HARI — CAUSAL CHAIN PROJECTION

# Circulation.

*Which Proposals for a Post-Wage Economy Will Hold*


AUTHOR Joseph Logan · josephlogan.com



DATE May 2026

VERSION 1.0

FRAMEWORK Hari Projection v3.0





SERIES Hari Projection · Standalone deep dive on the April 2026 displacement projection. Monthly cadence February–May 2026; framework debuted August 2025.


AUDIENCE Policy staff, investors, political figures, AI company leadership, investment bank strategy teams, and journalists covering displacement policy.

# Executive Summary.

Displacement is a given. Six months of monthly Hari projections, the academic literature on automation externalities, and the public-policy conversation now converge on the same point: the W2 wage—the dominant transaction mechanism for distributing purchasing power for roughly a century—is no longer carrying the weight alone. AI is the proximate driver, but robotics, geopolitical reordering, and the steady decline of Coase-era transaction costs are layered into the same structural shift. The question is not whether displacement is real. The question is what circulatory system replaces or supplements the wage as the primary mechanism through which capital flows back to the people who buy what companies sell.

This is not a death-spiral argument. A thriving economy is possible—arguably more resilient than the one being displaced—if the circulatory architecture is designed deliberately rather than left to emerge by default. *Circulation* takes that design problem seriously across six candidate architectures already credible in the public conversation: universal basic income, universal basic services, citizen-equity warrants, citizen-equity sovereign wealth fund, Pigouvian automation tax, and public ownership, with sectoral worker ownership as adjacent. Each is evaluated against seven structural criteria and three uniform stress scenarios, with leading indicators specified for falsification.

**Three findings.**

*First*, the public conversation has converged on capital-share architectures (warrants, SWF, dividend mechanisms) across ideologically incompatible actors—Varoufakis, Bruenig, Bores, Musk, the Trump administration. This convergence is real but is a signal about political viability under pressure, not a signal about which architecture corrects the underlying market failure. The two signals point in different directions and the brief separates them deliberately.

*Second*, the Hemenway Falk and Tsoukalas (2026) diagnostic reframes displacement as a market failure—a wedge between private and social returns to automation that produces over-automation at roughly twice the collectively efficient rate. By this diagnosis, UBI, UBS, and citizen-equity (warrants and SWF) are aimed at a different problem: they redistribute outcomes but do not change the marginal automation incentive. Only a Pigouvian automation tax acts on the per-task margin. This analysis stands by the claim against three named objections (Section 4.1) while acknowledging the empirical magnitude is contested.

*Third*, no architecture is robust across all three stress scenarios (sustained fiscal pressure, AI financial collapse, political reversal). The closest is the citizen-equity SWF variant, which is robust to fiscal pressure (asset-backed dividend), mostly robust to AI financial collapse (diversified portfolio), and robust to political reversal (asset trust structure). The cross-ideological convergence on the SWF mechanism, combined with this stress profile, makes SWF the most credible candidate.

**The position.** Among the architectures examined, the citizen-equity SWF variant is the best-positioned candidate to be designed, legislated, funded, and operating before the displacement window closes. Pigouvian elements should be layered on top to correct the underlying automation margin. UBI or warrants may serve as fast cushioning to buy time. Public ownership operates principally as framing pressure that shapes the other architectures rather than as a near-term legislative candidate.

This is a position, not a hedge. It rests on the analysis in Sections 1–13 and is held to the confidence levels named in Section 12.2.

**What would change the assessment.** Section 10 names the leading indicators that would shift the read. The single highest-leverage signal: whether the Trump administration's February 2025 SWF executive order acquires a dedicated seeding mechanism through legislation by mid-2027. If it does, the SWF lean strengthens substantially. If 18 months pass without it, the political-viability score should be revised down materially.

The displacement window is open. The architectural decision is structural and multi-year. The signals favor capital-share mechanisms politically and Pigouvian elements diagnostically. The strongest available combination layers them rather than choosing between them.

# Contents.



 

# 1. Frame.

The April 2026 Hari projection on AI-driven labor displacement treated displacement as a structural arc—a multi-year movement to be tracked, not a series of monthly events to be reacted to. The projection deliberately stopped at the arc. It did not recommend a policy response, because choosing a destination on the arc is a different exercise than identifying that the arc exists.

That choice now needs to be made. The projection raises a question it does not answer:

> *If displacement compounds faster than purchasing power can be reconstituted through ordinary labor markets, what transfer architecture can be designed, legislated, funded, and operating before the political and economic window closes?*

The question contains four hard constraints, each of which dominates a different candidate architecture:

- **Designed.** A specific mechanism, not a slogan. UBI is well-designed in this sense; "tax the AI companies" is not.
- **Legislated.** Capable of passing through actual political institutions, not just academic preference. Public ownership is well-theorized but legislatively distant.
- **Funded.** A revenue source that exists, that scales with the problem, and that does not depend on the same political coalition that resists the architecture itself.
- **Operating.** Standing infrastructure that pays out reliably under stress. Alaska's Permanent Fund Dividend has been operating for forty-four years; most candidate architectures have never operated at scale.

The window matters because of the asymmetry between AI capital accumulation and political organization. As AI companies hit displacement-scale valuations, they become the most politically powerful entities in the economy. Architectures that depend on legislating against them after that point are structurally weaker than architectures installed before. The clock on installation is shorter than the clock on payout.

This analysis examines six architectures across these four constraints. It does not commit to a recommendation. It commits to making the trade-offs visible and the leans contestable.

## 2. Convergence as elevated finding.

Before walking the architectures one by one, a finding this analysis flags in advance: actors operating from incompatible ideological starting points are arriving at overlapping mechanisms, without coordination. This is a *descriptive* claim, well-evidenced. This analysis is careful in this section to separate it from a stronger *interpretive* claim about what the convergence means, which is held as a hypothesis until Section 7.

The clearest cluster is around citizen-equity—variants in which citizens hold ownership claims on capital and receive dividend flows from those claims. The cluster includes:

- Yanis Varoufakis (left, Europe, 2016–present), proposing a Universal Basic Dividend funded by mandatory share contributions from public companies (interview, 2020).
- Matt Bruenig (left, US, 2018–present), proposing a Social Wealth Fund—recently rebranded as the American Solidarity Fund—that issues a non-transferable share to every American and pays an annual dividend from investment income (People's Policy Project; Money With Katie interview, 2025).
- Alex Bores (Democrat, NYC, 2026), proposing an AI Dividend funded by an AI consumption tax, equity warrants on major AI companies, and tax reforms to disincentivize automation (Axios, April 20, 2026; Klein interview, April 28, 2026).
- Elon Musk (right, capital, April 2026), advocating "universal high income" funded by AI productivity gains, mechanism-light but rhetorically aligned (X post, April 16, 2026).
- The Trump administration's federal sovereign wealth fund developments (right, US, 2025–present), institutionally aligned with the social wealth fund mechanism even if the political framing is opposite (OMFIF, December 2025).

The descriptive convergence is striking. The interpretation is genuinely open. Two readings are available, and this analysis does not commit to either at this point:

- **Convergence as signal.** The underlying problem—that AI substitution shifts income from labor to capital, and that a transfer mechanism that does not address ownership misses the structural fix—is shaped such that very different starting positions arrive at the same general answer. This would be evidence the diagnosis is correct, even if the prescriptions vary.
- **Convergence as coincidence.** Each actor is solving for their own constraint with the closest available tool. Varoufakis wants ownership socialized; Bruenig wants wealth diversified across the population; Bores wants insurance against AI displacement that is politically saleable in 2026 New York; Musk wants productivity gains distributed without disturbing ownership of his own companies; Trump wants a federal investment vehicle that looks like national strength. They land near each other because the tool—equity claims paying dividends—is the most flexible structure available, not because they share an analysis.

This analysis treats both readings as live, and Section 7 attempts to disentangle them after the architectural comparison and matrix have been worked through. Three tests apply there:

1. **Underlying problem analysis.** Do the actors share a diagnosis, or only a mechanism? Where the diagnosis differs but the mechanism aligns, convergence is weaker evidence.
2. **Structural similarity.** Are the mechanisms structurally similar (warrants vs. SWFs differ in concentration, contingency, and payout mechanics) or only superficially similar?
3. **Sequence vs. simultaneity.** Did the actors arrive at the mechanism through documented influence, or independently? Independent arrival is stronger evidence.

A further hypothesis worth surfacing here, but not yet earning treatment as a finding: the current public discussion of these architectures appears to treat them as camp-coded—Varoufakis and Bruenig as left, Bores as a Democratic insurance proposal, Musk as right-coded, Trump's SWF as a national-strength play. If the architectural overlap is real and the camp-coding is preventing cross-pollination, that would be analytically significant. The brief raises this question but does not answer it; whether it stands as a finding depends on what the architectural comparison and convergence analysis surface in Sections 5–7.

The convergence finding stays out of the executive summary. It is foregrounded here because it shapes how the architecture walkthroughs that follow should be read: the descriptive overlap is real, the interpretation is open, and the differences that matter most for analysis are at the *mechanism* layer, not the *ideology* layer.

# 3. The architectures (Part 1 of 2).

This section walks the first four architectures: UBI, UBS, citizen-equity (warrants variant), and citizen-equity (social wealth fund variant). Architectures 4 and 5—Pigouvian automation tax and public ownership, with sectoral worker ownership treated as adjacent—are covered in Section 4 (Gate 2).

Each architecture is treated with the same internal structure: diagnosis it answers, mechanism, lead proponents and current status, lineage, and stress-scenario behavior. The seven comparison criteria are scored in Section 5; the convergence/viability matrix in Section 6.

## 3.1 Universal Basic Income (UBI)

**Diagnosis.** Labor displacement produces an income shortfall for affected workers. The shortfall can be offset by direct, unconditional cash transfers to all citizens, removing the targeting and means-testing overhead that limits traditional welfare. UBI corrects the income side of the displacement equation. It does not redirect capital ownership.

**Mechanism.** Periodic cash payments to all adults, set at a level intended to provide a basic floor. Funding is typically proposed through a value-added tax, carbon tax, financial transaction tax, or general appropriation. The Andrew Yang 2020 presidential campaign proposed $1,000 per adult per month funded by a 10% VAT.

**Lead proponents and current status.** UBI is the most legible architecture to existing welfare-state thinking and has the longest pilot history in the US and globally. The most operationally informative current US data point is the Denver Basic Income Project, which has been running unconditional cash transfers ($1,000 per month for twelve months, $6,500 lump sum, or $50 control) to unhoused participants since 2022. First-year results found that nearly half of unhoused participants moved into housing within a year, and the project has continued to expand (Denver Basic Income Project research; State of Homelessness 2024 report). The Stockton SEED project (2019–2021) produced earlier US evidence that recipients used cash for basic needs and that employment among recipients rose, not fell. Internationally, the Finnish basic income trial (2017–2018) and the GiveDirectly experiments in Kenya are the largest comparable studies.

UBI's federal political status in the US remains theoretical. No major federal proposal is currently under serious legislative consideration. Yang's 2020 campaign moved the Overton window but produced no legislation. Modest precedents exist—pandemic stimulus checks, the expanded Child Tax Credit in 2021—that demonstrate the federal government can deliver direct payments operationally, but neither was framed or designed as UBI.

**Lineage.** Milton Friedman's negative income tax (1962) and the Family Assistance Plan proposed by Nixon (1969) are the conservative ancestors; the Speenhamland system in late-eighteenth-century England is the older institutional reference. The contemporary left-coded form runs from Philippe Van Parijs through the Basic

 

Income Earth Network to Yang. The Alaska Permanent Fund Dividend, treated separately as a citizen-equity variant in Section 8 (historical analogues), is sometimes cited as a UBI but is structurally a dividend on resource wealth—closer to citizen-equity in design than UBI.

**Stress-scenario behavior.** UBI's stress profile is the inverse of citizen-equity's. It is robust to AI financial collapse—the program does not depend on AI companies being solvent or successful—but vulnerable to fiscal pressure, because funding requires ongoing congressional appropriation against competing priorities. In a recession, the same political dynamic that suppresses welfare spending generally would suppress UBI levels. The 2026 Alaska PFD fight ($3,800 statutory benefit reduced to $1,500 in the House-passed budget) is a directional indicator: even institutionalized cash transfers degrade under fiscal stress when the funding mechanism allows annual appropriation override.

**Lean.** UBI is the architecture most likely to be passed under acute political pressure, because of legibility and pilot evidence. It is also the architecture least likely to correct the underlying capital-vs-labor split, because cash transfers funded by taxes do not change who owns the productive assets. Under the seven comparison criteria (Section 5), UBI scores high on clock speed and coalitional viability, low on structural correction depth and funding durability, and ambiguous on stress behavior depending on the funding source. Confidence: moderate-to-high on the directional claims; the operational details depend on pilot-to-program scaling questions that current evidence does not resolve.

## 3.2 Universal Basic Services (UBS)

**Diagnosis.** Labor displacement compounds because real income includes both cash and the cost of essential services. If healthcare, housing, transit, childcare, and education are decommodified—provided as universal entitlements rather than market goods—the income required to maintain a basic standard of living drops, even if cash income falls. UBS corrects the cost-of-living exposure side of the displacement equation. Like UBI, it does not redirect capital ownership; unlike UBI, it directly reshapes consumption costs.

**Mechanism.** Universal provision of a defined bundle of essential services, funded by general taxation or earmarked revenue. The defining design questions are: which services qualify as "basic," what quality level is guaranteed, who delivers (state, contracted private providers, hybrid), and how the services are paid for at the margin (free at point of use vs. heavily subsidized vs. universal voucher).

**Lead proponents and current status.** The fullest articulation comes from the UCL Institute for Global Prosperity (Anna Coote and Andrew Percy), which proposed in 2017 a UBS package including healthcare, education, transport, food, housing, and information services for the UK. The UK National Health Service is the most fully operating real-world UBS instance for a single category (healthcare). Most of the existing UBS literature is European; the US version is closer to Medicare for All plus housing and childcare guarantees, none of which is currently advancing federally.

**Lineage.** The post-war welfare states of Europe—particularly the Beveridge Report (1942) and the social democratic models of the Nordic countries—are the institutional ancestors. Decommodification as a concept goes back to Polanyi and was theorized in Esping-Andersen's three worlds of welfare capitalism. The contemporary UBS proposals are explicit about positioning themselves against UBI: services rather than cash, on the argument that markets in essential goods produce inflation that consumes UBI payments.

**Stress-scenario behavior.** UBS is the most robust architecture to political reversal once installed (the NHS has survived multiple Conservative governments) but the slowest to install, because it requires standing up service delivery infrastructure, not just a transfer mechanism. Under fiscal pressure, UBS degrades through service quality reduction and waiting times rather than through outright benefit cuts—a different and slower failure mode than cash-transfer architectures. UBS is not a fast response to acute displacement; it is a long-term restructuring of the cost base. Like UBI, it is robust to AI financial collapse because it does not depend on AI capital outcomes.

**Lean.** UBS is the slowest of the candidate architectures by a wide margin and the architecture least likely to be standing before the displacement window closes. It is included in the comparison because (a) it is a real candidate in the public conversation, (b) its different failure modes illuminate the others by contrast, and (c) the strongest political response to displacement may end up being a UBI/UBS hybrid where cash handles the acute shortfall and service expansion handles the structural cost base. Confidence: high on the directional claims; the operational details depend on US political conditions that are difficult to forecast.

### 3.3 Citizen-equity—warrants variant (Bores)

**Diagnosis.** Labor displacement, if it occurs at scale, will be a direct result of AI capital appreciation: AI companies become enormously valuable precisely because they are substituting for labor. The fix is to give the public a contingent claim on that appreciation—warrants that pay out only if AI companies hit displacement-scale valuations. Citizen-equity in this form corrects the capital-vs-labor split specifically in the success scenario for AI capital. It does not address scenarios where displacement happens but AI capital does not appreciate as expected.

**Mechanism.** Out-of-the-money equity warrants on major AI companies, granted to or purchased by the federal government at strike prices set today. If AI valuations rise dramatically, the warrants can be exercised and the proceeds fund direct payments to citizens. Bores has framed the mechanism as insurance: low cost in the scenario where AI does not displace labor at scale (warrants expire worthless), high payout in the scenario where it does. He has also suggested AI companies and their investors should support pre-success warrants on the grounds that the alternative—post-success political pressure for "seizing the means of production"—is structurally worse for them.

The full Bores AI Dividend plan combines the warrants mechanism with an AI consumption tax (described as "modest," measured in tokens), tax reforms to reduce the relative incentive to substitute capital for labor, and trigger conditions tied to objective economic indicators (Axios, April 20, 2026; Yahoo Finance / Bloomberg coverage; Gizmodo coverage).

**Lead proponents and current status.** Alex Bores, New York Assemblymember and 2026 Democratic House candidate, is the lead proponent and the only US political figure with a fully specified warrants proposal. The plan was unveiled on April 20, 2026, as part of a broader 43-idea framework Bores rolled out beginning February 12, 2026. AI super PACs are reportedly spending significant sums against Bores's congressional bid, treating the proposal as a direct industry threat (Politico, February 9, 2026).

The status of the proposal is: introduced, not legislated, contested. No federal legislation has been drafted. The earliest plausible path to law would be a Bores victory in 2026 followed by federal sponsorship in 2027–2028; the political headwinds from AI industry capital are substantial.

**Lineage.** Warrants as a public-finance instrument have a long history—TARP included warrants on major banks in 2008–2009, exercised at varying levels of profitability for the Treasury. The conceptual ancestor for citizen-equity via warrants is closer to Henry George's land-value taxation (capturing the public component of asset appreciation) than to traditional welfare or social democracy. Bores's proposal explicitly positions itself as conservative-compatible insurance rather than redistribution, which is a deliberate framing choice.

**Stress-scenario behavior.** This is where the warrants variant has its sharpest analytical exposure, and where it differs most from the SWF variant. Three failure modes apply, treated as a scenario tree rather than a single claim:

*Scenario 1—AI succeeds at scale, warrants pay out as designed.* This is the Bores success case. Warrants are exercised, proceeds fund the dividend, the architecture works. This scenario is probably the least likely outcome over the next decade based on current capital structures, but it is the scenario the mechanism is built for.

*Scenario 2—AI partially succeeds, but financial structure cracks.* Training costs continue compounding faster than revenue; hyperscaler capex commitments outrun customer monetization; circular financing arrangements between model labs, cloud providers, and chip manufacturers unwind. AI valuations decline meaningfully but the underlying labor-displacement effect continues, because the deployed capability is sufficient to substitute for many roles even if the companies producing it are not commercially successful at the level current valuations imply. In this scenario the warrants are worthless or near-worthless precisely when displacement is happening. The architecture is misaligned with the actual failure mode.

*Scenario 3—Distress event, warrant claims renegotiated.* Even if AI companies remain solvent in aggregate, distress events produce equity restructuring. Government warrants are junior contingent claims. The historical analogue is GM 2009: in the federal government's intervention, existing equity was wiped, bondholders took significant haircuts, and the government's preferred and common stock positions were exited at substantial losses despite being senior to the wiped-out equity. A government warrant on a distressed AI company would be a more junior claim than the GM TARP positions and would likely be renegotiated, diluted, or extinguished before senior creditors took their full losses.

These three scenarios are not exhaustive but cover the most likely failure surfaces. The deeper analytical claim: warrants are designed for one scenario (sustained AI success at scale, displacing labor while compounding valuations) and have poor coverage of the adjacent scenarios. The intermediate scenario—where displacement happens but AI capital does not appreciate as designed—may be the most likely outcome based on current capex math, and is exactly the case the warrants do not handle.

**Lean.** The warrants variant is the only candidate architecture whose clock speed could plausibly match displacement—installation requires one legislative event, payout is contractual and automatic. This advantage is contingent on three conditions: (a) installation in a political window that may already be closing; (b) AI issuers surviving financially through the payout window; (c) warrant claims surviving renegotiation pressure during distress. The combination of contingencies materially weakens the architecture against the headline clock-speed claim. The warrants variant is best understood as a high-variance bet—strong in the scenario it is designed for, weak across most adjacent scenarios. Confidence: moderate, with the clock-speed claim being the most fragile element of the brief's overall analysis.

### 3.4 Citizen-equity—social wealth fund variant (Bruenig / Varoufakis lineage)

**Diagnosis.** The same diagnosis as the warrants variant—labor displacement is fundamentally a capital-vs-labor split problem, and any transfer architecture that does not redistribute ownership treats a symptom rather than a cause. The SWF variant differs from warrants in *how* it redistributes ownership: through a diversified, ongoing equity stake in the broader economy rather than through a concentrated, contingent claim on a small number of AI companies.

**Mechanism.** A federal investment fund holds a diversified portfolio of public equity, fixed income, and other assets. Every adult citizen receives a single, non-transferable share of the fund. The fund pays an annual dividend to all shareholders from investment income; the dividend level is set by a formula based on fund returns, with smoothing across years to reduce volatility.

The fund is capitalized through some combination of: a tax on share issuances and corporate transactions (Bruenig's preferred mechanism, modeled after Sweden's Meidner Plan), reinvested government revenues, a wealth tax, or asset transfers from existing federal holdings. Bruenig's American Solidarity Fund proposal models the fund accumulating to roughly $10 trillion in assets over time, paying a per-citizen dividend in the low thousands of dollars annually at scale (People's Policy Project; Money With Katie interview, March 2025; In These Times, 2018; Next System Project SWF report).

Varoufakis's Universal Basic Dividend variant applies the same logic at a different scale: every public company contributes a fixed percentage of its share issuance into a public trust, which then pays dividends to all citizens. The Varoufakis design is more aggressive (mandatory dilution of existing shareholders) and theoretically anchored; Bruenig's design is more incremental and operationally specified.

**Lead proponents and current status.** Bruenig (People's Policy Project, US) and Varoufakis (DiEM25, Europe) are the lead theorists. The Norwegian Government Pension Fund Global is the closest existing real-world analogue—a sovereign wealth fund with broad equity holdings paying into the Norwegian fiscal account, though not paying direct citizen dividends. Alaska's Permanent Fund is the closest US analogue but is built on resource royalties rather than diversified equity, and its dividend mechanism is structurally vulnerable in ways covered in Section 8.

The Trump administration's federal sovereign wealth fund initiative (2025) has institutionally aligned the right with the SWF mechanism, even though the intended use is closer to industrial policy than citizen dividends (OMFIF analysis, December 2025). The institutional convergence—both major US political coalitions now have an SWF mechanism in their active policy frame—is one of the strongest convergence signals identified in this brief.

**Lineage.** Sweden's Meidner Plan (1976), which proposed transferring shares to industry-wide worker funds at a fixed annual rate until workers held majority stakes, is the most direct historical precursor. The plan was opposed by Swedish capital, partially implemented in weakened form, and abandoned in the 1990s. This analysis treats Meidner as a cautionary case, with an important qualification.

The maximal cautionary reading—*Meidner failed, therefore SWFs are politically unviable*—overstates. Advocates correctly note that Meidner used a mandatory-dilution design, transferring 20% of profits annually as new shares to worker funds, which directly diluted existing equity holders. Bruenig's American Solidarity Fund, by contrast, accumulates assets through new revenue streams (taxes on share issuance and corporate transactions, asset transfers, reinvested government revenue) rather than through dilution of existing equity. The political opposition to mandatory dilution is structurally different from, and stronger than, the opposition to tax-funded accumulation.

The accurate cautionary reading: mandatory-dilution SWF designs face durable structural opposition, demonstrated across multiple political contexts—Sweden in the 1970s–80s, the UK Labour Inclusive Ownership Funds proposal in 2019, which faced similar attacks and contributed to the party's election loss. Tax-funded SWF designs avoid that specific opposition vector but have not yet been tested at scale against the next-most-likely opposition: fiscal cost objections under conditions of compressed tax base. The brief's net assessment is that the Bruenig design is materially more politically tractable than Meidner was, but the political question has shifted rather than disappeared.

Varoufakis's lineage runs through European democratic socialism with influence from James Meade's *Agathotopia* and from the post-Keynesian tradition. Bruenig's design explicitly draws on the Alaska Permanent Fund as proof of concept, on Norway as proof of scale, and on Meidner as proof that the political coalition matters as much as the design.

**Stress-scenario behavior.** The SWF variant has the strongest stress profile of the citizen-equity architectures, by a meaningful margin:

*AI financial collapse.* Diversified across the economy, the fund's value is largely independent of AI sector performance. An AI-specific collapse damages the AI portion of the portfolio but leaves the broader fund intact. Compare warrants, where AI collapse is catastrophic.

*Distress events.* The fund holds equity claims pari passu with other shareholders, not contingent or junior claims. Restructuring of any individual portfolio company affects the fund proportionally to its holding, not asymmetrically. The Bruenig design holds non-controlling stakes; this insulates the fund from being treated as the residual claimant in distress workouts.

*Fiscal pressure.* This is the SWF variant's main stress exposure, but it is structurally lower than UBI's. The dividend is paid from investment income, not from annual appropriation. The fund itself is shielded from being raided by the same legal protections that insulate Norway's Government Pension Fund Global. The remaining political risk is at the level of reducing the formula or freezing dividends during recessions—closer to the Alaska 2026 dynamic, but at the federal level and against more diversified asset backing. The fund persists; the dividend level can degrade.

*Political reversal.* Once installed, the fund is hard to dismantle without a constitutional or super-majority political event. This is similar to the NHS's resilience pattern: the architecture survives but the level of provision can be eroded.

**Lean.** The SWF variant has the highest stress robustness of the candidate architectures and the deepest structural correction depth available within citizen-equity. Its weakness is clock speed: capitalizing a fund to scale takes years to decades, and the political mobilization required to install one is comparable to UBS's. The fund pays small or zero dividends in the early years; the displacement window may close before the fund matures into operational relevance.

The institutional convergence finding (Trump federal SWF + Bruenig SWF + Varoufakis UBD + Alaska/Norway operational analogues) is the strongest structural-agreement signal in the comparison. The mechanism is one that capital, labor, left, and right have all separately moved toward—for different reasons, at different scales, but converging on the same legal-financial form. Confidence: high on the architectural strengths and weaknesses; moderate on the convergence interpretation.

# 4. The architectures (Part 2 of 2).

This section completes the architectural walkthrough begun in Section 3, covering the Pigouvian automation tax and public ownership, with sectoral worker ownership treated as adjacent in the public-ownership chapter. The same internal structure applies: diagnosis, mechanism, lead proponents and current status, lineage, and stress-scenario behavior.

## 4.1 Pigouvian automation tax

**Diagnosis.** The Pigouvian framing rests on a different analytical foundation than UBI, UBS, or citizen-equity. It does not treat displacement as an income shortfall to be transferred away, or as an ownership inequity to be redistributed. It treats displacement as a *market failure*—specifically, a negative externality that firms impose on each other and on the broader economy when they automate.

The argument, formalized in Hemenway Falk and Tsoukalas, *The AI Layoff Trap* (2026) and applied in subsequent policy commentary (Westover, *AI Automation Paradox*, April 16, 2026), runs as follows: when a firm automates a task, it captures the full cost saving, but the resulting demand destruction (displaced workers spend less, including on this firm's products) is borne mostly by *other* firms. This wedge between private and social returns to automation persists under competition, intensifies as more firms enter and as AI capabilities improve, and produces over-automation at roughly twice the collectively efficient rate under illustrative parameters. The result is not redistribution; it is deadweight loss, harming both workers and shareholders relative to the cooperative optimum.

The diagnostic move matters because it forces a sharper test on the other architectures. If displacement is a market failure caused by a wedge at the per-task automation margin, then any architecture that does not act on that margin leaves the underlying distortion in place. By Hemenway Falk's logic, UBI and capital income taxation are structurally incapable of correcting the externality—they redistribute outcomes but do not change the marginal automation incentive. UBS narrows the demand-loss term but does not eliminate the wedge. Citizen-equity (warrants and SWF) operates on capital ownership and dividend flow, not on the automation margin. Only an automation tax, properly calibrated, internalizes the externality.

This is a strong claim. This analysis stands by the claim. It does not establish that the Pigouvian architecture is better than the others on all dimensions—clock speed, political viability, and operational design have their own constraints—but it does establish that the others are aimed at a different problem. A reader who accepts the Hemenway Falk diagnosis must accept that UBI/UBS/citizen-equity, however valuable on their own terms, do not correct the underlying market failure.

Three objections to the diagnostic claim deserve direct response, since each has been made in adjacent literatures:

*Objection 1—"Redistribution closes the demand loop, so UBI does correct the externality indirectly."* The argument: if displaced workers receive transfers, their consumption is partially restored, which reduces the demand-loss term in the externality. The response: this is true at the level of aggregate demand but not at the per-task automation margin. The wedge between private and social returns to *each automation decision* is unchanged by downstream redistribution, because the firm automating still captures the full cost saving while the demand loss is borne by rivals. UBI shifts where the demand reappears (concentrated in transfer recipients rather than displaced workers' prior employers), but does not change the marginal incentive to automate that produces the over-automation result. The Hemenway Falk model is explicit on this point: redistribution attenuates the *welfare loss* but does not correct the *marginal distortion*.

*Objection 2—"Citizen-equity gives workers a stake in the firm doing the automating, which internalizes the externality through ownership."* The argument: if displaced workers own warrants or SWF shares in AI firms, they capture the cost saving alongside the firm, eliminating the wedge. The response: ownership claims correct the *distributional* failure (workers share in capital returns) but not the *allocative* failure (over-automation relative to the cooperative optimum). A worker who owns 1% of AI firms via a citizen-equity structure still loses far more from displacement than they gain from the dividend, because the dividend reflects only their pro-rata share of cost savings while the displacement is concentrated. The math is in the Hemenway Falk model: ownership share would need to approach majority control to equalize private and social returns at the automation margin, and no citizen-equity proposal in serious circulation is at that scale.

*Objection 3—"The model is theoretical; we don't know if the wedge is empirically large enough to matter."* The argument: even granting the externality, the calibration matters. If the wedge is small, the over-automation result is small. The response: under the illustrative parameters in Hemenway Falk and Tsoukalas (2026), automation occurs at roughly twice the collectively efficient rate. This is one model under one set of assumptions; this analysis acknowledges the empirical fragility of the magnitude. But the *direction* of the result is robust to parameter choices in standard task-based automation models. Acemoglu and Restrepo's task-displacement framework (2022, MIT working paper) and Guerreiro, Rebelo, and Teles (2022, *Review of Economic Studies*) both reach qualitatively similar conclusions about over-automation under current incentive structures. The wedge's existence and direction are well-established; the magnitude is contested but plausibly material.

Where this analysis could be wrong: if the wedge turns out to be small under realistic parameters (Objection 3, strong form), or if labor markets reallocate displaced workers fast enough that the demand-loss term decays before the externality compounds. These are empirical questions, and this analysis is making a structural claim that depends on the empirical questions resolving in directions consistent with the existing models. If they resolve otherwise, the diagnostic claim weakens proportionally.

**Mechanism.** A per-task levy on automated work, set equal to the uninternalized demand loss the automation imposes on rivals. The rate depends on observable sector parameters: average wage, share of income spent in the sector, and the fraction of displaced labor income that gets replaced (through new jobs, transfers, or savings). The levy is collected on automation events, not on AI capital ownership or AI consumption.

Revenue can be deployed three ways: (a) lump-sum returned to firms (preserves the marginal incentive correction without redistributive effect), (b) transferred to displaced workers (combines correction with cushioning, but introduces moral hazard at the margin), or (c) channeled into retraining (transitional: raises the income replacement rate, which shrinks the future tax base). Hemenway Falk treats option (c) as the strongest pairing on theoretical grounds. The Danish flexicurity model is the closest existing real-world analogue—payroll taxes fund retraining at greater than 80% income replacement, producing an implicit wedge against rapid automation without an explicit Pigouvian levy.

**Lead proponents and current status.** Hemenway Falk and Tsoukalas (2026) provide the formal model. The Windfall Trust Policy Atlas treats automation taxes as one of a small number of "frontier" policy responses to AI displacement and tracks their policy status across jurisdictions. South Korea reduced tax incentives for automation in 2017 in what was widely framed as the first national automation tax, though the mechanism was subtractive (removing prior tax preferences) rather than additive (imposing a new levy).

The right-coded counterpoint comes from the Mercatus Center (October 2025), which argues for *removing* obsolete IRC restrictions to unleash market-driven adaptation rather than imposing new taxes. The Mercatus position rejects the Hemenway Falk externality diagnosis implicitly by treating displacement as a transitional adjustment that markets handle, given the right tax-code conditions. This analysis notes the disagreement but does not adjudicate it: the two positions disagree on whether the underlying market failure exists, not on how to respond to it.

No federal Pigouvian automation tax is currently under serious legislative consideration in the US. Bill Gates floated the concept in 2017 in widely circulated comments, which produced commentary but no policy traction. The architecture is academically well-developed and politically nascent—closer to UBI's status circa 2010 than to citizen-equity's current 2026 momentum.

**Lineage.** Pigou's 1920 framework for taxing negative externalities (carbon taxes, congestion pricing, alcohol and tobacco excise) is the analytical foundation. The application to automation specifically traces through Acemoglu and Restrepo's task-based models of automation (2022), Guerreiro, Rebelo, and Teles (2022) on robot taxation, and the Hemenway Falk and Tsoukalas formalization (2026). The political tradition closest to it is the carbon-tax movement: an externality-correction frame that has produced significant policy in Europe and Canada and persistently failed at the US federal level despite broad academic support.

**Stress-scenario behavior.**

*AI financial collapse.* The architecture is largely robust to AI sector financial conditions because the tax is on automation *events*, not on AI company valuations. If AI companies struggle financially but their tools are still deployed at scale, the levy continues to apply. If AI deployment slows because companies fail commercially, the levy collects less but the underlying displacement is also reduced.

*Fiscal pressure.* The Pigouvian architecture is not a fiscal cost—it generates revenue. This makes it structurally durable in fiscal-stress scenarios, and is the single strongest comparative argument for it relative to UBI/UBS. The political opposition to it comes from the firms whose marginal incentives it changes, not from fiscal hawks.

*Political reversal.* This is the architecture's main vulnerability. Once installed, an automation tax can be repealed by a future Congress, and the political coalition that opposes it (firms benefiting from the externality wedge) is concentrated, well-resourced, and persistent. The carbon-tax history is the cautionary precedent: jurisdictions that have installed carbon taxes have generally preserved them, but the US federal government has resisted installation for thirty-plus years against strong academic consensus, and incumbent industry opposition has been the proximate reason in nearly every case.

*Implementation challenges.* Two are unique to the Pigouvian architecture: observability (firms must report automation rates accurately) and offshoring (firms can avoid the levy by relocating automated work to jurisdictions without it). Hemenway Falk treats the first as tractable (AI deployment leaves digital records that can be audited) and the second as requiring multilateral agreement or border adjustment, neither of which is currently in motion.

**Lean.** The Pigouvian automation tax is the only architecture in the comparison that operates on the per-task automation margin, and on the Hemenway Falk diagnosis that is structurally significant. It is also fiscally self-funding, which materially distinguishes it from UBI/UBS in stress scenarios. Its weaknesses are political: incumbent industry opposition is concentrated and persistent, and the architecture has produced essentially no US legislative momentum in the decade since Bill Gates surfaced it. The brief's net assessment is that Pigouvian architecture is undervalued in the current displacement debate relative to its diagnostic strength, but its political clock speed is the slowest of the candidates examined. Confidence: high on the diagnostic claim (the Hemenway Falk model is well-formalized and increasingly cited); moderate on the political-feasibility claim (the carbon-tax precedent is informative but not deterministic).

## 4.2 Public ownership

**Diagnosis.** Public ownership treats displacement as a problem of *who owns the productive capital itself*, not as an income, cost-of-living, ownership-claim, or externality problem. Where citizen-equity gives citizens dividend claims on capital they do not control, public ownership puts capital under collective control directly. The diagnostic move is the deepest of the candidates: every other architecture, in this view, leaves the underlying ownership structure intact and intervenes downstream of it.

The fullest contemporary articulation is Zack Exley's four-part series *Why Capitalism Can't Survive AI* (New Consensus, February 2025 – April 2026). The argument across the series:

- **Part 1 ("This Time It's Different")** argues AI displacement is structurally distinct from prior automation waves because the displaced are not low-wage workers being absorbed into new sectors but the professional-managerial class being displaced from sectors where there is no obvious next destination (Part 1, February 2025).
- **Part 2 ("The Shape of the Collapse")** maps the demand-side cascade: displaced professionals reduce consumption, which propagates through service sectors, which produces second-order displacement.

 

- **Part 3 ("UBI Can't Save Us")** argues that even at adequate funding levels, UBI fails to address the underlying ownership shift and produces a society of permanent dependents on transfer payments from a small AI-owning class (Part 3, March 2026).
- **Part 4 ("Only One Door Left")** concludes that public ownership of the AI productive base is the only architecture that addresses the diagnosis at the right level.

The Exley argument is not incremental: it treats the other architectures as partial responses that buy time but do not solve the underlying problem.

**Mechanism.** Direct public ownership of AI infrastructure (compute, models, deployment platforms) and AI-derived productive capacity. The mechanism varies by proposal: nationalization of frontier AI companies, public-option AI development at federal scale (analogous to public broadcasting or the postal service), municipal or state-level public AI utilities, or worker-controlled AI cooperatives in specific sectors. The Exley framing emphasizes federal-scale public investment in AI development as a public good, with the productivity gains accruing to the public rather than to private equity holders.

**Lead proponents and current status.** Exley and the New Consensus team are the most visible US proponents of the public-ownership frame for AI specifically. Adjacent proposals exist in the democratic-socialist policy community (DSA, Roosevelt Institute) and in international contexts (UK Common Wealth, European public-AI initiatives). The legislative status in the US is essentially zero: no major proposal is under consideration, and the political coalition required is substantially smaller than the coalitions for UBI or citizen-equity.

**Lineage.** The lineage runs through twentieth-century democratic socialism (the Labour Party's pre-Blair tradition, the Swedish social democratic model in its strongest form, post-war French and Italian planning traditions) and through US progressive-era public-ownership experiments (TVA, public utilities, postal service). The contemporary application to AI specifically is newer and is led conceptually by Exley.

**Stress-scenario behavior.** Public ownership has the strongest theoretical correction depth and the slowest practical clock. Once installed, it is highly robust to political reversal (privatization of public utilities is politically expensive in most jurisdictions and rarely complete). It is also robust to AI financial collapse, since the public entity continues to operate at production cost regardless of private market valuations. Its main vulnerability is operational: public AI development at scale has limited precedent, and the institutional capacity required does not currently exist in US federal government.

**Lean.** Public ownership is the architecture least likely to be standing before the displacement window closes, and the architecture with the deepest structural correction if it were installed. Its analytical role in this analysis is as a *framing force*—the rhetorical and theoretical pressure of the public-ownership argument may shape what the other four architectures look like, even if public ownership itself does not advance legislatively. Exley's Part 3 critique of UBI, in particular, partially overlaps with the Hemenway Falk diagnostic critique despite reaching it from a very different starting point—a convergence the brief returns to in Section 7. Confidence: high on the directional claims; the operational details depend on political conditions that are not currently in motion.

### 4.3 Sectoral worker ownership (adjacent treatment)

Sectoral worker ownership is treated as adjacent rather than as a sixth architecture. The exclusion is a methodological choice and worth making explicit, because the alternative—silently ignoring it—would be a bigger analytical error than the choice to exclude.

**What it is.** Policy that gives workers in entire industries (not single firms) ownership stakes in the firms they work for. The canonical proposal is Sweden's Meidner Plan (1976), which would have transferred 20% of profits annually as new shares to industry-wide worker funds until workers held majority stakes (Wikipedia: Employee funds; PDXScholar working paper on Swedish Social Democracy and the Meidner Plan; nordics.info on wage earner funds). The UK Labour Party's 2019 Inclusive Ownership Funds proposal was a smaller-scale modern descendant: 10% of company shares transferred to worker trusts over ten years, with dividends partly distributed to workers and partly to government. Mondragon in the Basque Country is the firm-level cousin operating at scale. German codetermination provides governance rights without equity transfer.

**Why it could be relevant to AI displacement.** UBI, UBS, and the citizen-equity architectures all redistribute *after* the labor-vs-capital split has occurred. Sectoral worker ownership intervenes *in* the split: workers own a fraction of the productive capital, so as AI substitutes for their labor, the substitution flows partially back to them as capital income. It is a structural fix to the same problem citizen-equity addresses, located inside the firm rather than at the national-fund level. The Hemenway Falk policy table classifies it as "moderate" effectiveness—partially internalizing the externality but leaving residual demand leakage to other sectors.

**Why it is excluded from the main comparison.** Three reasons:

1. *Different intervention category.* Architectures 1–5 are transfer mechanisms (or correctives to the substitution incentive). Sectoral worker ownership modifies the firm's ownership structure. Comparing it on the same seven criteria is awkward—it answers different questions about who benefits from automation rather than how the costs are absorbed.
2. *Slower political clock than public ownership.* The Meidner Plan took two decades of political mobilization in a country with the world's strongest social democratic coalition, was implemented in significantly weakened form, and was abandoned in the 1990s. The UK Labour 2019 proposal contributed to the party's election loss. The political constituency for sectoral ownership in the contemporary US is a fraction of the constituency for cash transfers or citizen-equity dividends.
3. *Coverage problem.* Sectoral worker ownership only helps people who remain employed in covered firms. AI displacement is partly a problem of people *exiting* employment. The architecture handles part of the problem but does not address the displaced-worker case directly.

**Why it nonetheless belongs in the analysis.** Two arguments for inclusion as adjacent:

- It is the cleanest example of an architecture that intervenes *inside* the firm rather than at the transfer layer. This analysis implicitly treats the post-firm transfer layer as the right intervention point; sectoral ownership challenges that assumption. Naming and dismissing it is more honest than ignoring it.

- It connects to citizen-equity through the Bruenig and Varoufakis lineages—both have written about social ownership at multiple scales. Excluding sectoral ownership entirely cuts a thread the cited literature explicitly walks down.

**Lean.** Sectoral worker ownership is structurally adjacent to the citizen-equity SWF variant (both involve worker or citizen ownership claims on capital) but operationally and politically more difficult. The Meidner Plan's specific failure (mandatory dilution of existing equity) is the same failure mode flagged in Section 3.4 for the maximalist SWF designs. This analysis treats sectoral ownership as a useful analytical reference point—it sharpens the distinction between firm-level and economy-level interventions—but not as a candidate for the displacement window. Confidence: moderate; if the political environment shifts dramatically (a major recession plus organized labor revival), sectoral ownership could re-enter the conversation, but no current indicators point that direction.

# 5. Comparison across criteria.

This section scores each of the six architectures (UBI, UBS, citizen-equity warrants, citizen-equity SWF, Pigouvian automation tax, public ownership) against the seven comparison criteria from the scope. Sectoral worker ownership is included in the table for reference but flagged as adjacent rather than principal.

Scores are qualitative—high, moderate, low, or contested—with reasoning following the table. Confidence levels apply to the overall scoring exercise, not to each cell individually.

## 5.1 Comparison table

| CRITERION | UBI | UBS | CE—WARRANTS | CE—SWF | PIGOUVIAN TAX | PUBLIC OWNERSHIP | SECTORAL (ADJACENT) |
|---|---|---|---|---|---|---|---|
| **Structural correction depth** | Low | Low–moderate | Moderate (success scenario) | High | High (margin) | Highest | Moderate |
| **Clock speed** | High | Low | Highest (if window open) | Low | Low | Lowest | Low |
| **Legislative dependency** | High (annual) | Moderate (entitlement) | Low (autopilot) | Low (formula) | Moderate (rate-setting) | High (ongoing) | High |
| **Funding durability** | Low | Moderate | Contingent | Moderate–high | High (revenue-positive) | Moderate | Low–moderate |
| **Coalitional viability** | Moderate–high | Moderate | Contested | High (cross-camp) | Low | Low | Low |
| **Stress behavior** | Vulnerable (fiscal) | Robust (slow degrade) | Fragile (3 contingencies) | Robust (diversified) | Robust (revenue-side) | Robust (post-install) | Mixed |
| **Automation-incentive effect** | None (or perverse) | Slight negative | None | None | Strong negative | Strong negative | Moderate negative |

## 5.2 Reading the table

No architecture dominates across all criteria. This is the first non-obvious finding from the comparison: the public conversation often treats one architecture or another as *the answer*, but each architecture wins on different criteria and loses on others, and the patterns of strength and weakness are highly architecture-specific.

Four patterns are worth surfacing explicitly:

**Pattern 1—Speed and depth trade off.** The architecture with the highest clock speed (warrants) has the most fragile structural correction. The architecture with the deepest structural correction (public ownership) has the slowest clock. The Pigouvian tax is the only candidate that combines high diagnostic depth with revenue-positive funding, but its political clock speed is among the slowest. There is no candidate that combines high speed and high depth—the trade-off appears structural, not contingent.

**Pattern 2—Funding durability divides the candidates.** Pigouvian tax (revenue-generating), SWF (asset-backed), UBS (entitlement-protected) cluster on the durable side. UBI (annual appropriation), warrants (contingent on AI valuations), and public ownership (ongoing political contestation) cluster on the contestable side. This pattern is not aligned with structural correction depth—durability and depth are independent dimensions, and the brief should not collapse them.

**Pattern 3—Automation-incentive effect is the most under-discussed criterion.** Of the seven criteria, this one has received the least attention in public discussion of these architectures. Yet on the Hemenway Falk diagnosis, it is the criterion that determines whether an architecture corrects the underlying market failure or merely cushions its effects. UBI scores zero or worse here; the citizen-equity variants score zero; only the Pigouvian tax and public ownership directly disincentivize the per-task automation that generates the externality. If the Hemenway Falk diagnosis is correct, this row of the table is dispositive—and most current public proposals are aimed at architectures that score zero on the row.

**Pattern 4—Coalitional viability is anti-correlated with structural correction depth.** The architectures most likely to assemble a winning coalition (UBI, warrants in the right political moment) are not the architectures with the deepest structural correction. The architectures with the deepest correction (Pigouvian, public ownership) face the highest organized opposition. This is the political asymmetry that Section 1 named as the central difficulty: the displacement window favors what is fast and shallow over what is durable and deep.

## 5.3 What the table does not capture

Three important properties of the architectures sit outside the seven-criterion frame:

- *Hybridization.* Several pairs of architectures are partially compatible. UBI funded by a Pigouvian automation tax would combine the cushioning of the first with the marginal correction of the second. SWF dividends paid alongside UBS would address both income and cost-of-living exposure. The brief's analysis is principally on each architecture in isolation; hybridization opens design space the comparison does not yet explore. Section 9 (open questions) returns to this.

- *Sequencing.* This analysis treats each architecture as a candidate for installation now. In practice, multiple architectures can be installed in sequence: a fast cushioning mechanism first (UBI or warrants), followed by deeper structural mechanisms (SWF, Pigouvian tax) as political conditions allow. Sequencing changes the clock-speed analysis materially.
- *Subnational installation.* Several architectures can be installed at state or municipal level (UBI pilots, public-option AI utilities, automation taxes on specific sectors). This analysis is centered on federal-scale architectures because the displacement problem is national in scope; subnational pilots inform the operational record but do not substitute for federal installation.

These properties are flagged for the matrix in Section 6 and revisited in Section 9.

# 6. The convergence/viability matrix.

The matrix is the analytical centerpiece. It plots each architecture on two axes—degree of cross-ideological convergence (Meru output) and degree of viability (Hari output, weighted toward stress behavior, clock speed, and coalitional viability)—and shows what the four resulting quadrants mean.

The matrix is a 2×2 with gradations. Architectures are placed by qualitative judgment based on the comparison in Section 5 and the convergence evidence in Section 2. Exact placements are uncertain and this analysis is explicit about the gradations.

## 6.1 The matrix

↑ HIGH VIABILITY

← LOW CONVERGENCE HIGH CONVERGENCE →

| | |
|---|---|
| LOW CONV. / HIGH VIABILITY<br>**Pigouvian Automation Tax**<br>*Clears the diagnostic and viability tests; lacks cross-camp coalition.* | HIGH CONV. / HIGH VIABILITY<br>**Citizen-Equity SWF**<br>UBI *(boundary)*<br>*Strongest credibility position. SWF clearest; UBI on the boundary with moderate convergence and viability.* |
| LOW CONV. / LOW VIABILITY<br>**Public Ownership**<br>**UBS**<br>Sectoral Worker Ownership *(adjacent)*<br>*Present in the literature; weak on both axes in the US frame.* | HIGH CONV. / LOW VIABILITY<br>**Citizen-Equity Warrants**<br>*Cross-camp framings exist; viability constrained by installation window, issuer survival, and warrant claim under distress.* |

↓ LOW VIABILITY

## 6.2 Reading the quadrants

**High convergence / high viability—the strongest credibility position.** The citizen-equity SWF variant occupies this quadrant most clearly. Multiple ideologically incompatible actors have arrived at the SWF mechanism (Bruenig left-US, Varoufakis left-Europe, Trump administration right-US institutional, Norwegian and Alaskan operational analogues), and the architecture clears the seven viability criteria reasonably well—particularly on funding durability, stress behavior, and coalitional viability. The combination is what makes SWF the most credible candidate, even though it is not the fastest or the deepest. UBI sits closer to the boundary of this quadrant: moderate cross-camp interest (some right-coded support exists, particularly from libertarian quarters; left-coded support is the dominant strain), moderate viability scores. The brief's net assessment is that SWF is the architecture currently best-positioned by both signals.

**Low convergence / high viability—could work but only one camp wants it.** The Pigouvian automation tax sits here. The architecture clears the diagnostic and viability tests cleanly but has minimal cross-ideological support: incumbent industry opposes it, and there is no equivalent right-coded version that has gained traction comparable to the federal-SWF alignment on the citizen-equity side. This quadrant is vulnerable to political reversal; depends on coalitional persistence, which Pigouvian advocates have not assembled in three decades of academic prominence.

**High convergence / low viability—popular across camps but does not survive contact with reality.** The citizen-equity warrants variant sits in this quadrant despite its high clock-speed claim. The convergence is real (Bores has produced the most operationally specific proposal, Musk's UHI rhetoric provides right-coded resonance, the warrants mechanism aligns conceptually with the broader citizen-equity cluster), but the three contingencies established in Section 3.3—installation window, AI issuer survival, warrant claim survival under distress—produce a low viability score under realistic stress scenarios. This quadrant is rhetorically powerful and may force a weaker, more durable variant into existence (which would push the architecture toward SWF).

**Low convergence / low viability—present in the literature but not strong on either dimension.** Public ownership and UBS sit here, for opposite reasons. Public ownership has very deep structural correction but extremely low political viability and no significant cross-ideological convergence—it is essentially a left-coded position that draws rhetorical force from the radical end of the policy space. UBS has moderate viability but slow clock speed and limited cross-camp support in the US specifically; the architecture is more credible in European contexts. Both belong in the brief because their analytical contributions are real (public ownership provides the framing pressure that shapes the other architectures; UBS sharpens the distinction between cash and service approaches), but neither is a strong candidate on both axes simultaneously. Sectoral worker ownership is shown as adjacent.

## 6.3 What the matrix is and is not

The matrix is a *credibility heuristic*, not a ranking. It separates two signals—*people agree this is the answer* and *this is the answer*—that get conflated in current public discussion. An architecture in the high/high quadrant has the strongest case for both political feasibility and structural soundness; that does not establish that it is correct. An architecture in the low/low quadrant has the weakest case on both axes; that does not establish that it is wrong.

The matrix is also a snapshot. Positions can shift: warrants could move from high/low to high/high if the AI capital structure stabilizes and the installation window remains open longer than expected. Pigouvian could move from low/high to high/high if a right-coded version emerges with political traction. Public ownership could move from low/low to low/high if the framing pressure produces an unexpected legislative window. Section 9 (leading indicators) tracks the signals that would justify revising matrix positions.

## 6.4 The matrix's clearest finding

The matrix's clearest finding is that no architecture currently sits unambiguously in the high/high quadrant. Citizen-equity SWF is closest; UBI is at the boundary; the others are weaker on at least one axis. This is consistent with the comparison-table finding that no architecture dominates across all seven criteria. The brief's net read is that the strongest available combination is one in which a fast cushioning architecture (UBI, possibly warrants if the window holds) buys time for a deeper structural architecture (SWF) to mature into operational relevance, with Pigouvian elements layered in to correct the underlying margin. This is hybridization, and Section 9 returns to it as an open design question rather than a finding.

## 6.5 Where the matrix could be wrong

The matrix is built on two judgments—convergence (Meru) and viability (Hari)—that each have known failure modes. Naming them up front is more honest than retreating to ambiguity.

*Convergence judgments could overweight rhetorical resonance.* The SWF placement in the high-convergence quadrant rests on Bruenig (left), Varoufakis (left-Europe), Trump-administration interest (right-institutional), and the Alaska/Norway operational analogues. This is a real cross-camp pattern, but the right-coded interest in SWFs is partially extractive (a vehicle for state-directed investment) rather than redistributive (a dividend mechanism for citizens). If the right-coded variant diverges from the left-coded variant on the dividend question, the convergence is shallower than the matrix shows and SWF moves toward moderate-conv/high-viability.

*Viability judgments are weighted toward stress behavior and may underweight raw clock speed.* Warrants score low on viability primarily because of the three contingencies in Section 3.3. If the installation window is wider than current evidence suggests, or if AI capital structures stabilize sooner than the Hari projection assumes, warrants move toward the high-conv/high-viability quadrant and become competitive with SWF on the principal axis. Section 9's leading indicators are designed to detect these shifts.

*Public ownership could be miscategorized as low/low rather than low-conv/moderate-viability.* The placement reflects current US legislative conditions. In jurisdictions where public-AI infrastructure is being built (UK Common Wealth advocacy; European AI sovereignty initiatives), the viability score is materially higher. The brief's US frame is a methodological choice, not a global claim about the architecture.

The matrix is rough, not precise. Confidence: high on the rank ordering (SWF strongest, warrants high-conv/low-viability, public ownership weakest in the US frame); moderate on the specific quadrant placements; low on the precise position within each quadrant.

# 7. Clock-speed and stress analysis.

This section consolidates and extends the per-architecture stress analysis from Sections 3 and 4. The goal is to make the clock-speed and stress-vulnerability claims directly comparable across architectures, and to surface the specific scenarios under which each architecture fails.

## 7.1 Clock-speed comparison

Clock speed is decomposed into four components: *design* (mechanism specification), *legislation* (passage through political institutions), *funding* (revenue source operational), and *operational* (transfer or correction actually flowing).

| ARCHITECTURE | DESIGN | LEGISLATION | FUNDING | OPERATIONAL | TOTAL TIME-TO-OPERATION |
|---|---|---|---|---|---|
| **UBI** | Mature | Slow | Slow (new tax) | Fast (delivery infrastructure exists) | Years (5–10) |
| **UBS** | Mature | Slow | Slow (new tax) | Slowest (service infrastructure) | Decade+ |
| **CE—Warrants** | Mature (Bores) | Moderate (single legislative event) | Pre-funded | Automatic on trigger | Years to event; then fast |
| **CE—SWF** | Mature (Bruenig) | Slow | Slow (asset accumulation) | Fast once funded | Decade+ to scale |
| **Pigouvian Tax** | Mature (Hemenway Falk) | Slowest | Self-funding | Fast | Years (10+) |
| **Public Ownership** | Less mature | Slowest | Variable | Slowest | Decade+ |

The table reveals an asymmetry that matters: warrants are unique in that the *legislative* clock is moderate (single event), the *funding* clock is zero (warrants are pre-funded by their structure), and the *operational* clock is automatic (payout is contractual). This is the source of the H1 clock-speed claim. The other architectures all have at least one clock component that is slow; warrants compress all of them.

The asymmetry is contingent. If the installation window has closed (Section 3.3, contingency a), warrants no longer have the moderate-legislation property—the legislative clock becomes "impossible." If AI companies fail financially (contingency b) or if warrant claims are renegotiated under distress (contingency c), the operational clock becomes "never." Under any of these conditions, warrants lose the asymmetry and join the rest of the candidates in the slow tier.

## 7.2 Stress scenarios applied across architectures

This analysis applies three stress scenarios uniformly across architectures rather than allowing each architecture to define its own stress test:

**Stress scenario 1—Sustained fiscal pressure.** A multi-year period of compressed federal tax revenue, elevated mandatory spending, and political pressure to reduce discretionary outlays. Comparable to the 2010–2014 austerity period or the late-1970s/early-1980s fiscal squeeze.

*Architecture responses: - UBI*: Annual appropriation creates direct exposure. Benefit levels are reduced or program is suspended. The Alaska PFD 2026 dynamic generalizes. - *UBS*: Service quality and waiting times degrade rather than benefits being cut outright. Slower failure mode. - *Warrants*: Insulated from fiscal pressure since the architecture is contingent on AI valuations, not appropriation. - *SWF*: Insulated from fiscal pressure as the dividend is paid from investment income; political risk at the margin (formula adjustment) but architecture survives. - *Pigouvian Tax*: Revenue-positive; fiscal pressure may *strengthen* the political case for the tax, not weaken it. - *Public Ownership*: The public entity's operational budget is exposed to fiscal pressure; productivity and capacity degrade.

*Net*: Pigouvian, SWF, and warrants are most robust to fiscal pressure. UBI is most exposed. UBS and public ownership occupy the middle.

**Stress scenario 2—AI financial collapse.** A multi-year period of compressed AI sector valuations, capex contraction, and reduced AI deployment. Could be triggered by training-cost compounding, customer monetization shortfalls, or a circular-financing unwind across the hyperscaler/lab/chip-manufacturer ecosystem.

*Architecture responses: - UBI*: Robust. The architecture does not depend on AI capital outcomes. - *UBS*: Robust. Same reason. - *Warrants*: Catastrophic. Warrants become worthless; architecture exists but does not pay out. - *SWF*: Mostly robust. AI sector exposure damages the AI portion of the diversified portfolio but leaves the broader fund intact. Dividend levels reduce modestly. - *Pigouvian Tax*: Largely robust, with reduced revenue scaling with reduced AI deployment. If displacement is also reduced, the architecture is correctly responding to a smaller problem. - *Public Ownership*: Robust to AI sector collapse, since the public entity operates at production cost regardless of private market valuations.

*Net*: Warrants are catastrophically exposed. SWF is robust by design. The other architectures are largely insulated.

**Stress scenario 3—Political reversal.** A change in federal political control that is hostile to the installed architecture and seeks to dismantle, weaken, or de-fund it.

*Architecture responses: - UBI*: High exposure. Annual appropriation makes the program reversible by single-Congress majority. Pandemic stimulus precedent shows the program can be expanded or contracted readily. - *UBS*: Low exposure once installed. Service-delivery infrastructure is hard to dismantle (NHS precedent across multiple Conservative governments). Quality erosion is the more likely failure mode. - *Warrants*: Mixed. The pre-funded structure resists immediate dismantling, but a hostile administration could refuse to exercise warrants when triggered, requiring litigation to enforce. - *SWF*: Low exposure once installed. Asset trust structure resists raiding. Formula adjustment is the more likely failure mode. - *Pigouvian Tax*: High exposure. Tax can be repealed by single-Congress majority. Carbon-tax precedent shows incumbent industry persistence in seeking repeal. - *Public Ownership*: Mixed. Privatization is politically expensive but not impossible. Capability erosion through underfunding is the more likely failure mode.

*Net*: Pigouvian tax and UBI are most exposed to political reversal. SWF and UBS are most robust. Warrants and public ownership are mixed, with delayed or partial failure modes rather than immediate collapse.

## 7.3 Stress profile summary

No architecture is robust across all three stress scenarios. The closest is the citizen-equity SWF variant, which is robust to fiscal pressure (asset-backed dividend), mostly robust to AI financial collapse (diversified portfolio), and robust to political reversal (asset trust structure). This stress profile, combined with the cross-ideological convergence on the SWF mechanism, is the strongest combination identified in the brief.

among the architectures examined, the citizen-equity SWF variant is the best-positioned candidate to be designed, legislated, funded, and operating before the displacement window closes. Pigouvian elements should be layered on top to correct the underlying automation margin. UBI or warrants may serve as fast cushioning to buy time. Public ownership operates principally as framing pressure. This is a position, not a hedge, and the analysis in Sections 5–7 is what supports it. Where the position could be wrong is named in Sections 4.1 (the diagnostic claim could be empirically wrong about the wedge), 6.5 (the matrix placements could be miscalibrated), and 9 (the leading indicators are designed to detect the conditions that would falsify it).

The weakest stress profile is the warrants variant, which fails catastrophically in two of three scenarios (AI financial collapse; political reversal during a triggered payout) and is structurally robust only to fiscal pressure. The clock-speed advantage of warrants is real, but the stress profile is the source of the H1 contingencies and the matrix placement in the high-convergence/low-viability quadrant.

Net read on stress: an architecture that fails in *one* of the three stress scenarios may still be the right choice if the failure scenario is sufficiently unlikely. An architecture that fails in *two* of three scenarios is structurally fragile and should be installed only as part of a portfolio that includes more robust architectures. Warrants fall into the second category. SWF falls into the first.

# 8. The convergence finding, in full.

Section 2 set up convergence as a descriptive observation and held the interpretive claim for synthesis. This section delivers it.

## 8.1 What the convergence is and is not

The descriptive claim is well-evidenced and uncontroversial: actors from incompatible ideological starting points have arrived at structurally similar transfer-architecture proposals during 2024–2026. Bores (Axios, April 2026) on the Democratic side, Musk's UHI rhetoric and the Trump administration's SWF executive order (February 2025) on the right, Bruenig's People's Policy Project SWF design and Varoufakis's Universal Basic Dividend on the left, all converging on capital-share-of-output mechanisms rather than purely income-based transfers. This is not a coincidence and not a manufactured agreement; it is multiple actors responding to the same underlying structural pressure with mechanically similar designs.

The interpretive claim is sharper, and this analysis now states it: *cross-ideological convergence on capital-share architectures is a signal about political viability under pressure, not a signal about correctness.* Convergence tells us which architectures can survive contact with a political system that requires cross-camp tolerance. It does not tell us which architectures correct the underlying market failure (Section 4.1 settles that question separately) or which architectures will produce the deepest structural correction (Section 5 settles that). It tells us which architectures are most likely to be installed.

This distinction matters because the public conversation routinely conflates the two signals. An idea that draws cross-camp interest is treated as if it has been validated on the merits. An idea that draws only one camp's interest is treated as if it has been rejected on the merits. Neither inference holds. Convergence is a coalitional signal; correctness is an analytical signal. This analysis separates them deliberately.

## 8.2 Why capital-share architectures specifically

The convergence is not arbitrary. Capital-share architectures (warrants, SWF, dividend mechanisms tied to AI-driven productivity gains) have three properties that explain why ideologically incompatible actors land on them:

*They do not require new tax authority.* SWF dividends come from investment income on accumulated assets. Warrants come from contractual claims on AI firms. Both avoid the political third rail of new income or wealth taxes. This makes them passable in fiscal environments where UBI, UBS, and Pigouvian taxes face revenue-side opposition.

*They are framed as ownership rather than redistribution.* The right-coded framing is "every American becomes an owner." The left-coded framing is "workers share in the productivity gains they helped create." Both framings describe the same mechanism. This rhetorical bivalence is what permits the cross-camp coalition; it would not exist for a UBI funded by a new tax, where the redistributive frame is unavoidable.

*They map onto existing political traditions on both sides.* The right-coded tradition runs through ownership society rhetoric, sovereign wealth funds as state-as-investor, and the libertarian dividend tradition (Friedman's negative income tax has more in common with citizen-equity than with universal basic services). The left-coded tradition runs through social ownership, worker capital funds, and Meidner-style profit-sharing. The architectures sit at the intersection where these traditions can meet without either side abandoning its frame.

The convergence is therefore not philosophical agreement; it is a coincidence of structural fit. Both camps find a mechanism that solves their version of the problem without requiring concessions to the other camp's frame. This is exactly the property that makes an architecture politically durable.

## 8.3 What the convergence does not establish

The discipline on this point: convergence does not establish that capital-share architectures correct the underlying market failure. The Hemenway Falk diagnostic (Section 4.1) shows that capital-share architectures operate on capital ownership and dividend flow, not on the per-task automation margin. They redistribute the gains from automation; they do not change the marginal incentive that produces over-automation in the first place. A capital-share architecture installed without a Pigouvian element addresses the distributional consequences of displacement but leaves the allocative distortion in place.

This is the most uncomfortable analytical move in the synthesis and the place where it diverges most clearly from the public conversation. The political signals point toward capital-share architectures as most likely to be installed. The structural signals point toward Pigouvian elements as necessary to correct the underlying failure. The convergence and the diagnosis disagree about what the right answer is. This analysis takes both signals seriously and lands on a layered architecture (SWF as the political vehicle, Pigouvian elements layered on top) rather than choosing between them.

## 8.4 The cross-pollination question

Section 2 deferred the cross-pollination hypothesis (does the convergence reflect mutual influence, or independent convergence on the same structural attractor?) until synthesis. The available evidence in 2026 is mixed. The Bores design draws explicitly on the broader citizen-equity literature, which includes Bruenig and Varoufakis. The Trump administration's SWF executive order draws on the OMFIF tradition of sovereign-fund advocacy, which is institutionally adjacent to but not derivative of the left-coded SWF literature. Musk's UHI is functionally separate; it is a basic-income proposal with capital-share rhetoric layered on top.

The read: convergence is partly genuine (the structural pressure is real, and multiple traditions independently identify capital-share mechanisms as fitting it) and partly cross-pollinated (the Bores design and the broader 2026 dividend conversation are clearly drawing on prior left-coded SWF advocacy). The mix matters less than the destination. Whether actors arrived independently or by influence, they have arrived at structurally similar mechanisms, and the political viability signal is the same either way.

The one place where the cross-pollination question matters: if the convergence is mostly cross-pollinated (rather than independent), it is more fragile to political reversal. Cross-pollinated agreements unwind when one side disengages. Independent convergence persists because each side arrived for its own reasons. This analysis treats the SWF convergence as mostly independent (and therefore politically durable) and the warrants convergence as more cross-pollinated (and therefore more fragile). This judgment is moderate-confidence and worth tracking in Section 10.

# 9. Historical analogues.

Every analogue is imperfect. The discipline is to extract the structural lesson about *how the mechanism behaves under stress* without overclaiming the parallel. Four analogues earn space here.

## 9.1 Social Security (1935)

**Mechanism.** Federal payroll-tax-funded universal old-age insurance, passed during the Great Depression as the foundational pillar of the New Deal social safety net.

**Structural lesson.** Crisis legislation can install a transfer architecture in months when the political conditions align. The Social Security Act was transmitted to Congress on January 17, 1935 and passed the Senate on June 19, 1935 with a vote of 77–6—roughly five months from administration proposal to passage. The conditions were extreme: 25% unemployment, a Democratic supermajority, a President with electoral mandate, and visible pre-existing models (German social insurance, state-level pension experiments). When all four conditions are present, the legislative clock collapses from "years" to "months."

**What it tells us about the displacement case.** AI displacement may produce conditions analogous to 1935 if it compounds rapidly enough, and a fast cushioning architecture (UBI most plausibly, possibly warrants) could be installed on a Social Security timeline if those conditions arrive. Net read: the *capacity* exists for fast installation under crisis; what's missing in 2026 is the political-coalition piece. The fiscal-stress signal would need to align with a coherent administration proposal and a passable congressional majority. None of these is currently in motion.

**Where the analogue breaks.** Social Security was funded by a new payroll tax; the contemporary fiscal environment is more constrained, and a new tax of comparable scale would face structural opposition that did not exist in 1935. The 1935 program also had near-universal support among economists across the political spectrum; current displacement architectures do not have comparable academic consensus. The analogue suggests fast installation is *possible* in crisis conditions, not that it is *probable*.

## 9.2 Alaska Permanent Fund (1976/1982)

**Mechanism.** Constitutional amendment establishing a state sovereign wealth fund seeded by oil revenue (1976), with a dividend formula added six years later (1982), paying out a per-capita share of the fund's investment income to every Alaska resident annually.

**Structural lesson.** A two-stage installation. The fund was established by constitutional amendment in 1976 as a savings vehicle without a clear distribution rule. The dividend formula was created by the legislature in 1982—six years after the fund existed. This sequencing matters: the political fight over the *fund* and the political fight over the *dividend formula* were separable. The fund was passable on "saving for the future" rhetoric; the dividend was passable later, once the fund had visible assets.

**What it tells us about the displacement case.** A federal SWF could plausibly be installed in two stages: fund first (passable on diversified-investment-vehicle rhetoric, which the Trump administration's February 2025 executive order demonstrates), dividend formula second (passable later, once the fund has assets and political pressure for distribution exists). This sequencing reduces the legislative load at any single moment and is one of the principal reasons SWF clears the viability bar more easily than other architectures.

**Where the analogue breaks.** Alaska is small (population under 750,000), petroleum-rich, and constitutionally distinctive. The federal-scale equivalent would be substantially harder to fund (no equivalent of dedicated petroleum royalties), substantially harder to constitutionally entrench, and would face distribution-rule fights at much larger scale. The PFD has also been under sustained political pressure in recent years, with formula adjustments and underpayments showing the formula is robust but not invulnerable. The analogue confirms two-stage installation is possible and shows the failure modes (formula erosion under fiscal pressure) the federal version would face.

## 9.3 Norway Government Pension Fund Global (1990)

**Mechanism.** Sovereign wealth fund seeded by petroleum tax revenue, established by Norwegian Parliament in 1990, with no dividend distribution to citizens—accumulation-only, with returns supporting general government expenditure.

**Structural lesson.** A SWF can be politically durable for decades when it is framed as intergenerational saving rather than current redistribution. The Norwegian fund accumulated for 21 years before its first transfer to the government budget (2001), and its political durability rests on the *absence* of a current dividend. Voters who would object to redistribution support the fund because it preserves wealth for their children's generation.

**What it tells us about the displacement case.** This is the analogue most relevant to the right-coded SWF variant. The Trump administration's 2025 executive order is closer in design to the Norwegian model (intergenerational saving, no current dividend) than to the Bruenig model (current dividend distributed to citizens). The convergence between the two SWF variants is real but partial: both camps support the *fund*, but they disagree on whether it should pay a current dividend. This is the disagreement Section 6.5 flagged as a risk to the convergence claim.

**Where the analogue breaks.** Norway's fund is petroleum-funded; the seed capital came from a clearly identifiable resource rent. A US federal SWF would need to identify a comparable seeding mechanism (carried-interest taxation, AI-sector capital gains, federal land monetization, deficit-funded buy-in), each of which is politically more contested. The Norwegian model also assumes a high-trust government that voters allow to manage substantial assets without current return; US trust in federal asset management is materially lower.

## 9.4 Swedish Meidner Plan (1976–1991)

**Mechanism.** Mandatory annual transfer of 20% of large-firm profits to industry-wide worker-controlled funds, accumulating over decades until workers held majority stakes. Implemented in significantly weakened form (1984), abandoned (1991).

**Structural lesson.** Mandatory dilution of existing equity is the failure mode that distinguishes Meidner-style architectures from tax-funded SWF designs. The Meidner Plan failed politically because its mechanism—transferring shares away from existing owners—created a concentrated, well-resourced opposition that mobilized in unison (PDXScholar working paper on Swedish Social Democracy; Wikipedia: Employee funds). Tax-funded SWF designs (Bruenig) avoid this failure mode by purchasing assets at market prices rather than confiscating equity.

**What it tells us about the displacement case.** The cautionary value is real: maximalist designs that confiscate equity rather than purchase it produce coordinated opposition and tend to fail. The contemporary citizen-equity literature has largely absorbed this lesson—Bruenig, Varoufakis, and Bores all design around purchase or contractual claim rather than confiscation. The architectures that survive in 2026 are descendants that learned from Meidner's specific failure mode.

**Where the analogue breaks.** Sweden in 1976 was a different political environment from the contemporary US. Sweden had near-universal labor-union density, a dominant social democratic party, and a corporatist negotiation tradition that made even confiscatory proposals seem within the political mainstream. The contemporary US has none of these conditions. The Meidner Plan's failure tells us about the *mechanism* (mandatory dilution attracts coordinated opposition), not about a parallel political environment.

## 9.5 What the analogues collectively suggest

Reading the four analogues together: fast installation is possible under acute crisis conditions (Social Security), two-stage installation reduces legislative load (Alaska), durability is highest when the architecture is framed as saving rather than redistribution (Norway), and confiscatory mechanisms attract coordinated opposition that defeats them (Meidner). The composite picture: a federal SWF installed in two stages—fund first on intergenerational-saving rhetoric, dividend formula later under displacement-pressure conditions—has the strongest historical precedent of any candidate in the comparison. This is the principal historical-analogue support for the net read in Section 7.3.

Where the analogue composite could mislead: the four cases span different political environments, different scales, and different funding sources. Treating them as a single "lesson" risks smoothing over conditions that may not hold in the contemporary US. This analysis takes the analogues as suggestive structural evidence rather than as definitive guidance.

# 10. Leading indicators.

The matrix and stress analysis are snapshots. This section names the specific signals that would justify revising the assessment, with explicit thresholds where possible. The discipline: indicators must be *falsifiable*—naming a date and a condition under which the brief's read shifts.

## 10.1 Indicators that would strengthen the SWF lean

*Federal SWF legislation moves past committee in either chamber by end of 2026.* The Trump administration's February 2025 executive order created the executive scaffolding. Legislative codification is the next stage. If a federal SWF bill clears committee in either chamber by December 2026—even without floor passage—the architecture's clock-speed score moves up materially.

*Cross-camp endorsements of a specific dividend formula.* The convergence is currently on the *fund*, not on the *distribution rule*. If Bruenig-aligned and right-institutional advocates publicly endorse the same dividend formula (or formulas within close range) by mid-2027, the convergence is genuine rather than rhetorical, and the matrix placement strengthens.

*Visible operational learning from Alaska PFD pressure tests.* The 2026 Alaska legislative session produced the $1,500 PFD that compromised between fiscal pressure and dividend obligation. If the Alaska formula proves stable through additional fiscal-stress cycles (next two budget cycles), the federal-scale viability case strengthens.

## 10.2 Indicators that would weaken the SWF lean

*The Trump administration's SWF stalls without legislative momentum by mid-2027.* If the executive order produces no legislation, no dedicated revenue stream, and no operational fund within 18 months of issuance, the political-viability signal weakens substantially. The Norwegian fund took years to accumulate but had a clear seeding mechanism from the start. A federal SWF without a seeding mechanism by mid-2027 signals the political coalition is rhetorical rather than operational.

*Right-coded SWF advocates publicly reject the dividend mechanism.* If the right-institutional camp (think tanks, administration-aligned voices) explicitly rejects current-dividend distribution in favor of accumulation-only models, the convergence with the Bruenig/Varoufakis tradition collapses. This analysis currently treats the disagreement as latent; if it becomes explicit, SWF moves from high-conv/high-viability toward moderate-conv/high-viability.

*Bruenig-aligned advocates publicly reject the right-coded SWF as captured.* Symmetric to the above. If the left-coded SWF tradition disowns the federal SWF as a vehicle for state-directed investment rather than citizen distribution, the cross-camp coalition unwinds.

## 10.3 Indicators that would strengthen the warrants lean

*Bores wins his House race and uses the seat to advance the AI dividend.* Bores is a Democratic House candidate facing AI super PAC opposition. If he wins in November 2026 and introduces warrants legislation in the 120th Congress, the architecture's legislative clock moves materially.

*Co-sponsorship across party lines on the warrants design.* Bores's design has sufficient libertarian-friendly framing (every American becomes an AI owner) that cross-aisle co-sponsorship is conceivable. If the bill attracts even a small number of Republican co-sponsors, the convergence claim for warrants strengthens.

*AI capital structures stabilize.* Section 3.3 named three contingencies for warrants. If AI sector valuations stabilize (no major collapse, no extreme acceleration), the warrant-claim-survival contingency loosens. The relevant signal is multi-quarter pricing stability in the major AI public companies and continued private capital availability for frontier labs.

## 10.4 Indicators that would weaken the warrants lean

*A major AI hyperscaler experiences a capital-structure crisis.* If any frontier AI company or major hyperscaler enters distress (Chapter 11, forced restructuring, government bailout) during 2026–2028, the warrants architecture's robustness is tested in real time. Warrants on companies that fail or are renegotiated under distress lose value, which is the second contingency in Section 3.3.

*The installation window closes politically.* If the 120th Congress (January 2027) shows no movement on warrants legislation, and no hearings are held, the legislative clock-speed assumption that drove warrants' high-convergence/low-viability placement weakens further. The architecture becomes high-convergence/very-low-viability and effectively drops out of the candidate set.

## 10.5 Indicators that would strengthen the Pigouvian lean

*A jurisdiction installs a pure Pigouvian automation tax.* No US state or federal jurisdiction has installed an additive automation tax. If any does (California, New York, or a federal pilot), the architecture moves from academically-strong/politically-nascent to politically-active. Even sub-federal installation provides the operational record that has been missing.

*Carbon-tax-style international coordination begins.* The Hemenway Falk diagnosis implies an automation tax requires multilateral coordination to handle offshoring (Section 4.1). If OECD or G7 discussions on automation taxation emerge in 2026–2028, the structural clock-speed problem (offshoring) is plausibly addressable.

*Empirical confirmation of the wedge magnitude.* As noted earlier, the wedge's *magnitude* is empirically contested. If subsequent empirical work (using 2024–2026 displacement data) confirms over-automation at meaningful magnitudes, the diagnostic claim moves from theoretical to empirically grounded. This is the slowest indicator but the most foundational.

## 10.6 Indicators that would weaken the entire framework

*Labor markets reallocate displaced workers fast enough that displacement does not compound.* The April 2026 Hari projection assumes structural displacement on multi-year timelines. If the Conference Board Leading Economic Index and related signals show robust labor-market reallocation through 2026–2028—displaced professional-managerial workers finding new employment at comparable wages within reasonable timeframes—the projection's structural arc weakens, and the urgency for a transfer architecture weakens with it. This would not falsify the analytical framework but would reduce its policy relevance.

*AI capabilities plateau.* If frontier AI capability gains slow materially during 2026–2028 (no compounding from current frontier through agentic systems through full economic-task automation), the displacement projection's tail risk reduces and the displacement window may close on its own without a transfer architecture being needed.

*A transfer architecture is installed that is not in the candidate set.* The candidate set is bounded by the literature surveyed. If a fundamentally different architecture emerges (something not derivative of UBI/UBS/citizen-equity/Pigouvian/public-ownership families), the analysis would need to reincorporate it. This is the indicator most likely to be unexpected.

## 10.7 The most important single signal to watch

If if forced to name one signal: *whether the Trump administration's SWF gets a dedicated seeding mechanism through legislation in 2026–2027.* This is the single highest-leverage indicator because it tests (a) whether the right-coded SWF interest is operational rather than rhetorical, (b) whether cross-camp legislative alignment on a SWF is possible, and (c) whether the federal-scale two-stage installation pattern (fund first, dividend later) can begin. If the seeding mechanism passes, the SWF lean strengthens substantially. If 18 months pass without it, the political-viability score for SWF should be revised down materially.

# 11. Open questions.

The leans here are positions, not closed cases. Six open questions remain genuinely open and would benefit from additional analysis, empirical evidence, or political development. They are listed without attempt at resolution.

## 11.1 Hybridization design

This analysis lands on a layered architecture (SWF as political vehicle, Pigouvian on top, UBI/warrants as fast cushioning) without specifying how the layers integrate operationally. Hard design questions: should the Pigouvian tax revenue seed the SWF, fund retraining, return as dividends, or some mix? Should warrants be designed as an interim mechanism that converts to SWF shares at scale? Is there a clean way to combine Bores-style warrants with Bruenig-style SWF without producing two competing claim structures? This analysis flags hybridization as the strongest available combination and treats the design specifics as open.

## 11.2 Sequencing

The Alaska analogue suggests two-stage installation (fund first, dividend later). Other sequencing patterns are also plausible: warrants first (fast, but contingent), then SWF (slow, but durable), with Pigouvian elements layered in once political conditions allow. This analysis does not adjudicate between sequencing patterns. The choice depends on which architectures pass first under the prevailing political conditions of any given Congress.

## 11.3 Subnational pilots and federalism

Several architectures can be installed at state level (UBI pilots, public-option AI utilities, automation taxes). The relationship between subnational pilots and federal installation is unclear: do pilots build the operational record that enables federal scaling (the Wisconsin unemployment insurance precedent for Social Security), or do they create a fragmented patchwork that resists federal harmonization? This analysis is centered on federal architectures because the displacement problem is national, but subnational installation could shift the analysis.

## 11.4 International coordination

The Pigouvian architecture's offshoring problem (Section 4.1) is solvable only through multilateral agreement or border adjustment. Neither is currently in motion. The same problem applies in weaker form to SWF (capital flight to lower-tax jurisdictions if seeded by capital-gains taxation) and to warrants (AI firms relocating headquarters to avoid warrant obligations). The international coordination question is open and may determine which architectures are operationally viable at scale.

## 11.5 The civic-claim variant of citizen-equity

The scope considered Section 3.3 (warrants) and 3.4 (SWF) as the two principal citizen-equity variants. A third variant exists in the literature: *civic-claim* mechanisms where citizens have legally enforceable rights to AI-derived productivity gains without an underlying asset trust. The brief treated this as adjacent because no major proposal in 2026 is using this design. If it emerges (the closest contemporary precedent is the discussion around AI-training-data compensation), it could change the citizen-equity comparison.

## 11.6 The displacement projection itself

This analysis builds on the April 2026 Hari projection and treats its structural arc as the operative scenario. The projection's confidence intervals widen substantially beyond 2028. The transfer-architecture analysis is most reliable in the 2026–2028 window and progressively less reliable beyond it. Subsequent projections may revise the structural arc, which would propagate through the architecture analysis.

# 12. Methodological appendix.

## 12.1 The April-projection baseline and why it holds

This analysis builds from the April 2026 Hari projection rather than re-baselining to weekly or monthly data. The justification draws on Nassim Nicholas Taleb's signal-to-noise framing, articulated in *Antifragile* and *Fooled by Randomness*: at short observation timescales, the ratio of signal to noise approaches zero; at longer timescales, signal becomes recoverable.

Taleb's specific framing (*Fooled by Randomness*, the noise-bottleneck chapter of *Antifragile*) uses an illustrative numerical heuristic: at the scale of seconds, noise dominates almost completely; at the scale of a year, the signal-to-noise ratio improves to roughly 50/50; at the scale of a decade, signal dominates. The numbers are not empirical claims but illustrative—Taleb's argument is that observers who track variance at high frequencies systematically over-react to noise and miss structural change, while observers who track variance at the scale appropriate to the underlying process recover the structural signal.

The transfer-architecture decision is structural and multi-year. A monthly re-baselining would over-fit to noise and produce architectural recommendations that whipsaw with each labor-market data release. The April projection captures the structural arc at the resolution at which the architectural decision is useful. Subsequent monthly data is tracked as confirmation or disconfirmation of the arc, not as independent input to the analysis.

The Taleb framing is invoked once in this analysis (here) and is applied as a methodological principle, not as a primary analytical framework. The principal methods are causal-chain projection (Hari v3) and consensus archaeology (Meru), both of which are described elsewhere in the source material.

## 12.2 Confidence calibration

This analysis uses three confidence markers: *high*, *moderate*, and *exploratory*. They map roughly as follows:

- *High*: claim is supported by multiple independent sources, structural reasoning, and would require substantial new evidence to overturn. Examples: the Hemenway Falk wedge exists in direction; capital-share architectures exhibit cross-ideological convergence; warrants face three named contingencies under stress.
- *Moderate*: claim is supported by available evidence and reasoning but rests on judgments that could be revised with new evidence. Examples: SWF is the best-positioned candidate; the convergence is partly genuine rather than purely cross-pollinated; the magnitude of the Hemenway Falk wedge is plausibly material.
- *Exploratory*: claim is logically coherent and worth considering but rests on conditions not yet observed. Examples: the layered hybridization is operationally implementable; subnational pilots can scale to federal installation; cross-camp legislative alignment on a specific dividend formula will materialize.

This analysis uses these markers throughout rather than collapsing into uniform hedging. Claims with high confidence are stated directly; claims with moderate confidence carry the qualifier; claims with exploratory status are flagged as open.

## 12.3 Source disclosure

This analysis draws on policy proposals (Bores, Bruenig, Varoufakis, Exley, Hemenway Falk and Tsoukalas), institutional documents (Trump administration executive order, Norwegian and Alaskan fund histories), academic literature (Acemoglu and Restrepo, Guerreiro et al., Hemenway Falk and Tsoukalas, Westover), think-tank publications (People's Policy Project, Mercatus, Roosevelt Institute, Windfall Trust), and journalistic coverage (Axios, Politico, Klein interview). Each citation appears inline at the point of use.

This analysis does not draw on confidential sources or unattributed claims. Where this analysis makes a structural inference (e.g., the SWF as best-positioned candidate), the inference is stated as the brief's own judgment based on the cited evidence, not attributed to any of the cited authors.

# 13. Financial-stress sensitivity appendix.

The stress analysis in Section 7 applied three scenarios uniformly. This appendix explores the *AI financial collapse* scenario in greater depth because it is the scenario most specific to AI displacement and least covered in the broader transfer-architecture literature.

## 13.1 Scenario specification

AI financial collapse refers to a multi-year period in which AI-sector valuations contract substantially, capex slows, and AI deployment slows or reverses. The scenario is not specified by a single trigger; plausible triggers include training-cost compounding outrunning monetization, customer-acquisition shortfalls at the major labs, circular-financing unwinds across the hyperscaler-lab-chip ecosystem, regulatory restrictions on AI deployment, or a broader tech-sector contraction analogous to the 2000–2002 dotcom unwind.

The scenario does not specify whether AI capabilities continue to improve (collapse is purely financial) or whether capability gains also slow (collapse is financial and technical). The architecture-level effects differ between the two sub-scenarios.

## 13.2 Architecture behavior in pure financial collapse (capability gains continue)

If AI capabilities continue to improve while financial conditions deteriorate, displacement continues but the AI-sector capital base does not. This is the most damaging scenario for warrants: the architecture is contingent on AI firm valuations, which are contracting precisely when the displacement they were designed to address is intensifying. Warrants fail at the moment they are most needed.

SWF behavior in this sub-scenario is mostly robust. The fund's exposure to AI-sector assets contracts in line with the sector but the broader diversified portfolio insulates the dividend. Dividend levels reduce modestly rather than collapsing.

UBI and UBS are insulated because they do not depend on AI capital outcomes. Their funding sources (general taxation) may face fiscal pressure if the AI-sector contraction reduces corporate tax receipts, but the architecture-level mechanism survives.

Pigouvian taxation continues to function: the levy is on automation events, and continued AI deployment despite financial distress means the levy continues to apply. Revenue may reduce if deployment slows, but the architecture is correctly responding to a smaller problem rather than failing.

Public ownership operates at production cost regardless of private valuations and is therefore robust to this sub-scenario by design.

### 13.3 Architecture behavior in combined financial and technical collapse

If AI capabilities also plateau, displacement intensity reduces alongside the financial contraction. In this sub-scenario, the urgency for a transfer architecture decreases, but the architectures already installed face different challenges.

Warrants still fail because the underlying capital base is contracting. UBI and UBS continue to function but face general fiscal pressure from a tech-sector recession. SWF dividends reduce in proportion to portfolio performance. Pigouvian revenue contracts but so does the displacement it was designed to correct. Public ownership is robust.

Net read on this sub-scenario: it reduces the urgency of the displacement problem and creates an asymmetric architecture-test in which warrants fail catastrophically while other architectures degrade gracefully. The right policy response is to slow the warrants installation and accelerate the SWF and Pigouvian elements, on the grounds that warrants are fragile to a scenario the others survive.

### 13.4 The probability question

This analysis does not assign a probability to AI financial collapse during the displacement window. The analysis is conditional: *if* the scenario occurs, *then* the architecture-level behavior is as described. Probability assessment requires inputs (capital-structure dynamics, monetization timelines, capability progression rates) that are tracked elsewhere in the Hari projection but are not duplicated here.

Structural read: AI financial conditions are a non-trivial risk during the 2026–2028 window, and any architecture installed during this window should be evaluated for behavior under collapse. Architectures that fail catastrophically under this scenario (warrants) should not be the principal pillar of a transfer-architecture portfolio. Architectures that degrade gracefully (SWF, Pigouvian, UBI/UBS) are preferable as principal pillars even when they have lower clock-speed scores.