

Table Of Contents

Contents

1. Headline results	3
2. What the audit passes changed	3
3. The problem this solves	4
What actually happens in 2034?	4
The cash burden through exhaustion (\$bn nominal)	4
Why does the plan hold the special issues rather than cancel them?	5
4. Benefits and the contributory link	5
5. Distribution — who pays and who gains	5
The largest progressive mechanism is not the rate	6
What a middle-class retiree faces	7
Is this the flat tax the right has been asking for?	7
6. Single taxation — the flow rule	7
7. The Senior Income Supplement — the most heavily revised element	8
What the material error was, and what it cost	8
v62: the floor is flat per person, and there is no household test at all	9
Why no household test at all	9
Why 200%, and why flat	9
v51: the floor is an old-age floor, and the headline is relabeled	10
Lifted to the floor is not lifted out of poverty	10
v53: which tier of benefit taxation is repealed	10
v56: Medicare sits entirely outside the proposal	11
The floor does not cost anyone their Medicare assistance	11
Interest no longer paid on retired federal debt	11
Stress testing the margin	12
v51: citizens aged 65 and over with no earnings record	12
What the v50 and v51 changes did to the results, and where they landed	13
How income is measured for the floor	13
Valuing imputed rent without valuing every house	13
8. Investment and return assumptions	14
What the return assumption actually changes	14
9. What is deliberately not claimed	15
Effects excluded from every scored figure	15
Baseline choices that make the result look worse	15
Base construction where the lower figure was taken	16
The one assumption that is not conservative	16
The honest headline	16
The rate sits sixteen basis points above a delivery cliff	16
10. Known limitations	17

11. Version history 18

SAVING SOCIAL SECURITY — A PROPOSAL

Methodology, Assumptions and Known Limitations

Companion to model v91, deck v91 and reconciliation v91 · 23 September 2026 · New Frontiers in Public Policy

This document exists so that a reviewer does not have to discover our assumptions by reverse-engineering the model. Every methodological choice that materially affects the result is stated here, together with what would change if the choice were made differently. Where a figure rests on an estimate rather than a measurement, it is labeled as an estimate and the range is given.

Figures are in real 2026 dollars and are modeled against the 2026 OASDI Trustees Report intermediate assumptions. The plan is scored as a delta against the published long-range actuarial deficit of 4.42% of taxable payroll.

1. Headline results

Measure	Result
Levy rate	5.95% on seventeen sources, replacing the 12.4% payroll tax
Actuarial balance — current-law benefits only	+2.56% of taxable payroll
Actuarial balance — after all commitments	+0.74% of taxable payroll
Improvement on current law	+5.16%
Trust fund depletion	Never; reserves (the trust fund's bonds plus 75% of the fund) never fall below one year of benefits
Senior Income Supplement	200% of the guideline, flat per person, all ages 65+. Funded in full in 75 of 75 years
Medicare made whole	Entirely outside the plan; no payment flows in either direction
Publicly-held debt retired by 2100	\$6.6tn real
Trust fund bonds	\$2.1tn held as the reserve at cutover; retired, non-cash, from 2033 to 2086 as the fund replaces them
National fund by 2100	\$16.6tn real; peak ownership 4.0%, average 1.37%

The number to quote externally is +0.74%. The +2.56% figure is the balance before the Senior Income Supplement is charged, and should never be presented on its own.

2. What the audit passes changed

Version 48 is an audit pass across all twenty tabs of the workbook and all four external documents. It is a correction pass: no design decision was revisited and no published figure moved. Coverage is stated rather than asserted, because a claim that a model is clean means nothing unless it says what was examined.

Finding	Effect
Poverty tab C28 divided a count of households by a count of people	The only figure in the workbook that moved. Displayed 16.4% where 30% belonged. Not printed in any external document.
Summary Ledger 4 pointed at households under a label reading 'Recipients'	Now points at the person count. The published figure was always correct; the traceability was not.
Six cells described a 'seven levy' design	The model has seventeen sources. The headline narrative also quoted +2.33% without the caveat its own ledger requires.

Four notes on the poverty tab argued for multiples the model no longer uses	Superseded v43 commentary quoting counts computed on pre-v44 units. Relabeled as history.
The solving-rate sensitivity grid is v27 vintage	RESOLVED AT v78. The grid was retired and re-run live from nine explicit rate runs of the current workbook. Doing so exposed the delivery cliff described in section 9.
Two cells claimed a live link they do not have	The bare-solvency and never-deplete rates are stored records of explicit runs, not formulas. Now labeled as such.
The change log stopped at v45 while the workbook reached v47	v46 and v47 entries are reconstructed on the Changes tab and marked as a reconstruction.

The last two items belong to the same family as the fertility-lever note that survived four versions: commentary asserting a wiring the formulas do not have. One instance is now closed. The sweep is not complete, and a reviewer who finds a second incorrect provenance note is entitled to doubt the rest.

3. The problem this solves

What actually happens in 2034?

This is the single most misunderstood point in the whole debate, and getting it wrong undermines the case rather than strengthening it. Trust fund exhaustion does not trigger a bill to the Treasury. It triggers an automatic benefit cut.

The trust fund is a claim, not a pot of money. There are two distinct phases, and they work in opposite directions. Between now and 2034, outgo exceeds dedicated revenue, so Social Security redeems its special-issue Treasury securities and Treasury borrows from the public to find the cash — roughly \$3.2tn over that window, being \$2.6tn of redeemed securities plus the interest owed along the way. From 2034 onward, the Social Security Act contains no borrowing authority, so outgo is clipped to incoming revenue: a cut of roughly 17% across the board, deepening to about 35% by 2100. At that moment Treasury's obligation to the program goes to zero and the unified deficit actually improves.

The uncomfortable implication is that even the \$3.2tn understates the case. No Congress will vote for a 17% cut to 71 million people, so the realistic path is that the shortfall returns to the general fund indefinitely. The honest measure of what is at stake is not \$2.6tn; it is the 4.42%-of-payroll actuarial deficit in perpetuity.

Any presentation that says 'the Treasury will have to find \$2.1tn in 2034' invites a correction from the first actuary who reads it. The stronger and more accurate framing is that the bill is being paid right now, at \$271bn this year and rising past \$460bn by 2034, and that what happens in 2034 is that it converts into a permanent cut.

The cash burden through exhaustion (\$bn nominal)

Year	Cost	Tax revenue	Treasury finds	of which interest	Redeemed	Fund left
2026	1,696	1,425	271	102	169	2,392
2027	1,777	1,489	288	97	191	2,201
2028	1,867	1,556	311	90	221	1,980
2029	1,962	1,626	336	82	254	1,726
2030	2,056	1,700	356	72	283	1,443
2031	2,153	1,777	376	61	314	1,129
2032	2,261	1,857	404	49	355	773
2033	2,378	1,942	436	34	402	371
2034	2,498	2,030	468	16	371	0

Cumulatively \$3,246bn nominal, or \$2,926bn in real 2026 dollars. Two features are worth drawing out. The cash gap and the redemption are different numbers: Treasury funds both the interest it owes the fund and the

securities it redeems, and as the balance shrinks the interest component collapses from \$102bn to \$16bn, so redemptions accelerate faster than the underlying shortfall. And the cumulative burden is larger than the \$2.6tn usually quoted, because that figure is the trust fund balance rather than the Treasury's cash requirement.

Why does the plan hold the special issues rather than cancel them?

Through v90 the plan cancelled them: the claim was written off at cutover and Treasury paid a \$600bn settlement. From v91 the bonds are held. New revenue funds benefits from the first year, so they are rarely called; they stay on the trust fund's books as its reserve, keep earning the special-issue rate, and pay any shortfall before anything else — \$0.2tn toward the Senior Income Supplement in 2027–32, before the fund exists. Each year the bonds not needed to hold one year of benefits in reserve, counting 75% of the fund's value, are retired as a non-cash entry, but only as far as the fund has already replaced them. Retirement starts in 2033, the holding peaks at \$2.4tn in real terms in 2048, and the last bond is retired in 2086. No surplus is spent redeeming internal debt, so every debt-retirement dollar still goes to publicly-held debt; paydown begins in 2033, once the swept surplus exceeds the Supplement.

Holding rather than cancelling costs nothing in cash, keeps at least a year of benefits in reserve from the first day — which the Trustees' short-range adequacy test requires and a write-off would fail — and keeps the protection beneficiaries have in a debt-limit standoff, which depends on the trust fund holding bonds it can redeem. Retirement is still an accounting event and is presented as one: it reduces gross federal debt and leaves debt held by the public unchanged. Under the Office of the Chief Actuary's convention a retirement is a transfer out of the trust fund. Its present value is \$1,908bn, 0.29% of payroll, and the model shows it as a memo (Summary row 70) rather than charging it, because the same convention would also count the investment fund as a trust fund asset, which the model does not.

4. Benefits and the contributory link

'Benefits as defined by current law' is the right policy answer but not a sufficient specification. Three questions have to be answered explicitly, because each changes the score. Income above the current taxable maximum does not earn benefit credit. Non-wage income — capital gains, dividends, business cash flow — earns no credit at all. W-2 wage reporting to SSA continues unchanged. So the primary insurance amount formula, the bend points, wage indexing, the COLA and insured-status rules are all untouched. The taxable maximum survives, not as a revenue parameter but as a benefit-computation parameter.

Raising the maximum benefit for those now paying on uncapped income was considered and rejected. The objection is that it is purely regressive in effect: a broad flat levy falling on the entire economy would fund an upward transfer to people earning above \$192,252. The estimated cost was a quarter to a half point of payroll, taken out of the surplus that funds the Senior Income Supplement and debt retirement.

Three technical corrections worth carrying: the benefit formula uses the best 35 years, not 40; each year is capped at that year's taxable maximum, not the current one; and the maximum benefit varies substantially by claiming age, exceeding \$5,000 at 70 and falling below \$3,000 at 62. 'The maximum benefit' is not a single number.

Current-law benefits are unchanged, so nobody receives less than the existing formula provides. But the framing needs care. If the proposal states its purpose as ensuring people receive benefits 'from what they paid into the system', that contradicts the design — the whole point is that capital, inheritance and business cash flow now contribute without earning benefit credit. The defensible framing is that the contributory link governs what people receive, while the funding base is broadened to whatever the economy can bear.

The Senior Income Supplement is treated as a Social Security benefit, and that is the stricter choice. Section 201 of the Social Security Act dedicates trust fund money strictly to benefits and administration. Either all levy revenue flows into OASDI — in which case poverty and Medicare spending are scored as OASDI expenditures, which is what the model does — or revenue is split at source into separate funds, in which case the OASDI-only balance becomes the headline. What is not available is claiming the higher balance while also spending the surplus elsewhere.

5. Distribution — who pays and who gains

The plan levies one flat rate with no brackets, no thresholds and no carve-outs. That structure reads as regressive on its face and is not. The progressivity is real but sits in two places a reader does not immediately look: the concentration of the bases, and the cap on what contributions buy.

Base	Revenue \$bn	Top-decile share	\$bn from top decile
Personal income — uncapped	805	42%	338
Capital income — incl. retirement accounts	274	85%	233
Business cash flow	587	70%	411
Inheritances, death benefits and charitable transfers at source	158	88%	139
Foreign holders of US-source income	64	not US households	0
Employer benefits ex-health	42	35%	15
Net rent	16	80%	13
Non-profit gross operating surplus	16	not US households	0
US persons working abroad	7	70%	5
Government enterprises	7	not US households	0
Nonresident-alien student wages	1	not US households	0
Residential property disposal gains	11	60%	6
Insurance and annuity inside buildup	7	not allocated	—
Section 1031 rollover gains	3	not allocated	—
Lottery prizes and gambling winnings	8	not allocated	—
Irrevocable trust retained gains	4	not allocated	—
Royalties to individuals	3	not allocated	—
TOTAL	2,014		1,160

Roughly 58% of all revenue is borne by the top decile, before any share of Sources 13–17 is allocated, so the figure is a floor. The payroll tax it replaces is capped, and therefore falls hardest as a share of income on the middle.

The business cash flow line is the contestable one. It is 29% of revenue and the estimate assumes 70% of its incidence falls on owners. Economists divide on how corporate-level taxation splits between owners, workers and consumers; at a 50/50 split the top-decile share falls from 58% to about 52%. Still progressive, but expect this to be the point of attack.

The largest progressive mechanism is not the rate

It is that contributions are uncapped while benefits are not. Every dollar above the taxable maximum is levied and earns no benefit credit whatsoever, so the entire contribution above that line is a transfer.

Wage income	Levy paid	Earns credit on	Earning nothing
\$184,500	\$10,978	\$184,500	0%
\$250,000	\$14,875	\$184,500	26%
\$500,000	\$29,750	\$184,500	63%
\$1,000,000	\$59,500	\$184,500	82%
\$5,000,000	\$297,500	\$184,500	96%

Over a 35-year career a \$1m earner contributes approximately \$1.7m that generates no additional benefit. The principle extends further: capital income, business cash flow and inheritance earn no benefit credit at any level, which means 51% of total revenue is contribution against which no claim can ever be made, and 60% once every other non-wage source is added. This is a direct consequence of holding the benefit ceiling at current law, and it is the substantive reason to decline raising the maximum benefit rather than merely a cost argument.

Flat rate on the way in, capped return on the way out. The tax side is clean enough to satisfy a flat-rate advocate; the redistribution happens on the benefit side.

What a middle-class retiree faces

Take a retired couple with a \$48,000 benefit and \$48,000 of investment income. The benefit is entirely outside the tax system. The investment income carries the 5.95% levy, \$2,856, which is 3.0% of their total income; adding the federal income tax that still applies to the investment portion brings the all-in federal burden to roughly \$4,336, or 4.5%. Half of a middle-class retiree's income becomes permanently invisible to the tax system, and the tax torpedo disappears — today each additional dollar earned can drag benefit income into taxation at effective marginal rates above 50%.

What should NOT be claimed is a large cash windfall for this household. Against current law the same couple pays about 4.3% today, because the levy on their investment income roughly offsets the value of the benefit exemption. The honest argument is about the structure and the marginal rate, not the level of their disposable income.

The tax-free provision does not help the poorest recipients, and the plan is better described honestly on this point. Benefit taxation only applies above \$25,000 for single filers and \$32,000 for couples, so the bottom two deciles already pay nothing and gain nothing from repeal. Those 11.3 million people are reached by the Senior Income Supplement instead. The two instruments are complementary rather than overlapping: the floor reaches the bottom two deciles and nobody else; the exemption reaches deciles three through ten and nobody else.

Is this the flat tax the right has been asking for?

In rate structure, yes: one rate, no brackets, no deductions, no carve-outs, nothing to plan around. Expect the counter to be about the base rather than the rate. Most flat-tax proposals exempt or preferentially treat investment returns; this one reaches capital harder than the current system does, because capital currently contributes nothing to Social Security at all. The argument to have ready is that this is the removal of an exemption rather than the imposition of a new tax.

6. Single taxation — the flow rule

Is a dollar levied twice if it is earned, saved and then bequeathed? Yes, and that is the rule. 'Taxed once and only once' means each FLOW is levied once as it moves; it does not mean a dollar is levied once in its lifetime. A dollar earned is levied as wages, and if it is later bequeathed the bequest is a new flow and is levied again. This is how the present system already works — income tax, then estate tax.

The distinction matters because the alternative is not available. A wealth-based rule that never touches already-taxed money would exempt most of the inheritance base, since the majority of a bequest is accumulated savings out of taxed income and only the step-up-basis appreciation is genuinely untouched. Base 2 would collapse from \$158bn to roughly \$60-70bn. The two principles cannot both be claimed, and the flow rule is the one the model implements.

What the levy never does is reach the SAME flow twice. Three corrections were required and all three are in the model. Outbound — a US company paying a foreign holder — is levied in Source 8, withheld at source under IRC 1441, and the same dividend is therefore deducted from Base 3 (\$374bn). Inbound — a foreign company paying a US person — has no US withholding agent, so it is reached at the recipient in Base 4 (\$372bn), and domestic dividends are excluded from Base 4 entirely. Corporate net interest is deducted from Base 3 (\$44bn) because it is levied at the recipient.

Flow	Payor	Recipient	Where levied	Mechanism
Outbound	US company	Foreign holder	Source 8 (\$1,080bn)	Withheld at source, IRC 1441
Inbound	Foreign company	US person	Base 4 (\$372bn)	Levied at the recipient
Domestic	US company	US person	Base 3 only	Levied at the paying company

The asymmetry is not a design choice. The US can withhold at source only where the US IS the source; it has no jurisdiction over a foreign payor, so inbound income must be reached at the recipient.

A fourth correction, made in v37, removed \$1,070bn of imputed financial services (FISIM) from personal interest income. More than half of the BEA personal interest series is imputed depositor and insurance services that no

household actually receives; levying it would have produced roughly \$102bn a year of phantom revenue. Base 3 nets to \$9,860bn after both deductions.

Does the base really exceed GDP? Yes, and this is not an error. GDP measures value added once. The levy reaches seventeen distinct flows, several of which are transfers of existing wealth rather than current production — inheritances at \$2.6tn being the clearest example. A dollar of corporate cash flow and a dollar of the wage it later pays are two different flows in the same year, and both are inside the base. The design constraint is not that the base be smaller than GDP; it is that no single dollar be levied twice.

7. The Senior Income Supplement — the most heavily revised element

This is where a reviewer should look hardest, because it is the part of the model that changed most and the part that still rests on interpolation rather than measurement.

The decile incomes come from Congressional Research Service report R47341, which publishes household income quintile upper bounds for householders aged 65 and over: \$24,132, \$43,252, \$68,480 and \$116,252. These are inflated from 2021 to 2026 dollars using CPI-U, fitted to a lognormal distribution, and integrated to give the mean income within each decile. The SSA publication 'Income of the Population 55 or Older' was examined first and rejected as the primary source because it was discontinued after the 2014 edition.

The lognormal fit has an R-squared of 0.9994. It is strongest through the middle of the distribution, where the bounds anchor it, and weakest in the two tails, where nothing does. The top decile figure produced by the fit is an artifact of an unbounded lognormal tail and should not be quoted. The bottom decile carries more model risk than the middle deciles, and it is the decile that drives most of the cost.

An external benchmark validates the rebuild. Published sources put roughly one in three Americans aged 65 and over below 200% of the federal poverty level; tested on the same basis, the rebuilt distribution puts 32% of senior households below that line. The same test is what condemned the previous hardcoded figures, which implied only about 21%.

What the material error was, and what it cost

The decile incomes are household income. They were being tested against the one-person federal poverty guideline, which compares a retired couple's pooled income to a single person's needs. The HHS guideline is linear in household size — roughly \$15,650 for one person plus about \$5,500 for each additional person. Census and BLS data put average household size at 2.0 for ages 65 to 74 and 1.6 for 75 and over, which weights to 1.83 persons. Evaluating the guideline at 1.83 persons gives \$20,226, and the result is insensitive to the cohort weighting.

A second error ran the other way. Both the Poverty tab and the Allocation projection multiplied a HOUSEHOLD income shortfall by a count of BENEFICIARIES. Since the decile incomes are household income, the shortfall must be multiplied by HOUSEHOLDS — 71mn beneficiaries divided by 1.83 gives 38.8mn households, not 71mn. That overstated the cost by exactly 1.83x. With both corrected independently, the 200% floor costs \$186bn a year and is affordable.

Flat floor multiple	Cost, year one	Balance after commitments	Years funded	Debt retired by 2100
160% of the guideline	\$240bn	+1.44%	75 of 75	\$43.8tn
175%	\$292bn	+1.20%	75 of 75	\$28.5tn
190%	\$349bn	+0.94%	75 of 75	\$11.7tn
195%	\$370bn	+0.84%	75 of 75	\$9.1tn
200%	\$390bn	+0.74%	75 of 75	\$6.6tn
210%	\$432bn	+0.55%	75 of 75	\$2.3tn

The original v41 figure of \$164bn at 200% was close to correct BY ACCIDENT — two errors of opposite sign were roughly cancelling. A reviewer who checked only the headline would have found nothing wrong; the errors were only visible once the units were forced to agree. A third instance of the same unit failure was found in the v48 pass, in a cell that reports the share of beneficiary households reached.

v62: the floor is flat per person, and there is no household test at all

The floor has been through three structures. Through version 49 a single threshold of \$40,452 applied to every household, obtained by evaluating the federal poverty guideline at an average beneficiary household size of 1.83 persons — too generous to the beneficiary living alone and too austere to two beneficiaries sharing a home. Version 50 replaced it with a household-composition rule: 2.00 times the one-person guideline for the first beneficiary at an address and 1.25 times for each additional one. Version 62 replaced that in turn, and the current design is the third and simplest of the three.

From version 62 the floor is FLAT PER PERSON. Every beneficiary aged 65 or over faces the same threshold — 200% of the one-person poverty guideline, \$31,920 in 2026. There is no advanced-age step: the 205% tier added at v64 was retired at v78. The floor does not depend on who a person lives with, how many beneficiaries share an address, or whether they are married. The live constants are FloorFirstMultiple, FloorAddlMultiple and Floor85Multiple, all 2.00x, on the Constants tab. The two multiples are equal by design: that equality is what makes the floor flat, and setting them apart would reintroduce the household determination this structure exists to remove.

Beneficiary	Floor	Per person
Aged 65-84	\$31,920	\$31,920
Aged 85 and over	\$31,920	\$31,920
Two beneficiaries sharing an address	\$63,840	\$31,920
Any household size	2.00x the one-person guideline each	unchanged by composition

Why no household test at all

The v50 rule rested on an address: the Social Security Administration already records a payment address for every beneficiary, so two beneficiaries at one address could be treated as one household without asking anyone a question. That was administratively sound and it is why the rule survived a dozen versions. What it could not do was survive its own premise. Whether two people genuinely share a household is a fact that cannot be falsified from the outside, so the rule created a determination that could be misstated — a fraud surface — and it did so in order to fine-tune a payment. Eliminating the test removes that surface structurally rather than mitigating it, which is a stronger position than any verification regime.

It also removes a penalty. Under the composition rule, two beneficiaries sharing a home received less each than either would living alone. That is a marriage penalty in substance whatever it is called in form, and it is imposed on a population in which isolation is already a measured harm and in which grey divorce is rising. Married couples, unmarried partners, siblings and housemates are now treated identically, because they are not distinguished at all. The cost of the change is stated rather than argued away: a single beneficiary faces a lower floor than the 2.00x the composition rule gave them, and single beneficiaries are the poorer group.

Why 200%, and why flat

The multiple is chosen rather than derived, and it is stated as such. What it is not is arbitrary: it is the point at which the model funds the floor in all 75 years while leaving a working margin, and the margin is itself part of the benefit.

The reasoning runs through the waterfall. The Senior Income Supplement is the FIRST claim on the surplus, which means it is the thing that fails if the buffer is gone. A margin of safety is therefore not a luxury sitting beside the floor; it is a property of the floor. Recalculated on v91, the floor costs \$390bn in year one and \$289bn a year averaged across the 75 years, real 2026 dollars, and is met in full in 75 of 75 years.

The sensitivity is worth stating, and it is looser than earlier versions of this paper implied. Re-run in full on v91, a flat floor at 160% costs \$240bn in year one and leaves a balance of +1.44%; at 175%, \$292bn and +1.20%; at 190%, \$349bn and +0.94%; at the adopted 200%, \$390bn and +0.74%. At 210% it costs \$432bn and is funded at +0.55%, and it is now delivered in 75 of 75 years — the trust fund’s bonds carry the early shortfall that cost three years at v81 — but reserves dip to 0.85 years of benefits and no bond is ever retired. The adopted multiple has real headroom below it, and that is the honest description of it. Earlier versions of this paper described a 200%

floor as unaffordable; on the v91 revenue path it is funded at +0.74% in 75 of 75 years, and that claim is withdrawn.

v51: the floor is an old-age floor, and the headline is relabeled

Through version 50 the floor population was the full OASDI beneficiary count of 71 million. That count was divided by 1.83, a household size measured for ages 65 and over, and tested against Congressional Research Service decile incomes for households with a householder aged 65 or over. Three different populations. Roughly 13 million beneficiaries under 65 — disability recipients, children of retired and deceased workers, early claimants and young survivors — were being converted into households and assigned the income profile of retired ones.

From version 51 onward the floor covers everyone aged 65 or over, whatever benefit they receive, and no one under 65, whatever benefit they receive. The line is age, not benefit type: an early retiree at 63 is outside it exactly as a disability beneficiary at 45 is. A disability beneficiary aged 65 or 66 is inside it, and keeps the floor under a different label at full retirement age, when the benefit converts automatically and the amount does not change. Nobody applies for anything and nobody is reclassified.

The share is taken from the Social Security Administration's Monthly Statistical Snapshot for July 2026: 56,809,000 people aged 65 or over receiving Social Security only, plus 1,428,000 receiving both Social Security and Supplemental Security Income, against 71,343,000 OASDI beneficiaries in total — 81.63%. The household count moves from 38.80 million to 35.78 million and the base floor to \$220bn a year.

What this costs is real and should be stated rather than buried. A person of 50 with a severe disability and \$14,000 of income receives nothing from the floor for fifteen years. The plan funds their benefit in full — 1.8 of the 12.4 points being replaced are Disability Insurance — and they enter the floor at 65. A Senior Income Supplement for working-age disability is a coherent proposal and probably a good one. It is a different proposal, resting on a different income distribution, and this model does not contain it.

Lifted to the floor is not lifted out of poverty

The floor pays up to 200% of the poverty guideline, flat across all ages 65+, not Being below that line is not being in poverty. Versions through 50 labeled that count as people lifted out of poverty, in the workbook and in every external document. That was wrong, and it was the kind of wrong that gets found: 18.2 million against roughly 59 million beneficiaries would imply a senior poverty rate near 31%, which no official source supports and which a reviewer would falsify in one step.

Version 53 computes both. Of the ten decile means, only the first sits below 100% of the guideline — \$13,581 against the 2026 guideline of \$15,960 for one person and \$21,640 for two. The floor lifts 26.0 million people to the floor, and 6.5 million of those were below the poverty line outright. The remaining 19.6 million were between the poverty line and twice it: near-poor, not poor. Both figures are now separate cells, and Summary row 47 carries a note stating which one the words out of poverty may be attached to.

The smaller figure is the stronger claim, because it is the one that survives being checked. Six and a half million people is close to the published senior poverty count, which is what a model measuring the right thing ought to produce.

v53: which tier of benefit taxation is repealed

Current law taxes benefits in two tiers and credits the revenue to two different funds. Tax on the first 50 percent of included benefits goes to the Social Security trust funds; tax on included benefits above 50 percent goes to Medicare's Hospital Insurance fund. Through version 52 this provision repealed both, and the plan then repaid Medicare for its share -- \$51.1bn in 2027, \$624bn over the first decade, \$3.70tn in present value.

That was a loss-making trade on its own terms. Repealing the Medicare tier hands beneficiaries \$38.8bn a year of relief while costing the plan \$51.1bn a year, because Medicare must be made whole on a gross basis while the taxpayer sees only the net. From version 53 the provision repeals the Social Security tier only. Beneficiaries below the second-tier thresholds of \$34,000 single and \$44,000 joint -- unindexed since 1993 -- stop paying income tax on their benefits entirely. Above those thresholds the taxable share falls from up to 85 percent to up to 35 percent. Medicare loses nothing, because nothing of Medicare's is taken.

The honest headline is that benefits become tax-free for most beneficiaries rather than for all. The alternative -- repealing both tiers and simply not repaying Medicare -- is worth 0.570 %, and it is the one change considered

anywhere in this model that would have been a deliberate misrepresentation. The Hospital Insurance fund is projected to deplete in the second quarter of 2033 holding \$256bn of reserves; removing \$39bn a year would pull that date forward. Repairing Social Security by worsening Medicare, and presenting the result as a solvency gain, is precisely the claim this proposal exists in order not to make.

Through version 53 one part of the add-back survived: the spillover onto income-related Part B and D premiums, about \$10.4bn in 2027. Version 54 moves that to the general fund as well, for the reasons set out below.

v56: Medicare sits entirely outside the proposal

Version 53 still paid Medicare about \$10bn a year. Making benefits partly tax-free lowers adjusted gross income, and the income-related premiums retirees pay on Medicare Parts B and D are keyed to that income, so some beneficiaries move into a lower premium bracket and Medicare collects less. Version 53 reimbursed that from the Social Security surplus.

Version 54 does not, and the reason is consistency rather than economy. Provision 1 lowers adjusted gross income by exactly the same mechanism, and its own effect on those premiums was always charged to the general fund, sitting quietly among the AGI-linked spillovers alongside the ACA premium credits, the earned income credit and the child credit. Charging one to the trust fund and the other to the general fund reflected the order the tabs were written in, not a principle. Both now sit in the income-tax ledger.

The distinction that matters is between a tax and a premium. The Hospital Insurance tier of benefit taxation is a tax earmarked to Medicare, and repealing it would have taken Medicare's money -- which is why version 53 stopped repealing it. The income-related premium is Medicare charging people according to their income. When income falls for any reason the premium follows, by design, and no one reimburses Medicare. That happens every time Congress changes any provision affecting adjusted gross income.

The result is that no payment flows between this proposal and Medicare in either direction. Medicare's tax rates, its thresholds, its trust fund, its benefits and its premium rules are all outside the design, and the deck no longer carries a Medicare line at all.

The floor does not cost anyone their Medicare assistance

Placing Medicare outside the design does not by itself prevent the design from reaching into it, because the help Medicare gives poor beneficiaries is income-tested and the floor raises income. The 2026 limits for a single person are \$16,200 a year for the Qualified Medicare Beneficiary programme, \$19,392 for Specified Low-Income Medicare Beneficiary and \$21,792 for Qualifying Individual. The floor pays \$31,920. Every enrollee it reached would cross every limit.

What they would lose is substantial: the Part B premium at \$202.90 a month, or \$2,435 a year; for QMB enrollees the Part A and Part B deductibles and all coinsurance, worth roughly \$2,500; and the Part D low-income subsidy, worth about \$5,700. A QMB enrollee stands to lose around \$10,635 a year and an SLMB or QI enrollee around \$8,135.

Left unaddressed, this would hollow out the floor precisely where it matters most. A household in the first decile, lifted from \$13,581 to \$31,920, would gain \$18,339 in cash and lose \$10,635 in kind, netting 42 cents on the dollar. A household just below the QI limit would net 20 cents. The person the floor helps most on paper would be helped least in fact.

The floor payment is therefore disregarded in determining eligibility for the Medicare Savings Programs and the Part D low-income subsidy. Nobody's countable income changes for those purposes, nobody loses in-kind support, and Medicaid and Part D continue to pay exactly what they pay today. This is ordinary drafting rather than a novel device: the earned income credit is disregarded for means-tested benefits on the same reasoning.

The disregard costs the levy nothing. It does forgo something: had the floor been allowed to disqualify these households, that is a saving deliberately not taken and is not claimed anywhere in this model. A floor that pays for itself by stripping poor people of their medical assistance would not be a floor.

Interest no longer paid on retired federal debt

Retiring \$6.6tn of publicly held debt by 2100 means the Treasury stops paying interest on it. In constant 2026 dollars that saving averages \$50bn a year across the projection and totals \$3.7tn; on a present-value basis it is

\$1,079bn, which offsets 6.8% of the plan's general-fund cost. The debt figure is net of the \$0.2tn of bonds cashed for the Supplement in 2027–32, which Treasury funds by borrowing. It is a general-fund effect, it is computed in the model at Allocation rows 102 to 105 from the annual series rather than asserted, and it is credited nowhere in the actuarial balance.

Stress testing the margin

The 75-year balance of +0.74 percent is not the interesting number. What matters is how far the result can be pushed before a commitment fails. Three of the plan's revenue lines rest on author estimates with published ranges, and the honest test is to run them at the unfavorable end of those ranges simultaneously.

Scenario	Revenue	Balance	Senior Income Supplement met
Central — live constants (earnings share 69%, Base 6 mortgage deduction 22%)	\$2,014bn	+0.74%	75 of 75
Adverse — earnings share 62%, deduction 24%	\$2,006bn	+0.66%	75 of 75
Severe — earnings share 55% (range floor), deduction 26%	\$1,998bn	+0.58%	75 of 75

Reserves never deplete in any scenario and the Senior Income Supplement is met in every year of all three; in the severe case the reserve dips to 0.95 years of benefits. Demographics are not the vulnerability: moving fertility from the Trustees' 1.75 to the actual 1.60 changes the result by nothing at all, because that lever feeds a stated sensitivity rather than the projection. The vulnerability is the revenue estimates, which is why two of them were replaced with derivations at version 53 rather than left as judgements.

v51: citizens aged 65 and over with no earnings record

A Senior Income Supplement built on beneficiary status cannot reach the people least likely to be beneficiaries. The Social Security Administration calls them never beneficiaries: about 3.3% of the population aged 60 or over, some 2.6 million people, of whom 54.3% live below 100% of the poverty line against 5.8% of current and future beneficiaries. A floor that excludes the group nine times more likely to be poor is not a Senior Income Supplement.

Version 53 brings US citizens aged 65 and over with no earnings record inside it. Citizenship is recorded on the Numident when the number is issued, so it is held data in the same way the payment address is: no new question, no verification, no redetermination. Non-citizens are excluded by the same condition federal law already applies — the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 requires lawful permanent residents to have 40 qualifying quarters, which is the same forty credits that qualify a person for Social Security. The floor imports a line Congress drew rather than drawing a new one, and that is the honest description of it: it is a line that falls on lawful permanent residents, most of whom arrived too late to accumulate ten years of covered work.

The addition costs about \$23bn a year and reaches roughly 880,000 households. It lifts 1,147,000 people to twice the guideline and 681,000 out of poverty outright, of whom about 429,000 are women. It is the most expensive part of the floor measured per person reached — roughly \$20,400 against \$14,000 for the main floor — because these households start at zero and the whole distance to the floor has to be paid. That is an argument for it rather than against it: ranking anti-poverty spending by cost per person reached funds the near-poor first every time and the destitute never.

What exists today for this group is Supplemental Security Income, whose 2026 federal benefit rate is \$994 a month, or \$11,928 a year — 76% of the one-person poverty guideline — enforced with a \$2,000 asset limit unchanged since 1989. From v83 aged SSI is absorbed into the Supplement entirely: about \$18.5bn a year, \$16.7bn federal and \$1.8bn state, ceases and is carried inside the Supplement's cost. The federal saving lands in the general fund while the floor is paid from the swept surplus, so it sits in a different account from the cost and is not claimed anywhere in the actuarial balance. It belongs with the other unclaimed offsets in Section 9.

The population figure is the weakest input in the model and is flagged as such in its own constant. It is derived, not published. Version 51 put it at 780,000 by applying SSA's 38.9% infrequent-worker share as if it were a citizenship split; it is not, since SSA's four categories describe why forty credits are missing rather than nationality. Version 53 corrects this by applying a naturalisation rate of about 50% to the late-arrival category, giving a citizen

share of 75.6% and a population of 1.53 million. That reconciles with the Administration's own count of 1,071,267 people aged 65 or over receiving SSI and no Social Security, of whom about 12.1% are noncitizens — roughly 850,000 citizens already enrolled, before counting anyone above the \$2,000 asset limit, anyone with income between the SSI cutoff and the floor, or anyone who never applied. The range is 1.20 to 1.75 million on a 40-60% naturalisation rate. A Current Population Survey microdata run would settle it, along with the household-composition share and the housing-tenure question already outstanding.

What the v50 and v51 changes did to the results, and where they landed

This paragraph is history and is retained so the chain is auditable. At v50, on the pre-v51 all-ages base, households reached fell from 11.64 million to 11.25 million, because single-beneficiary households faced a lower floor and those in the third decile cleared it; people reached fell from 26.0 million to 18.2 million on the same base. The annual cost rose from \$186bn to \$198bn. Both figures were then superseded by the v51 age scope, and the composition rule itself was superseded at v62. On the current flat floor the model reaches 26.0 million people, of whom 6.48 million were below the poverty line outright.

A unit error was caught during the rebuild and is recorded because it is the fourth instance of the same failure in this model. The person count was initially left as households multiplied by an derived average size of 1.77. With composition explicit that is wrong — it applies a blended average to a blended count. It now derives from composition directly, and the implied average size below the floor comes out at 1.62 rather than 1.77.

One approximation remains. The share of households containing a single beneficiary is 38%, derived as 2 minus BenefPerHH (1.62) from the average number of beneficiaries per beneficiary household of households below the floor. That figure described below-floor households while the share is applied to all beneficiary households, so the two are not the same quantity. A Current Population Survey microdata run keyed to Social Security address records would measure both directly, and is the same run that would settle the housing-tenure question below.

How income is measured for the floor

The floor is an income test, not a means test. There is no asset verification, no bank statement, no application and no redetermination. Eligibility is computed from records the Social Security Administration already holds, in the same way the Medicare income-related surcharge is already applied automatically using Internal Revenue Service data from two years prior. That is what keeps the guarantee free of the administrative cost that makes Supplemental Security Income expensive to run, at an agency now operating at roughly 1,429 beneficiaries per employee against 400 in 1967.

One adjustment is made to the income concept, and it is made for consistency rather than convenience. A household that owns its home outright receives a real flow of housing services that a renter with identical cash income must buy. Counting only cash treats those two households as equally poor when they are not. The model therefore measures income for the floor as cash income plus NET IMPUTED RENT on owner-occupied housing.

Net imputed rent is gross rental value less the costs of producing it: property tax, homeowners association fees, insurance and maintenance, and mortgage interest where a mortgage exists. That is not an expense deduction introduced for this purpose. It is the identical rule the model already applies to Base 6, where rent is levied net of property taxes, loan interest and a fixed maintenance allowance. A landlord is taxed on net rental income; an owner-occupier is measured on net imputed rental income. Same income type, same treatment, whether the tenant is a stranger or the owner.

Four tenure states cover substantially all cases. A renter has no imputed income, because housing is bought out of cash income already counted. An outright owner has the full net rental value. A mortgaged owner has the net rental value less mortgage interest, which correctly approaches zero for a recent purchase, since a household that has just bought is not living rent-free. A household whose housing is paid by someone else is credited with the amount paid on its behalf, which includes rental subsidy: a housing voucher is a transfer received and is counted as income. Residents of institutional care are excluded from the imputation entirely.

Valuing imputed rent without valuing every house

The natural objection is that this requires a property valuation for sixteen million households, which would rebuild exactly the administrative apparatus the design avoids. It does not, because a published federal series already exists for the purpose.

Gross rental value is taken from the Department of Housing and Urban Development's Fair Market Rents, published annually for every county and metropolitan area in the United States, by unit size, at the fortieth percentile of area rents. Three properties recommend it. It is a federal standard already used to administer housing assistance, so the valuation is not this proposal's to defend. It is geographically adjusted by construction, so a home in rural Michigan and one in San Diego are not treated as equivalent. And the fortieth percentile is deliberately conservative, which errs toward paying the floor rather than denying it.

The alternatives were considered and rejected. User cost — capital cost plus depreciation less expected appreciation — is the economically pure measure and is unusable, because it depends on price expectations and turns negative during booms. A fixed percentage of assessed value, the method used in the Netherlands, fails in the United States because assessment practice is not uniform: California's Proposition 13 holds long-held homes far below market value while other states assess near it, so identical houses would be imputed wildly different income. Fair Market Rent requires only county and unit size, both of which follow from records already held.

This measurement rule is stated but NOT yet modeled. Applying it requires the joint distribution of income, housing tenure and unit size among beneficiaries, which cannot be derived from the four published quintile bounds the decile distribution currently rests on. It needs the same Current Population Survey microdata run that would replace the derived household size. Until that is done, the floor is costed on cash income alone, which OVERSTATES the number of households below the floor and therefore overstates the cost. The published figures are the conservative case.

One boundary is drawn deliberately and should be stated plainly. Out-of-pocket medical spending is not deducted. The Census Bureau's Supplemental Poverty Measure does deduct it, and on that basis the share of Americans over 65 below 200% of poverty rises from 29.4% to 42.2% — from 17.0 to 24.4 million people. The omission is therefore material and it is not an oversight. The rule applied throughout is that income is measured net of the cost of producing it and never net of consumption. Property tax is a cost of holding the asset that generates imputed rent; medical spending is consumption, however involuntary. Admitting consumption deductions would convert an automatic income test into a case-by-case means test, which is the administrative machinery this design exists to avoid. That the line falls here is a choice, not a derivation, and it is the one place in the floor's construction where the reader is asked to accept a judgement rather than follow an argument.

8. Investment and return assumptions

Is a 7.25% real return defensible? It is an assumption, not a hindcast, and it should be presented as one. The historical geometric real return on US equities is approximately 6.8% — the S&P 500 total return with dividends reinvested compounded at roughly 10.0% nominal from 1928 to 2025 against average CPI inflation of about 3% (Damodaran, NYU Stern). The fund's mandate is a 70/30 screened global portfolio rather than the S&P, and a 70/30 blend at historical rates lands nearer 6.2% to 6.4% real. So 7.25% sits roughly 0.85 to 1.05 points above a market-weight historical benchmark. The case for it rests on the capital share of national income rising as automation proceeds, which is a growth argument rather than a risk-premium argument, and a reviewer is entitled to discount it.

What the return assumption actually changes

Less than one might expect, and nothing that matters. Moving from 7.00% to 7.25% leaves the actuarial balance unchanged to the basis point at +0.74%, and both commitments are met in 75 of 75 years under either. What moves is the size of the fund and the debt it retires, and almost all of that movement is post-2060 compounding. Re-run at v91: the fund ends 2100 at \$14.9tn against \$16.6tn, and debt retired at \$6.3tn against \$6.6tn. The fund moves proportionally further than debt does, and the ownership cap binds under neither return: peak ownership is 4.0% at 7.25% and 3.6% at 7.00%. At a 4.00% real return the Supplement is still met in all 75 years and reserves never deplete; the fund ends at \$4.9tn.

The architectural point behind this is worth stating explicitly, because a reviewer skimming could easily mistake the proposal for an equity-investment plan and stop there. The fund sits downstream of the trust fund: benefits are paid first and the fund is built from what remains, while the trust fund itself continues to earn the special-issue rate. The solvency claim therefore requires no change whatever to trust fund investment authority, and can be scored entirely under current-law investment rules.

The 70/30 US-to-global split is a stated deviation from market weight. The model's own constants put the US at roughly 64% of the investable universe, so 70/30 represents a six-point home bias. An 80/20 split was tested and rejected. Re-run at v91, the ownership cap binds under neither split, so the fund and the debt retired do not change at all; what changes is government ownership of domestic equity, which rises from 4.0% to 4.6% of the eligible US market at peak. It makes the central political vulnerability worse and buys nothing.

The fund does not reach the 10% ownership cap. Peak ownership of the eligible universe in the binding sleeve is 4.0% at the base-case return, and average ownership across the projection is 1.37%. The 10% figure is the Section 16 insider threshold and functions as a hard guardrail rather than a target; on the current design it is a guardrail that is never tested. At a 7.00% return peak ownership is 3.6%. Earlier versions of this paper stated that the cap bound at exactly 10.00%; that was true of the pre-v56 design, in which the fund absorbed the cap overflow, and it stopped being true when the rebate was introduced and again when it was retired.

One asymmetry is disclosed rather than defended. If the automation argument that justifies the return premium holds, then the assumed real market-cap growth of 2.6% should rise with it. It has not been raised. The favorable half of the argument is in the model and the offsetting half is not — which is a conservative combination, since a faster-growing market would lower the fund's ownership share further still.

9. What is deliberately not claimed

A reviewer's first instinct on seeing a result this favorable is to look for the assumption that produced it. This section exists to make that search short. Set out below is every material choice where the model could have been made to look better and was not. The stack runs one way, and it is the conservative way.

Effects excluded from every scored figure

Each of these is real. None of them appears in the balance, the depletion test, the Senior Income Supplement or the debt path. The plan is solvent before any of them is counted.

Excluded effect	Approximate worth
The growth dividend. Removing a per-head tax on employment and a double tax on capital should raise trend growth. Half a point of extra real growth is worth roughly +2.6% of payroll to this design — around 60% of the entire deficit — because 60% of revenue comes from bases that earn no benefit credit and therefore carry no offsetting benefit cost.	+2.6% of payroll at +0.5%
The consumption multiplier on the Senior Income Supplement. Money delivered to the lowest income deciles is spent almost in full, whereas the same money retained at the top is largely saved.	Real, unquantified
Interest saved on retired debt. \$6.6tn of publicly-held debt is retired by 2100, with paydown beginning in 2033, and the interest never paid on it compounds for sixty-seven years. The model computes this at Allocation rows 102 to 105: an average of \$50bn a year in real 2026 dollars, \$3.7tn cumulatively, \$1,079bn in present value. It accrues to the general fund rather than to OASDI, so not a dollar of it appears in the actuarial balance.	\$50bn/yr average — offsets 6.8% of the plan's general-fund cost
Reduced public cost of senior poverty — Medicaid, SSI, housing assistance and emergency care avoided when 26.0 million people are lifted above the line.	Real, unquantified
Income and corporate tax revenue from higher employment once hiring is no longer taxed per head.	Real, unquantified

Keeping these outside the scored case is not modesty. It means the case cannot be attacked through them. An opponent who disputes the growth dividend has disputed nothing that matters, because nothing in the result depends on it.

Baseline choices that make the result look worse

Three decisions about what to score against, each of which costs the plan measurable ground.

Choice	Cost to the result
Scored against the Trustees' published deficit of -4.42% rather than the model's own from-scratch projection of current law, which summarises to about -4.74%. The published figure is the easier baseline to beat and was chosen anyway, because it is the one the Office of the Chief Actuary uses.	0.32 %
Fertility held at the Trustees' ultimate assumption of 1.75, when the actual US rate fell to roughly 1.60 in 2025. On the actual rate the current-law baseline is about -4.77%, and the plan's improvement is correspondingly larger. The adjustment is not claimed.	0.35 %

The year-by-year depletion test runs on the stricter internal path rather than the published one, so the 'never depletes' result is proved against the harder baseline while the balance is scored against the easier.	Not quantified
--	----------------

Taken together the first two give away roughly two thirds of a point of payroll before any behavioral argument is made.

Base construction where the lower figure was taken

- Imputed financial services (FISIM) removed from personal interest income — \$1,070bn of base and about \$102bn a year of revenue, forgone because more than half of the BEA personal interest series is imputed depositor and insurance services that no household actually receives. Levying it would have been phantom revenue.
- Inbound dividends levied on what arrives, net of foreign withholding at source, rather than on the gross declaration.
- The retirement earnings share is set at 69%, derived at v52 from aggregate compounding rather than chosen as a midpoint of the 55% to 70% range. This is stated here for completeness and it is NOT a conservative choice: 69% sits near the top of the range, and the low end of the range would reduce the base. It is listed in this section only so that a reader is not left to discover it elsewhere.
- A ceiling of 95 million on the projected beneficiary count — a constraint that can only ever raise projected cost.
- The maximum benefit held at current law despite contributions being uncapped, so the plan generates no new benefit obligation against the newly reached income.

The one assumption that is not conservative

The assumed real fund return of 7.25% sits roughly 0.85 to 1.05 points above a market-weight historical benchmark for a 70/30 blend. It is the single above-benchmark assumption in the model, and it is stated as an assumption rather than a hindcast.

It also affects nothing that is scored. At 7.00% the actuarial balance is identical to the basis point and both commitments are still met in 75 of 75 years; what moves is the size of the fund and the debt retired. The only non-conservative assumption in the model touches no result on which the case rests.

The honest headline

The plan reaches +0.74% of taxable payroll after every commitment it makes, and +2.56% before the Senior Income Supplement is charged, with no behavioral response counted, against a baseline deliberately chosen to be easier to beat than the evidence supports, and with the largest excluded effect worth more than half the deficit on its own.

The rate sits sixteen basis points above a delivery cliff

This is the most important thing version 71 found, and it was invisible before. The solving-rate sensitivity grid in the workbook had stood unchanged since version 27 — it predated four of the seventeen sources, two rounds of poverty corrections and the current fund treatment — and it reported the Senior Income Supplement as met in full at a levy of 5.85%, with a fund of \$20.0tn. Retired and re-run live on the current workbook from nine explicit rate runs, that is not what happens.

Re-run on v91, with the trust fund's bonds held as the reserve: at 5.78% the Supplement is met in 67 of 75 years. At 5.79% it is met in 75 of 75, but only by cashing bonds, so publicly-held debt ends \$3.7tn higher rather than lower and the fund reaches \$0.8tn. At 5.85% it is met in 75 of 75 with a fund of \$2.0tn; at 5.90%, \$7.7tn; at the adopted 5.95%, \$16.6tn. The central claim of this proposal — the Supplement delivered in every year of the projection — therefore appears at about 5.79%, and the chosen rate clears it by 0.16% of payroll, up from 0.05% at v90 because the bonds now carry the early shortfall the fund could not.

The steepness is mechanical rather than mysterious. The fund is fed by whatever residual is left after benefits and the Supplement have been paid, and that residual is then compounded for seventy years. Close to the threshold

the residual is near zero, so a small movement in the rate moves the terminal fund by an order of magnitude. The corollary is that the fund figure — \$16.6tn — is the single most rate-sensitive number in the model and the least suitable one to lead with. The figure worth defending is the margin over the delivery threshold, which is about 0.16% of payroll.

Two consequences follow and both are accepted here rather than argued away. First, solvency is not the binding constraint: the system is solvent on current-law benefits from about 5.19% and the reserves never deplete from about 5.70% (a test that, from v91, includes Supplement draws on the bonds), but it does not deliver the Supplement in all 75 years until about 5.79%. Quoting a bare-solvency rate understates what the proposal actually requires. Second, any adverse movement worth roughly 0.16% of payroll — in the fertility path, in one of the larger revenue bases, or in the cost of the Supplement itself — puts the headline delivery claim at risk in a way the fund total does not reveal. The sensitivities in sections 8 and 9 are run against that margin, not against the fund.

It is stated here because it is the first sensitivity any actuary will run, and it is better read in this paper than discovered in a review.

10. Known limitations

The following are stated because a reviewer will find them, and it is better that they are disclosed than discovered.

Retirement accounts are levied on withdrawal, and the base is now a published series.

Through v48 the model levied gains realized INSIDE a retirement account and assumed \$1,025bn a year — the midpoint of a \$500bn to \$1,550bn range with no published measurement behind it, because such realizations generate no 1099-B and are reported nowhere. That was the dominant uncertainty in the model and it has been removed.

Version 49 levies distributions instead. The base is the earnings portion of retirement distributions net of rollovers: IRS SOI Table 1.4 reports \$912bn of taxable pensions and annuities and \$438bn of taxable IRA distributions for tax year 2022, grown to 2026 and increased by an estimate of qualified Roth distributions, giving \$1,771bn of distributions. At the live 69% earnings share the levyable base is \$1,222bn. Contributions are returned free because they already bore the levy on the way in, as wages in Base 1 and employer contributions in Base 5.

Three features of that treatment are worth drawing out. It closes a double levy: under the previous rule a fund that sold one holding at a gain, bought another and sold that at a gain was levied twice on a single pool of money that never left the wrapper. It excludes rollovers automatically, because a rollover reports zero in the taxable box of the published series — so the roughly \$800bn a year of people changing jobs or providers is not levied, and no new rule was needed to achieve that. And it uses a pro-rata exclusion ratio in the manner of IRC section 72, not basis-first ordering: basis-first would defer almost all revenue into the 2050s and beyond, hollowing out precisely the decades in which the trust fund is thinnest.

Roth accounts are levied on the same terms as traditional accounts. Roth earnings are tax-free under current law, so this is a real change to an affluent and financially attentive group, and it is stated as a deliberate decision rather than left to be discovered. The alternative — exempting Roth — would open an avoidance channel that conversions would exploit immediately, since under this plan money entering either wrapper has already been levied.

What remains an estimate is the earnings share, set at 69% and derived at v52 from aggregate compounding. Across a range of 55% to 70% the levyable base runs from \$974bn to \$1,240bn, which straddles the figure it replaces, and the plan is solvent and never depletes across the whole of it. That is a materially narrower band than the \$500bn to \$1,550bn it supersedes. Two smaller estimates sit alongside it: the 5% annual growth carrying the 2022 figures to 2026, and \$130bn of qualified Roth distributions, which do not appear in the SOI series because they are not taxable under current law.

The SOI figures were verified against Tax Foundation, Sources of Personal Income, Tax Year 2022 (January 2026), which cites IRS SOI Table 1.4. They have not been read directly from Publication 1304. Anyone relying on them should confirm against the primary table.

The decile distribution is interpolated, not measured.

A CPS microdata run keyed to actual household size would replace the lognormal interpolation with measured data and remove the tail uncertainty in a single step. This is the highest-value refinement available to the model and it has not been done. The average number of beneficiaries per beneficiary household is set at 1.62

(Constants BenefPerHH), an estimate; the count of people reached moves proportionally with it, while the cost of the guarantee does not move at all, because the cost is computed on households.

Provenance notes have been partially swept.

The fertility lever carried a note stating it was wired into the projection when it drives two display cells and nothing else, and that note survived four model versions. The v48 pass found and closed a second instance: two cells reporting the bare-solvency and never-deplete levy rates carried a note claiming they read live off a sensitivity table, when both are hardcoded records of explicit runs and the table contains neither rate. The remaining 'this is now wired' notes have still not been individually traced against the formulas that would have to reference them.

The sensitivity grid has been retired and re-run.

The v27 solving-rate table was retired at v78 and replaced by explicit rate runs of the current workbook (section 9). The full-delivery and never-deplete rates were run again at v91, because holding the trust fund's bonds changes both. They are stored results of those runs rather than live solves, and must be re-run whenever the waterfall or the Supplement changes.

Only the intermediate assumptions have been tested.

The plan is scored against the Trustees' intermediate assumptions. The Office of the Chief Actuary publishes low-cost and high-cost variants alongside and scores against all three. A balance that holds under the high-cost variant is a claim very few reform proposals can make; a balance that does not will be quoted at high-cost by opponents permanently. This has not been run, and it is the most consequential remaining gap in advance of a formal score.

Fertility is aligned to the Trustees, not to the actual rate.

The model uses the Trustees' ultimate assumption of 1.75. The actual US rate fell to roughly 1.60 in 2025, which would imply a current-law baseline of about -4.77% rather than -4.42%. That adjustment is deliberately not claimed. Scoring against the Trustees' own published baseline is the conservative choice when seeking their score, and net migration is the offset that could close a fertility gap on a timetable no one can forecast.

Two items remain unresolved as policy questions.

Foreign holders of US-source income, worth \$64bn of revenue, collides with approximately 60 bilateral tax treaties; the model includes it, and whether it survives negotiation is a separate question. Inheritances levied at source and gross, worth \$158bn, is the largest political attack surface in the design for 7.9% of revenue. Withholding at the estate removes the administrative objection entirely — there is no heir-level filing and no valuation dispute between an heir and the government — but it does not remove the rhetorical one, and the history of estate-tax politics in the United States runs largely independent of the arithmetic.

11. Version history

Version	What changed
v43	Fertility aligned to the Trustees ultimate 1.75. IRMAA reimbursement funded. Poverty decile incomes rebuilt from CRS R47341. Poverty guideline moved to a household-size basis.
v44	The unit bug found and fixed; 200% FPL restored. Person-level reach made native to the model.
v45	Inbound dividends levied on the net receipt rather than the gross declaration. Life insurance death benefits added to Base 2. The 7.29% real return claim withdrawn as an arithmetic-mean artifact.
v46 / v47	Not logged contemporaneously; reconstructed in the v48 pass. Sources 9, 10 and 11 added. Return sensitivity refreshed, establishing that the return lever does not touch solvency. Balance moved from +0.34% to +0.48%.
v48	Full twenty-tab audit pass. One live formula error corrected, one mislabeled cell repointed, stale commentary and a stale sensitivity grid flagged, two provenance claims corrected, and the change-log gap closed. Every hardcoded number lifted onto a Constants tab with a source; every year and discount period rewired to derive from the year column. The robot duty cycle, once visible as a named assumption, was corrected from 84 hours a week to 22 hours a day over 50 weeks.
v49	Retirement accounts moved to a withdrawal basis: distributions net of rollovers, earnings only, pro-rata rather than basis-first, Roth levied on the same terms. Replaces the model's dominant

	uncertainty with a published IRS series and closes a double levy on churn inside the wrapper. Balance moved from +0.48% to +0.53%; debt retired by 2100 from \$18.7tn to \$22.0tn.
v50	The Senior Income Supplement is set by household composition rather than by a single average-household threshold: 2.00x the one-person guideline for the first beneficiary at an address, 1.25x for each additional. Households reached fall to 11.25mn; people reached rise to 18.2mn. Balance moved from +0.53% to +0.45%; debt retired by 2100 from \$22.0tn to \$15.2tn. A fourth instance of the household/person unit error was caught during the rebuild.
v51	The floor is scoped to ages 65 and over, so that the household count and the CRS 65+ income distribution measure the same population; the previous count divided an all-ages beneficiary total by a 65+ household size. A 100% poverty test is added alongside the 200% floor, and the headline is relabeled: the floor lifts 15.5mn people to twice the guideline, of whom 5.0mn were below the poverty line outright. Earlier versions described the 200% count as lifted out of poverty, which it is not. US citizens aged 65+ with no earnings record are brought inside the floor. Balance moved from +0.45% to +0.70%; debt retired by 2100 from \$15.2tn to \$27.0tn.
v52	The household divisor is corrected. The model had divided the beneficiary count by 1.83 persons per household while assuming a composition of 55/45 that implies 1.77 BENEFICIARIES per household -- two positions that imply 31.8mn and 40.2mn households and cannot both hold. Multiplied out, the composition accounted for only 45.9mn of 58.2mn beneficiaries. Share1PersonHH is now DERIVED as 2-BenefPerHH so the two can never disagree again. The never-beneficiary population is corrected from 0.78mn to 1.53mn. Base 6 stops deducting mortgage principal and the retirement earnings share moves to 69%.
v53	Provision 3 repeals only the tier of benefit taxation credited to OASI and DI, leaving in place the tier credited to Medicare HI. Medicare loses nothing and there is nothing to make whole beyond the IRMAA spillover, which falls from \$51.1bn to \$10.4bn. The surplus split moves to 55/45 in favor of the fund. Balance +0.77%, debt retired by 2100 \$35.3tn, 26.05mn lifted to twice the guideline and 6.5mn out of poverty.
v56	Medicare is placed entirely outside the proposal: no payment flows to it or from it. The IRMAA spillover, previously charged to the trust fund, moves to the general fund where Provision 1's identical effect already sat. The floor is disregarded for Medicare Savings Program and Part D low-income subsidy eligibility, so no beneficiary loses in-kind support. The wage levy is remitted by the employee and withheld by the employer, uncapped, with the pre-tax exclusion left capped at the old wage base. Audit fixes: Share1PersonHH and PovertyGuideline were orphaned constants carrying superseded values and now mirror the live cells. Balance +0.70%, debt retired \$7.7tn, average annual interest saved \$102bn.
v59	Sensitivity grid re-run on the post-rebate base.
v62	THE FLOOR BECOMES FLAT PER PERSON. FloorFirstMultiple and FloorAddMultiple are set equal, removing the household-composition determination and the address test with it. The advanced-age tier is added at 205% from age 85. THE LEVY REBATE IS RETIRED: RebateSwitch set to 0, the effective rate holds at 5.95% throughout, and the annual draw stays in the pool — trading the declining-rate claim for a larger fund and more debt retired.
v64	The flat multiple is settled at 195% of the one-person guideline.
v66	State and local fiscal ledger added; reconciliation merged into the workbook as a tab and made live; sensitivity grid unfrozen and the base-case column wired to read the model.
v67	Correction pass across the workbook and all four documents. Summary D89 was reading the PEAK ownership cell under an AVERAGE label and now reads Allocation C108 (2.25% against 8.5%) — the seventh instance of a pointer or unit error in this model. Summary D101 was adding the never-beneficiary annual cost to a 75-year average, printing \$262bn where \$415bn belongs. Stale labels quoting a 200% floor and a live rebate were swept from the workbook and from this paper. The floor-multiple and revenue-estimate sensitivity tables were re-run in full; both had gone non-monotonic on superseded bases. No scored figure moved: the balance after all commitments is +0.70%.
v68	Model carried forward; companion documents unchanged.
v69	SOURCE 7 MOVED ONTO BEA GROSS OPERATING SURPLUS. The non-profit base was struck on net income, after depreciation and interest, while Base 3 and Source 10 both levy gross operating surplus. Rebuilt from BEA's non-profit institutions serving households account: gross value added of \$1,631.9bn less compensation of \$1,339.6bn less taxes on production of \$29.8bn leaves consumption of fixed capital of \$262.4bn (2025), the components exhausting gross value added because BEA assigns the sector no net operating surplus by construction. Base \$120bn to \$275bn; revenue \$7.1bn to \$16.4bn. No investment income was added — all dividend income is levied at source in Base 3.

v70	EIGHT AUTHOR ESTIMATES REPLACED BY PUBLISHED SERIES. State and local wages \$1,150bn to \$1,550bn and non-profit wages \$700bn to \$1,170bn, both from BEA wage and salary accruals; federal wages from the BEA civilian and military series; Base 1 wages re-anchored to the published 2025 actual of \$12,958.7bn; US market capitalisation \$62tn to the measured \$75.3tn, which lowers the fund's ownership share and pushes the per-issuer cap further out. Sector wage bills are attribution only — corporate taxable payroll is the residual — so total revenue does not move.
v78	THE v27 SENSITIVITY GRID RETIRED AND RE-RUN LIVE, AND A DELIVERY CLIFF FOUND FIVE BASIS POINTS BELOW THE PLAN RATE (section 9). Separately, the state and local cap ratio was corrected for Social Security COVERAGE as well as the cap: about 5 million state and local employees are not covered by the programme at all, so no employer contribution is paid on them today. Ratio 0.9192 to 0.6968; state and local employer share abolished \$88.3bn to \$67.0bn, which cuts the windfall the v70 wage-bill correction had just increased. Bare solvency re-solved to 5.28% and the after-commitments crossing to 5.34%; the never-depletes rate stored since v60 was wrong by 0.17% and is corrected to 5.28%.
v83-v86	2026 poverty guidelines updated; the floor rises to \$31,920. Aged-65+ SSI absorbed into the Supplement. Base 12 maintenance allowance set at 1% of the CPI-indexed purchase price per year of ownership. Sources 13 to 17 added: insurance and annuity inside buildup, Section 1031 rollover gains, lottery and gambling winnings, irrevocable trust retained gains, and royalties to individuals.
v87	PROPERTY TAX IS NOT DEDUCTIBLE IN BASE 12. Excluding imputed rent is generous to owner-occupiers, and non-deduction is the partial offset. The tiers were then re-derived from the model rather than assumed: over 15 years moves from 0% to 19.36%, 2 to 15 years from 5% to 7.10%, under 2 years unchanged at 30%. Base 12 rises from \$77.63bn to \$177.98bn. There is no longer a break-even tenure.
v88	THE RESERVE DOUBLE-COUNT REMOVED. The model had been spending the trust fund reserves and cancelling them at the same time: Allocation cancelled the special issues while Solvency continued to roll the reserve forward and release it over ten years. The legacy release is disabled. Reserves restated to \$2,102.2bn, the 2026 Trustees Report Table VI.G2 projection for end-2027, on a 2028 cutover. The special issues are settled for \$600bn, \$300bn in each of the first two years, which restores the Supplement to 75 of 75 in every year.
v89	CHARITABLE TRANSFERS AT DEATH ADDED TO BASE 2, \$62.19bn (Giving USA 2026, bequest giving 2025). They were never inside the \$2,600bn: the Cerulli figure the model cites is the flow to HEIRS, and charity is separate. No charitable exclusion applies. The IRC s2055 deduction covers every 501(c)(3), so bequests into private foundations, donor-advised funds and charitable trusts are already within this figure and are not added again. The Cerulli charity:heirs ratio was rejected as the basis because it includes lifetime giving, which is already levied under the gifts-at-source switch. Base 2 rises to \$2,662bn and total revenue to \$2,014bn.
v91	THE TRUST FUND'S BONDS ARE HELD, NOT CANCELLED. v88 wrote the special issues off at cutover and paid a \$600bn settlement. v91 holds them as the reserve: they earn the special-issue rate, pay any shortfall first, and are retired non-cash only above one year of benefits (bonds plus 75% of the fund) and only as far as the fund has replaced them — from 2033 to 2086. Reserves never fall below one year of benefits. The second \$300bn settlement payment had been added to Supplement delivered on top of a need already met; removing it moves the balance from +0.70% to +0.74%. Without the first payment seeding the fund, the fund reaches \$16.6tn by 2100 and \$6.6tn of publicly-held debt is retired. Starting reserves corrected to the Trustees' \$2,561bn at end-2025; current-law depletion now 2034. Full delivery from 5.79%. The stress test now applies its shock to revenue, as its label says, rather than to the swept surplus, about one-sixth the base: at equities down 60% and revenue down 10% for five years the Supplement is met in 75 of 75.

All figures trace to model v91. Where this document and the deck disagree, the model governs. Section 9 lists every material choice where the model could have been made to look better and was not.

v56: audit and correction pass

Figures quoted inside this version history are as at the version under which they appear and are superseded by the v91 headline set in section 1. They are left unrestated so the record of what each version found remains intact.

Nine corrections, one design decision, and no change to the design itself. The largest was a Medicare charge that should not have existed. From version 53 this proposal repeals the Social Security tier of benefit taxation only, so Medicare keeps its own tier and loses nothing — a position stated on the provision tab, in both Allocation

switches, in this paper and in the reconciliation. The model nevertheless continued to book the entire Medicare share as a general-fund loss: \$38.8bn a year and \$3,079bn in present value. It was a leftover from the pre-53 design. Removing it improves the general-fund ledger from -\$19,477bn to -\$16,316bn in present value and does not touch the actuarial balance, which sits in a different ledger.

One genuine second-order effect remains and is deliberately not charged. Repealing the Social Security tier drops benefit inclusion from up to 85 percent to up to 35 percent, so Medicare's surviving band sits lower in the taxpayer's bracket stack and yields modestly less — on the order of \$1 to \$2bn a year. Pricing it properly requires the provisional-income distribution the model does not have, and the amount is immaterial against a \$2.0tn revenue base. It is held at zero in a named constant rather than omitted silently, so a reader can see the decision. Four errors were corrected in the presentation of results. The fund stake per recipient in 2100 was printing \$305 because two pointers still referenced a cell that version 54 had repurposed to hold average interest saved; the correct figure is \$229,666. The share of beneficiary households reached by the floor was dividing a beneficiary-based household count by a persons-per-household figure, printing 36.6 percent where 32.4 percent belongs — which is also the figure that matches the published benchmark that roughly one in three of those aged 65 and over live below twice the poverty guideline. Sources 9 through 11 were computed and then omitted from the projection total, with the transition bridge silently standing in for them, so the projection was still collecting \$159bn of an abolished payroll tax in 2100; the total was unaffected but the attribution was wrong. And the distribution table double-counted the employer share while omitting four sources, reporting the plan raising \$70bn less than current law on the same tab where another cell reported \$607bn more.

The general-fund interest credit was also restated. It previously added interest saved in 2060, in nominal dollars, to an annual cost stated in 2026 dollars, which flattered the plan by roughly a factor of two. Both are now 75-year present values on the same discount factors as every other figure on the scorecard. On that basis the interest no longer paid on retired debt is \$1,832bn and covers 11.2 percent of the plan's general-fund cost. (Restated at v59: the \$6,986bn stated here through v58 predated the v56 rebate, which cut debt retired from \$38.6tn to \$7.7tn and the interest with it.) It does not eliminate it, and the paper does not claim it does.

Two Control Panel levers were found to drive no calculation and are now labeled display-only, and one shared constant was doing duty as an unrelated phase-in weight. The headline set after all of this, as at v56: the balance after commitments is +0.61 percent of taxable payroll, an improvement of 5.03 percent on the published deficit, closing 113.8 percent of the gap; \$8.9tn of publicly-held debt retired and a \$23.1tn fund by 2100 in real terms; the floor met in all 75 years and the fund never depleted.

The levy rebate, and why it was retired

Version 56 added a rebate: the annual fund draw returned to payers as a credit against the return they already file, so that the levy a payer actually bore fell every year while the headline rate stayed at 5.95%. It was scoped to people, US business and non-profits — estates, foreign holders and government enterprises paid in and received no credit — and it ranked below the Senior Income Supplement, so the floor was never at risk from it. It was the most distinctive single claim in the design: a tax whose effective rate declines by construction rather than by promise.

Version 62 retired it. RebateSwitch is set to 0 and every rebate row in the model evaluates to zero; the effective levy rate is 5.95% in 2100 and 5.95% on a present-value basis, flat throughout. This paper states the trade rather than burying it. Routing the annual draw to payers cost reach: with the draw retained in the pool, the fund reaches \$16.6tn by 2100 and \$6.6tn of publicly-held debt is retired (v91), against a materially smaller result when the draw was being paid out. The self-amortizing rate decline was given up for a larger fund, more debt retired, and a retained surplus that functions as protection against a black-swan event, with supplementary distributions contemplated once the reserve passes a threshold.

What is claimed now is therefore narrower and easier to check: one rate, 5.95%, in statute and in fact, for all 75 years. There is no declining-rate mechanism to defend and no rebate arithmetic to audit.

The precedent argument that accompanied the rebate is recorded here because it may be wanted again if the provision is ever revived. Most US states already select their unemployment insurance tax rate from a schedule triggered by the trust fund balance; Montana enacted House Bill 210 in 2025 creating an automatic trigger to cut rates when the fund exceeds 2.8 percent of total wages, and Missouri and Delaware run comparable balance-triggered methodologies. Roughly forty legislatures have passed a version of this.

The distinction that made the rebate different from those schemes also stands on the record. A state unemployment fund is a RESERVE: it rises and falls with the business cycle, so its rate reductions are cyclical and reverse in the next recession. The rebate rested on an ENDOWMENT compounding above the growth of its own base, so the reduction would have been secular rather than cyclical. That was the claim, and it is no longer being made.

What the rebate was, and what it was not

It was a structural claim, not a distributional one, and it was never sold as the latter. The credit was pro rata to levy paid, so of the eligible base 51.2 percent sat with capital and business and 48.4 percent with labor. At the

2100 credit rate a worker earning \$60,000 would have received about \$528 a year while a holder of \$600,000 of capital income would have received about \$5,282. That is proportional, which is coherent for a proportional levy, but it is not progressive and this paper never claimed it was.

The progressive content of this proposal lies where it always has: the wage levy falling from 12.40 percent to 5.95 percent, and the Senior Income Supplement lifting 26.0 million people to the floor with 6.5 million taken out of poverty outright. Retiring the rebate removed a distinctive claim; it removed nothing progressive.

For the record, the rebate was small where it was measurable: it returned 0.83 percent of what payers handed over in the first decade, and 5.1 percent in present value across the 75 years. Its value was that the rate declined without further legislation, not that the relief was large — which is also why giving it up cost less than the headline suggests.

The real fund return assumption

The 7.25 percent real return is an author estimate against a historical geometric real return of about 6.8 percent for US equities (Damodaran, NYU Stern, 1928-2025). The basis for the premium is the rising capital share from automation, modeled on the AI scenario tab: a shift of income from labor to capital raises the return to capital directly. Arguments from the fund's own market impact are deliberately NOT relied on. A large price-insensitive buyer raises asset prices and therefore lowers expected forward returns, and a lower long-term interest rate lowers the required return on equities for the same reason. Both raise valuations once rather than raising returns permanently, and citing them would be an error. The sensitivity analysis shows the actuarial balance is identical at 7.00 and 7.25 percent, so nothing that determines solvency rests on this figure.

Who actually pays: the employer share by sector

Every employer in the United States stops remitting the 6.2% employer payroll contribution. Only some of them pick up a levy on a base of their own. That asymmetry decides who gains and who pays, and it is not the split the debate assumes.

The decision rule is a single ratio. The employer contribution is 6.2% of taxable payroll; the levy is 5.95% of the base. A payer therefore remits more only if its levied base exceeds 1.042 times its taxable wage payroll. Below that ratio it remits less. The denominator must be taxable payroll rather than the full wage bill, because the employer contribution stops at the 2026 cap of \$184,500 while the levy does not; using full wages understates the ratio for any sector with above-cap earners and reverses the corporate verdict.

On that basis the corporate and private business sector runs a ratio of about 1.28 and remits roughly \$141bn a year more, once the Base 5 charge on non-health benefits is included. The non-profit sector runs a ratio of about 0.27 and remits roughly \$43bn a year less. State and local government has almost no levied base — a ratio of about 0.07, from the operating surplus of its enterprises — and stops writing a cheque worth roughly \$58.1bn a year net. The federal government saves about \$22.4bn, which is a transfer from itself to itself and is carried as a memo only.

The state and local figure is the largest and least discussed. Around 20 million people work for states, counties, cities, school districts, police and fire departments and transit authorities. Their employers pay the employer contribution today and would pay essentially nothing under the levy, because government has no cash flow, no operating surplus, no inheritances and no rent. It is close to a pure fiscal transfer to the level of government most constrained by pension obligations.

The non-profit result is narrower than it first appears, and narrower in the sector's favor. Non-profit net investment income was removed from Source 7 during the v48 audit, because dividends are already levied at the paying company under Base 3 and interest at the recipient under Base 4; counting them again at the institution would have been double-levying. Endowments are therefore untouched at the institution level. A university or a teaching hospital is exposed only on its operating surplus, while receiving the full payroll saving on a very large and very labor-intensive staff.

The self-employed are a separate case and a larger population. Roughly 16 million people remit the whole 12.40% under SECA today, both halves themselves. Under the levy their earnings are reached once at 5.95%, with their business return reached through Base 3.

An earlier version of this model netted the entire economy-wide employer share against the corporate levy alone, which credited corporations with the contribution that non-profits, governments and households actually remit. That comparison read as a \$11.4bn reduction in corporate remittance; corrected, it is an increase of about \$110bn before Base 5 and \$141bn after. The aggregate was never affected — the WHO PAYS total has always reconciled to 2026 OASDI receipts of \$1,368bn — but the distributional claim was wrong in the direction that flattered the proposal, and it is corrected here.

The sector wage bills are no longer author estimates. State and local and non-profit wages are BEA wage and salary accruals (Constants C103:C104, v70), federal wages are derived from the BEA civilian and military series (C105), and the remaining sector splits are derived at Constants rows 138-146. Nothing else in the model depends on them: replacing them changes who the SectorIncidence tab says wins, and no other result.

NOTE ON THE VARIANT TABLES ABOVE. The floor-multiple table and the revenue-estimate table were both re-run in full on the v81 workbook and supersede the versions carried through v66, which had been computed on earlier revenue bases and an earlier floor structure and were no longer monotonic. Central case at v81: +0.61 percent of taxable payroll after all commitments.