# A checked federal income-tax schedule in Lean

**Matha Labs | Research prototype v0.1.0 | November 15, 2026**

## Abstract

This paper presents an executable formal model of a bounded tax calculation: the 2026 federal ordinary taxable-income rate schedules for Single and Married Filing Jointly. Official threshold constants are transcribed from Revenue Procedure 2025-32 and cross-checked against IRS Publication 505. A certificate gate requires the selected source, revision, tax year, filing status, and reviewed scope. The arithmetic uses natural-number cents, an exact integer numerator, and one explicit final rounding operation. Fourteen named theorems establish acceptance preconditions, monotonicity, bounds, and a local continuity property. Thirty synthetic regression fixtures check schedule bases, representative amounts, rounding boundaries, and rejection cases. The contribution is an inspectable schedule illustration, not a complete Internal Revenue Code model or a Form 1040 computation.

## 1. A manageable first formalization

Tax law connects definitions, elections, classifications, phaseouts, procedural rules, exceptions, and annual changes. A useful first implementation should identify a small executable rule and state its dependencies clearly. Here the input is already reviewed ordinary taxable income. The model begins after income classification, adjustments, deductions, and filing-status determination have been completed elsewhere.

This boundary creates a meaningful experiment without pretending that a bracket table solves a tax return. A future system could propose facts from records and attach reviewed evidence, while a deterministic kernel checks their relationships and computes the represented rule. A model's confidence would not replace the scope checks. The present repository includes no learned model, document extraction, taxpayer records, filing integration, or automatic policy-update process.

## 2. Primary-source selection

The source for the 2026 constants is Revenue Procedure 2025-32, section 4.01, Tables 1 and 3 [1]. The IRS also publishes that procedure in Internal Revenue Bulletin 2025-45 [2]. Publication 505 provides an additional IRS cross-check for the 2026 rate schedules and projected-tax worksheets [3]. These sources were reviewed on November 15, 2026.

The statutory context is 26 U.S.C. section 1, including subsection (j) [4]. Reading only the historical dollar tables near the beginning of that section would produce the wrong annual threshold constants. A source-selection layer must connect statutory authority with the applicable inflation-adjusted publication and tax year. This project records that relationship through compiled identifiers and a source manifest, rather than attempting a general legal applicability engine.

The represented statuses are Single and Married Filing Jointly. The latter uses the joint schedule only; this prototype does not determine eligibility for joint filing or extend the implementation to qualifying surviving spouses. Other statuses deliberately require review, even if a published schedule happens to share some constants with a supported one.

## 3. Reviewed input and certificate

`ReviewedFacts` contains the tax year, status, ordinary taxable income in cents, a review attestation, and flags describing whether the case is ordinary-only, involves credits, requires AMT, or needs a special computation. The income is a natural number, so negative values cannot be represented. An ingestion adapter would need to reject invalid numeric representations explicitly rather than silently converting a loss or an unknown value to zero.

`Certificate` supplies a source URL, policy revision, statute identifier, tax year, and status. `Conditions` requires exact agreement with the compiled source and revision, the represented statute identifier, tax year 2026, matching supported status, positive review and ordinary-only attestations, and negative flags for credits, AMT, and special computations. `verify` decides these conditions; `evaluate` returns an optional amount only after successful verification.

The policy and bracket constants are code-owned. A proposal cannot replace them by supplying a different table. Source-string equality establishes consistency with the selected bundle, not authentication of the IRS publication or independent verification of the taxpayer's facts. Review flags likewise record external attestations; setting a Boolean is not itself a professional review procedure.

## 4. Exact arithmetic representation

Let x be ordinary taxable income in cents and let t1 through t6 be the six increasing threshold amounts in cents. Natural-number subtraction is truncated at zero. The model calculates the exact numerator:

```
N(x) = 10*x
     +  2*(x-t1) + 10*(x-t2) + 2*(x-t3)
     +  8*(x-t4) +  3*(x-t5) + 2*(x-t6)
```

The coefficients are increases in marginal percentage points. Once the first threshold is exceeded the total marginal rate becomes 12%; the subsequent totals become 22%, 24%, 32%, 35%, and 37%. Below each threshold its excess term is zero. This representation is equivalent to summing the occupied marginal bands, while making monotonicity particularly straightforward to prove.

The numerator's unit is one hundredth of a cent. `illustrationCents` divides N by 100 once, using natural-number division. Consequently, the implemented convention floors the final result to cents after exact accumulation. It does not round each marginal contribution separately and does not use binary floating point. This convention is chosen for the research artifact; it is not asserted to be an IRS filing-rounding instruction.

## 5. Examples and bracket boundaries

For reviewed ordinary taxable income of exactly $100,000, the schedule illustration is $16,712.00 for Single and $11,504.00 for Married Filing Jointly. These values depend on taxable income rather than gross wages. No deduction should be subtracted again merely because the example began with an amount named income.

The fixture suite also checks every cumulative bracket base in both supported schedules. At a threshold, the newly activated excess term is zero. Raising income by one cent activates only the additional marginal contribution rather than repricing all prior income at the higher rate. The universal one-cent bound formalizes that behavior after the chosen final rounding. It prevents a discontinuous whole-cent jump greater than the one-cent income change anywhere in the represented schedules.

At extremely small inputs, the final floor can yield zero. At one cent of income the exact first-band contribution is one tenth of a cent, so the illustrated output is zero cents. These cases document the implementation's numeric contract and are not tax-return instructions.

## 6. Theorem inventory

Fourteen named theorems are checked in `IRSTax/Model.lean`:

- `verification_iff` connects Boolean success with the complete gate proposition.
- `accepted_source_matches`, `accepted_year_is_2026`, and `accepted_status_matches` establish identity and applicability requirements.
- `accepted_requires_review` and `accepted_is_ordinary_only` establish reviewed-scope requirements.
- `unsupported_status_refused` and `rejected_has_no_illustration` establish refusal behavior.
- `numerator_monotone` and `illustration_monotone` prove that increasing taxable income cannot decrease either the exact numerator or the rounded illustration.
- `numerator_bounded_by_top_rate` proves that N(x) is at most 37 times x, and `illustration_le_income` proves the illustrated amount is at most income.
- `zero_income_zero_illustration` proves the zero-input case.
- `one_cent_increase_bounded` proves that one additional cent of input increases the illustration by at most one cent.

The arithmetic results quantify over all natural-number inputs and threshold structures. The concrete official constants are separately checked by fixtures. The project uses Lean 4.35.0 and its bundled standard library, with no proof placeholders, unsafe declarations, admitted results, or project-specific axioms.

## 7. Regression evaluation

The executable reports thirty named fixtures. Positive cases cover both accepted statuses, zero income, the $100,000 examples, a higher-income example, all Single bracket bases, the final joint base, and the complete joint base list. Cent-level cases test immediately below and above the first threshold, a final-floor boundary, and a sub-cent contribution.

Negative cases change the source, revision, statute identifier, fact tax year, certificate tax year, status agreement, supported status, review attestation, ordinary-only flag, credit flag, AMT flag, and special-computation flag. The entire fixture list is also evaluated through a compile-time `example`. An additional compile-time check covers all six Single bracket bases.

The pinned toolchain successfully built the project and the executable reported 30/30 passing fixtures. These are authored regression cases and transcribed schedule checks, not an external tax-preparation benchmark or evidence of complete return correctness. The universal proofs establish properties beyond the examples, but

only within the stated model.

## 8. Why the output is an illustration

The currently published 2025 Form 1040 instructions distinguish the Tax Table for taxable income below $100,000 from the Tax Computation Worksheet at or above that amount, with exceptions requiring other methods [5]. This older-year source is cited to explain the difference between rate schedules and return mechanics, not to supply 2026 constants. This project implements neither those tax tables nor the complete 2026 filing workflow.

Capital gains, qualified dividends, credits, alternative minimum tax, special worksheets, other federal taxes, withholding, payments, refunds, and state taxation are unimplemented. The model also cannot establish that a taxpayer's ordinary-only attestation is accurate. A filing system must resolve those issues before producing a legally applicable return amount. Deterministic repetition of an incomplete calculation does not remove its omissions.

## 9. Research extensions

The next useful additions would formalize taxable-income derivation and supported filing-status eligibility as separate modules, introduce the applicable IRS tables and worksheets, and connect source snapshots to independently reviewed policy translations. An independent reference implementation should compare amounts and abstentions across boundary cases and held-out reviewed examples. Each added rule should retain its source, tax year, dependency set, and explicit failure path. Those are proposed extensions and have not been delivered by this initial schedule model.

## References

1. IRS. Revenue Procedure 2025-32, section 4.01. https://www.irs.gov/pub/irs-drop/rp-25-32.pdf
2. IRS. Internal Revenue Bulletin 2025-45. https://www.irs.gov/irb/2025-45_IRB
3. IRS. Publication 505 (2026). https://www.irs.gov/publications/p505
4. Office of the Law Revision Counsel. 26 U.S.C. section 1. https://uscode.house.gov/us/usc/t26/s1/
5. IRS. Instructions for Form 1040 (2025). https://www.irs.gov/instructions/i1040gi