Water and the regulation of monopoly utilities - White Paper

A proposal to fix the rules that govern England and Wales's water monopolies so they hold whoever owns the pipes - and to put the two genuine value choices, who owns the system and how far and fast to invest, honestly to the public.

Discussion draft · version 1.0 · June 2026. Prepared under *The Pragma Method* - an approach for turning long-unsolved problems into implementation-ready policy on graded evidence and across the political spectrum. It is presented for development and public deliberation, not as a finished government position: it **sets out options and the evidence for them, and does not advocate their adoption.** Whether the country changes who owns its water companies, and how ambitious it wants to be in fixing the system, is for the public and Parliament to decide. Because government is reforming water right now, this is a living proposal kept current as the reform moves - see Living updates.

How to read the evidence grades. Each factual claim below is graded for the strength of the evidence behind it: **A** robust / audited (an audited regulatory account, official statistic or determination) · **B** strong official projection (a large official dataset or a government, National Audit Office or Environment Agency projection) · **C** single estimate, modelled or contested (one source, a residual, or a figure under active dispute) · **D** weak or absent (no agreed primary figure; an extreme-case estimate carried only to mark the edge of a range). Every claim carries a grade; full citations are in the Evidence Annex.

One number guard, stated once and held throughout. Two investment figures recur and must never be added together. The five-year programme of **£104 billion (2025-2030)** is the first five years *inside* the long-term **£290 billion (2025-2050)** programme. They are the same money seen over different windows. Anywhere both appear, the five-year figure is a slice of the larger one, never an addition to it.

Executive summary

Water is the one essential service nobody can shop around for. Wherever you live, a single company owns the pipes, and that is the only company you can buy from - so the competition that disciplines price and quality in other markets does not exist here. The job of substituting for that missing competition falls to regulation, and for decades it has not done it. The companies were allowed to load up on debt and pay money out to their owners while the pipes, sewers and reservoirs fell behind. The result is the most salient fact in the whole debate: the average household bill is now **£603 a year (2025/26, cash)**, after a rise of about **26% in real terms in a single year** (Grade A). Sewage spills rose by roughly **27% over the five years to 2025** despite a promised fall (Grade B for the trend). And the largest company, **Thames Water - which serves about a quarter of the population - loaded up debt to the point of near-collapse and has faced the prospect of special administration** (Grade B).

This is a problem the state has to solve, because it is the product of three things a market cannot fix on its own: a **natural monopoly** (no competition is possible, Grade A), a **regulatory failure** (the system meant to substitute for competition permitted high gearing and large distributions while investment lagged, Grade B), and an **externality** (a company that discharges sewage does not bear the cost it imposes on rivers, bathers and downstream communities, Grade A).

Government is already mid-reform. The Independent Water Commission (Cunliffe) reported in mid-2025 with 88 recommendations; the government has accepted **abolishing the existing economic regulator (Ofwat)** and replacing it with a new one (Grade A). This proposal builds on that settlement rather than starting over - and concentrates on what the Commission left unresolved or was *barred* from considering. Crucially, the Commission's terms of reference **forbade it from recommending public ownership** (Grade A) - so there is a live, official gap on exactly the question the public most wants answered.

The proposal has two parts.

An ownership-neutral core - the substance that holds whoever owns the pipes. A successor regulator with the powers and the in-house expertise Ofwat lacked; a hard financial-resilience regime (a binding cap on borrowing, conditions that stop owners being paid while a company is failing, and a ring-fence so group debt cannot drain the company that runs the network); a pre-defined rescue route for a failing company that keeps the water flowing while owners and lenders - not the public - bear the losses; enforceable environmental and supply targets with penalties big enough to bite; and transparent governance of the long-term investment programme. None of this depends on who owns the companies, and all of it is designed directly on evidence.

Two genuine choices, routed honestly to the public. *Choice 1 - who owns the system, and who pays for the past:* keep the companies private but properly regulated, take them into public ownership, or run them as not-for-profit companies with no shareholders - together with how the legacy and forward bill is split between bill-payers, investors and taxpayers. *Choice 2 - how far and how fast to invest:* an ambition ladder costing roughly **+£19 / +£28 / +£38 per household per year** for a slow, steady or accelerated pace.

The headline numbers, carried honestly. The investment needed to put the system right and future-proof it is about **£290 billion over 2025-2050** (Grade B), of which the bigger half is securing future supply, not only cleaning up. The one-off cost of *changing* who owns the companies has **no agreed figure**: credible estimates run from about **£50 billion to £107 billion**, with a contested argument that a failed company could be transferred for far less (Grade C, and D at the extreme). That enormous spread is **not** a measurement dispute - it is a choice about how generously existing owners are compensated.

That last point is the discipline of this whole paper. The investment need, the sector's finances and the bill trajectory are **empirical questions**, settled here on evidence. *Who should own the pipes, who should pay for the years of neglect, and how far and how fast to invest are **value questions*** - and the cardinal error this institute exists to avoid is presenting a value choice as a technical necessity. So those are routed to the public, with no option labelled "recommended". The rest of this paper sets out the problem, why markets and the current regulation failed, the ownership-neutral core, the investment need, the two choices, the honest costing, an adversarial three-perspective review, the measures of success, and the route to implementation.

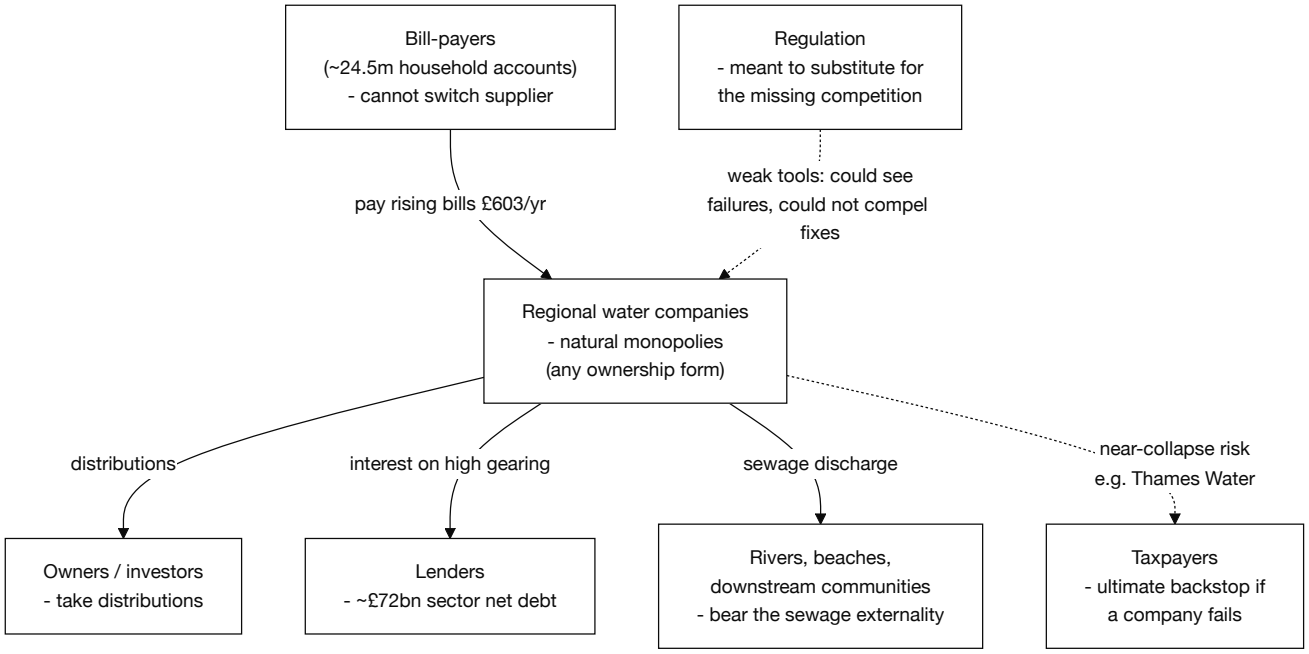

1. The problem

The water and sewerage industry in England and Wales is a set of **regional monopolies**. A household cannot switch its water company the way it switches a phone or energy supplier - wherever it lives, one company owns the pipes, and that company is the only one it can buy from. There is therefore no competitive discipline on price, quality or investment, and the rules are supposed to do that job instead. They have not.

The facts that are not really in dispute:

- **The average household bill is £603 a year (2025/26, cash), after a rise of about 26% in real terms in a single year**, with further increases planned to 2030 (Grade A). This is the single most salient fact for the public. (The cash bill and the regulator's real-terms trajectory are different bases and should not be mixed; the £603 is what households actually pay now.)
- **Around 2 million customers are in water poverty, and roughly 40% said the planned rises would be hard to afford** (Grade C, survey-based - carried as context, not a precise count).
- **Sewage spills rose by roughly 27% over the five years to 2025**, despite a promised reduction (Grade B for the trend). The *direction* is robust; the precise level is less certain, because monitoring coverage expanded over the period - more storm overflows were fitted with monitors, which improves measurement and partly inflates the apparent rise. The trend is the relevant fact.
- **Thames Water, which serves about a quarter of the population, is in acute financial distress:** gearing around 80%, debt of roughly £20-23 billion against a regulated asset value of about £20 billion, in cash lock-up and major restructuring, having faced the prospect of special administration (Grade B). It is the live illustration of how a financial structure the rules permitted can turn into a national supply risk.

Why this is a problem the state must address. Three distinct market insufficiencies operate together:

- **Natural monopoly.** The fixed cost of regional pipe networks makes competition impossible; customers cannot switch, so there is no market discipline (Grade A - a structural fact of the industry).

- **Regulatory failure.** The regulatory model meant to *substitute* for that missing competition permitted high gearing and shareholder distributions while network and environmental investment lagged (Grade B, National Audit Office April 2025). This is not a market failure in the textbook sense; it is a *government-created* shortfall in the substitute for competition, and is labelled as such here rather than disguised as inevitable.
- **Externality.** Sewage discharge imposes costs on rivers, bathers and downstream communities, not on the company that discharges - a classic externality the market has no reason to price (Grade A).

2. Why markets and the current regulation have not resolved it

A natural monopoly cannot be fixed by competition, so the only options are to regulate it well or to own it differently - and the existing regulation governs all three ownership forms today. The honest diagnosis is that the regulatory model itself broke, in ways that are now measurable.

- **The companies are geared far above the level the rules assume.** Sector regulatory gearing - net debt as a share of the regulated asset base - is **67.9%**, against a notional **55%** the regulator funds the sector as if it were at (Grade A). The regulator funds returns as though the companies were geared at 55%, but they are geared well above that, and the surplus debt is a return-on-equity engineering choice the model permitted. That **12.9-percentage-point gap is the single clearest measure of the financial-resilience failure.** Thames Water, at around 80%, is the gap taken to its extreme (Grade B).
- **Large distributions flowed while investment lagged.** The sector has paid out something on the order of **£85 billion in distributions since privatisation in 1991** (Grade C - an order-of-magnitude figure, deflator- and definition-dependent; the *direction* is robust, the exact total is not). Sector net debt now stands at about **£72 billion** (Grade B).
- **The regulator could see the failures but could not compel the fixes.** Ofwat published a financial-resilience report showing the 67.9%-versus-55% gap but lacked the binding tools to close it (Grade A for the figures). And it has been out-resourced and out-advised by the companies it regulates - the "intelligent client" weakness, an asymmetry of expertise and data between regulator and regulated (Grade C; both the National Audit Office and the Independent Water Commission criticised it). A regulator that cannot out-think the firms it polices will be captured by them regardless of its formal powers - a problem that links directly to the wider question of the state's capacity to deliver.
- **The official reform leaves the central question open.** The Independent Water Commission reported in mid-2025 with 88 recommendations; the government has accepted abolishing Ofwat and replacing it with a new regulator (Grade A). But the Commission was **barred by its terms of reference from recommending public ownership** (Grade A) - so the live, official process has deliberately left untouched the very question the public cares most about.

The lesson is not that regulation cannot work, but that the *substance* of it - financial resilience, enforceable targets, the expertise to enforce them, and a ready plan for a failing company - has to be far stronger, and has to hold whoever owns the companies. That substance is the next section.

3. The ownership-neutral core - the substance that holds whoever owns the pipes

This is the part the evidence can largely settle: how to make monopoly water regulation work. It is designed directly on the graded evidence, it builds on the accepted abolition of Ofwat, and it is built so that it is the

same machinery whether the pipes are privately owned, publicly owned or run by a not-for-profit. The full operational detail - the legislation, the powers in licence-condition form, the sequencing and the exit conditions - is in the Delivery Design; the core has five elements.

A successor regulator with teeth and expertise. The body that replaces Ofwat keeps the economic remit (setting the five-year price control that determines bills and allowed investment) and adds three duties Ofwat lacked the tools to discharge: a **financial-resilience duty** (to *enforce*, not merely report on, resilience), an **environmental duty** (exercised jointly with the Environment Agency, to drive discharge reduction and river-quality improvement), and an **"intelligent client" / capability duty** (to build and keep the in-house financial, engineering and data expertise to hold sophisticated, well-advised monopolies to account). The capability duty is what stops the new powers being paper tigers. The regulator is funded to a capability standard, ring-fenced in statute so it cannot be quietly out-resourced by the companies, and given a data and analytics function as a first-class part of the body rather than a back office.

A financial-resilience regime that bites - four interlocking instruments. This is where the system actually broke, so it is the most detailed part. (1) A **binding gearing cap** - a ceiling on borrowing as a share of the regulated asset base - closing the gap between the 55% the rules assume and the 67.9% the sector actually runs at, with a **de-gearing glide-path** for companies above the cap so an abrupt limit does not itself precipitate the distress it exists to prevent. Welsh Water's own history shows de-gearing is achievable over time (gearing fell from about 93% at its 2001 inception to 58% by 2021-22, Grade B). (2) **Conditions on distributions** - owners can be paid only when the company is within its gearing cap, meeting its targets and holding defined liquidity; the lock-up that Thames Water reached *under crisis* becomes a *standing rule applied before crisis*. This is ownership-neutral by construction: it bites on extraction that undermines resilience, whoever the owner is, and does not prevent a no-shareholder company reinvesting surpluses or returning a customer dividend. (3) A **ring-fence** insulating the regulated operating company from its corporate group, so group-level debt cannot drain the entity that has to keep water flowing. (4) The special-administration route below, as the ultimate sanction.

A pre-defined rescue route for a failing company. A credible, *pre-specified* special-administration route is part of the core because the alternative - improvising under crisis pressure with a company serving a quarter of the population - is how a financial failure becomes a national supply emergency. It exists in outline in the Water Industry (Special Administration) Regulations 2024 but has never been operated at scale. The route, written down *in advance*: a trigger (when a company breaches its resilience conditions to the point of insolvency, *before* supply is at risk); a guarantee of **continuity of supply** as the first, non-negotiable objective (a special administrator runs the company as a going concern; customers see no change); and **creditor and investor treatment stated honestly - equity is the first loss, creditors take losses by seniority, and the taxpayer backstops continuity of supply, not investor returns.** Public money keeps the water flowing during administration and is recovered from the restructured entity; it does not exist to make whole the owners or lenders of a company that failed its resilience conditions. The route stabilises and restructures; it does not pre-decide who ends up owning the pipes.

Enforceable environmental and supply targets with penalties that bite. The defect today is not the absence of targets but the absence of *consequences* - penalties have been small enough to be absorbed as a cost of doing business. The regime sets time-bound, binding targets across **both** environmental quality (discharge reduction *and* river-quality outcomes, so a company cannot meet a spill count while a watercourse stays degraded) **and** supply security (leakage reduction on the existing –50%-by-2050 trajectory; closing the projected supply gap; drought and climate headroom). Penalties are calibrated to *exceed the money a company saves by under-performing* - or they are simply a fee for pollution - and are linked to the distribution conditions, so a company failing its targets cannot pay its owners while in breach.

Penalty proceeds are directed to remediation where possible, so the externality is at least partly internalised at source.

Transparent governance of the investment programme. Because customers cannot switch, the only discipline on a water monopoly is **visibility**. The regulator must collect and publish, comparably across companies: gearing and distributions; allowed versus *delivered* investment by programme; target performance; and the bill impact attributable to each. The forward programme is specified through the five-year price control, monitored against **delivered outputs** (kilometres of mains renewed, overflows eliminated, megalitres of new supply) not just spend, and enforced through claw-back - a company cannot bank an allowance for investment it did not make and distribute the difference. The structural backlog this must track is stark: at the current mains-replacement pace, full network renewal would take roughly **700 years** (Grade B).

4. The investment need

The amount the country needs to invest to put the system right and future-proof it is about **£290 billion over 2025-2050** (Grade B - a Water UK / Independent Water Commission / NAO aggregate). This is an empirical figure, needed **whoever owns the companies** - it is the cost of fixing a system that was left to run down, not a cost of any particular owner. It decomposes into named sub-programmes that reconcile back to the total:

Sub-programme	Central	Range	Grade	What it covers
Maintenance & renewal	**£178bn**	£150-220bn	C	Renewing the worn-out network; the residual of the £290bn less the itemised programmes
Environmental / sewage	**£60bn**	£56-120bn+	A (for £60bn)	The Storm Overflows Discharge Reduction Plan to 2050
Supply resilience	**£52bn**	£40-70bn	B	30 strategic projects to ~2050 - reservoirs, transfers, desalination, demand reduction
Leakage reduction	**£0 discrete**	-	D	Funded *within* maintenance & renewal - carried at zero to avoid double-counting, not because it is free
Climate adaptation	**£0 discrete**	-	D	Embedded in supply resilience and the environmental programme - carried at zero for the same reason
Total (central)	**£290bn**			Reconciles to the aggregate anchor ✓

Three things are worth drawing out. **First, the bigger half is the future, not the past.** Maintenance and renewal plus supply resilience (£230 billion) outweighs the environmental clean-up (£60 billion) - securing enough water as the population grows and the climate dries out summers is the larger task, not only cleaning up sewage. The projected **supply gap is about 4,000 megalitres a day by 2050** (Grade B), of which the Environment Agency attributes a large share to climate and drought - which is *why* climate adaptation is carried as embedded rather than as a separate line.

Second, leakage and climate are carried at £0 *discrete* as a deliberate double-count guard, flagged as such - not a claim they are free. Leakage (running at about 2,156 megalitres a day, with a –50%-by-2050 target, Grade A) is reduced *within* the maintenance-and-renewal programme; costing it separately

would double-count spend already in the £178 billion. Climate resilience is embedded in the supply and environmental programmes for the same reason.

Third, the £104 billion AMP8 programme (2025-2030; AMP8 is the water industry's current five-year investment plan set by the regulator, the eighth such plan) is the first five years *inside* this £290 billion - never an addition to it (Grade A for the £104 billion). Of that £104 billion, about £24 billion is environment (including roughly £12 billion on storm overflows) and about £8 billion is supply resilience - each the first tranche *inside* the relevant long-term programme above.

5. The two public choices

The core above is built to hold under any answer to the two questions below. They are routed to the public as neutral, costed menus, presented in full plain language in the Public Choices companion; no option is labelled "recommended".

Choice 1 - who owns the system, and who pays for the past

Two parts, tangled together: who owns the companies, and who pays for the decades of neglect and the forward bill.

Who owns it - three honest options. The decisive empirical finding, established by the two real British non-private comparators, is that **ownership form is *not* a reliable predictor of customer or environmental outcomes** (Grade A on the structural point, mixed on outcomes):

- **Keep them private, but regulate them properly.** The companies stay in private hands under the tougher core above. Quickest and cheapest route - nothing has to be bought out. But it leaves the same owners in place who oversaw the decline, and private owners ultimately need a return, which is a cost in the bill.
- **Public ownership** (the Scottish Water model - a public corporation). Scottish Water has **among the lowest bills in the UK - about £369.55 (2024-25), roughly £70/yr below the England & Wales average** (Grade A), and the state can borrow more cheaply with no outside owners taking a cut. But it costs a large, disputed sum to buy out (below), and it is no magic fix: Scottish Water has **materially worse leakage (~80 l/p/d vs 46.4 in England & Wales, Grade A)**, and Scotland's wetter, emptier geography makes its job easier - so its lower bills are not cleanly attributable to public ownership.
- **A not-for-profit with no shareholders** (the Glas Cymru / Welsh Water model - a company limited by guarantee). No outside owners take a profit; surpluses go back to the system or customers (Welsh Water returned a customer dividend of about £47m in 2018/19, Grade B), and it borrows more cheaply. But it is also no guarantee: Welsh Water has the **highest bills in England & Wales (~£485, year to March 2023, Grade B)** and **one of the worst environmental records** (two-star, 2023-24), and was fined £40m in March 2024 for misleading customers and regulators.

The honest truth running through all three: cheaper capital and no shareholders did *not* reliably deliver lower bills or cleaner rivers. That is precisely why the regulatory core must hold whoever owns the pipes, and why ownership is a value choice - about who you trust to run a vital service and who should profit from it - not a technical question with one right answer.

Why the cost of changing ownership is so uncertain (and must be carried as a range). If the country did take the companies out of private hands, the one-off acquisition cost has **no agreed figure** - every estimate is Grade C, or D at the extreme, *by construction*:

Basis	Figure	Who argues it	On what basis	Grade
Regulated value (gross)	**~£107bn**	Defra (2025)	Buy out equity at the £106.7bn regulated asset value *and* assume the existing ~£72bn of debt	C
Market / enterprise	**~£85bn**	Social Market Foundation (2018)	Enterprise / takeover value; the equity-only slice is far lower given high gearing	C
Statutory / historic	**~£50bn**	PSIRU / University of Greenwich (2019)	A historic invested-capital basis	C
Special administration	**~£0**	McGaughey / Common Wealth (2025)	A failing company taken through the rescue route could transfer for near-zero compensation	D (extreme low)

The ~£107 billion spread between top and bottom is almost entirely a function of a **legal and political choice about what owners are owed** - not a measurement of any underlying asset value. It must be presented as such, never as a technical input with a single right answer. (And the honest creditor-and-investor treatment in the special-administration route, §3, is exactly *why* the near-zero end is a live legal argument.)

Who pays for the past - a distributional choice. Whoever ends up owning the companies, the legacy of under-investment plus the forward bill has to be paid, and there are only three groups it can come from: bill-payers, investors, or taxpayers. How the same burden is split differs by option - illustratively (Grade C, to make the choice *visible*, not to recommend a column):

Option	Bill-payers	Investors	Taxpayers
Regulated-private (status quo)	£80.0bn	£16.0bn	£10.7bn
Public ownership	£42.7bn	£0.0bn	£64.0bn
Not-for-profit	£64.0bn	£10.7bn	£32.0bn

Loading the burden onto bills protects taxpayers but hits households directly, including those who can least afford it. Loading it onto investors feels fair to many but the companies argue it makes the money the system needs harder and dearer to raise. Loading it onto taxpayers spreads it by means but competes with every other call on public money. Each column moves the *same* burden between the three groups; which split is fair is a question for the public. Presenting any one column as the "efficient" or "correct" answer would disguise a distributional choice as a technical necessity - the cardinal error.

Choice 2 - how far and how fast to invest

The £290 billion is what it takes to put the system right and future-proof it. But *how fast* is a genuine choice, and going faster costs more on the bill. The honest ladder - three speeds, with the real impact on the average household bill, **on top of** the rises already happening:

- **Slow - about +£19 a year.** Cheapest on bills; the work still gets done, but rivers and beaches stay dirtier for longer and the supply gap closes more slowly, leaving more exposure to drought in the meantime.

- **Steady - about +£28 a year.** A middle pace: cleaner rivers and more drought headroom sooner, for a bigger addition to the bill.
- **Accelerated - about +£38 a year.** The fastest: cleanest rivers and strongest supply protection soonest, but the biggest addition to the bill.

There is no free version: going faster genuinely costs more each year; going slower genuinely leaves rivers dirtier and supply less secure for longer. How much that is worth is a value judgement, and it belongs to the public.

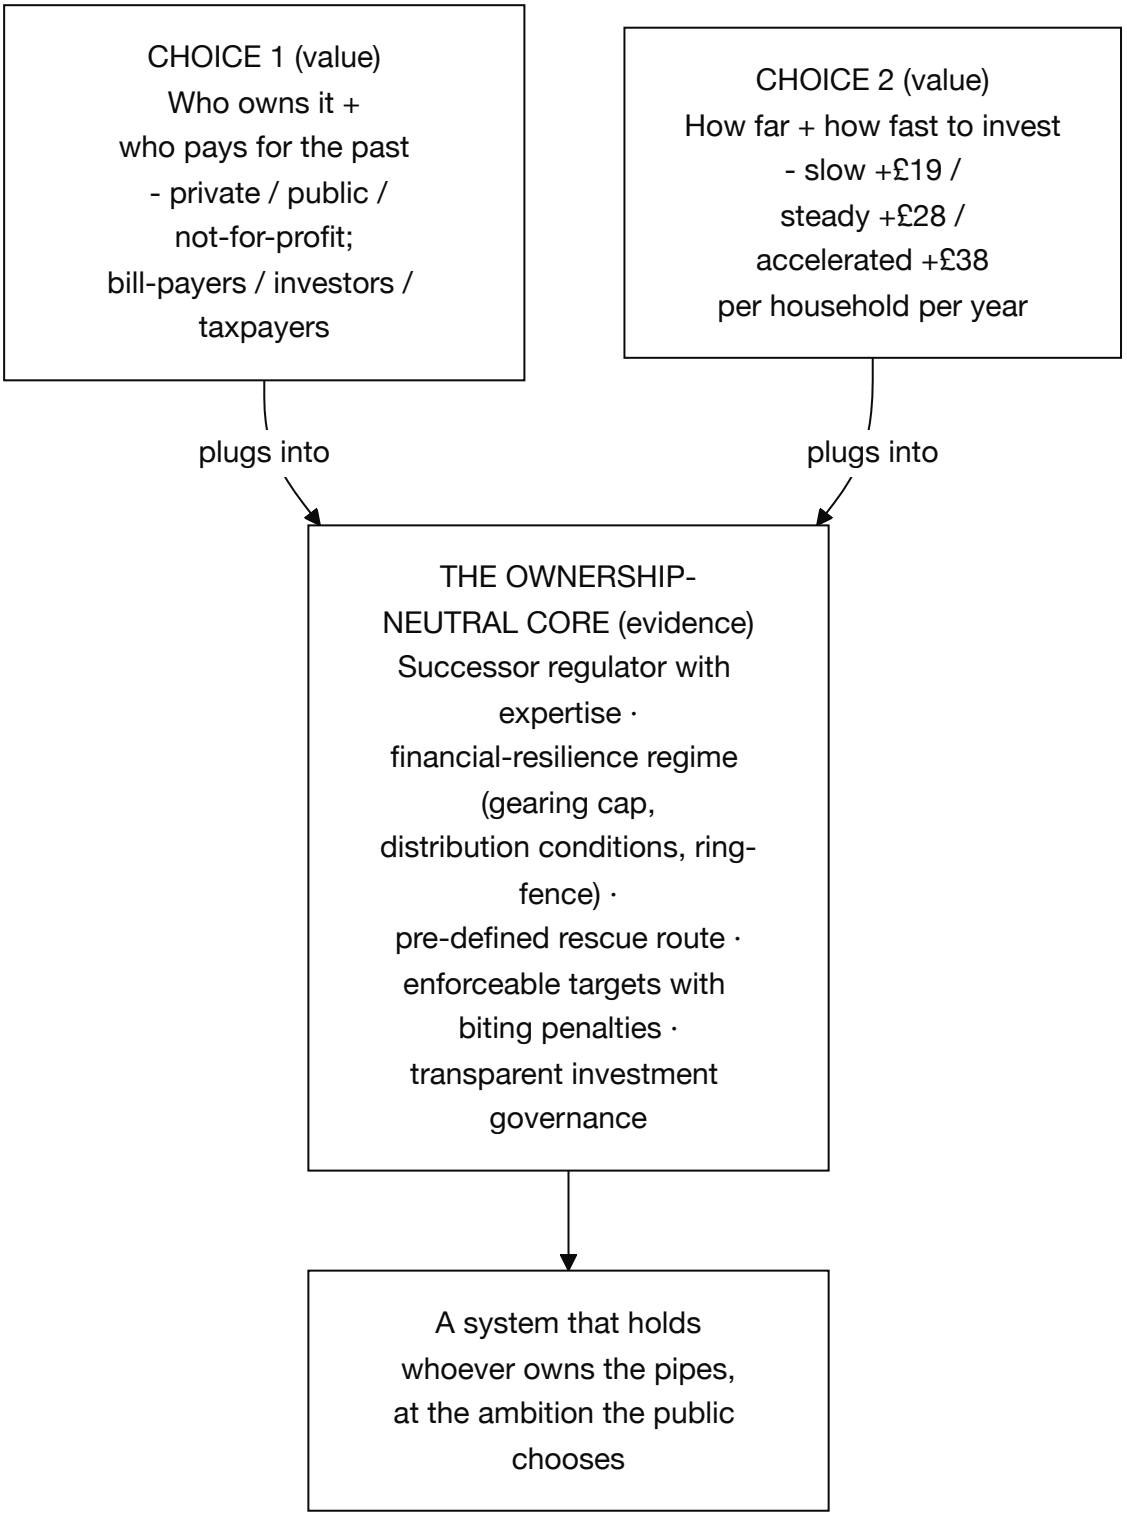

6. The honest costing

This section is where honesty is proven. The full model, its parameters and graded data are published openly (see the costing model and its results); every figure here matches the Evidence Annex at the grade it records.

The investment programme: £290 billion (2025-2050), Grade B, decomposed as in §4 (maintenance & renewal £178bn Grade C; environmental/sewage £60bn Grade A; supply resilience £52bn Grade B; leakage and climate carried at £0 discrete, Grade D, as a double-count guard). The decomposition ties exactly to the £290 billion aggregate (178 + 60 + 52 = 290). The £104 billion AMP8 programme (Grade A) is the first five years *inside* it, never an addition.

The bill-impact ladder (Choice 2): +£19 / +£28 / +£38 per household per year for slow / steady / accelerated, on top of the rises already in train. These are model outputs at the grade the annex records.

The ownership-transition cost: a range, never a single headline. About **£50 billion to £107 billion** depending on the valuation basis and legal choice, with a contested near-zero argument at the extreme (Grade C, and D at the extreme) - see the table in §5. The spread is a value choice, not a measurement, and is reported as a range with the basis attributed to who argues it.

The recurring cost-of-capital effect - modest on the model basis, and genuinely contested. Separately from the one-off acquisition cost, public ownership could change the *annual* cost of financing the same asset base, because public borrowing at the gilt rate is cheaper than the private allowed return. **The model's central estimate is about £0.5 billion a year (Grade C, contested)**, computed as the regulated asset value times the spread between a nominal-equivalent private cost of capital (~6.0%, itself a Grade-C re-basing of the 4.03% real allowed return that the regulator set in its 2024 price review, known as PR24) and the ~5.55% nominal gilt (the interest rate on UK government borrowing) - a spread of only about 0.45 percentage points, hence the modest figure (a rough rule of thumb: about £1 billion a year per percentage point of spread on the asset base). This is **genuinely disputed in both directions**, and the paper carries the counter-arguments rather than banking a saving:

- **PSIRU puts it far higher, at about £3 billion a year** (Grade C, contested) - but it does so by additionally banking the *removal of shareholder distributions*, which bundles a distributional choice into what is presented as a financing-cost saving.
- **Cunliffe judged the regulator's cost-of-capital allowance too low** - if the allowed private return is already below the true cost of capital, the private "excess" over public borrowing is small or nil, shrinking any saving.
- **The Treasury "no free lunch" critique** - the gilt rate understates the true social cost of capital because public ownership transfers risk from investors to taxpayers; the apparent saving is partly just risk shifted onto the public balance sheet, not a real efficiency gain.

The verdict is therefore that the recurring effect is a **scenario parameter, not a settled saving** - somewhere between roughly £0.5 billion a year (modest spread basis) and ~£3 billion a year (distribution-inclusive basis), with credible arguments that the true figure is lower still once risk transfer is priced. It is presented as a contested range with the basis explicit, **never as a guaranteed dividend of public ownership.**

The discipline throughout: settle the empirical figures (the £290 billion programme, the sector's finances, the bill trajectory) on evidence; carry the contested figures (the transition cost, the recurring effect, the legacy split) as ranges with the basis attributed; and never present a value or ownership choice as

technical. Nothing here is banked that cannot be shown, and where the evidence is weak it is graded as such.

7. Adversarial review - three perspectives and the strongest case against

This is the section a sceptical reader should check hardest, and ownership is where partisanship is fiercest - so the discipline here is scrupulous even-handedness. Each perspective is given its genuine reading of **both** the ownership-neutral core **and** each of the three ownership options, reporting where they agree (robust) and where they diverge (shown, not dropped). No perspective is allowed to win.

On the ownership-neutral core (§3)

Fiscal-conservative reading. Broadly supportive: a gearing cap and distribution conditions impose the financial discipline that competition would in a normal market; the ring-fence protects the operating company; the honest special-administration route means owners and lenders - not taxpayers - bear the losses of a failed financial structure, which is exactly market discipline restored. Its caution: the capability duty must not become an excuse for an empire-building regulator, and penalties must be predictable enough not to deter the investment the sector needs.

Social-democratic reading. Also broadly supportive, for different reasons: enforceable targets with biting penalties finally put the environment ahead of distributions; the distribution conditions stop owners being paid while rivers are polluted; transparency lets the public hold a monopoly to account. Its caution: a core that merely *regulates* private extraction better may not be enough if the deeper objection is to extraction at all - which is why it pushes hardest on Choice 1.

Libertarian reading. More sceptical, but finds real value: clear, rules-based limits known in advance are preferable to discretionary regulatory intervention; the ring-fence and honest creditor waterfall respect the consequences of private risk-taking rather than socialising losses. Its caution: a binding gearing cap and distribution conditions are genuine constraints on private property and contract, justified here *only* because the monopoly removes the market discipline that would normally apply - and it would want them no broader than that justification.

Where the three agree on the core (robust): financial resilience must be enforced, not merely monitored; owners must not be paid while a company is failing; a failing company's losses fall on its owners and lenders, not the public; targets need penalties that bite; and a monopoly that cannot be escaped must be held to account by transparency. These are the load-bearing commitments, and all three perspectives can sign them - which is what makes the core robust enough to build first, ahead of the contested ownership question.

On the three ownership options (Choice 1)

Option A - keep private, regulate properly.

- *Fiscal-conservative:* the natural home - keeps private capital raising its own investment off the public balance sheet, avoids a contested and expensive buy-out, and trusts the strengthened core to do the disciplining.
- *Social-democratic:* the weakest option - it leaves in place the owners who oversaw the decline and keeps a permanent return-on-equity cost in every bill; it trusts that *this time* the rules will hold, when they did not before.

- *Libertarian:* acceptable in principle (private ownership, regulated only as far as the monopoly requires), but uneasy that the regulation needed to make it safe is itself heavy.

Option B - public ownership.

- *Fiscal-conservative:* the least attractive - a large, contested acquisition cost largely on taxpayers, the state's borrowing exposed to a £100bn-plus asset base, and (on the Treasury critique) an apparent cheaper-capital saving that is partly just risk shifted onto the public.
- *Social-democratic:* often the preferred option - no shareholders taking a cut, cheaper public borrowing, and the encouraging Scottish Water comparator on bills; the legacy burden shifts toward taxpayers, who pay by means.
- *Libertarian:* the least attractive - the largest expansion of the state, public money tied up in operating a utility, and the buy-out cost falling on taxpayers who did not choose it.

Option C - not-for-profit, no shareholders.

- *Fiscal-conservative:* a middle path it can accept - no permanent shareholder return, capital raised through bonds rather than the public purse, no nationalisation bill - though wary that without equity discipline a company-limited-by-guarantee can drift (Welsh Water's high bills and weak environment are the warning).
- *Social-democratic:* attractive - surpluses to customers and the system rather than investors - but it notes the same Welsh Water warning that no-shareholder ownership is no guarantee of good outcomes.
- *Libertarian:* tolerable - it is a private, contractual form rather than state ownership - but it observes the model still depends on the regulatory core to perform, just as the others do.

Where the three agree on Choice 1 (robust): ownership form does *not* reliably determine outcomes (the comparators settle this); the transition cost is a contested range driven by a choice, not a measurement; and whatever is chosen, the regulatory core must hold. **Where they genuinely disagree (the live value question):** whether private owners should remain in a monopoly essential service and take a return; how generously, if at all, existing owners should be compensated on any transition; and how the legacy burden should fall across bill-payers, investors and taxpayers. These disagreements are real, reasonable and not resolvable by evidence - which is exactly why Choice 1 is routed to the public rather than settled here. The menu deliberately removes the *mechanism* disagreement (the core is the same under any answer) so the public can argue the genuine value question cleanly.

The strongest case against the whole proposal

Put as forcefully as a serious opponent would: *Government is already abolishing Ofwat and acting on 88 recommendations - so a further proposal is at best redundant and at worst muddies a reform in flight. The ownership-neutral core is just "better regulation", which is what was promised last time and failed; the contested £50-107 billion transition range and the disputed £0.5-3 billion recurring effect show the ownership question cannot be costed honestly enough to put to a public vote; and routing the hardest choices to the public is an abdication - voters cannot be expected to adjudicate a regulated-asset-value valuation dispute. Better to let the Commission's reform bed in and judge it on results.*

The answer is in the design, not a deflection. The proposal is **not** redundant: it concentrates precisely on what the Commission left unresolved or was *barred* from considering - the financial-resilience regime in enforceable detail, a pre-defined rescue route, targets with real teeth, and above all the ownership question the Commission's terms of reference forbade it from touching. "Better regulation failed last time" is true and is the reason the core specifies *binding* instruments with *biting* penalties and the in-house *capability* to enforce them, rather than the monitoring tools that failed. The costing's honesty is the point, not a

weakness: the transition cost *is* a contested range because it is a value choice, and saying so plainly - rather than inventing a single headline number - is what lets the public choose on a true picture. And routing the choice to the public is not an abdication but the opposite of the cardinal error: the public are asked the *value* questions (who should own a monopoly, who should pay), presented with the costed consequences in plain language, while the *technical* design is settled on evidence and the same whatever they decide. What the case correctly establishes is that the country should not pre-commit to one contested ownership answer on present evidence - which is exactly why ownership is a public choice and the core is built to hold under any answer to it.

8. Measures of success

System-level, evidence-graded, and built so that no single number can be gamed - covering financial resilience, environment, supply security and value to bill-payers:

- **Financial resilience:** sector and per-company gearing trending toward the cap (the 67.9% → cap trajectory, against the 55% notional, Grade-A baseline); distributions made only within the conditions; companies in breach; companies in or near special administration.
- **Environment:** discharge reduction against the time-bound targets; **river-quality outcomes, not just spill counts**; whether penalties exceed the savings from under-performance.
- **Supply security:** leakage against the –50%-by-2050 trajectory (from ~2,156 Ml/day, Grade A); progress closing the ~4,000 Ml/day supply gap (Grade B); drought and climate headroom.
- **Investment delivery:** allowed-versus-delivered investment by programme; mains-renewal pace against the structural backlog (the 700-year finding, Grade B); the £290 billion programme tracked in five-year tranches, never conflated with the £104 billion AMP8 window.
- **Value to bill-payers:** the bill trajectory and what it buys (the £603 cash bill for 2025/26, Grade A, as the baseline the public can check), published comparably across companies.

The two value-choice outcomes (ownership and ambition) are **not** success measures of the core - they are the public's choices, and the core is judged on resilience, environment, supply and value regardless of how they are answered. An **independent evaluator**, resourced and reporting independently of the regulator and the companies with a statutory protected budget, assesses the regime against these measures. Because government is mid-reform, the evaluation is **standing, not one-off**: it tracks the forthcoming water legislation and the new regulator standing up, and reports as the regime beds in.

9. Implementation summary

The full operational detail - the legislation, the successor regulator's powers in licence-condition form, the financial-resilience instruments, the special-administration route step by step, the sequencing and critical path, the first hundred days, and the explicit exit conditions - is in the Delivery Design. In outline:

- **Legislation:** the successor-regulator statute (already accepted in principle) with the financial-resilience, environmental and capability duties on the face of the Act; the financial-resilience powers (gearing cap, distribution conditions, ring-fence); the standing special-administration regime built out from the 2024 Regulations; and the enforceable-targets power. Any change of ownership (Choice 1) and the investment-ambition trajectory (Choice 2) carry their own, separate legislation *after* the public decides - off the critical path for the core.

- **Sequencing:** stand up the regulator with its capability function first (until it can out-think the companies, the powers are paper); bring the resilience regime live (caps with glide-paths, distribution conditions, ring-fence, the pre-defined rescue route, biting targets); then operate the investment governance and transparency regime. The two public choices follow on their own track once the core is standing.
- **Exit and correction conditions:** a company that fails its resilience conditions goes through the *defined* rescue route, not an improvised bailout; whatever the public decides on ownership and ambition, the core transfers intact and keeps protecting customers and rivers. The resilience core stands alone under any answer - which is what makes it a robust core rather than a bet on one contested question.

10. Open questions

Honesty requires naming what is not yet settled:

- **Legal structure is undecided** and is the founder's call; nothing here presumes charitable status, and the "present options, do not advocate" posture is what keeps that route open.
- **Items flagged for a primary-source upgrade before headline use:** the ~£85 billion cumulative-distributions-since-1991 figure (order-of-magnitude, deflator-dependent - Grade C); Scottish Water's cited "+30-35% investment per household" (advocacy-sourced - to be upgraded via the Scottish regulator's primary data); and the precise sewage-spill *level* (monitoring coverage grew over the period - the direction is robust, the level less so). A full primary-source check precedes publication, as for the social-care and employment-service annexes.
- **The recurring cost-of-capital effect remains a contested scenario parameter** (~£0.5-3 billion a year depending on basis), not a settled saving, and is carried as such.
- **The value-question channel** - a Pragma-hosted deliberation platform with authenticated users able to weigh the two choices and see how others weigh them - is the intended route; its build timing is a dependency (publish the costed choices now, wire the voting later).

Living updates

Under the Pragma Method, a proposal is a living document: as government publishes reports, consultations and legislation, the institute publishes a short non-partisan note on how the proposal bears on them, and revises the proposal where the new evidence improves on its own. Every version is retained and each change recorded.

- **June 2026 (v1.0) - born into a live reform.** Unlike most areas the Register covers, water is being reformed right now. The Independent Water Commission (Cunliffe) reported in mid-2025 with 88 recommendations, and the government has accepted abolishing Ofwat and replacing it with a new regulator (Grade A). This proposal builds on that settlement and concentrates on what it left unresolved or was barred from considering - above all the ownership question the Commission's terms of reference forbade it from recommending on. The forthcoming Water (regulation) legislation and the new regulator's standing-up will be assessed against this proposal on publication, with any revisions recorded here.

Annexes and sources

- **Evidence Annex** - every claim above, with A-D grades and full citations; the contested ownership-transition and cost-of-capital figures are carried there as ranges with the basis attributed.
- **Delivery Design** - the implementation-ready detail behind §3, §8 and §9: legislation, the successor regulator, the financial-resilience instruments, the special-administration route, sequencing and exit conditions.
- **Public Choices** - the plain-language presentation of the two choices (ownership + who-pays, and the investment-ambition ladder).
- **Costing model** and its **results** - the re-runnable appraisal behind §4 and §6.

Key sources (full list in the Evidence Annex): Ofwat PR24 Final Determinations and Monitoring Financial Resilience Report 2024-25; Ofwat Thames Water enforcement decision (May 2025); the Independent Water Commission (Cunliffe) report and the government response; National Audit Office (April 2025); Environment Agency National Framework for Water Resources (2025) and event-duration-monitoring data; Defra Storm Overflows Discharge Reduction Plan; Water UK; Consumer Council for Water; Glas Cymru / Welsh Water and Scottish Water as the two real British non-private comparators; Social Market Foundation (2018) and PSIRU / University of Greenwich (2019) on transition cost; House of Commons Library sewage-discharges briefing CBP-10027. Cross-references: Problem Register entry 7 (water and the regulation of monopoly utilities) and entry 11 (state capacity - the "intelligent client" link in §3).