glad of it?
