



Optimal Wealth Lab (Owl): Mathematical Foundations of a Retirement Planning Optimizer

*A mixed-integer linear programming approach to
US federal tax-aware retirement wealth management*

2026 Edition

Martin-D. Lacasse

March 10, 2026

© 2025–2026 Martin-D. Lacasse

Contents

1	Introduction	1
2	US Retirement Tax Rules and Their Implementation in Owl	3
2.1	Overview	3
2.2	Savings Accounts	3
2.3	Household Financial Profile and Cash Flow	5
2.4	Social Security Benefits	7
2.5	Pension Income	8
2.6	Roth Account Rules	9
2.7	Health Savings Accounts	10
2.8	ACA Marketplace Premiums (Pre-65)	13
2.9	Medicare and IRMAA Premiums	14
2.10	Long-Term Capital Gains Tax	15
2.11	Net Investment Income Tax	16
2.12	Ordinary Income Tax	17
3	Model Architecture and Solution Approach	19
3.1	Overview	19
3.2	Planning Horizon and Time Steps	19
3.3	Account Dynamics	20
3.4	Asset Allocation and Glide Paths	21
3.5	Spending Profile and Objectives	21
3.6	The Self-Consistent Iteration Loop	22
3.7	Filing Status and the Death of a Spouse	23
3.8	Solvers	23
4	Indices, variables, and parameters	24
4.1	Indices	24
4.2	Variables	25

4.3	Parameters	29
4.4	Intermediate variables	38
5	Formulation with imposed asset allocation ratios	42
5.1	Constraints	43
5.1.1	Required minimum distributions (RMDs)	43
5.1.2	Income tax brackets	43
5.1.3	Standard exemption	43
5.1.4	Withdrawal limits	44
5.1.5	Posthumous account activities	44
5.1.6	Roth conversions	44
5.1.7	Minimum taxable balance	45
5.1.8	Initial balances	45
5.1.9	Roth 5-year maturation and conversion ladder	45
5.1.10	Cash-flow surplus	46
5.1.11	Account balances	46
5.1.12	Net spending	49
5.1.13	LTCG bracket constraints	50
5.1.14	Taxable ordinary income	51
5.1.15	Social Security taxability	51
5.1.16	Medicare brackets and costs	53
5.1.17	ACA marketplace premium constraints	57
6	Mapping of decision variables	59
6.1	Reverse mapping of indices	63
7	Building constraint matrices	64
7.1	Inequality constraints	64
7.2	Equality constraints	68
7.3	Other considerations	72
8	Objective Functions and Their Formulation	76
8.1	Maximizing net spending	76
8.2	Maximizing bequest	78
8.3	Lexicographic tie-breaking	78
9	Rate models for return scenarios	79
9.1	Historical data	79
9.2	Constant rate methods	80

9.3	Deterministic varying rates: <code>historical</code>	80
9.4	Stochastic models	81
9.4.1	Multivariate normal: <code>histogaussian</code> and <code>gaussian</code>	81
9.4.2	Log-normal models: <code>lognormal</code> and <code>histolognormal</code>	81
9.4.3	Bootstrap sequence-of-returns: <code>bootstrap_sor</code>	83
9.4.4	Vector autoregression: <code>var</code>	84
9.4.5	DCC-GARCH(1,1): <code>garch_dcc</code>	85
9.5	User-supplied rates: <code>dataframe</code>	86
9.6	Monte Carlo and stress testing	87
A	Source File Reference	88

Abstract

This document describes the mathematical foundations of Owl (Optimal Wealth Lab), a Python application that optimizes US retirement planning using mixed-integer linear programming. Owl models federal tax laws, Medicare and ACA marketplace premiums, Social Security rules, pension income, Roth conversions, Health Savings Accounts, and account dynamics jointly within a single optimization framework. The self-consistent iteration loop and MIP reformulations handle nonlinear tax interactions. Chapters 1–3 present the conceptual model for a general audience; Chapters 4–9 provide the formal mathematical formulation, constraint structure, and rate models for implementers and extenders.

1. Introduction

This document describes the mathematical model underlying the optimization algorithms implemented in Owl, which is a Python application optimizing retirement planning using mixed-integer linear programming. Owl is designed for US retirees as it considers US federal tax laws, Medicare premiums, rules for 401k including required minimum distributions, maturation rules for Roth accounts and conversions, Social Security rules, etc. The goal of these calculations is to optimize the financial aspects of retirement planning, considering the types of savings accounts, federal income tax, contributions, return rates, Medicare premiums, Roth conversions, and desired income among many other things.

The approach is described here mathematically and the Python implementation follows the structure and notation presented in this document. The intent of this document is to provide a guide to the source code for any individual desiring to extend the model to other cases. The mathematical description is not tied to any specific mathematical programming language so that it remains as generic as possible and can be solved using different linear solvers. The source code, installation instructions, and usage examples are available on the project’s GitHub repository [1].

This document is intended for two audiences. Chapters 1 through 3 require only basic financial literacy: familiarity with retirement account types, tax-deferred vs. tax-free savings, Social Security, and Medicare is sufficient to follow the discussion. Chapters 4 through 9 require mathematical maturity: comfort with linear algebra, index notation, and the basics of linear and mixed-integer programming is assumed. Chapter 9 additionally assumes familiarity with probability distributions and time-series models. Readers interested only in the conceptual model may stop after Chapter 3; readers seeking to extend or verify the implementation will need the later chapters.

The remainder of this document is organized as follows. Chapter 2 surveys the US retirement tax rules that Owl models—Social Security benefits, pension income, Roth account rules, ACA marketplace premiums, Medicare premiums, capital gains tax, the Net Investment Income Tax, and ordinary income tax—with notes on how

each is implemented and where known limitations apply. Chapter 3 describes the overall model architecture: how account balances evolve over time, how asset allocation glide paths are specified, how spending objectives are formulated, and how the self-consistent iteration loop resolves the nonlinear tax interactions that arise within a linear programming framework. Chapters 4 through 7 present the formal mathematical formulation: the indices, variables, and parameters; the constraint structure; the mapping of decision variables into the solver’s one-dimensional array; and the construction of the constraint matrices. Chapter 8 defines the objective functions. Chapter 9 describes the rate models available for generating annual return sequences, from simple fixed-rate assumptions to fully stochastic simulations. Appendix A provides a reference guide to the Python source files.

2. US Retirement Tax Rules and Their Implementation in Owl

2.1 Overview

Retirement income in the United States comes from four broad account types: taxable brokerage accounts, tax-deferred accounts such as traditional 401(k)s and IRAs, tax-free Roth accounts, and Health Savings Accounts (HSAs). Withdrawals from these accounts interact with multiple simultaneously computed taxes and premiums: ordinary income tax (with bracket stacking), long-term capital gains (LTCG) tax, the Net Investment Income Tax (NIIT), ACA Premium Tax Credit, Medicare Part B premium surcharges, and the taxability fraction of Social Security (SS) benefits. Because all these quantities depend on the same underlying income variables, optimizing any one strategy in isolation—such as Roth conversion timing or Social Security claiming age—can inadvertently worsen another.

Owl models all of these rules *jointly* within a single linear program (LP) or mixed-integer program (MIP) so that strategies such as Roth conversions, SS claiming age, and withdrawal order are optimized together rather than sequentially. Many rules are nonlinear (piecewise, min/max); Owl handles them within a linear framework using a self-consistent (SC) iteration loop for quantities that can be lagged by one iteration, and MIP reformulations for quantities that require exact simultaneous computation.

2.2 Savings Accounts

US retirement savings are held in four fundamentally different account types, each with its own tax treatment, contribution rules, and withdrawal constraints. Owl models all four types and optimizes across them jointly. For married couples, each spouse holds their own independent set of accounts; balances, contribution histories, claiming ages, and Required Minimum Distribution (RMD) schedules are tracked

separately for each individual while the optimizer considers the combined household portfolio.

Taxable brokerage accounts

Taxable accounts hold securities purchased with after-tax dollars. Dividends and interest are taxed as ordinary income (or at preferential qualified-dividend rates) in the year received, regardless of whether the cash is withdrawn. Capital gains are recognized and taxed only when securities are sold. There is no tax deduction for contributions, no contribution limit from a federal tax standpoint, and no RMD requirement.

Owl models the taxable account as generating a dividend and interest yield each year proportional to the balance. Capital gains arise when the optimizer elects to sell holdings; all realized gains are treated as long-term (see the limitation noted in Section 2.10).

Tax-deferred accounts (traditional IRA, 401(k), and 403(b))

Tax-deferred accounts accept pre-tax contributions that reduce current taxable income, subject to annual IRS limits. Growth inside the account is not taxed until withdrawal. Every dollar withdrawn is taxed as ordinary income in the year taken. Required Minimum Distributions (RMDs) apply starting at age 70–75 depending on the account owner’s birth year (see Section 2.12). A 10% early withdrawal penalty applies to withdrawals taken before age 59½, subject to IRS exceptions not modeled here.

Owl supports contributions to traditional 401(k) and IRA accounts as entered in the Wages and Contributions table described in Section 2.3. The optimizer determines the optimal annual withdrawal amount from these accounts subject to RMD floors and other constraints. The pro-rata rule (IRS Form 8606) is not modeled: Owl treats the entire tax-deferred balance as pre-tax, so every dollar converted or withdrawn is fully taxable (see the limitation in Section 2.6).

Tax-free Roth accounts (Roth IRA and Roth 401(k))

Roth accounts accept after-tax contributions. Qualified withdrawals of both principal and earnings are entirely tax- and penalty-free. Roth accounts are not subject to RMDs during the original owner’s lifetime. Roth conversions—transfers from a tax-deferred account to a Roth account—are fully taxable as ordinary income in the

year of conversion but produce tax-free growth thereafter. The detailed five-year seasoning and age-59½ rules governing Roth withdrawals and conversions are described in Section 2.6.

Owl supports contributions to Roth 401(k) and Roth IRA accounts, and optimizes Roth conversion amounts annually subject to user-specified constraints and tax-bracket considerations.

Health Savings Accounts (HSA)

HSAs are the only triple tax-advantaged savings vehicle in the US tax code: contributions are pre-tax (reducing ordinary income and MAGI), growth is tax-free, and qualified medical withdrawals are entirely tax-free at any age. After age 65 (Medicare enrollment), non-qualified withdrawals are taxed as ordinary income, effectively making the HSA a second traditional IRA without Required Minimum Distributions. A surviving spouse inherits the account intact, preserving its full tax-advantaged status. Crucially, Medicare Parts B and D premiums, IRMAA surcharges, Medigap premiums, and out-of-pocket medical costs all qualify as HSA withdrawals, so the optimizer can fund these expenses entirely tax-free. Annual contribution limits are set by the IRS (approximately \$4,400 for self-only and \$8,750 for family coverage in 2026) and must cease at Medicare enrollment. See Section 2.7 for the detailed rules and implementation.

2.3 Household Financial Profile and Cash Flow

The primary input to Owl is the **Household Financial Profile (HFP)**, which captures the year-by-year financial data for each individual and the household as a whole. The HFP consists of three tables: the *Wages and Contributions* table (one tab per individual), the *Fixed Assets* table, and the *Debts* table.

Wages and Contributions table

Each individual has their own sheet with the following year-indexed columns:

- *year* — calendar year;
- *anticipated wages* — gross earned income expected for that year;
- *other inc* — other ordinary income not captured elsewhere (optional; defaults to zero);
- *net inv* — rent and trust distributions that count as net investment income for NIIT (optional; defaults to zero);

- *taxable ctrb* — contributions to the taxable brokerage account;
- *401k ctrb* — traditional 401(k) contributions;
- *Roth 401k ctrb* — Roth 401(k) contributions;
- *IRA ctrb* — traditional IRA contributions;
- *Roth IRA ctrb* — Roth IRA contributions;
- *HSA ctrb* — Health Savings Account contributions (optional; defaults to zero; automatically zeroed at Medicare enrollment age; entries past age 65 are ignored);
- *Roth conv* — planned or historical Roth conversion amounts; the five years preceding the plan start date are used to implement the five-year lookback rule (Section 2.6);
- *big-ticket items* — large planned expenses such as a home purchase or college tuition that must be funded from the portfolio in a specific year.

Fixed Assets table

The Fixed Assets table describes assets held outside investment accounts: real estate (including a primary residence), fixed annuities, collectibles, precious metals, or individual stocks. Each row specifies the asset name, type, acquisition year, cost basis, current value, expected return rate, year of disposal, and sales commission. The optimizer accounts for sale proceeds and any resulting capital gains as fixed assets are disposed of during the plan horizon.

Debts table

The Debts table lists outstanding loan and mortgage obligations. Each row specifies the debt name, type (loan or mortgage), start year, term, original principal, and interest rate. Annual debt-service payments enter the cash-flow constraint as fixed outflows.

Cash-flow balance

In every plan year, Owl enforces an annual cash-flow balance. Spending plus all taxes, Medicare premiums, debt payments, and account contributions must equal income from wages, Social Security benefits, account withdrawals, fixed-asset sale proceeds, and other income. Big-ticket expenses shift additional spending into specific years. The optimizer determines the combination of account withdrawals and

Roth conversions—across all account types and both spouses—that satisfies this constraint while maximizing the chosen objective.

2.4 Social Security Benefits

Rules

Monthly Social Security retirement benefits equal the Primary Insurance Amount (PIA) multiplied by an age-based adjustment factor. The Full Retirement Age (FRA) depends on birth year: 65 for those born before 1938; 66 for those born between 1943 and 1954; and 67 for those born in 1960 or later, with two-month-per-year interpolation for the transition cohorts (born 1938–1942 and 1955–1959). A special Social Security Administration (SSA) rule treats individuals born on January 1st as attaining age on the last day of the prior month, shifting their effective FRA by one month. Reference: SSA POMS RS 00615.003 [2].

Benefits claimed before FRA are permanently reduced: by 5/9 of 1% per month for the first 36 months before FRA, and by 5/12 of 1% per month for any additional months beyond that, yielding a minimum of approximately 70% of PIA at age 62 when FRA is 67. Benefits deferred beyond FRA earn delayed retirement credits of 8% per year up to age 70. Reference: SSA POMS RS 00615.015 [2].

The spousal benefit is up to 50% of the higher-earning spouse’s PIA, reduced by the claiming spouse’s own PIA (floored at zero), with an additional reduction of 25/36 of 1% per month for claiming before the spousal FRA (which mirrors the own FRA schedule) and no bonus for deferral beyond FRA.

The survivor benefit is the greater of the deceased spouse’s actual benefit and 82.5% of the deceased’s PIA, reduced if the survivor claims before their own survivor FRA (minimum 71.5% at age 60). Reference: CFR § 404.338 [2].

The fraction of SS benefits subject to federal income tax ranges from 0% to 85% based on “provisional income” (PI), defined as modified adjusted gross income (MAGI) minus half of total SS benefits. Thresholds are \$25,000/\$34,000 (single) and \$32,000/\$44,000 (MFJ), frozen since 1983 and 1994 respectively. Reference: IRS Publication 915 [3].

Owl also accepts a `trim_pct` and `trim_year` to model a projected reduction in SS benefits, for example to reflect a potential trust fund depletion scenario.

Implementation in Owl

SS benefit amounts, FRA schedules, spousal benefits, and survivor benefits are computed from the PIA inputs and each individual's date of birth. SS taxability is computed iteratively using the provisional income formula, updated at each pass until it converges; a more exact method is also available via the option `withSSTaxability="optimize"` [3].

Limitations and assumptions

The SSA family maximum benefit cap (CFR § 404.403) is not modeled; this may overstate benefits when multiple beneficiaries claim on one worker's record. The SSA earnings test is not modeled; users are assumed to be fully retired before their FRA. Survivor benefits are assumed to be claimed in the year of the spouse's death; the strategy of deferring survivor benefits to a later age is not optimized. The Windfall Elimination Provision (WEP) and Government Pension Offset (GPO) were repealed by the Social Security Fairness Act (signed January 5, 2025) and therefore require no modeling.

2.5 Pension Income

Overview

A pension is a fixed periodic payment from an employer-sponsored defined-benefit plan, paid for the life of the retiree. Unlike Social Security, pension amounts are set by the plan sponsor and are not subject to federal phase-in rules; they begin in full at the elected commencement age. Some plans offer cost-of-living adjustments (COLAs) that partially or fully index payments to inflation; many do not.

Implementation in Owl

Each individual may have a monthly pension amount and a commencement age. Owl converts the monthly amount to an annual figure and determines the first plan year in which payments begin, using the individual's birth year and birth month to compute the age reached in each calendar year. In the first year of receipt, the annual amount is prorated by the fraction of the year remaining after commencement. For subsequent years the full annual amount (twelve times the monthly benefit) is credited.

When indexed, the annual pension payment grows with inflation each year so that its purchasing power remains constant in real terms. When not indexed, the nominal dollar amount is fixed for the entire duration of the pension. In both cases no pension income is credited before the commencement year. Pension income enters the cash-flow constraint as ordinary income alongside wages, other income, and Social Security benefits.

Joint-and-survivor option

When a pension holder elects a joint-and-survivor (J&S) option, the surviving spouse receives a fraction (e.g., 50%, 75%, or 100%) of the primary's pension after the primary's death. The user specifies the *actual* monthly amount received (already actuarially reduced if J&S was elected) and the survivor fraction as a number between 0 and 1. Starting in the year after the first spouse passes away, the surviving spouse receives that fraction of the deceased's pension each year for the remainder of their horizon. If the primary's pension is inflation-indexed, the survivor benefit follows the same cost-of-living adjustment, so both the primary's pension (while alive) and the survivor benefit grow at the same rate over time. The resulting survivor benefit is added to the surviving spouse's pension income and enters the cash-flow as ordinary income.

Limitations and assumptions

Early-commencement reductions and late-commencement credits are not modeled.

2.6 Roth Account Rules

Rules

Roth IRA contributions are made with after-tax dollars; qualified withdrawals are entirely tax- and penalty-free. Reference: IRC § 408A [4]. Contribution *principal* may be withdrawn at any time without tax or penalty. Roth *earnings* on contributions require both (a) the account must be at least five years old and (b) the owner must be at least age $59\frac{1}{2}$ (or meet another IRS exception) for a penalty-free withdrawal.

Each Roth *conversion* carries its own independent five-year clock. Conversion principal withdrawn before age $59\frac{1}{2}$ is subject to a 10% early withdrawal penalty if taken within five years of that conversion. Roth conversions are fully taxable

as ordinary income in the year of conversion. Reference: IRC § 408A(d)(3); IRS Publication 590-B [5].

Roth IRA accounts are not subject to Required Minimum Distributions (RMDs) during the original owner’s lifetime. Roth 401(k) RMDs were eliminated by SECURE 2.0 Act § 325 effective for tax years beginning after December 31, 2023. A 10% early withdrawal penalty also applies to tax-deferred account withdrawals before age $59\frac{1}{2}$ (subject to exceptions not modeled here).

Implementation in Owl

The five-year maturation rule is enforced as a set of minimum-balance constraints that retain recent conversions at compounded value, plus gains-only on recent contributions, as a floor on the Roth account balance. The age- $59\frac{1}{2}$ threshold is computed per individual from their birth month. Conversion amounts are added to taxable ordinary income in the year of conversion. Conversion behavior can be further controlled via the options `maxRothConversion`, `noRothConversions`, `startRothConversions`, and `swapRothConverters` [4, 5].

Limitations and assumptions

A single unified five-year lookback is applied conservatively: recent conversions at compounded value plus gains-only on contributions are retained as a minimum balance floor. Contribution principal is not separately tracked, so some valid early withdrawals of contribution principal are conservatively disallowed. The pro-rata rule (IRS Form 8606) is not modeled: Owl treats the entire tax-deferred balance as pre-tax, so every dollar converted or withdrawn is fully taxable ordinary income. Users with significant IRA after-tax (nondeductible) basis should be aware of this limitation. Roth conversions are also disallowed in the last two years of the plan horizon.

2.7 Health Savings Accounts

Rules

A Health Savings Account (HSA) is a tax-advantaged account available to individuals enrolled in a High-Deductible Health Plan (HDHP). HSAs offer a unique triple tax advantage: contributions are pre-tax (above the line, reducing AGI), investment growth is tax-free, and withdrawals used for qualified medical expenses are

tax-free at any age. After age 65, non-qualified withdrawals are taxed as ordinary income—making the HSA functionally equivalent to a traditional IRA for general use, but without Required Minimum Distributions. Reference: IRC § 223 [6]; IRS Publication 969 [7].

Eligibility requires enrollment in a High-Deductible Health Plan (HDHP) and no disqualifying coverage (e.g., Medicare, non-HDHP insurance). Enrollment in any part of Medicare (Part A, B, C, or D) immediately disqualifies one from contributing; existing HSA balances may still be withdrawn for qualified expenses. Annual contribution limits for 2026 are \$4,400 for self-only coverage and \$8,750 for family coverage, with an additional \$1,000 catch-up for individuals age 55 and older (each spouse with family coverage may add \$1,000 to their own account if both are 55+). Contributions must cease once the account owner enrolls in Medicare, which typically occurs at age 65. Reference: IRC § 223(b)(7) [6]; IRS Publication 969 [7].

A surviving spouse who is the named beneficiary inherits the HSA intact and may treat it as their own HSA, preserving all three tax advantages. A non-spouse beneficiary must include the full account balance in taxable income in the year of inheritance. Reference: IRC § 223(f)(8)(B) [6].

Implementation in Owl

Owl models the HSA as a fourth savings account type alongside taxable, tax-deferred, and Roth accounts. Because the dominant retiree use is funding qualified medical expenses, all HSA withdrawals are treated as tax-free (simplified model); see Limitations below. The key modeling choices are as follows.

- *Contributions reduce taxable income.* Annual HSA contributions are deducted from ordinary taxable income, from provisional income used to compute Social Security taxability, and from the modified adjusted gross income (MAGI) used in the two-year Medicare IRMAA lookback. Contributions are zeroed automatically at Medicare enrollment age; HFP entries past age 65 are ignored.
- *Withdrawals are tax-free and fund any expense.* HSA withdrawals enter the annual cash-flow balance at full dollar value, identical to Roth account withdrawals, and are excluded from taxable income. Because Medicare Parts B and D premiums, IRMAA surcharges, Medigap premiums, and other out-of-pocket medical costs are all qualified expenses under IRS Publication 969, the optimizer automatically routes HSA funds to cover these costs tax-free whenever doing so improves the household's after-tax outcome.
- *No Required Minimum Distributions.* Unlike tax-deferred accounts, HSAs are never subject to mandatory annual distributions.

- *Cannot exceed account balance.* As with taxable and Roth accounts, the optimizer may not withdraw more than the current HSA balance in any year.
- *Bequest at heirs' tax rate.* The HSA balance remaining at the end of the plan is valued after the heirs' assumed marginal income tax rate, identical to the treatment of tax-deferred accounts. A surviving spouse inherits the HSA intact with no tax (the full balance transfers to their own HSA), but a non-spouse beneficiary must include the full amount in ordinary income. The reported estate value applies the same discount.
- *Asset allocation.* In account-specific allocation mode, the HSA defaults to the same allocation as the Roth account, reflecting that both grow tax-free. Users may also specify a separate HSA allocation.
- *Spousal transfer at death.* At the death of the first spouse, the full HSA balance transfers to the surviving spouse's HSA account, preserving its tax-advantaged status.

Limitations and assumptions

Withdrawals. All HSA withdrawals are treated as qualified medical expenses and therefore entirely tax-free. Non-medical withdrawals after age 65—which would be taxable as ordinary income—are not separately tracked. Before age 65, non-qualified withdrawals incur ordinary income tax plus a 20% penalty; Owl does not model this penalty because all withdrawals are assumed qualified. This simplification is conservative for the dominant retiree use case: funding Medicare premiums and out-of-pocket health costs. Future work may introduce a more detailed model that distinguishes medical from non-medical withdrawals.

Contributions. The annual IRS contribution limit (\$4,400 self-only and \$8,750 family for 2026) is not enforced as an optimizer constraint; users are responsible for not exceeding it in the *HSA ctrb* column of the Household Financial Profile. The catch-up contribution (\$1,000 for age 55+) is not automatically applied; users aged 55+ must add it manually if desired. Self-only vs. family coverage is not distinguished; the applicable limit is the user's responsibility. The mandatory contribution stop at Medicare enrollment is enforced by zeroing contributions from that year forward. `setHSA()` is invoked either explicitly or automatically when loading from configuration (e.g., TOML). HFP entries in the *HSA ctrb* column for years when the individual is age 65 or older are overwritten with zero. If Medicare enrollment occurs mid-year, the IRS allows prorated contributions for months before enrollment [8]; Owl zeroes the entire year, which is conservative.

Eligibility. Owl assumes the user is HSA-eligible (enrolled in an HDHP with no

disqualifying coverage). HDHP status, disqualifying coverage, and dependent status are not verified.

Bequests. Non-spouse beneficiary inheritance is modeled by applying the heirs' marginal income tax rate to the terminal HSA balance, identical to the treatment of tax-deferred accounts. The simplified qualified-withdrawal model (all HSA withdrawals tax-free during the plan horizon) is not extended to non-qualified withdrawals after age 65; those remain deferred to a future path.

2.8 ACA Marketplace Premiums (Pre-65)

Rules

Before Medicare eligibility at age 65, individuals and couples may obtain health coverage through the Health Insurance Marketplace (Exchange) established by the Affordable Care Act. The Premium Tax Credit (PTC) reduces the net premium cost for Marketplace plans based on household modified adjusted gross income (MAGI) and the Federal Poverty Level (FPL). Reference: IRC § 36B [9].

The applicable percentage—the share of income a household is expected to contribute toward the benchmark plan—varies with the ratio of MAGI to FPL. The IRS publishes the applicable percentage table annually. For 2025, Rev. Proc. 2024-35 provides the table under the Inflation Reduction Act (IRA), which suspends indexing and caps the contribution at 8.5% for income above 400% FPL. Reference: Rev. Proc. 2024-35 [10].

For 2026 and beyond, Rev. Proc. 2025-25 provides the applicable percentage table when the IRA enhanced subsidies have expired: indexing resumes, contribution percentages are higher, and there is no subsidy for income above 400% FPL (the “subsidy cliff” returns). Reference: Rev. Proc. 2025-25 [11].

Federal Poverty Level guidelines are published annually by the Department of Health and Human Services (HHS). For 2026, the FPL for the 48 contiguous states and DC is \$15,960 (single) and \$21,640 (couple). Reference: HHS Poverty Guidelines [12].

Household size for ACA purposes is the number of individuals in the tax household who are under age 65 and require coverage (Medicare-eligible individuals are excluded). Below 138% of FPL, individuals in Medicaid expansion states generally qualify for Medicaid rather than Marketplace subsidies.

Implementation in Owl

Owl models ACA net premium costs for pre-65 years when the user supplies the annual benchmark Silver plan (SLCSP) premium via `setACA()` or the `aca_settings` configuration. Two computation modes are available.

In the default `withACA="loop"` mode, ACA costs are computed after each solver pass using `acaCosts()`, which applies the IRS applicable-percentage table to MAGI and FPL for each plan year. Rules are year-aware: plan years before 2026 use the 2025 table; 2026 and later use the 2026 table. Below 138% FPL, Owl returns the full SLCSP (no PTC), reflecting Medicaid eligibility in expansion states.

In `withACA="optimize"` mode, ACA bracket selection and cost are folded into the optimization. The formulation uses 7 MAGI brackets (6 intervals up to 400% FPL, plus one above) with SOS1 binary selection, MAGI decomposition, and a piecewise cost: brackets 0–5 use proportional contribution rates; bracket 6 (above 400% FPL) imposes the full SLCSP. ACA uses current-year MAGI (no two-year lag like Medicare IRMAA).

Limitations and assumptions

The SLCSP benchmark premium is user-supplied and inflated by the plan's inflation factor. FPL and applicable-percentage tables are updated annually in the code; manual updates are required when HHS or IRS publish new values. If Congress extends IRA subsidies (e.g., the 8.5% cap above 400%), parameters must be updated accordingly. Below 138% FPL, Owl assumes Medicaid-eligible and returns full SLCSP; expansion-state rules only.

2.9 Medicare and IRMAA Premiums

Rules

Medicare Part B eligibility begins at age 65. The standard 2026 monthly premium is \$202.90 (\$2,434.80/year). Reference: Centers for Medicare & Medicaid Services (CMS) 2026 Medicare Part B premium announcement [13].

The Income-Related Monthly Adjustment Amount (IRMAA) adds a surcharge based on MAGI from two years prior. In 2026, the five incremental per-person per-month surcharges are \$81.20, \$121.70, \$121.70, \$121.70, and \$40.70, applied at MAGI thresholds of \$109k, \$137k, \$171k, \$205k, and \$500k (single); these thresholds are doubled for married filing jointly (MFJ). Reference: CMS 2026 IRMAA tables [13].

MAGI for IRMAA purposes includes all ordinary taxable income, LTCG, dividends, and the *full* amount of Social Security benefits (not just the taxable fraction). Reference: IRS Publication 915 [3]. For couples, premiums reflect both individuals when both are Medicare-eligible. Filing status transitions from MFJ to single at the first spouse's death.

Implementation in Owl

Owl determines which IRMAA bracket applies to each year's MAGI and computes the exact cumulative premium for that bracket. Two computation modes are available. The default, `withMedicare="loop"`, computes IRMAA costs after each solver pass. Setting `withMedicare="optimize"` folds bracket selection directly into the optimization for an exact result. The two-year MAGI lookback is fully implemented. Because IRMAA in plan year n depends on MAGI from year $n - 2$, the first two plan years require MAGI values from before the plan start. These are supplied by the user via the `previousMAGIs` option as a list of the two MAGI values from two and one years before the plan start. When not supplied, both default to zero, which may underestimate IRMAA surcharges in the first two plan years for higher-income households. Reference: [13].

Limitations and assumptions

Medicare Part D (both base premiums and IRMAA surcharges) is not modeled. CMS 2026 Part D monthly surcharges range from \$14.50 to \$91.00 per person for MAGI above \$109k, resulting in a modest underestimate of Medicare costs for higher-income households. IRMAA thresholds in Owl are assumed to scale with the general inflation rate; in practice, CMS adjusts brackets annually but not necessarily by the Consumer Price Index (CPI).

2.10 Long-Term Capital Gains Tax

Rules

Long-term capital gains (LTCG) and qualified dividends are taxed at preferential federal rates of 0%, 15%, or 20% based on total taxable income (ordinary income plus LTCG combined). Reference: IRS Topic 409 [14]; IRS Publication 550. LTCG “stacks on top of” ordinary income: the portion of gains that pushes combined taxable income above each threshold is taxed at the corresponding higher rate. For

2026, the 0%/15% threshold is \$49,450 (single) or \$98,900 (MFJ); the 15%/20% threshold is \$545,500 (single) or \$613,700 (MFJ). These thresholds are indexed for inflation annually.

Implementation in Owl

LTCG bracket allocation is always active (no option required). Total gains are split across the three rate brackets; because the rates increase (0%, 15%, 20%), the optimizer naturally uses the lowest-rate brackets first. The upper limit of each bracket depends on how much ordinary income already occupies the corresponding threshold; these limits are refined at each solver pass until they converge. The total LTCG tax is derived from the bracket allocation after solving. Reference: [14].

Limitations and assumptions

Short-term capital gains (assets held one year or less) are taxed as ordinary income. Owl does not separately model short-term versus long-term gains within the taxable account; all realized gains are treated as long-term.

2.11 Net Investment Income Tax

Rules

The Net Investment Income Tax (NIIT) is a 3.8% surtax imposed by the Affordable Care Act on net investment income (NII) for taxpayers whose MAGI exceeds \$200,000 (single) or \$250,000 (MFJ). Reference: IRC § 1411 [15]; IRS Form 8960. The thresholds are *not* indexed for inflation. The tax equals 3.8% of the lesser of net investment income and the amount by which MAGI exceeds the applicable threshold. NII includes dividends, interest, LTCG, rents, and royalties, but excludes wages, Social Security benefits, tax-deferred retirement distributions, and Roth withdrawals.

Implementation in Owl

The NIIT is computed annually from MAGI and net investment income (dividends, interest, LTCG, and any rent or trust income entered via the optional `netinv` column; when absent, that column defaults to zero. Rent and trust income from `netinv` is included in NII and also flows into cash-flow and taxable-income as ordinary income. The tax uses the IRS formula described above. Because the NIIT depends on the

solution, it is held fixed during each solver pass and updated from the new income values until it converges. The threshold is fixed in nominal dollars; its real impact over time is correctly reflected in the model. Reference: [15].

Limitations and assumptions

None.

2.12 Ordinary Income Tax

Rules

The US federal income tax is a progressive tax applied to taxable ordinary income (adjusted gross income, or AGI, minus deductions). Reference: IRS Publication 17 [16]. Under the One Big Beautiful Bill Act (OBBBA, signed July 4, 2025), seven marginal brackets apply at rates of 10%, 12%, 22%, 24%, 32%, 35%, and 37%. Standard deductions for 2026 under the OBBBA are \$16,100 (single) and \$32,200 (MFJ). An additional deduction of \$2,000 (single) or \$1,600 (MFJ) per eligible individual applies at age 65 or older.

OBBBA § 1002 provides an extra \$6,000 per-person senior bonus deduction for taxpayers aged 65 or older, phasing out at \$6 per \$100 of MAGI above \$75,000 (single) or \$150,000 (MFJ) and fully phased out at \$175k/\$250k. This bonus expires after tax year 2028. Reference: [17].

Required Minimum Distributions (RMDs) are mandatory annual withdrawals from tax-deferred accounts, beginning at age 70 for those born before 1949, age 72 for those born 1949–1950, age 73 for those born 1951–1959, and age 75 for those born in 1960 or later (effective 2033), per SECURE 2.0 Act § 107 [18]. RMD fractions are taken from the IRS Uniform Lifetime Table III (effective 2022), as published in IRS Publication 590-B [5]. Roth accounts are exempt from RMDs. Filing status transitions from MFJ to single in the calendar year of the first spouse’s death.

Owl can also model a future reversion to pre-OBBBA tax rates (10%, 15%, 25%, 28%, 33%, 35%, 39.6%) from a user-specified calendar year (`yOBBBA` parameter), enabling scenario analysis of potential future tax law changes.

Implementation in Owl

Ordinary income is allocated across the seven tax brackets, with bracket widths and rates updated each year for inflation and for the OBBBA sunset schedule. RMD

fractions are computed per individual and year from the IRS Uniform Lifetime Table. The OBBBA 65+ bonus deduction phaseout is computed from the current MAGI and refined at each solver pass. Reference: [16, 18, 17].

Limitations and assumptions

State income taxes are not modeled. Itemized deductions are not modeled; only the standard deduction (plus the 65+ and OBBBA bonus deductions) is applied. The OBBBA 65+ bonus deduction phaseout is computed outside the LP (as a post-solution update) because it is a nonlinear function of MAGI. The IRS Uniform Lifetime Table III is used for standard cases. When the spouse is the sole designated beneficiary and more than ten years younger than the account owner, IRS Table II (Joint and Last Survivor, Pub. 590-B) is applied instead. Inherited IRA and beneficiary RMD rules are not modeled.

3. Model Architecture and Solution Approach

3.1 Overview

Owl formulates retirement planning as a linear program (LP) or, when binary variables are required, a mixed-integer program (MIP). The plan is divided into annual time steps covering the entire household planning horizon. In each year, the optimizer simultaneously determines withdrawal amounts from all accounts, Roth conversion amounts, and spending levels, subject to tax laws, account balance dynamics, and user-specified constraints. The resulting solution determines the optimal sequence of financial decisions across the full horizon in a single solve—not year by year.

A complication arises because several tax quantities depend nonlinearly on income, yet also influence the cash-flow balance that income must satisfy. Owl resolves this circularity through a self-consistent (SC) iteration loop: nonlinear quantities are fixed at their previous-iteration values during each LP solve, then recomputed from the new solution, and the process repeats until convergence. For quantities that require exact simultaneous treatment, binary variables and MIP reformulations replace the SC loop entirely.

3.2 Planning Horizon and Time Steps

The plan horizon runs from the current calendar year (year 0) through the year following the last individual’s projected death. Each year index corresponds to one calendar year; all decision variables are resolved at annual granularity.

For a single individual, the horizon spans the years from the plan start through the assumed year of death, which is derived from a user-supplied life expectancy or target age. For married couples, the shorter-lived spouse’s death marks the transition from a joint plan to a survivor plan. Both death years are inputs to Owl; they are

treated as certain within a given scenario. Monte Carlo and historical stress-test simulations run the optimizer repeatedly across many return scenarios; stochastic mortality is not modeled.

3.3 Account Dynamics

Each individual maintains up to four savings accounts: a taxable brokerage account, a tax-deferred account, a Roth account, and a Health Savings Account (HSA). A married couple therefore has up to eight accounts in total, tracked and optimized jointly.

At the start of each plan year each account balance grows by the year’s investment return for that account’s asset allocation. During the year the following flows are applied:

- contributions from the Wages and Contributions table are deposited (modeled at mid-year so they earn half a year of return);
- withdrawals are made to fund spending, pay taxes, and cover other expenses including debt service and big-ticket items;
- Roth conversions transfer a chosen amount from the tax-deferred account to the Roth account, adding the converted amount to that year’s ordinary taxable income;
- the taxable account generates a dividend and interest yield proportional to its balance, which is taxed as ordinary income or at preferential qualified-dividend rates;
- Required Minimum Distributions enforce a minimum withdrawal from each tax-deferred account once the owner reaches the applicable RMD starting age;
- any spending surplus—income exceeding expenses in a given year—is deposited back into the taxable account;
- HSA contributions reduce ordinary income (pre-tax) and cease at Medicare enrollment; HSA withdrawals fund qualified medical expenses entirely tax-free and are not included in taxable income;
- at the death of the first spouse, a user-specified fraction of each account is transferred to the surviving spouse’s corresponding account, reflecting beneficiary designations; the surviving spouse inherits the HSA intact at full tax-advantaged status.

The balance at the start of the following year is the result of all of these flows. All accounts are assumed to be rebalanced to their target asset allocation at the end of each year.

3.4 Asset Allocation and Glide Paths

The target allocation of each account across asset classes (equities, corporate bonds, Treasury notes, and cash) is specified by the user. Allocation can be prescribed at three levels of granularity: per individual per account, per individual across all accounts combined, or identically for the entire household.

When the user specifies starting and ending allocation values, Owl interpolates smoothly over the plan horizon using either a linear glide path or an S-curve (hyperbolic tangent) glide path. The S-curve produces a gradual, natural transition that accelerates through the midpoint and tapers at both ends, which is often a more realistic model of how investors de-risk their portfolios over time. Allocation values can also be set to step down at the death of the first spouse, reflecting the survivor's changed risk profile or time horizon.

3.5 Spending Profile and Objectives

Spending profile

Net spending in each year follows a user-specified *spending profile* that scales a common first-year basis. The profile encodes the desired shape of spending over time; the optimizer determines the overall scale—the basis—that is achievable given the household's resources and constraints.

Common profiles include a flat profile (constant real spending throughout retirement), a *smile* profile (higher spending in early and late retirement, lower in the middle years), and a linearly decreasing profile. At the death of the first spouse, spending is multiplied by a user-specified survivorship factor to reflect the reduction in household expenses for a single person.

Primary objectives

Two primary optimization objectives are supported:

Maximize net spending.

Subject to a desired bequest—which may be zero—the optimizer maximizes the sum of inflation-adjusted annual spending over the full plan horizon. This is the most common objective for retirees who want to sustain the highest possible standard of living.

Maximize bequest.

Subject to a fixed annual spending level, the optimizer maximizes the after-tax

estate value delivered to heirs at the end of the horizon. The heir’s marginal income tax rate on inherited tax-deferred balances can be specified by the user.

Lexicographic weight

A small secondary weight can be added to the objective to break ties among solutions with the same primary objective value. By default, this weight mildly discourages unnecessary Roth conversions (preferring fewer conversions when spending is otherwise unaffected) and mildly prefers drawing withdrawals from the first individual before the second when both choices are equivalent. The weight is small enough that it never overrides the primary objective.

3.6 The Self-Consistent Iteration Loop

Several quantities that enter the LP constraints are nonlinear functions of the income variables being optimized. The most important are:

- the *Social Security taxability fraction*—the share of SS benefits subject to ordinary income tax—which depends on provisional income, which in turn includes SS benefits themselves and all other income sources;
- *Medicare IRMAA costs*, which depend on MAGI from two years prior and interact with the current year’s spending and tax calculations;
- the *Net Investment Income Tax (NIIT)*, which depends on MAGI and on investment income;
- the *LTCG bracket upper bounds*, which depend on how much of each threshold is already occupied by ordinary income;
- the *OBBBA 65+ senior bonus deduction phaseout*, which is a nonlinear function of MAGI.

Owl handles these nonlinearities through the SC loop. At each iteration, all nonlinear quantities are held fixed at their values from the previous iteration, making the problem an LP (or MIP). The solver produces a new income solution; the nonlinear quantities are recomputed from that solution and compared to their previous values. If the maximum change across all years falls below a convergence tolerance, the loop terminates; otherwise the updated values are substituted and the process repeats.

To prevent oscillation near the SS taxability thresholds—where small changes in income can flip the taxability fraction back and forth—the updated SS taxability estimate is blended with the previous estimate using a damped update. The loop typically converges in fewer than fifteen iterations for well-posed plans.

Two optional MIP formulations replace portions of the SC loop with exact binary constraints, eliminating the corresponding nonlinearity entirely: `withSSTaxability="optimize"` models SS taxability as an exact piecewise-linear MIP; `withMedicare="optimize"` models IRMAA bracket selection within the optimization itself. These exact formulations are more computationally demanding but guarantee the globally optimal bracket assignment without relying on iterative convergence.

3.7 Filing Status and the Death of a Spouse

For married couples, the year of the first spouse’s death is a known input to the plan. Beginning in that year, Owl automatically transitions all relevant quantities from married-filing-jointly (MFJ) status to single status:

- the standard deduction drops from the MFJ amount to the single-filer amount;
- tax bracket widths narrow to the single-filer schedule;
- Social Security taxability thresholds shift to the single-filer values;
- Medicare premiums reflect only the surviving spouse;
- the survivor begins receiving the survivor benefit computed at the year of the spouse’s death;
- the spending profile is multiplied by the user-specified survivorship factor.

Because the death year is fixed at plan-construction time, this transition requires no binary variables. Separate parameter vectors are constructed for the MFJ period and the single-survivor period; the LP selects the correct values automatically for each year.

3.8 Solvers

Owl’s optimization model is assembled in a solver-agnostic matrix form: inequality constraints, equality constraints, variable bounds, and an objective vector. Any LP or MIP solver that accepts this standard input format can be used.

The default solver is **HiGHS**, a high-performance open-source LP and MIP solver accessed directly via the `highspy` Python package. An optional interface to **MOSEK**, a commercial solver requiring a separate license, is also supported and generally provides faster solution times for large MIP problems.

4. Indices, variables, and parameters

In the next sections, the indices, variables, and parameters are described in detail. Then the model constraints are introduced. For implementation in a linear programming solver, index mapping functions are introduced to map all variables into a single one-dimensional array that is optimized subject to inequality and equality constraints expressed in matrix form. Finally, the constraint matrices are built and so are some useful objective functions.

4.1 Indices

For all indices, we will follow the C array style (starting at 0), rather than the traditional mathematical standard starting at 1. This will facilitate the final sequential mapping of all the variables into a single one-dimensional array, and serve as a direct reference for better understanding the code implementation.

The indices used and their range are defined here, while we also introduce the characteristics and dimensions of the problem. Upper bounds on indices are indicated by the letter N , with the index name as a subscript, e.g., N_i for index i . In other instances, the subscript will indicate that the array depends on subscripts, for example, b_{ikn} is a multidimensional array depending on subscripts i , k , and n .

- i Individual. i runs from 0 to $N_i - 1$ where $N_i = 2$ for couples, or $N_i = 1$ for single individuals. The first individual to pass is denoted by i_d while the survivor is i_s .
- j Type of savings account. j goes from 0 to $N_j - 1$, for taxable ($j = 0$), tax-deferred ($j = 1$), tax-free Roth ($j = 2$), and Health Savings Account ($j = 3$) respectively. Therefore $N_j = 4$.

k	Type of asset class. k goes from 0 to $N_k - 1$, for S&P 500, Baa corporate bonds, Treasury notes, and cash, respectively, and therefore $N_k = 4$. More asset classes could be considered at the cost of increasing the complexity of the problem while not generating many more insights.
n	Index of the year being modeled. The period being modeled runs from the beginning of year 0 to the first day of the year following year $N_n - 1$, and therefore $N_n + 1$ years are actually considered. Year N_n is the first year following the passing of all individuals in the plan. The time period for all decision variables is annual. For spouses, the end of year $n_d - 1$ is when the first individual is assumed to pass while the survivor is assumed to die at the end of year $N_n - 1$ of the plan.
p	Index for long-term capital gains (LTCG) tax rate brackets. $p \in \{0, 1, 2\}$ corresponds to the 0%, 15%, and 20% LTCG rates, respectively, giving $N_p = 3$ brackets.
q	Index for income brackets related to Medicare premium adjustments. q goes from 0 to $N_q - 1$, from low to high. At the time of writing, there are 5 Medicare premium step adjustments in the federal tax code separating $N_q = 6$ income brackets.
N_{aca}, q (ACA con- text)	N_{aca} : number of ACA-eligible plan years ($n < N_{aca}$ when at least one individual is under 65 and within their horizon). When discussing ACA, q denotes the ACA MAGI bracket with $N_q^{aca} = 7$ brackets (6 FPL-based intervals up to 400% FPL, plus one above).
t	Federal income tax bracket. t goes from 0 to $N_t - 1$, from low to high. At the time of writing, there are $N_t = 7$ federal income tax brackets in the current federal tax code.

4.2 Variables

We will use lowercase roman letters to represent variables. This is done to separate variables from parameters, which will be represented using Greek letters or caligraphic fonts. All variables are assumed to take only non-negative values (≥ 0 inequality). Variables are resolved by the optimizer while parameters are prescribed by the user, historical data, or tax laws.

b_{ijn}	Balance for individual i in savings account j at the beginning of year n . When we consider each asset class k , the variable b_{ijkn} is used instead.
d_{in}	Deposit of year- n net spending surplus in taxable account of individual i . These deposits are coming from the surplus s_n , distributed to spousal taxable accounts depending on parameter η .
e_n	Adjusted standard exemption for year n . This is a variable as the taxable income can sometimes be less than the standard exemption $\bar{\sigma}_n$, leading to a negative taxable income if the inflation-adjusted standard exemption is simply subtracted from the gross taxable income G_n in years of low income.
f_{tn}	Amount of taxable ordinary income falling into tax bracket t . The ordinary taxable income is

$$G_n = \sum_t f_{tn}, \quad (4.1)$$

with f_{tn} bound between 0 and the bracket width $\bar{\Delta}_{tn} = \bar{\Gamma}_{tn} - \bar{\Gamma}_{(t-1)n}$. For interpretation, one can introduce the fractional fill f'_{tn} with $f_{tn} := f'_{tn} \bar{\Delta}_{tn}$ and

$$G_n = \sum_t f'_{tn} \bar{\Delta}_{tn}, \quad (4.2)$$

$$0 \leq f'_{tn} \leq 1. \quad (4.3)$$

When considering taxes, the product $\bar{\Delta}_{tn} \theta_{tn}^x$ need not be ordered monotonically, so we cannot rely on the cost structure to fill lower brackets first when minimizing. It is therefore more appropriate to optimize directly in f_{tn} , as the tax rates are progressive. Given that the rates θ_{tn}^x on tax brackets t are increasing monotonically with income, the lower brackets will be filled in first when optimizing. Definitions of θ_{tn}^x , Γ_{tn} and Δ_{tn} are in the section describing the parameters below.

g_n	Net spending in year n .
h_{qn}	Segment of Medicare bracket q in which the modified adjusted gross income (MAGI) in year $n - 2$ is contained. This variable is used to model the income-related monthly adjustment amounts (IRMAA) brackets using

special ordered sets of type 1 (SOS1) for the selection on the bracket indicators (as opposed to using a big- \mathcal{M} method). One of the brackets, say q' , will contain the MAGI from two years before $\mathbb{G}_{(n-2)}$, and therefore

$$h_{qn} = \delta_{q,q'} \mathbb{G}_{(n-2)}, \quad (4.4)$$

where $\delta_{q,q'}$ is the Kronecker delta (1 if $q = q'$, 0 otherwise), such that

$$\mathbb{G}_{(n-2)} = \sum_q h_{qn}. \quad (4.5)$$

The role of h will become clearer as the formulation of the MAGI calculation is exposed below.

m_n	Medicare costs in year n , including Part B premiums and IRMAA income-related adjustments. This does not include Part D premiums and adjustments.
h_{qn}^{aca}	Segment of ACA bracket q in which the current-year modified adjusted gross income (MAGI) \mathbb{G}_n is contained for ACA-eligible year n . Used only when <code>withACA="optimize"</code> ; index n runs over the first N_{aca} plan years when at least one individual is under age 65 and within their horizon. One bracket q' contains MAGI, so $\sum_q h_{qn}^{\text{aca}} = \mathbb{G}_n$. ACA uses <i>current-year</i> MAGI (no two-year lag like Medicare IRMAA).
m_n^{aca}	ACA net premium cost in year n (after Premium Tax Credit). Used only when <code>withACA="optimize"</code> ; zero for $n \geq N_{\text{aca}}$ (Medicare-eligible years). In <code>withACA="loop"</code> mode, ACA cost is computed via the SC loop and treated as a parameter.
s_n	Surplus of funds during year n , most likely caused by required minimum distributions (RMDs) or influx of money from big-ticket items (inheritance, gifts, etc.), or from the disposition of fixed assets (sale of residence, restricted stocks, etc.).
w_{ijn}	Withdrawal from account j belonging to individual i at the beginning of year n . For the ($j = 1$) tax-deferred savings account, w_{i1n} is referred to as a distribution for tax purposes as it is a taxable withdrawal, and will always satisfy required minimum distributions. For the ($j = 3$) HSA, w_{i3n} represents qualified medical withdrawals, which are entirely tax-free and excluded from taxable income (simplified model; see Section 2.7).

x_{in} Roth conversion performed by individual i during year n . These events are taxable as ordinary income. Roth conversions are forced to zero in the last two years of each individual's horizon (see Limitations in Section 2.6).

$p_n^{\sigma, \text{lo}}, p_n^{\sigma, \text{hi}}$ Excess provisional income above the lower and upper Social Security taxability thresholds, respectively:

$$p_n^{\sigma, \text{lo}} = \max(0, \Pi_n - \mathcal{P}_n^{\text{lo}}), \quad (4.6)$$

$$p_n^{\sigma, \text{hi}} = \max(0, \Pi_n - \mathcal{P}_n^{\text{hi}}). \quad (4.7)$$

These continuous non-negative variables are present only when the associated solver option `withSSTaxability="optimize"`. See Eq. (5.33) for the formal constraint formulation.

$p_n^{\sigma, \text{min}}$ 50%-zone Social Security taxability allocation, bounded in $[0, \Delta\mathcal{P}]$, where $\Delta\mathcal{P} := \mathcal{P}_n^{\text{hi}} - \mathcal{P}_n^{\text{lo}}$ is the 50%-zone width (see Parameters):

$$p_n^{\sigma, \text{min}} = \min(\Delta\mathcal{P}, p_n^{\sigma, \text{lo}}). \quad (4.8)$$

Only present when `withSSTaxability="optimize"` is selected. Together with $p_n^{\sigma, \text{lo}}$ and $p_n^{\sigma, \text{hi}}$, all SS taxability LP variables use the letter p for consistency.

q_{pn} Portion of LTCG and qualified dividends allocated to bracket p in year n , where $p = 0, 1, 2$ correspond to the 0%, 15%, and 20% LTCG rate brackets, respectively. These $N_p = 3$ continuous non-negative LP variables are always present and together partition the total LTCG:

$$q_{0n} + q_{1n} + q_{2n} \geq Q_n, \quad (4.9)$$

$$q_{0n} \leq \max(0, \bar{\mathcal{B}}_n^{15} - G_n), \quad (4.10)$$

$$q_{0n} + q_{1n} \leq \max(0, \bar{\mathcal{B}}_n^{20} - G_n), \quad (4.11)$$

where $\bar{\mathcal{B}}_n^{15}$ and $\bar{\mathcal{B}}_n^{20}$ are the inflation-adjusted total taxable income thresholds above which the 15% and 20% rates apply (see Parameters), and $G_n = \sum_t f_{tn}$ is the ordinary taxable income (formally defined in the Intermediate Variables section). Since the cost function is strictly convex ($0 < 0.15 < 0.20$), no binary variables are required: the LP naturally fills the lowest-cost bracket first. See Eqs. (5.27)–(5.28) for the full constraint formulation.

t_n^σ	Taxable Social Security amount in year n , bounded in $[0, 0.85 \sum_i \bar{\zeta}_{in}]$. In the default SC-loop mode this quantity equals $\Psi_n \sum_i \bar{\zeta}_{in}$ and is treated as a fixed parameter; with the solver option <code>withSSTaxability="optimize"</code> it becomes an LP decision variable computed exactly from the IRS formula without iteration.
z_*^*	Binary variables are all designated by the letter z , but a superscript is used to distinguish the different families. For example, z_{qn}^m represents binary variables associated with Medicare calculations, z_{qn}^{aca} with ACA bracket selection (when <code>withACA="optimize"</code>), z_{nz}^x are used for the exclusion constraints, and $z_{0n}^\sigma, z_{1n}^\sigma$ encode the two $\min(\cdot, \cdot)$ operations in the SS taxability MIP formulation (see the Social Security taxability paragraph in the Constraints section).

4.3 Parameters

For more easily distinguishing parameters from variables, all parameters are expressed either in Greek letters or uppercase roman letters using caligraphic fonts. Parameter values are either set by the user, historical data, or by the tax code.

β_{ij}	Initial balances in savings accounts. These amounts are used to initialize b_{ij0} .
τ_{kn}	Annual rate of return for asset class k in year n . A time series of annual return rates for each asset class. Here, inflation and the rate of return of cash ($k = 3$) are assumed to be the same. In other words, investing in cash yields constant dollars as the return perfectly matches inflation, equivalent to a Treasury inflation-protected security (TIPS) having a real yield of 0%.
γ_n	Cumulative inflation at the beginning of year n computed as the product

$$\gamma_n = \prod_{n'=0}^{n-1} (1 + \tau_{3n'}), \quad (4.12)$$

with $\gamma_0 := 1$, and where n' is a dummy index. As the time span of interest goes from the beginning of the first year to the beginning of the year following the last year, variable γ_n will have $N_n + 1$ elements. Parameters indexed for inflation will be indicated by a bar on top as in $\bar{\sigma}_n$. See the next entry for a specific example.

σ_n Standard deduction. It can be adjusted for inflation as follows

$$\bar{\sigma}_n = \sigma_n \gamma_n, \quad (4.13)$$

and can be modified for additional exemptions after age 65, for example. It is a simple time series which can include any foreseeable changes in the tax code, or change in filing status due to the passing of one spouse for $n \geq n_d$. The value of $\bar{\sigma}_n$ is an upper bound for variable e_n .

ξ_n Spending profile. This is a time series that multiplies a basis for the desired net spending amount. It is $\xi_n = 1, \forall n$ for a flat profile, or can be a *smile* profile allowing for more money at the start of retirement and modulating it over retirement. Parameter ξ_n can also contain spending adjustments typically made at the passing of one spouse. The *smile* can be implemented using a cosine superimposed over a gentle linear increase such as in

$$\xi_n = 1 + a_1 * \cos(2n\pi/(N_n - 1)) + a_2 n/(N_n - 1), \quad (4.14)$$

and then normalized by factor $N_n/(\sum_n \xi_n)$ to be sum-neutral with respect to a flat profile. Values of $a_1 = 15\%$ and $a_2 = 12\%$ provide curves that are similar to realistic spending profiles reported in the literature. See Fig. 4.1 for an example. At the passing of one spouse, both profiles are reduced by a factor χ for $n \geq n_d$, and the normalizing factor is adjusted accordingly. The inflation-adjusted profile is $\tilde{\xi}_n = \gamma_n \xi_n$.

χ Factor to reduce spending profile after the passing of one spouse. It is typically assumed to be 0.6. That is, we are assuming that the surviving spouse can live with 60% of the net spending amount that was available to the couple.

n_d Year index at which the first spouse (i_d) is assumed to pass. For single individuals, $n_d = N_n$. For couples, n_d is derived from the life-expectancy inputs so that the end of year $n_d - 1$ coincides with the assumed date of death of i_d . Filing status, tax brackets, Social Security thresholds, and LTCG thresholds all change at $n = n_d$ (from married filing jointly to single).

n_m Year index at which Medicare first applies (i.e., when the youngest individual in the plan turns 65). For $n < n_m$ no Medicare cost m_n is incurred. Used in the Medicare constraints; see the Medicare paragraph in the Constraints section.




Figure 4.1: Example of a spending profile with 15% cosine factor and a 12% linear profile.

- $\mathcal{N}_i^{\text{hsa}}$ Year index at which individual i enrolls in Medicare and HSA contributions must cease (typically when individual i turns 65). For $n \geq \mathcal{N}_i^{\text{hsa}}$, the contribution κ_{i3n} is zero; values in the *HSA ctrb* column for those years are ignored. $\mathcal{N}_i^{\text{hsa}}$ is set by `setHSA()`, which is invoked explicitly or automatically when loading from configuration.
- ρ_{in} Required minimum distribution for individual i in year n . Expressed in fractions determined from IRS tables. The IRS Uniform Lifetime Table III (Pub. 590-B) is used in standard cases. When the spouse is the sole designated beneficiary and more than ten years younger than the account owner, the Joint and Last Survivor Table II is used instead.
- Γ_{tn} Bounds for federal income tax brackets. We define $\Gamma_{(-1)n} := 0$, so that Γ_{0n} is the upper bound for the 10% tax bracket in year n . As the filing status can change for couples, and so can the tax code, Γ_{tn} will be changing over

n . Inflation-adjusted brackets are $\bar{\Gamma}_{tn} = \gamma_n \Gamma_{tn}$.

Δ_{tn} Difference between upper bound Γ_t and lower bound Γ_{t-1} of a federal income tax bracket,

$$\Delta_{tn} = \Gamma_{tn} - \Gamma_{(t-1)n}. \quad (4.15)$$

Inflation-adjusted bracket widths are $\bar{\Delta}_{tn} = \gamma_n \Delta_{tn}$. Once adjusted for inflation, the taxable income can be expressed as in Eq. (4.1). These data are 7 time series. The filing status changes after the passing of one spouse ($n \geq n_d$) and income tax brackets and differences are adjusted accordingly.

θ_{tn}^x Tax rate for ordinary income tax bracket t in year n . Using N_t time series allows one to adjust income tax rates in the foreseeable future. For example, as of 2026 the rates (in decimal) are .10, .12, .22, .24, .32, .35, and .37. While these rates were extended indefinitely by Congress in 2025, they could still revert in the future to 2017 or similar higher rates (.10, .15, .25, .28, .33, .35, and .396). See Eq. (4.31) for its use.

$\mathcal{B}^{15}, \mathcal{B}^{20}$ Base (2026 nominal) total taxable income thresholds above which the 15% and 20% long-term capital gains (LTCG) tax rates apply, respectively. LTCG is taxed at 0% when total taxable income (ordinary + LTCG) is below $\bar{\mathcal{B}}_n^{15}$, at 15% between $\bar{\mathcal{B}}_n^{15}$ and $\bar{\mathcal{B}}_n^{20}$, and at 20% above $\bar{\mathcal{B}}_n^{20}$, where $\bar{\mathcal{B}}_n^{15} = \gamma_n \mathcal{B}^{15}$ and $\bar{\mathcal{B}}_n^{20} = \gamma_n \mathcal{B}^{20}$ are the inflation-adjusted thresholds. In 2026, $(\mathcal{B}^{15}, \mathcal{B}^{20}) = (\$49,450, \$545,500)$ for single filers and $(\$98,900, \$613,700)$ for married filing jointly. Filing status changes at $n = n_d$ for couples. Used in the LTCG bracket constraints; see Eqs. (5.27)–(5.28).

α_{ijkn} Desired asset allocation for savings account j of individual i in asset class k during year n . Allocation ratios come in many flavors as they could be specified globally between individuals and accounts as α_{kn} , for example. When specified by the user, allocation ratios are given two values, one at the beginning of the plan α_{ijk0} and the other at the end $\alpha_{ijkN_{n-1}}$, or α_{ijkn_d} for a spouse passing before the other. Then, intermediate values are interpolated either using a linear relation,

$$\alpha_{ijkn} = a + \frac{n}{N-1}(b-a), \quad (4.16)$$

where N is either n_d or N_n , or using an s-curve as in

$$\alpha_{ijkn} = a + \frac{(b-a)}{2}(\tanh((n-n_1)/n_2) + 1), \quad (4.17)$$

where n_1 is the number of years ahead when the inflection point will occur, and n_2 is the width (in years) of the transition. Constants n_1 and n_2 can be adjusted by the user. Default values are $n_1 = 15$, and $n_2 = 5$, meaning that the transition center will occur in 15 years, taking place from 15 – 5 years to 15 + 5 years from now. Using $a = \alpha_{ijk_0}$ and $b = \alpha_{ijk(N-1)}$ is an approximation as values of ± 1 are only reached at $\pm\infty$ for a hyperbolic tangent. More precise bounds a' and b' for matching the desired start and end values can be determined by solving a 2×2 system of equations leading to

$$\begin{aligned} a' &= (a - k_{12}b')/k_{11} \\ b' &= (b - (k_{21}/k_{11})a)/(k_{22} - (k_{21}/k_{11})k_{12}), \end{aligned} \quad (4.18)$$

where

$$\begin{aligned} k_{11} &= \frac{1}{2}(1 + \tanh(n_1/n_2)) \\ k_{12} &= \frac{1}{2}(1 - \tanh(n_1/n_2)) \\ k_{21} &= \frac{1}{2}(1 - \tanh((N-1-n_1)/n_2)) \\ k_{22} &= \frac{1}{2}(1 + \tanh((N-1-n_1)/n_2)). \end{aligned} \quad (4.19)$$

These interpolation functions allow the allocation ratios to gradually change or *glide* during retirement. Fig. 4.2 provides an example of an *s-curve* gliding allocation ratios.

It is also possible to have a coarser granularity on the portfolio by having an asset allocation scheme defined on a sum of accounts. For example, allocation can be coordinated between accounts leading to α_{ikn} , or even between spouses as α_{kn} . For any of these cases, it is assumed that weights are always properly scaled so that

$$\begin{aligned} \sum_k \alpha_{ijkn} &= 1, \\ \text{or} \quad \sum_k \alpha_{ikn} &= 1, \\ \text{or} \quad \sum_k \alpha_{kn} &= 1, \end{aligned} \quad (4.20)$$




Figure 4.2: Example of an allocation portfolio with 60/40% stocks/bonds transitioning to 70/30% using an s-curve.

depending on the scheme selected.

\mathcal{T}_{ijn} When the allocation ratios α_{ijkn} are prescribed, it is sometimes more convenient to express the return rates as

$$\mathcal{T}_{ijn} = \sum_k \alpha_{ijkn} \tau_{kn}. \quad (4.21)$$

Λ_{in}^{\pm} Big-ticket item requested by individual i in year n . These are large expenses or influx of money that can be planned. Therefore, Λ^{\pm} can be positive (e.g., gift received, inheritance) or negative (e.g., buy a house, large gifts).

λ Allowed deviation from the desired net spending profile during one year. Parameter λ can be better understood as a percentage. If $\lambda = 0.10$, then

the net spending amount is allowed to vary by up to 10% from the prescribed profile. This parameter is mainly provided for educational purposes.

- π_{in} Sum of pension benefits for individual i in year n . These amounts are typically specified along with the ages at which these benefits begin. Pensions can optionally be indexed for inflation and then represented as $\bar{\pi}_{in}$.
- ζ_{in} Social Security benefits for individual i in year n . Starting age and the passing of one individual for spouses will determine the time series. $\bar{\zeta}_{in}$ is the same series adjusted for inflation.
- $\mathcal{P}_n^{\text{lo}}, \mathcal{P}_n^{\text{hi}}, \Delta\mathcal{P}$ Lower and upper provisional income thresholds for Social Security taxability, following IRS Publication 915. Below $\mathcal{P}_n^{\text{lo}}$ no Social Security is taxable; between $\mathcal{P}_n^{\text{lo}}$ and $\mathcal{P}_n^{\text{hi}}$ up to 50% of benefits enter taxable income; above $\mathcal{P}_n^{\text{hi}}$ up to 85% of benefits are taxable. We define $\Delta\mathcal{P} := \mathcal{P}_n^{\text{hi}} - \mathcal{P}_n^{\text{lo}}$ (the 50%-zone width; \$9,000 for single, \$12,000 for MFJ). For single filers $(\mathcal{P}^{\text{lo}}, \mathcal{P}^{\text{hi}}) = (\$25,000, \$34,000)$; for married filing jointly $(\mathcal{P}^{\text{lo}}, \mathcal{P}^{\text{hi}}) = (\$32,000, \$44,000)$. Unlike income-tax brackets, these thresholds are *not* indexed for inflation; the time subscript n reflects only the change in filing status at $n = n_d$.
- \mathbb{G}^{NIIT} Statutory modified adjusted gross income (MAGI) threshold above which the net investment income tax (NIIT) of 3.8% applies. \mathbb{G}^{NIIT} is not indexed for inflation. In 2026, $\mathbb{G}^{\text{NIIT}} = \$200,000$ for single filers and \$250,000 for married filing jointly, with filing status changing at $n = n_d$. See the definition of J_n in the Intermediate Variables section.
- ϵ_{N_n} Desired amount to leave as a bequest at the end of the final year of the plan, $N_n - 1$, which is the beginning of year N_n . This amount is the after-tax value of the estate for the heirs expressed in today's dollars. The inflation-adjusted target is $\bar{\epsilon}_{N_n} = \gamma_{N_n} \epsilon_{N_n}$. See parameter ν for the heirs tax rate.
- κ_{ijn} Sum of contributions to savings account j made by individual i during year n . We assume that contributions are made at half-year to better represent periodic contributions made throughout the year. In practice, a contribution amount κ_{ijn} is specified, in which case the contribution to each asset class is

$$\kappa_{ijkn} = \alpha_{ijkn} \kappa_{ijn}. \quad (4.22)$$

as savings account balances are assumed to be rebalanced periodically. For the ($j = 3$) HSA, contributions κ_{i3n} are pre-tax: they reduce ordinary

taxable income, provisional income for Social Security taxability, and MAGI for Medicare IRMAA purposes in year n . HSA contributions are zeroed for $n \geq \mathcal{N}_i^{\text{hsa}}$, when individual i enrolls in Medicare (typically at age 65).

ω_{in}	Sum of wages earned by individual i during year n . Do not confuse wages ω with withdrawals w .
v_{in}	Other ordinary income for individual i in year n , beyond wages, pension, and Social Security. This is a user-supplied time series (e.g., rental income, alimony) loaded from the household financial plan spreadsheet.
ν_{in}	Net investment income from rent or trust distributions for individual i in year n . User-supplied time series loaded from the <code>netinv</code> column of the household financial plan spreadsheet. It enters cash-flow and taxable-income constraints as ordinary income and is also included in the net investment income total for NIIT computation. Defaults to zero when the column is absent.
\mathcal{A}_n^*	Amounts resulting from the proceeds of liquidation of fixed assets. The liquidation of these fixed assets can generate a taxable income, marked with a superscript x as \mathcal{A}_n^x , long-term capital gains, marked with superscript c as \mathcal{A}_n^c , or tax-free proceeds marked with a superscript f as \mathcal{A}_n^f . These are used to capture the sale of a residence or real estate, cashing restricted stocks, or receiving a lump-sum annuity.
\mathcal{C}_q	Cost of Medicare for bracket q of modified adjusted gross income (MAGI) \mathbb{G} . This includes Medicare Part B premiums and any additional income-related monthly adjusted amount (IRMAA). When adjusted for inflation, this becomes $\bar{\mathcal{C}}_{qn} = \gamma_n \mathcal{C}_q$. There are $N_q = 6$ brackets for IRMAA and therefore 5 step adjustments forming a piecewise constant function.
\mathcal{D}_n	Debt payment in year n to be considered in the cash flow. This allows the cash flow to account for a mortgage or a car loan, for example.
\mathcal{L}_q	Upper bounds for brackets used to determine Medicare adjustments based on the modified adjusted gross income (MAGI) \mathbb{G} . When adjusted for inflation, this becomes $\bar{\mathcal{L}}_{qn} = \gamma_n \mathcal{L}_q$. While there are 6 brackets for IRMAA adjustments forming a piecewise constant function, \mathcal{L}_q has $N_q - 1$ distinct thresholds as the last element is an arbitrarily large number bounding the last bracket. See Fig. 5.1 for a visual representation. Note that $\bar{\mathcal{L}}_{qn}$ needs

to be adjusted for inflation and marital status, including adjustments due to the passing of one spouse.

$\bar{\mathcal{L}}_{qn}^{\text{aca}},$ $\bar{\mathcal{C}}_q^{\text{aca}},$ \bar{S}_n^{aca}	ACA parameters (only when <code>withACA="optimize"</code> and $n < N_{\text{aca}}$). $\bar{\mathcal{L}}_{qn}^{\text{aca}}$: inflation-adjusted FPL bracket thresholds for year n ; $N_q^{\text{aca}} = 7$ brackets (6 FPL-based intervals up to 400% FPL, plus one above). $\bar{\mathcal{C}}_q^{\text{aca}}$: applicable contribution percentage for bracket q ($q = 0, \dots, 5$); contribution is proportional to MAGI within each bracket. \bar{S}_n^{aca} : inflation-adjusted annual benchmark Silver plan (SLCSP) premium; for bracket 6 (above 400% FPL), net cost equals \bar{S}_n^{aca} (no subsidy).
μ	Dividend return rate for equities in taxable accounts. Average is about 1.7% for S&P 500.
ν	Heirs' income tax rate to be applied on the tax-deferred portion of the estate. This is not an estate tax but rather the federal income marginal tax rate that heirs would have to pay on inherited tax-deferred accounts.
ϕ_j	Fraction of savings account j that is left to surviving spouse i_s as a beneficiary at the death of individual i_d , the first spouse to pass.
Ψ_n	Effective fraction of Social Security benefits subject to income tax in year n , bounded in $[0, 0.85]$. It is computed by the self-consistent loop from the IRS provisional income formula (IRS Publication 915) and enters the LP as a fixed parameter each iteration. See the Social Security taxability paragraph in the Constraints section.
η	Spousal ratio for surplus deposits, which goes from 0 to 1, as the fraction that goes to the $i = 1$ spouse's account. Therefore, a surplus s_n in year n will result in a deposit d in the taxable account of individual i as

$$\begin{aligned} d_{0n} &= (1 - \eta)s_n \\ d_{1n} &= \eta s_n. \end{aligned} \tag{4.23}$$

This choice is such that we can set a value depending on the surviving individual $\eta = i_s$ for $n \geq n_d$, after the passing of i_d . Default value is $(N_i - 1)/2$, i.e., 0.5 for couples and 0 for single individuals. When the beneficiary of the savings accounts is not the other spouse, i.e., when $\phi_j \neq 1, \forall j$, it is recommended that η be set to i_d so that all surplus gets deposited to i_d 's accounts, thus avoiding loopholes when optimizing for the final bequest.

\mathcal{M} This large constant is used in the so-called big- \mathcal{M} method to implement binary constraints. As this is mainly used around MAGI, a value of about 10^8 should be adequate for most cases. We also use \mathcal{M} to represent the upper bounds of top tax brackets.

4.4 Intermediate variables

We use intermediate variables for conciseness or clarity, but they are ultimately replaced in the final formulation. Intermediate variables are represented in roman uppercase letters, or in double stroke uppercase letters.

G_n Taxable ordinary income in year n . Proceeds from the liquidation of fixed assets taxable as ordinary income, the sum of wages, other ordinary income, pensions, taxable Social Security benefits, all withdrawals from tax-deferred accounts, including Roth conversions, and gains from securities (i.e., all gains except those from the $k = 0$ equities, which are taxed as capital gains) in the ($j = 0$) taxable account, including contributions κ , minus the standard deduction e_n ,

$$\begin{aligned}
G_n = & \mathcal{A}_n^x + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \Psi_n \bar{\zeta}_{in} + \bar{\pi}_{in}] - e_n \\
& + \sum_i [w_{i1n} + x_{in}] \\
& + \sum_{i,k \neq 0} [(b_{i0n} - w_{i0n} + d_{in} + 0.5\kappa_{i0n})\alpha_{i0kn}\tau_{kn}]
\end{aligned} \tag{4.24}$$

Social Security is indexed for inflation. The taxable portion $\Psi_n \sum_i \bar{\zeta}_{in}$ is determined by the self-consistent loop; see the Social Security taxability paragraph in the Constraints section. Pensions can optionally be indexed for inflation. We exclude $k = 0$ to capture only non-equity gains in taxable accounts. These gains are all taxed as ordinary income. Here, we assume that withdrawals and deposits in the taxable account are taking place at the beginning of the year, while contributions, if any, are taking place in mid-year.

G_n was also already defined in Eq. (4.1) as

$$G_n = \sum_t f_{tn},$$

and equating both equations links these variables together.

Q_n Qualified dividends and long-term capital gains obtained in year n . They only involve dividends occurring in the taxable savings accounts ($j = 0$) that were obtained from equities ($k = 0$), or sales of stocks due to withdrawals from taxable savings accounts. For simplicity, we assume that all equity sales only generate long-term capital gains and that all dividends are qualified, resulting in

$$Q_n = \sum_i \alpha_{i0n} [(b_{i0n} - w_{i0n} + d_{in} + 0.5\kappa_{i0n})\mu + w_{i0n}\max(0, \tau_{0(n-1)} - \mu)]. \quad (4.25)$$

A formulation where only a fraction of dividends are qualified can easily be implemented with the addition of another parameter. Notice that we are using return rates from the previous year. The first terms on the right-hand side represent dividends generated by equities ($k = 0$) in the ($j = 0$) taxable savings account plus half the yearly contributions. The second term accounts for withdrawals w of equities assumed to have been purchased a year ago. Capital gains are calculated as price appreciation only (total return minus dividend rate) to avoid double taxation of dividends, which are already included in the first term. It does not account for losses, but a market drop would most likely result in stock purchase rather than sale. For withdrawals, we make the assumption of selling the most recent stocks which would not be accurate in situations where the taxable savings account is being depleted slowly. An implementation keeping track of stock purchases and sales is beyond the scope of providing a guide for retirement decisions.

\mathbb{G}_n Modified adjusted gross income or MAGI for year n . When non-taxable municipal bonds are ignored, the MAGI can be expressed as

$$\mathbb{G}_n = G_n + Q_n + e_n + (1 - \Psi_n) \sum_i \bar{\zeta}_{in}. \quad (4.26)$$

Because G_n already contains the taxable Social Security term $\Psi_n \sum_i \bar{\zeta}_{in}$, adding back $(1 - \Psi_n) \sum_i \bar{\zeta}_{in}$ restores the full benefit in MAGI, consistent with the IRS definition. This approach ignores additional IRS rules around tax-free interests which are insignificant in most cases.

Π_n Provisional income for Social Security taxability in year n (IRS Publication 915): $\Pi_n = \mathbb{G}_n - \frac{1}{2} \sum_i \bar{\zeta}_{in}$. Used to determine the taxable portion of

Social Security; see Eq. (5.31) and the Social Security taxability paragraph in the Constraints section.

\mathcal{S}_n Amount of Social Security benefits subject to federal income tax in year n . \mathcal{S}_n is computed from the IRS formula in Eqs. (5.32)–(5.35), and $\Psi_n = \mathcal{S}_n / \sum_i \bar{\zeta}_{in}$.

\mathbb{I}_n Interest and dividend income from the taxable account, plus rent and trust income from the `netinv` column. Used as the base for NIIT computation (together with Q_n).

$$\mathbb{I}_n = \max \left(0, \sum_{i,k \neq 0} [(b_{i0n} - w_{i0n} + d_{in} + 0.5\kappa_{i0n})\alpha_{i0kn}\tau_{kn}] \right) + \sum_i \nu_{in}. \quad (4.27)$$

Only positive returns are taxable; the interest/dividend component is clamped to zero before adding ν_{in} .

P_n Amount of 10% early withdrawal penalty in year n . The penalty applies only to tax-deferred ($j = 1$) withdrawals before the IRS threshold of age $59\frac{1}{2}$:

$$P_n = 0.10 \sum_i (1 - \mathcal{H}(n - n_{i,59\frac{1}{2}})) w_{i1n}. \quad (4.28)$$

Here, $\mathcal{H}(n - n_{i,59\frac{1}{2}})$ is a Heaviside step function which is 0 or 1, depending on the sign of its argument:

$$\mathcal{H}(x) := \begin{cases} 0 & x < 0 \\ 1 & x \geq 0. \end{cases} \quad (4.29)$$

The parameter $n_{i,59\frac{1}{2}}$ is the year index when individual i turns $59\frac{1}{2}$, or 0 if the individual is already past $59\frac{1}{2}$ at the beginning of the plan. Because the model operates at yearly granularity:

$$n_{i,59\frac{1}{2}} = \max(0, 59 - Y_0 + y_{\text{obs},i} + \mathbf{1}[m_{\text{obs},i} > 6]), \quad (4.30)$$

where Y_0 is the current calendar year, $y_{\text{obs},i}$ is the birth year of individual i , and $m_{\text{obs},i} \in \{1, \dots, 12\}$ is the birth month. The indicator $\mathbf{1}[m > 6]$ adds one year for individuals born in July–December, whose $59\frac{1}{2}$ birthday falls in January–June of the following year. Tax-free ($j = 2$) Roth withdrawals are not subject to this penalty because the 5-year maturation constraints already restrict withdrawals to penalty-free amounts; see the maturation paragraph below.

T_n Amount of income tax paid on taxable ordinary income G_n in year n . This is the taxes paid on ordinary income expressed as the sum of the amounts paid in each tax bracket as

$$T_n = \sum_t f_{tn} \theta_{tn}^x. \quad (4.31)$$

Notice how f_{tn} also defines G_n in Eq. (4.1), and that optimal values of f_{tn} have to minimize T_n regardless of whether the bequest or the desired net spending is being maximized.

U_n Amount of income tax paid on long-term capital gains (LTCG) and qualified dividends in year n . The LTCG tax is computed exactly by the LP bracket variables q_{pn} :

$$U_n = 0.15 q_{1n} + 0.20 q_{2n}. \quad (4.32)$$

We assume that qualified dividends and long-term capital gains are taxed at the same preferential rate, which is the case for most situations.

J_n Net investment income tax (NIIT) paid in year n . The NIIT is 3.8% and applies when modified adjusted gross income \mathbb{G}_n exceeds the statutory threshold \mathbb{G}^{NIIT} (formally defined in the Parameters section; not indexed for inflation). We compute

$$J_n = 0.038 \max(0, \min(\mathbb{G}_n - \mathbb{G}^{\text{NIIT}}, \mathbb{I}_n + Q_n)), \quad (4.33)$$

where \mathbb{G}^{NIIT} is \$200k for single filers and \$250k for married filing jointly, with filing status changing after the passing of one spouse ($n \geq n_d$).

5. Formulation with imposed asset allocation ratios

We first present the case where the sums of assets in each savings account b_{ijn} are known and for which we assume prescribed asset allocation ratios. The amount in each asset class k for b_{ijkn} is simply obtained from $\alpha_{ijkn}b_{ijn}$ in this case. This formulation assumes that the accounts are always balanced. This is a reasonable assumption given the auto-balancing feature offered by many financial service providers and robo-advisers.

The benefit of this approach is that it has fewer variables and that only the sums of all asset classes in each savings account need to be considered. The rate of return of the account is then simply the product of the account balance with the sum of the rates of return weighted according to the desired allocation ratio. This approach allows us to eliminate k by summing over it and rewrite equations such as

$$\sum_k b_{ijk(n+1)} = \sum_k b_{ijkn}(1 + \tau_{kn}) + \dots, \quad (5.1)$$

for the annual evolution of account balances from year n to year $n + 1$ as the simpler expression

$$\begin{aligned} b_{ij(n+1)} &= b_{ijn} \sum_k \alpha_{ijkn}(1 + \tau_{kn}) + \dots, \\ &= b_{ijn}(1 + \mathcal{T}_{ijn}) + \dots, \end{aligned} \quad (5.2)$$

where

$$\mathcal{T}_{ijn} := \sum_k \alpha_{ijkn} \tau_{kn}. \quad (5.3)$$

It follows that the allocation ratios are normalized to unity, i.e.,

$$\sum_k \alpha_{ijkn} = 1. \quad (5.4)$$

In this formulation where the α_{ijkn} are prescribed, we will use \mathcal{T}_{ijn} to add the market returns to the savings balances.

5.1 Constraints

5.1.1 Required minimum distributions (RMDs)

Withdrawals from the ($j = 1$) tax-deferred savings accounts must be greater than or equal to the required minimum distributions, and therefore,

$$w_{i1n} - \rho_{in} b_{i1n} \geq 0. \quad (5.5)$$

As b_{ijn} are the balances at the beginning of year n , they are also the balances at December 31 of the previous year, which is the amount on which the IRS bases the RMDs. Eq. (5.5) has to hold for each year n and each individual i , and therefore, there are $N_i \times N_n$ such equations (although trivial when $\rho_{in} = 0$). These constraints avoid paying up to 25% penalty on amounts not withdrawn when RMDs are required. Note that aggregate rules need to be considered separately as this approach only considers the sum of assets in a class with similar tax treatment (e.g., IRA and 401k).

5.1.2 Income tax brackets

Taxable ordinary income is divided into tax brackets as defined in Eq. (4.1), and therefore

$$0 \leq f_{tn} \leq \bar{\Delta}_{tn}, \quad (5.6)$$

where $\bar{\Delta}_{tn}$ is the inflation-adjusted income tax bracket width.

5.1.3 Standard exemption

The standard exemption is constrained by

$$0 \leq e_n \leq \bar{\sigma}_n. \quad (5.7)$$

Variable e_n is required for accounting for years when the taxable ordinary income is smaller than the standard exemption.

5.1.4 Withdrawal limits

We introduce another set of constraints that might look unnecessary, but can help convergence, and prevent overdrafts during the year of passing of spouse i_d . As withdrawals and conversions are at the beginning of the year we impose that

$$w_{ijn} + \delta_{j,1}x_{in} \leq b_{ijn}, \quad (5.8)$$

which just states that account balances need to be at least as large as withdrawals and possible Roth conversions.

5.1.5 Posthumous account activities

For cases with spouses, no withdrawal, Roth conversion, or deposit should occur in the accounts of the passed individual i_d :

$$\begin{aligned} w_{i_d j n} &= 0, \\ d_{i_d n} &= 0, \\ x_{i_d n} &= 0, \\ &\quad \forall j \in \{0, \dots, N_j - 1\} \\ &\quad \forall n \in \{n_d, \dots, N_n - 1\}. \end{aligned} \quad (5.9)$$

5.1.6 Roth conversions

Roth conversions x_{in} are forced to zero in the last two years of each individual's horizon (see Section 2.6). Otherwise, Roth conversions cannot be larger than the balance at the beginning of the year in the account:

$$x_{in} \leq b_{i1n}. \quad (5.10)$$

This constraint, however, is naturally satisfied when $b_{ijn} \geq 0$ non-negativity bounds are enforced. Additional maximum Roth conversion constraints x_{max} can be imposed by the user. For a single individual, the previous equation becomes

$$x_{in} \leq \min(b_{i1n}, x_{max}). \quad (5.11)$$

For couples, the cap applies to each individual's conversions:

$$x_{in} \leq x_{max}, \quad i \in \{0, 1\}. \quad (5.12)$$

As these equations involve a variable and a parameter in the min function, they are coded as separate constraints.

5.1.7 Minimum taxable balance

To maintain liquidity throughout the plan, the user may impose a lower bound β_i^{\min} (in today's dollars) on the taxable account balance of each individual. The inflation-adjusted floor is enforced from year 1 through the end of the individual's life horizon:

$$b_{i0n} \geq \beta_i^{\min} \gamma_n, \quad n \in \{1, \dots, \text{horizon}_i - 1\}. \quad (5.13)$$

This is controlled by the `minTaxableBalance` option. The constraint is omitted when `minTaxableBalance` is not specified.

5.1.8 Initial balances

The initial balances β_{ij} are one of the main inputs of the model. The initial savings account balances are imposed through the constraints

$$b_{ij0} = \beta_{ij}. \quad (5.14)$$

At this point, we assume that all accounts are balanced according to the desired allocation ratios α_{ijk0} .

5.1.9 Roth 5-year maturation and conversion ladder

Two IRS rules govern penalty-free access to Roth funds. (1) Roth *contributions* can always be withdrawn tax- and penalty-free. (2) Roth *conversions* can be withdrawn penalty-free only after a 5-year holding period, with each conversion carrying its own clock; *earnings* additionally require age 59½. Because of rule (2), any Roth withdrawal allowed by the maturation constraint is already penalty-free, so no 10% early-withdrawal penalty is assessed on w_{i2n} .

The maturation constraint (row group I_8 in the matrix appendix) enforces this rule:

$$b_{i2n} - w_{i2n} \geq \sum_{k=1}^5 \Gamma_k x_{i,n-k} + \sum_{k=1}^5 (\Gamma_k - 1) \kappa_{i2,n-k}, \quad \Gamma_k = \prod_{\ell=1}^k \max(1, 1 + \mathcal{T}_{i2,n-\ell}), \quad (5.15)$$

forcing immature conversions (those made within the last 5 years) and recent contribution gains to remain in the account. Historical amounts ($n - k < 0$) use a conservative 10%/yr assumption.

This constraint enables the *Roth conversion ladder*, a strategy popular in the FIRE (Financial Independence, Retire Early) community. An early retiree converts

tax-deferred funds in years $n, n+1, \dots$ and, 5 years later, withdraws the corresponding Roth amounts penalty-free. The optimizer discovers this strategy automatically once the Roth penalty is removed and the AMO constraint is relaxed for pre-59½ years (see the AMO discussion below).

5.1.10 Cash-flow surplus

When both spouses are alive, surplus s_n gets deposited in the taxable accounts according to variable η as described in Eq. (4.23),

$$d_{in} = [\delta_{i,0}(1 - \eta) + \delta_{i,1}\eta]s_n. \quad (5.16)$$

Otherwise, for $n \geq n_d$, variable η gets redefined as $\eta = \delta_{1,i_s}$. Surplus can be caused by large influx of money coming from big-ticket items, compulsory RMDs, or the disposition of fixed assets. The `noLateSurplus` option pins surplus to zero in the final two plan years. This is useful because during market downturns, large sums can otherwise be recycled through the taxable account in those years with little tax or growth consequence, since the long-term compounding benefit of such deposits vanishes near the end of the plan.

5.1.11 Account balances

Contributions are assumed to be made at half-year to better represent periodic contributions made throughout the year. As we already mentioned, the account balance at the end of a year is the same as the balance at the beginning of the following year. Changes include contributions κ , distributions and withdrawals w , conversions x , surplus deposits d , and growth τ on the account through the year. For each spouse i , we track each savings account j separately, and tax-deferred accounts are coupled with the corresponding tax-free account through Roth conversions.

The timing of Roth conversions, withdrawals, and deposits brings additional coupling between these variables, and is worth a detailed discussion. First, the financial aspects, and then the algorithmic ones. For the former, some financial advisors would recommend making Roth conversions at the beginning of the year, while making withdrawals at the end. Obviously, financial simulators would always yield higher numbers when using this scenario, as the money needed to pay the regular bills stayed in the bank until the end of the year. More realistically, however, it would be more accurate to assume withdrawals at mid-year, to better represent evenly distributed withdrawals. So, financially, conversions at the beginning of the year, and withdrawals at mid-year make good sense. Conversions are also typically best when

timed with market downturns, which are obviously not always at the beginning of the year.

Now, let's look at the optimization side of these transactions. During years of positive returns, a direct withdrawal from the tax-deferred account at mid-year will always be unfavorable when compared to a Roth conversion at the beginning of the year, followed by a tax-free withdrawal later in the same year. This is because the second scenario involves gains which are tax-free over the half-year, while the first one does not. Moving account withdrawals at the beginning of the year, and the conversions in mid-year can only solve part of this artificial bias under specific conditions.

Moving all transactions at the beginning of the year can also cause undesired transactions. To solve these spurious scenarios, it would be desirable to make the following at-most-one (AMO) exclusions between surplus s_n , withdrawals w_{ijn} , and conversions x_{in} :

$$\left(\sum_i w_{i0n} + \sum_i w_{i2n} \right) \text{ AMO } s_n, \\ \sum_i x_{in} \text{ AMO } \sum_i w_{i2n},$$

for $j \neq 1$, i.e., for all withdrawals except those from tax-deferred accounts. The AMO relation can also be represented by a logical NAND operator. The first exclusion prevents surplus deposits when withdrawals from taxable or tax-free accounts occur from either spouse, while the second exclusion prevents simultaneous Roth conversions and tax-free withdrawals from either spouse. To favor tax-deferred withdrawals in most reasonable situations, we implement these exclusions by introducing binary variables $z_{n\ell}^x \in \{0, 1\}$ with $\ell \in \{0, 1, 2, 3\}$ for each year n (shared across both spouses). The first pair of binary variables (z_{n0}^x, z_{n1}^x) enforces the surplus-withdrawal exclusion:

$$\begin{aligned} \epsilon z_{n0}^x &\leq \sum_i (w_{i0n} + w_{i2n}) \leq \mathcal{M} z_{n0}^x, \\ \epsilon z_{n1}^x &\leq s_n \leq \mathcal{M} z_{n1}^x, \\ 0 &\leq z_{n0}^x + z_{n1}^x \leq 1, \end{aligned} \tag{5.17}$$

where ϵ is a small positive number (e.g., 0.001) to ensure that when a binary variable is active, the corresponding continuous variable is non-zero. However, the AMO exclusion does not require the lower bound on the variables as both values can be

zero. The second pair of binary variables (z_{n2}^x, z_{n3}^x) enforces the Roth conversion-tax-free withdrawal exclusion:

$$\begin{aligned} \epsilon z_{n2}^x &\leq \sum_i x_{in} \leq \mathcal{M} z_{n2}^x, \\ \epsilon z_{n3}^x &\leq \sum_i w_{i2n} \leq \mathcal{M} z_{n3}^x, \\ 0 &\leq z_{n2}^x + z_{n3}^x \leq 1. \end{aligned} \tag{5.18}$$

Here, \mathcal{M} is a large number such as 10^9 , just slightly larger than what x , w , and s can possibly be. The lower bound can be ignored here as well for the same reason. Note that the Roth conversion/tax-free withdrawal AMO (Eq. (5.18)) is *relaxed* for years $n < \max_i(n_{i,59\frac{1}{2}})$: before any individual reaches age $59\frac{1}{2}$, it is legitimate to simultaneously withdraw mature 5-year-old conversions and initiate new ladder rungs, so enforcing the AMO in those years would block the Roth conversion ladder. The AMO is reinstated once all individuals have passed $59\frac{1}{2}$, preventing artificial round-trips.

Another approach could be to perform Roth conversions at mid-year, while withdrawals could be made at the beginning of the year, and surplus deposits, if needed due to RMDs or receiving large sums of money, could be made at the end of the year. Let's formulate this approach in more detail and investigate for potential problems. Timing controls which terms get multiplied by the rate of return $(1 + \mathcal{T}_{ijn})$. Therefore, our current choice would yield

$$\begin{aligned} b_{ij(n+1)} &= [b_{ijn} - w_{ijn} + 0.5\kappa_{ijn}](1 + \mathcal{T}_{ijn}) + [\delta_{j,2} - \delta_{j,1}]x_{in}(1 + \mathcal{T}_{ijn}/2) \\ &\quad + \delta_{j,0}d_{in} + 0.5\kappa_{ijn}, \end{aligned} \tag{5.19}$$

where we use discrete Kronecker δ functions for selecting the specific accounts involved in Roth conversions. These conversions are made such that asset allocation ratios in the sending and receiving accounts are unchanged.

Bringing all variables to the left-hand side, this gets rewritten as

$$\begin{aligned} &b_{ij(n+1)} - (b_{ijn} - w_{ijn})(1 + \mathcal{T}_{ijn}) \\ &- [\delta_{j,2} - \delta_{j,1}]x_{in}(1 + \mathcal{T}_{ijn}/2) - \delta_{j,0}d_{in} = \kappa_{ijn}(1 + \mathcal{T}_{ijn}/2). \end{aligned} \tag{5.20}$$

When $j = 0$, this equation introduces a path to shelter negative returns by performing an over-withdrawal from the taxable account at the beginning of the year followed by a deposit in the same account at the end of the year. This can be removed by using another binary variable, thus making these events exclusive by using the same strategy as Eqs. (5.17) and (5.18).

A much simpler approach, while not so natural, is to move all transactions to be synchronous at the beginning or at the end of the year thus avoiding undesirable movements of funds. If we select the beginning of the year, except for contributions, this leads to

$$b_{ij(n+1)} - [b_{ijn} + \delta_{j,0}d_{in} - w_{ijn} + (\delta_{j,2} - \delta_{j,1})x_{in}](1 + \mathcal{T}_{ijn}) = \kappa_{ijn}(1 + \mathcal{T}_{ijn}/2).$$

This is the current approach used in Owl, coupled with binary variables excluding simultaneous overwithdrawals and deposits, and simultaneous Roth conversions and withdrawals from tax-free accounts.

5.1.12 Net spending

For calculating the net spending g_n , we consider the cash flow of all withdrawals, wages, other ordinary income, Social Security and pension benefits, proceeds from liquidation of fixed assets, and big-ticket items. Then we subtract potential surplus s_n and all taxes, penalties, debts, Medicare premiums, and ACA premiums paid:

$$\begin{aligned} g_n = & \sum_i [\omega_{in} + v_{in} + \nu_{in} + \bar{\zeta}_{in} + \bar{\pi}_{in}] + \sum_{i,j} w_{ijn} + \sum_i \Lambda_{in}^{\pm} + \sum_{*=x,c,f} \mathcal{A}_n^* \\ & - s_n - P_n - T_n - U_n - J_n - m_n - m_n^{\text{aca}} - \mathcal{D}_n. \end{aligned} \quad (5.21)$$

Here m_n^{aca} is the ACA net premium cost: a decision variable when `withACA="optimize"`, or a parameter (from the SC loop) when `withACA="loop"`. Notice how big-ticket items Λ^{\pm} contribute directly to the cash flow. Replacing intermediate variables and bringing all variables to the left-hand side, we get

$$\begin{aligned} g_n - \sum_{i,j} w_{ijn} + 0.1 \sum_i (1 - \mathcal{H}(n - n_{i,59\frac{1}{2}})) w_{i1n} \\ + m_n + m_n^{\text{aca}} + s_n + \sum_t f_{tn} \theta_{tn}^x + U_n = & \sum_{*=x,c,f} \mathcal{A}_n^* + \sum_i \Lambda_{in}^{\pm} \\ & + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \bar{\zeta}_{in} + \bar{\pi}_{in}] \\ & - \mathcal{D}_n - J_n. \end{aligned} \quad (5.22)$$

When `withACA="loop"`, m_n^{aca} is not an LP variable; the ACA cost is computed in the SC loop and subtracted on the RHS (omitted from the equation above for clarity). Notice that we do not consider market losses or tax-loss harvesting, and we do not track individual stock purchases over the years. We also treat J_n as a parameter as it is computed through a self-consistent loop.

We want the net spending to be predictable and smooth. For that purpose, we use

$$g_n/\bar{\xi}_n = g_0/\bar{\xi}_0, \quad (5.23)$$

where the net spending is adjusted for inflation and where we use the time series of parameter ξ_n , allowing for additional adjustments to the overall desired spending. Note that $\bar{\xi}_0 = \xi_0$ as $\gamma_0 = 1$. This profile is used to lower the desired net spending amount by a reduction factor χ after the passing of one spouse and/or to allow for more realistic spending profiles, such as the *smile* profile described above. Eq. (5.23) can be rewritten as

$$g_n \bar{\xi}_0 - g_0 \bar{\xi}_n = 0, \quad (5.24)$$

for the constraints to be enforced. Once g_0 is determined, the whole time series of net spending is determined. When using slack variable λ , the spending profile is then implemented as inequality constraints

$$\begin{aligned} g_n \bar{\xi}_0 - g_0(1 - \lambda) \bar{\xi}_n &\geq 0, \\ -g_n \bar{\xi}_0 + g_0(1 + \lambda) \bar{\xi}_n &\geq 0. \end{aligned} \quad (5.25)$$

In this case, the objective function will need to consider the sum over all years when optimizing net spending.

5.1.13 LTCG bracket constraints

The LTCG tax $U_n = 0.15 q_{1n} + 0.20 q_{2n}$ is determined by three LP bracket allocation variables q_{pn} ($p \in \{0, 1, 2\}$). The partition constraint

$$q_{0n} + q_{1n} + q_{2n} \geq Q_n \quad (5.26)$$

ensures the allocations cover the total LTCG Q_n . The bracket upper-bound constraints prevent overflow into higher-rate brackets. Because LTCG is stacked on top of ordinary taxable income $G_n = \sum_t f_{tn}$, the room available in each bracket depends on how much of the threshold is already occupied by ordinary income. The constraints can be expressed as

$$q_{0n} \leq \max(0, \bar{\mathcal{B}}_n^{15} - G_n), \quad (5.27)$$

$$q_{0n} + q_{1n} \leq \max(0, \bar{\mathcal{B}}_n^{20} - G_n), \quad (5.28)$$

where $\bar{\mathcal{B}}_n^{15}$ and $\bar{\mathcal{B}}_n^{20}$ are the inflation-adjusted total taxable income thresholds above which the 15% and 20% LTCG rates apply. In these constraints, G_n is taken from

the solution of the previous SC loop iteration (initialized to zero on the first pass). These constraints converge together with the SC loop. Since the cost is convex ($0\% < 15\% < 20\%$), no binary variables are needed: the LP naturally allocates gains to the cheapest bracket first.

5.1.14 Taxable ordinary income

We connect the two definitions for G_n stated above in Eqs. (4.1) and (4.24),

$$\begin{aligned} \sum_t f_{tn} &= \mathcal{A}_n^x + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \Psi_n \bar{\zeta}_{in} + \bar{\pi}_{in}] \\ &\quad + \sum_i [w_{i1n} + x_{in}] \\ &\quad + \sum_{i,k \neq 0} [(b_{i0n} - w_{i0n} + d_{in} + 0.5\kappa_{i0n})\alpha_{i0kn}\tau_{kn}] - e_n, \end{aligned} \quad (5.29)$$

and rearrange to move variables to the LHS as follows

$$\begin{aligned} e_n + \sum_t f_{tn} - \sum_i [w_{i1n} + x_{in}] \\ - \sum_{i,k \neq 0} [(b_{i0n} - w_{i0n} + d_{in})\alpha_{i0kn}\tau_{kn}] &= \mathcal{A}_n^x + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \Psi_n \bar{\zeta}_{in} + \bar{\pi}_{in}] \\ &\quad + 0.5 \sum_{i,k \neq 0} \alpha_{i0kn}\tau_{kn}\kappa_{i0n}. \end{aligned} \quad (5.30)$$

5.1.15 Social Security taxability

Under IRS Publication 915, the fraction of Social Security benefits subject to income tax is determined by the *provisional income* (PI), defined as

$$\Pi_n = \mathbb{G}_n - \frac{1}{2} \sum_i \bar{\zeta}_{in}. \quad (5.31)$$

The IRS piecewise formula gives the taxable SS amount as

$$\mathcal{S}_n = \begin{cases} 0 & \Pi_n < \mathcal{P}^{\text{lo}}, \\ 0.5 (\Pi_n - \mathcal{P}^{\text{lo}}) & \mathcal{P}^{\text{lo}} \leq \Pi_n < \mathcal{P}^{\text{hi}}, \\ \min\left(0.85 \sum_i \bar{\zeta}_{in}, 0.5 (\mathcal{P}^{\text{hi}} - \mathcal{P}^{\text{lo}}) + 0.85 (\Pi_n - \mathcal{P}^{\text{hi}})\right) & \Pi_n \geq \mathcal{P}^{\text{hi}}, \end{cases}$$

$$(5.32)$$

so that $\Psi_n = \mathcal{S}_n / \sum_i \bar{\zeta}_{in} \in [0, 0.85]$.

Introducing the excess provisional income above each threshold,

$$p_n^{\sigma, \text{lo}} = \max(0, \Pi_n - \mathcal{P}_n^{\text{lo}}), \quad p_n^{\sigma, \text{hi}} = \max(0, \Pi_n - \mathcal{P}_n^{\text{hi}}), \quad (5.33)$$

the three cases of Eq. (5.32) collapse into the single expression

$$\mathcal{S}_n = \min\left(0.85 \sum_i \bar{\zeta}_{in}, 0.50 \min\left(\sum_i \bar{\zeta}_{in}, \min(\Delta\mathcal{P}, p_n^{\sigma, \text{lo}})\right) + 0.85 p_n^{\sigma, \text{hi}}\right), \quad (5.34)$$

where $\Delta\mathcal{P} = \mathcal{P}^{\text{hi}} - \mathcal{P}^{\text{lo}}$ is the bracket width (\$12,000 for MFJ; \$9,000 for single filers). The inner $\min(\sum_i \bar{\zeta}_{in}, \cdot)$ guards against the unusual case $\sum_i \bar{\zeta}_{in} < \Delta\mathcal{P}$, in which SS benefits are too small to fill the 50%-bracket entirely. Because the clipped excess $\min(\Delta\mathcal{P}, p_n^{\sigma, \text{lo}}) \leq \Delta\mathcal{P}$, whenever $\sum_i \bar{\zeta}_{in} \geq \Delta\mathcal{P}$ the inner min is automatically satisfied and Eq. (5.34) simplifies to

$$\mathcal{S}_n = \min\left(0.85 \sum_i \bar{\zeta}_{in}, 0.5 p_n^{\sigma, \text{min}} + 0.85 p_n^{\sigma, \text{hi}}\right), \quad p_n^{\sigma, \text{min}} = \min(\Delta\mathcal{P}, p_n^{\sigma, \text{lo}}). \quad (5.35)$$

This condition holds for virtually all retirees, since combined annual SS benefits nearly always exceed \$12,000 (MFJ) or \$9,000 (single).

Because Ψ_n enters the LP only as a constant multiplier of the fixed series $\bar{\zeta}_{in}$ on the right-hand side of Eq. (5.30), it is treated as a *parameter* rather than a decision variable. It is updated between LP solves in the self-consistent (SC) loop: after each solve, \mathbb{G}_n is computed from the current solution, Π_n is evaluated, and Ψ_n is refreshed via Eq. (5.32). A 30% damping blend $\Psi_n \leftarrow 0.3 \Psi_n^{\text{new}} + 0.7 \Psi_n$ prevents oscillation near the bracket boundaries. To avoid unnecessary LP rebuilds, Ψ_n is only updated when $\max_n |\Delta\Psi_n| > 10^{-3}$. The SC loop terminates when the change in the objective function (net spending or bequest) falls below a prescribed tolerance.

Alternatively, when the solver option `withSSTaxability="optimize"` is set, all max/min operations in Eqs. (5.33)–(5.35) are encoded exactly within the LP as a mixed-integer program (MIP) using the continuous variables $p_n^{\sigma, \text{lo}}$, $p_n^{\sigma, \text{hi}}$, $p_n^{\sigma, \text{min}}$, and t_n^σ , together with two binary variables z_{0n}^σ and z_{1n}^σ per year n for which $\sum_i \bar{\zeta}_{in} > 0$; in years with no Social Security income all four continuous variables and both binaries are trivially fixed to zero and excluded from the MIP enumeration. The non-negativity of $p_n^{\sigma, \text{lo}}$ and $p_n^{\sigma, \text{hi}}$ combined with the lower-bound constraints $p_n^{\sigma, \text{lo}} \geq \Pi_n - \mathcal{P}_n^{\text{lo}}$ and $p_n^{\sigma, \text{hi}} \geq \Pi_n - \mathcal{P}_n^{\text{hi}}$ enforce the $\max(0, \cdot)$ operations; because maximizing spending minimizes taxes, the optimizer pushes these variables to their lower bounds at the

optimum. The two $\min(\cdot, \cdot)$ operations in Eq. (5.35) are formulated via big- \mathcal{M} lower-bound constraints. For $p_n^{\sigma, \min} = \min(\Delta\mathcal{P}, p_n^{\sigma, \text{lo}})$, the upper bounds

$$p_n^{\sigma, \min} \leq \Delta\mathcal{P}, \quad p_n^{\sigma, \min} \leq p_n^{\sigma, \text{lo}} \quad (5.36)$$

are complemented by binary-controlled lower bounds that force the optimizer to the tighter constraint:

$$p_n^{\sigma, \min} - \mathcal{M} z_{0n}^\sigma \geq \Delta\mathcal{P} - \mathcal{M}, \quad p_n^{\sigma, \min} - p_n^{\sigma, \text{lo}} + \mathcal{M} z_{0n}^\sigma \geq 0. \quad (5.37)$$

Analogously, for $t_n^\sigma = \min(0.85 \sum_i \bar{\zeta}_{in}, 0.5 p_n^{\sigma, \min} + 0.85 p_n^{\sigma, \text{hi}})$:

$$t_n^\sigma \leq 0.85 \sum_i \bar{\zeta}_{in}, \quad t_n^\sigma \leq 0.5 p_n^{\sigma, \min} + 0.85 p_n^{\sigma, \text{hi}}, \quad (5.38)$$

$$t_n^\sigma - \mathcal{M} z_{1n}^\sigma \geq 0.85 \sum_i \bar{\zeta}_{in} - \mathcal{M}, \quad t_n^\sigma - 0.5 p_n^{\sigma, \min} - 0.85 p_n^{\sigma, \text{hi}} + \mathcal{M} z_{1n}^\sigma \geq 0. \quad (5.39)$$

In this formulation t_n^σ is an LP variable rather than the fixed parameter $\Psi_n \sum_i \bar{\zeta}_{in}$ used in SC-loop mode. Moving t_n^σ to the left-hand side, Eq. (5.30) becomes

$$\begin{aligned} e_n + \sum_t f_{tn} - t_n^\sigma - \sum_i [w_{i1n} + x_{in}] \\ - \sum_{i,k \neq 0} [(b_{i0n} - w_{i0n} + d_{in}) \alpha_{i0kn} \tau_{kn}] = \mathcal{A}_n^x + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \bar{\pi}_{in}] \\ + 0.5 \sum_{i,k \neq 0} \alpha_{i0kn} \tau_{kn} K_{i0n}, \end{aligned} \quad (5.40)$$

and $\Psi_n = t_n^\sigma / \sum_i \bar{\zeta}_{in}$ is recovered from the LP solution for reporting. There is no circular dependency: provisional income Π_n is constructed from non-SS ordinary income, Roth conversions, Q_n dividends, and the fixed half of SS ($0.5 \sum_i \bar{\zeta}_{in}$), so t_n^σ does not appear in Π_n .

5.1.16 Medicare brackets and costs

Annual Medicare costs m_n include the income-related monthly adjusted amount commonly known as IRMAA. As this additional adjustment is a piecewise constant function, it can be computed using binary variables and mixed-integer linear programming. In the current tax code, this adjustment depends on the modified adjusted gross income (MAGI) \mathbb{G} from 2 years earlier. For the MAGI, we use

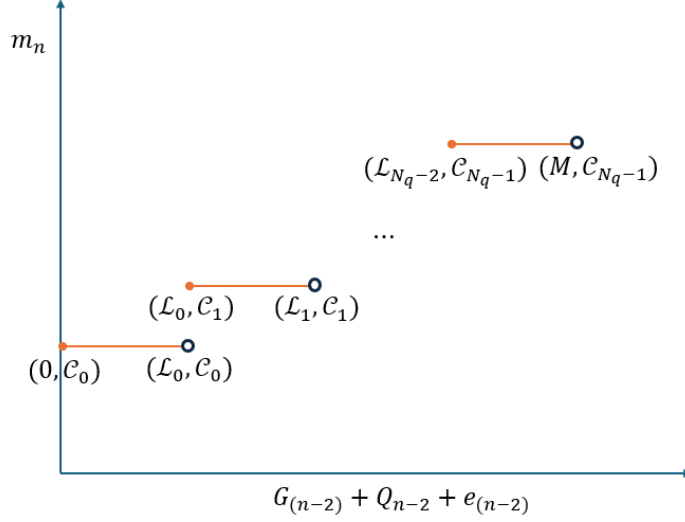


Figure 5.1: Modeling a piecewise constant monotonically increasing function. The x -axis represents the MAGI from two years ago, while m_n is the Medicare adjusted premiums. A binary variable z_q^m is associated with each segment q .

$\mathbb{G}_{(n-2)} = G_{(n-2)} + Q_{(n-2)} + e_{(n-2)} + (1 - \Psi_{(n-2)}) \sum_i \bar{\zeta}_{i(n-2)}$, i.e., gross taxable ordinary income plus dividends, plus standard exemption and the non-taxable portion of Social Security from 2 years ago. If the plan has its oldest individual currently either 64 or 65 years old, MAGI values for previous years must be provided by the user to complete the calculations.

Including regular premiums, there are $N_q = 6$ brackets of MAGI indexed for inflation defined with annual cumulative Medicare costs of $\bar{\mathcal{C}}_{qn} = \gamma_n \mathcal{C}_q$ also adjusted for inflation. These segments are separated by $N_q - 1$ thresholds $\bar{\mathcal{L}}_{qn} = \gamma_n \mathcal{L}_q$. Points $(\bar{\mathcal{L}}_{qn}, \bar{\mathcal{C}}_{qn})$ form a piecewise constant function that could be modeled using a special ordered set of type 1 (SOS1). Points $(\mathcal{L}_q, \mathcal{C}_q)$ are shown in Fig. 5.1. When $\mathbb{G}_{(n-2)} \leq \bar{\mathcal{L}}_{0n}$, the value $\bar{\mathcal{C}}_{0n}$ represents the (inflation-adjusted) base premium for Medicare Part B in year n , while higher q values include the IRMAA additional costs.

To model the IRMAA brackets we introduce binary variables $z_{qn}^m \in \{0, 1\}$ (one per bracket) and continuous variables $h_{qn} \geq 0$ for the MAGI assigned to the selected bracket q . Big- \mathcal{M} is used only for the upper bound of the top bracket (see below). We impose an SOS1 selection on binary z_{qn}^m as

$$\sum_q z_{qn}^m = 1, \quad (5.41)$$

and the disaggregation constraints

$$\sum_{q=0}^{N_q-1} h_{qn} = \mathbb{G}_{(n-2)}, \quad (5.42)$$

$$\begin{aligned} 0 \leq h_{0n} &\leq \bar{\mathcal{L}}_{0n} z_{0n}^m, \\ \bar{\mathcal{L}}_{(q-1)n} z_{qn}^m \leq h_{qn} &\leq \bar{\mathcal{L}}_{qn} z_{qn}^m, \quad q \in \{1, \dots, N_q - 2\}, \\ \bar{\mathcal{L}}_{(N_q-2)n} z_{(N_q-1)n}^m \leq h_{(N_q-1)n} &\leq \mathcal{M} \gamma_n z_{(N_q-1)n}^m. \end{aligned} \quad (5.43)$$

for all $n \in \{n_m, \dots, N_n - 1\}$, where n_m is the index year when Medicare starts for any individual i in the plan. The upper bound on the last bracket is necessary: when $z_{(N_q-1)n}^m = 0$ (the top bracket is not selected), the lower bound alone only gives $h_{(N_q-1)n} \geq 0$, so the solver could assign positive MAGI to an inactive bracket. The constraint $h_{(N_q-1)n} \leq \mathcal{M} \gamma_n z_{(N_q-1)n}^m$ forces $h_{(N_q-1)n} \leq 0$ when $z_{(N_q-1)n}^m = 0$; with $h_{qn} \geq 0$ we obtain $h_{(N_q-1)n} = 0$ for the inactive bracket. When $z_{(N_q-1)n}^m = 1$, $\mathcal{M} \gamma_n$ is chosen large enough to not bind (e.g. of order 10^8 in real units). As this applies to each individual eligible for Medicare, values $\bar{\mathcal{L}}_{qn}$ and $\bar{\mathcal{C}}_{qn}$ can include the accounting details for when both spouses are eligible and/or alive. For each year $n \geq n_m$, we can then express the Medicare costs m_n as

$$m_n = \sum_{q=0}^{N_q-1} z_{qn}^m \bar{\mathcal{C}}_{qn}, \quad (5.44)$$

as $z_{qn}^m = \delta_{q,q'}$ can only select a unique bracket q' . For $n \geq 2$, Eq. (5.42) can be

expanded using Eqs. (4.24) and (4.25) as

$$\begin{aligned}
\sum_{q=0}^{N_q-1} h_{qn} &= \mathcal{A}_{(n-2)}^x + \mathcal{A}_{(n-2)}^c + \sum_i [\omega_{i(n-2)} + v_{i(n-2)} + \nu_{i(n-2)} \\
&\quad + \bar{\zeta}_{i(n-2)} + \bar{\pi}_{i(n-2)}] \\
&\quad - e_{(n-2)} - \sum_i [w_{i1(n-2)} + x_{i(n-2)}] \\
&\quad - \sum_i [(b_{i0(n-2)} - w_{i0(n-2)} + d_{i(n-2)}) \\
&\quad \cdot \left(\sum_{k \neq 0} [\alpha_{i0k(n-2)} \tau_{k(n-2)}] + \mu \alpha_{i00(n-2)} \right)] \\
&\quad + \sum_i [w_{i0(n-2)} \alpha_{i00(n-2)} \\
&\quad \cdot \max(0, \tau_{0 \max(0, n-3)} - \mu)] \\
&\quad + 0.5 \sum_i \left[\kappa_{i0(n-2)} \left(\sum_{k \neq 0} [\alpha_{i0k(n-2)} \tau_{k(n-2)}] + \mu \alpha_{i00(n-2)} \right) \right].
\end{aligned}$$

Collecting variables on the left-hand side yields

$$\begin{aligned}
&\sum_{q=0}^{N_q-1} h_{qn} + e_{(n-2)} + \sum_i [w_{i1(n-2)} + x_{i(n-2)}] \\
&\quad + \sum_i [(b_{i0(n-2)} - w_{i0(n-2)} + d_{i(n-2)}) \\
&\quad \cdot \left(\sum_{k \neq 0} [\alpha_{i0k(n-2)} \tau_{k(n-2)}] + \mu \alpha_{i00(n-2)} \right)] \\
&\quad - \sum_i [w_{i0(n-2)} \alpha_{i00(n-2)} \max(0, \tau_{0 \max(0, n-3)} - \mu)] \\
&= \mathcal{A}_{(n-2)}^x + \mathcal{A}_{(n-2)}^c + \sum_i [\omega_{i(n-2)} + v_{i(n-2)} + \nu_{i(n-2)} + \bar{\zeta}_{i(n-2)} + \bar{\pi}_{i(n-2)}] \\
&\quad + 0.5 \sum_i \left[\kappa_{i0(n-2)} \left(\sum_{k \neq 0} [\alpha_{i0k(n-2)} \tau_{k(n-2)}] + \mu \alpha_{i00(n-2)} \right) \right].
\end{aligned} \tag{5.45}$$

Eq. (5.42) is expanded using Eqs. (4.24) and (4.25) to yield a linear equality in the decision variables; it is implemented directly in the constraint matrix. For indices

at the beginning of the time sequence, we use $\tau_{0,\max(0,n-x)}$, when $x > n$. Note that $G_{(n-2)}$ already includes the taxable portion of Social Security (via $\Psi_n \bar{\zeta}_{in}$), and the additional $(1 - \Psi_{(n-2)}) \bar{\zeta}_{i(n-2)}$ term restores the non-taxable portion so that MAGI uses total Social Security benefits.

For plans where the year index when Medicare first starts $n_m < 2$, we will request and use user-provided values for the first years' MAGIs as $\mathbb{G}_{(-1)}$ and possibly $\mathbb{G}_{(-2)}$ and set

$$\sum_{q=0}^{N_q-1} h_{qn} = \mathbb{G}_{(n-2)}, \quad n \in \{n_m, \dots, 1\}, \quad (5.46)$$

where $n_m \geq 0$. Using Eq. (5.43), we can also check for bounds and add constraints directly on z_{q0}^m and possibly z_{q1}^m .

5.1.17 ACA marketplace premium constraints

When `withACA="optimize"` and ACA is enabled ($\text{SLCSP} > 0$, $N_{\text{aca}} > 0$), ACA bracket selection and net premium cost are modeled within the optimization. ACA uses *current-year* MAGI \mathbb{G}_n (no two-year lag like Medicare). Let $N_q^{\text{aca}} = 7$ (brackets 0–5: FPL-based; bracket 6: above 400% FPL).

SOS1 selection:

$$\sum_{q=0}^{N_q^{\text{aca}}-1} z_{qn}^{\text{aca}} = 1, \quad n \in \{0, \dots, N_{\text{aca}} - 1\}. \quad (5.47)$$

MAGI decomposition:

$$\sum_{q=0}^{N_q^{\text{aca}}-1} h_{qn}^{\text{aca}} = \mathbb{G}_n, \quad (5.48)$$

where \mathbb{G}_n is the current-year MAGI (same structure as in Eq. (5.45) but with n instead of $n - 2$).

Bracket bounds (analogous to Eq. (5.43) with Big- \mathcal{M} for the last bracket):

$$\begin{aligned} 0 \leq h_{0n}^{\text{aca}} &\leq \bar{\mathcal{L}}_{0n}^{\text{aca}} z_{0n}^{\text{aca}}, \\ \bar{\mathcal{L}}_{(q-1)n}^{\text{aca}} z_{qn}^{\text{aca}} &\leq h_{qn}^{\text{aca}} \leq \bar{\mathcal{L}}_{qn}^{\text{aca}} z_{qn}^{\text{aca}}, \quad q \in \{1, \dots, N_q^{\text{aca}} - 2\}, \\ \bar{\mathcal{L}}_{(N_q^{\text{aca}}-2)n}^{\text{aca}} z_{(N_q^{\text{aca}}-1)n}^{\text{aca}} &\leq h_{(N_q^{\text{aca}}-1)n}^{\text{aca}} \leq \mathcal{M} \gamma_n z_{(N_q^{\text{aca}}-1)n}^{\text{aca}}. \end{aligned} \quad (5.49)$$

Cost constraint: For brackets 0–5, net cost equals the applicable percentage times MAGI in that bracket. For bracket 6 (above 400% FPL), no subsidy applies and net cost equals the full SL CSP:

$$m_n^{\text{aca}} = \sum_{q=0}^5 \bar{C}_q^{\text{aca}} h_{qn}^{\text{aca}} + \bar{S}_n^{\text{aca}} z_{6n}^{\text{aca}}, \quad n \in \{0, \dots, N_{\text{aca}} - 1\}. \quad (5.50)$$

For $n \geq N_{\text{aca}}$, $m_n^{\text{aca}} = 0$ (Medicare-eligible; no ACA cost). The variable m_n^{aca} is added to the cash-flow constraint Eq. (5.22) alongside m_n when `withACA="optimize"` is active.

6. Mapping of decision variables

At this point, one can use one of the many algebraic modeling languages such as AMPL, GAMS, MOSEK, AIMMS, and Gurobi, and code the equations above using that language, but most of these applications are proprietary and require a license and additional software installation. These languages allow the problem to be stated at a high level and steps to cast the problem in a form suitable for solution are performed automatically. There are also object-oriented language extensions, such as Python's Pyomo, that can ease the process of solving these problems. For completeness, however, we present here a simple index mapping approach that allows solving this problem using a generic linear programming solver.

Using a simple interface for mapping sparse objects to dense ones, the approach described here has been successfully tested with both the HiGHS open-source solver and the MOSEK proprietary solver. To cast the problem in a form suitable for a linear programming solver, we will use a single block vector represented by the array $y[q()]$ with index-mapping functions $q()$. While this process can be achieved using slicing and reshaping in some programming languages, we will present a generic approach suitable for most programming languages. The detailed approach presented here also allows us to determine the size of the problem to solve. We proceed alphabetically for all variables, and continue to use the convention of having index 0 for representing the first element.

To bring all variables in a single block vector, we will simply use two generic index mapping functions defined as

$$q_*(C, \ell_1, \ell_2, \ell_3; N_1, N_2, N_3) := C + \ell_1 N_2 N_3 + \ell_2 N_3 + \ell_3, \quad (6.1)$$

and

$$q_C(C, N_1, N_2, N_3) := C + N_1 N_2 N_3, \quad (6.2)$$

with the constraint that $0 \leq \ell_i < N_i$.

Account balances (b) For storing the savings account balances appropriately, variable b_{ijn} needs to have one more entry ($N_n + 1$) to store the end-of-life estate value. Therefore, we use

$$y[q_b(i, j, n)] = b_{ijn}, \quad (6.3)$$

where

$$q_b(i, j, n) = q_*(C_b, i, j, n; N_i, N_j, N_n + 1), \quad (6.4)$$

and where n exceptionally runs from 0 to N_n inclusively, and therefore q_b runs from $C_b = 0$ to $C_d - 1$, where

$$C_d = q_C(C_b, N_i, N_j, N_n + 1) = N_i N_j (N_n + 1).$$

Surplus deposits (d) For surplus deposits in the taxable savings accounts d_{in} we will use

$$y[q_d(i, n)] = d_{in}, \quad (6.5)$$

where

$$q_d(i, n) = q_*(C_d, i, n, 0; N_i, N_n, 1), \quad (6.6)$$

with q_d running from C_d to $C_e - 1$, where

$$C_e = q_C(C_d, N_i, N_n, 1) = N_i(N_j(N_n + 1) + N_n).$$

Standard exemption (e) For the standard exemption e_n we will use

$$y[q_e(n)] = e_n, \quad (6.7)$$

where

$$q_e(n) = q_*(C_e, n, 0, 0; N_n, 1, 1) = C_e + n, \quad (6.8)$$

with q_e running from C_e to $C_f - 1$, where

$$C_f = q_C(C_e, N_n, 1, 1) = N_i(N_j(N_n + 1) + N_n) + N_n.$$

Tax bracket amounts (f) For tax bracket amounts f_{tn} we will use

$$y[q_f(t, n)] = f_{tn}, \quad (6.9)$$

where

$$q_f(t, n) = q_*(C_f, t, n, 0; N_t, N_n, 1) \quad (6.10)$$

with q_f running from C_f to $C_g - 1$, where

$$C_g = q_C(C_f, N_t, N_n, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 1)N_n.$$

Net spending (g) For net spending g_n we will use

$$y[q_g(n)] = g_n, \quad (6.11)$$

where

$$q_g(n) = q_*(C_g, n, 0, 0; N_n, 1, 1) = C_g + n, \quad (6.12)$$

with q_g running from C_g to $C_h - 1$, where

$$C_h = q_C(C_g, N_n, 1, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 2)N_n.$$

MAGI bracket portions (h) For h_{qn} we will use

$$y[q_h(q, n)] = h_{qn}, \quad (6.13)$$

where

$$q_h(q, n) = q_*(C_h, q, n, 0; N_q, N_n - n_m, 1), \quad (6.14)$$

with q_h running from C_h to $C_m - 1$, where

$$C_m = q_C(C_h, N_q, N_n - n_m, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 2)N_n + N_q(N_n - n_m).$$

Medicare costs (m) For Medicare costs m_n we will use

$$y[q_m(n)] = m_n, \quad (6.15)$$

where

$$q_m(n) = q_*(C_m, n, 0, 0; N_n, 1, 1) = C_m + n, \quad (6.16)$$

with q_m running from C_m to $C_s - 1$, where

$$C_s = q_C(C_m, N_n, 1, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 3)N_n + N_q(N_n - n_m).$$

LTCG bracket allocations (q) For LTCG bracket allocation variables q_{pn} we will use

$$y[q_q(p, n)] = q_{pn}, \quad (6.17)$$

where

$$q_q(p, n) = q_*(C_q, p, n, 0; N_p, N_n, 1), \quad (6.18)$$

with q_q running from C_q to $C_s - 1$, where

$$\begin{aligned} C_q &= q_C(C_m, N_n, 1, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 3)N_n + N_q(N_n - n_m), \\ C_s &= q_C(C_q, N_p, N_n, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 3 + N_p)N_n + N_q(N_n - n_m). \end{aligned}$$

Surplus (s) Surplus can be generated if big-ticket items are received (inheritance, sale of a house, etc.) or due to RMDs. Surplus s is then deposited to taxable savings accounts according to variable η . We will use

$$y[q_s(n)] = s_n, \quad (6.19)$$

where

$$q_s(n) = q_*(C_s, n, 0, 0; N_n, 1, 1) = C_s + n, \quad (6.20)$$

with q_s running from C_s to $C_w - 1$, where

$$C_w = q_C(C_s, N_n, 1, 1) = N_i(N_j(N_n + 1) + N_n) + (N_t + 4 + N_p)N_n + N_q(N_n - n_m).$$

Withdrawals (w) For withdrawals w_{ijn} we will use

$$y[q_w(i, j, n)] = w_{ijn}, \quad (6.21)$$

where

$$q_w(i, j, n) = q_*(C_w, i, j, n; N_i, N_j, N_n), \quad (6.22)$$

with q_w running from C_w to $C_x - 1$, where

$$C_x = q_C(C_w, N_i, N_j, N_n) = N_i(N_j(2N_n + 1) + (N_t + 4 + N_p)N_n) + N_q(N_n - n_m).$$

Roth conversions (\mathbf{x}) Finally, for Roth conversions x_{in} we will use

$$y[q_x(i, n)] = x_{in}, \quad (6.23)$$

where

$$q_x(i, n) = q_*(C_x, i, n, 0; N_i, N_n, 1), \quad (6.24)$$

with q_x running from C_x to $C_* - 1$, where

$$C_* = q_C(C_x, N_i, N_n, 1) = N_i(N_j(2N_n + 1) + (N_t + N_i + 3 + N_p)N_n). \quad (6.25)$$

With $N_i = 2, N_j = 3, N_k = 4, N_t = 7, N_p = 3$ we have $(27 + N_p)N_n + 6 = 30N_n + 6$ variables. The $N_p N_n = 3N_n$ additional variables are the LTCG bracket allocations q_{pn} . When IRMAA is optimized, h_{qn} adds $N_q(N_n - n_m)$ more variables. If the time resolution is increased to months, the number of variables scales accordingly. Adding binary variables z_{nz}^x with $z \in \{0, 1, 2, 3\}$ can add $4N_n$ more decision variables, while adding z_{qn}^m adds $N_q(N_n - n_m)$ more, resulting in $(30N_n + 6) + N_q(N_n - n_m) + 4N_n + N_q(N_n - n_m)$ variables.

6.1 Reverse mapping of indices

The inverse functions for the index-mapping functions will be derived for the most complex case encountered in this paper. If we have

$$z = q_*(C, i, j, k, n; N_i, N_j, N_k, N_n) := C + iN_jN_kN_n + jN_kN_n + kN_n + n, \quad (6.26)$$

then $(i, j, k, n) = q_*^{-1}(z; N_i, N_j, N_k, N_n, C)$ is obtained from

$$\begin{aligned} n &= \text{mod}(\text{mod}(\text{mod}(z - C, N_jN_kN_n), N_kN_n), N_n), \\ k &= \text{mod}(\text{mod}(z - C - n, N_jN_kN_n), N_kN_n)/N_n, \\ j &= \text{mod}(z - C - n - kN_n, N_jN_kN_n)/(N_kN_n), \\ i &= (z - C - n - kN_n - jN_kN_n)/(N_jN_kN_n). \end{aligned} \quad (6.27)$$

While this holds for all cases presented in the previous section, this can be easily simplified for cases having fewer active indices. However, some modern languages can accomplish this mapping rather easily by providing `reshape()` functions.

7. Building constraint matrices

Let's first define generic index-mapping functions I and J as

$$\begin{aligned} I_l(n) &= C_l + n, \\ I_l(i, n; N_n) &= C_l + iN_n + n, \\ I_l(i, j, n; N_j, N_n) &= C_l + iN_jN_n + jN_n + n, \\ \dots &= \dots \end{aligned} \tag{7.1}$$

and so on, which would cumulatively increase row count C_l at each new instance l , similar to how we proceeded in the previous section. This allows us to build rectangular matrices by iteratively adding rows. These constraint matrices have C_* (defined in Eq. (6.25)) columns but can have fewer rows, forming an underdetermined system to be optimized using linear programming. Function J is defined similarly for equality constraints, while I is used for building the rows of the matrix containing the inequality constraints. Additional indices found in columns and not in rows imply the existence of multiple elements on that row.

7.1 Inequality constraints

Inequality constraints can be upper or lower bounds expressed in matrix form as $\ell \leq A_u y \leq u$. When ℓ is not specified it is assumed 0, and an unspecified u implies ∞ . Most if not all inequalities will be expressed as upper bounds.

Required minimum distributions (RMDs) We rewrite the inequality constraint on required minimum distributions Eq. (5.5) using matrix $A_u y \leq u$ starting

with the following $N_i N_n$ rows,

$$\begin{aligned}
A_u[I_0(i, n), q_w(i, 1, n)] &= -1 \\
A_u[I_0(i, n), q_b(i, 1, n)] &= \rho_{in}, \\
u[I_0(i, n)] &= 0, \\
\forall i \in \{0, \dots, N_i - 1\}, \\
\forall n \in \{0, \dots, N_n - 1\},
\end{aligned} \tag{7.2}$$

and all other elements in the same rows of A_u being 0. Notice that while b has $N_n + 1$ elements, the constraints for b go from 0 to $N_n - 1$ as there is no RMD required in the last year of the plan N_n , when all individuals have passed. Years before RMDs where $\rho_{in} = 0$ just add a trivial constraint. See Eq. (6.4).

Ordinary income tax brackets Similarly, we add $N_t N_n$ more rows to matrix $A_u y \leq u$ to express the inequality constraint in Eq. (5.6) setting an upper limit on amounts f_{tn} ,

$$\begin{aligned}
A_u[I_1(t, n), q_f(t, n)] &= 1, \\
u[I_1(t, n)] &= \bar{\Delta}_{tn}, \\
\forall t \in \{0, \dots, N_t - 1\}, \\
\forall n \in \{0, \dots, N_n - 1\},
\end{aligned} \tag{7.3}$$

and all other elements in the same rows of A_u being 0.

Standard exemption For $e_n \leq \bar{\sigma}_n$ we add

$$\begin{aligned}
A_u[I_2(n), q_e(n)] &= 1, \\
u[I_2(n)] &= \bar{\sigma}_n, \\
\forall n \in \{0, \dots, N_n - 1\}.
\end{aligned} \tag{7.4}$$

Non-negativity of e_n provides the lower bound.

Withdrawal limits For $w_{ijn} + \delta_{j,1}x_{in} \leq b_{ijn}$ we add:

$$\begin{aligned}
A_u[I_3(i, j, n), q_w(i, j, n)] &= 1, \\
A_u[I_3(i, j, n), q_x(i, n)] &= \delta_{j,1}, \\
A_u[I_3(i, j, n), q_b(i, j, n)] &= -1, \\
u[I_3(i, j, n)] &= 0, \\
&\forall i \in \{0, \dots, N_i - 1\}, \\
&\forall j \in \{0, \dots, N_j - 1\}, \\
&\forall n \in \{0, \dots, N_n - 1\} \text{ } (n < n_d \text{ for } i = i_d).
\end{aligned}$$

For a couple, the upper index for the first-to-pass individual i_d is $n_d - 1$ rather than $N_n - 1$, since posthumous constraints zero out all account activities.

Posthumous account activities For $n \geq n_d$ and the first-to-pass individual i_d , we set

$$\begin{aligned}
A_u[I_4(j, n), q_w(i_d, j, n)] &= 1, \\
u[I_4(j, n)] &= 0, \\
A_u[I_5(n), q_d(i_d, n)] &= 1, \\
u[I_5(n)] &= 0, \\
A_u[I_6(n), q_x(i_d, n)] &= 1, \\
u[I_6(n)] &= 0, \\
&\forall j \in \{0, \dots, N_j - 1\}, \\
&\forall n \in \{n_d, \dots, N_n - 1\}.
\end{aligned}$$

Roth conversion limit If a maximum conversion x_{max} is specified, we add the bound $\sum_i x_{in} \leq x_{max}$ for all n with

$$\begin{aligned}
A_u[I_7(n), q_x(i, n)] &= 1, \\
u[I_7(n)] &= x_{max}, \\
&\forall i \in \{0, \dots, N_i - 1\}, \\
&\forall n \in \{0, \dots, N_n - 1\}.
\end{aligned}$$

This maximum is applied to the combined conversions of both individuals in each year.

Exclusion constraints The at-most-one (AMO) exclusions of Eqs. (5.17) and (5.18) are implemented with big- \mathcal{M} constraints using binary variables z_{nz}^x , adding the corresponding rows for $\sum_i w_{i0n} + \sum_i w_{i2n}$, s_n , and $\sum_i x_{in}$ with upper bounds $\mathcal{M}z_{nz}$, along with the bounds $0 \leq z_{n0}^x + z_{n1}^x \leq 1$ and $0 \leq z_{n2}^x + z_{n3}^x \leq 1$.

$$\begin{aligned}
A_u[I_9(n), q_w(i, 0, n)] &= 1, \\
A_u[I_9(n), q_w(i, 2, n)] &= 1, \\
A_u[I_9(n), q_{z^x}(n, 0)] &= -\mathcal{M}, \\
u[I_9(n)] &= 0, \\
A_u[I_{10}(n), q_s(n)] &= 1, \\
A_u[I_{10}(n), q_{z^x}(n, 1)] &= -\mathcal{M}, \\
u[I_{10}(n)] &= 0, \\
A_u[I_{11}(n), q_{z^x}(n, 0)] &= 1, \\
A_u[I_{11}(n), q_{z^x}(n, 1)] &= 1, \\
u[I_{11}(n)] &= 1, \\
A_u[I_{12}(n), q_x(i, n)] &= 1, \\
A_u[I_{12}(n), q_{z^x}(n, 2)] &= -\mathcal{M}, \\
u[I_{12}(n)] &= 0, \\
A_u[I_{13}(n), q_w(i, 2, n)] &= 1, \\
A_u[I_{13}(n), q_{z^x}(n, 3)] &= -\mathcal{M}, \\
u[I_{13}(n)] &= 0, \\
A_u[I_{14}(n), q_{z^x}(n, 2)] &= 1, \\
A_u[I_{14}(n), q_{z^x}(n, 3)] &= 1, \\
u[I_{14}(n)] &= 1, \\
\forall i \in \{0, \dots, N_i - 1\}, \\
\forall n \in \{0, \dots, N_n - 1\}.
\end{aligned}$$

Medicare brackets Medicare's MAGI brackets use binary variables z_{qn}^m and the disaggregated MAGI variables h_{qn} . We set upper bounds for each bracket using

Eq. (5.43) as

$$\begin{aligned}
A_u[I_{15}(q, n), q_h(q, n)] &= 1, \\
A_u[I_{15}(q, n), q_{z^m}(q, n)] &= -\bar{\mathcal{L}}_{qn}, \\
u[I_{15}(q, n)] &= 0, \\
&\forall q \in \{0, \dots, N_q - 2\}, \\
&\forall n \in \{n_m, \dots, N_n - 1\},
\end{aligned} \tag{7.5}$$

the upper bound for the last bracket ($q = N_q - 1$) using big- \mathcal{M} as

$$\begin{aligned}
A_u[I_{15}(N_q - 1, n), q_h(N_q - 1, n)] &= 1, \\
A_u[I_{15}(N_q - 1, n), q_{z^m}(N_q - 1, n)] &= -\mathcal{M}\gamma_n, \\
u[I_{15}(N_q - 1, n)] &= 0, \\
&\forall n \in \{n_m, \dots, N_n - 1\},
\end{aligned} \tag{7.6}$$

and the lower bounds as

$$\begin{aligned}
A_u[I_{16}(q, n), q_h(q, n)] &= -1, \\
A_u[I_{16}(q, n), q_{z^m}(q, n)] &= \bar{\mathcal{L}}_{(q-1)n}, \\
u[I_{16}(q, n)] &= 0, \\
&\forall q \in \{1, \dots, N_q - 1\}, \\
&\forall n \in \{n_m, \dots, N_n - 1\}.
\end{aligned} \tag{7.7}$$

7.2 Equality constraints

Account balances For the equality constraints on account balances expressed in Eq. (5.21), we define an equality constraint matrix $A_e y = v$ starting with $N_i N_j N_n$ rows as

$$\begin{aligned}
A_e[J_0(i, j, n), q_b(i, j, n + 1)] &= 1, \\
A_e[J_0(i, j, n), q_b(i, j, n)] &= -(1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_x(i, n)] &= -(\delta_{j,2} - \delta_{j,1})(1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_w(i, j, n)] &= (1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_d(i, n)] &= -\delta_{j,0}(1 + \mathcal{T}_{ijn}), \\
&\forall i \in \{0, \dots, N_i - 1\}, \\
&\forall j \in \{0, \dots, N_j - 1\}, \\
&\forall n \in \{0, \dots, N_n - 1\},
\end{aligned} \tag{7.8}$$

where v is

$$v[J_0(i, j, n)] = \kappa_{ijn}(1 + \mathcal{T}_{ijn}/2). \quad (7.9)$$

The initial account balances expressed in Eq. (5.14) are imposed through

$$\begin{aligned} A_e[J_1(i, j), q_b(i, j, 0)] &= 1, \\ v[J_1(i, j)] &= \beta_{ij}, \end{aligned} \quad (7.10)$$

$$\begin{aligned} \forall i &\in \{0, \dots, N_i - 1\}, \\ \forall j &\in \{0, \dots, N_j - 1\}, \end{aligned} \quad (7.11)$$

leading to $N_i N_j$ additional rows to A_e .

Medicare bracket selection The SOS1 selection of Eq. (5.41) is encoded as

$$\begin{aligned} A_e[J_2(n), q_{zm}(q, n)] &= 1, \\ v[J_2(n)] &= 1, \\ \forall q &\in \{0, \dots, N_q - 1\}, \\ \forall n &\in \{n_m, \dots, N_n - 1\}. \end{aligned} \quad (7.12)$$

MAGI disaggregation The equality $\sum_q h_{qn} = \mathbb{G}_{(n-2)}$ is expanded using Eqs. (4.24) and (4.25) and implemented as

$$\begin{aligned}
A_e[J_3(n), q_h(q, n)] &= 1, \\
A_e[J_3(n), q_e(n-2)] &= -1, \\
A_e[J_3(n), q_w(i, 1, n-2)] &= -1, \\
A_e[J_3(n), q_x(i, n-2)] &= -1, \\
A_e[J_3(n), q_b(i, 0, n-2)] &= -\mu\alpha_{i00(n-2)} \\
&\quad - \sum_{k \neq 0} \alpha_{i0k(n-2)} \tau_{k(n-2)}, \\
A_e[J_3(n), q_d(i, n-2)] &= -\mu\alpha_{i00(n-2)} \\
&\quad - \sum_{k \neq 0} \alpha_{i0k(n-2)} \tau_{k(n-2)}, \\
A_e[J_3(n), q_w(i, 0, n-2)] &= \mu\alpha_{i00(n-2)} + \sum_{k \neq 0} \alpha_{i0k(n-2)} \tau_{k(n-2)} \\
&\quad - \alpha_{i00(n-2)} \max(0, \tau_{0 \max(0, n-3)} - \mu), \\
v[J_3(n)] &= \mathcal{A}_{n-2}^x + \mathcal{A}_{n-2}^c + \sum_i [\omega_{i(n-2)} + \nu_{i(n-2)} + \nu_{i(n-2)} \\
&\quad + \bar{\zeta}_{i(n-2)} + \bar{\pi}_{i(n-2)}] \\
&\quad + 0.5 \sum_i \left[\kappa_{i0(n-2)} \left(\mu\alpha_{i00(n-2)} + \sum_{k \neq 0} \alpha_{i0k(n-2)} \tau_{k(n-2)} \right) \right], \\
&\quad \forall i \in \{0, \dots, N_i - 1\}, \\
&\quad \forall k \in \{0, \dots, N_k - 1\}, \\
&\quad \forall q \in \{0, \dots, N_q - 1\}, \\
&\quad \forall n \in \{\max(n_m, 2), \dots, N_n - 1\}. \tag{7.13}
\end{aligned}$$

When $n_m < 2$, we instead set $v[J_3(n)] = \mathbb{G}_{(-1)}$ and $\mathbb{G}_{(-2)}$, which are user-provided MAGI values for the years before the plan started.

Medicare costs Binary variables z_{qn}^m express the active IRMAA bracket. The costs for Medicare are obtained using Eq. (5.44):

$$\begin{aligned}
A_e[J_4(n), q_m(n)] &= 1, \\
A_e[J_4(n), q_{zm}(q, n)] &= -\bar{C}_{qn}, \\
v[J_4(n)] &= 0, \\
\forall q \in \{0, \dots, N_q - 1\}, \\
\forall n \in \{n_m, \dots, N_n - 1\}.
\end{aligned}$$

Net spending For the equality constraint on net spending expressed in Eq. (5.22), we add N_n more rows to $A_e y = v$ as

$$\begin{aligned}
A_e[J_5(n), q_g(n)] &= 1, \\
A_e[J_5(n), q_m(n)] &= 1, \\
A_e[J_5(n), q_s(n)] &= 1, \\
A_e[J_5(n), q_f(t, n)] &= \theta_{tn}^x, \\
A_e[J_5(n), q_w(i, j, n)] &= -1 + 0.1 \delta_{j,1} (1 - \mathcal{H}(n - n_{i,59\frac{1}{2}})), \\
A_e[J_5(n), q_q(1, n)] &= 0.15, \\
A_e[J_5(n), q_q(2, n)] &= 0.20, \\
\forall t \in \{0, \dots, N_t - 1\}, \\
\forall i \in \{0, \dots, N_i - 1\}, \\
\forall j \in \{0, \dots, N_j - 1\}, \\
\forall n \in \{0, \dots, N_n - 1\},
\end{aligned}$$

where $q_q(p, n)$ is the flat index of q_{pn} in the solution vector, and v is

$$v[J_5(n)] = \sum_{*=x,c,f} \mathcal{A}_n^* + \sum_i \Lambda_{in}^\pm + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \bar{\zeta}_{in} + \bar{\pi}_{in}] - \mathcal{D}_n. \quad (7.14)$$

The condition of having a predictable net spending expressed as an equality in Eq. (5.24) adds $N_n - 1$ more rows to $A_e y = v$ as

$$\begin{aligned}
A_e[J_6(n), q_g(0)] &= -\bar{\xi}_n, \\
A_e[J_6(n), q_g(n)] &= \xi_0, \\
v[J_6(n)] &= 0, \\
\forall n \in \{1, \dots, N_n - 1\}.
\end{aligned} \quad (7.15)$$

Taxable ordinary income For the equality constraint in Eq. (5.30) establishing taxable ordinary income, we add N_n rows to $A_e y = v$ as follows. The taxable Social Security term $\Psi_n \sum_i \bar{\zeta}_{in}$ appears on the right-hand side as a parameter updated each SC-loop iteration.

$$\begin{aligned}
A_e[J_7(n), q_e(n)] &= 1, \\
A_e[J_7(n), q_f(t, n)] &= 1, \\
A_e[J_7(n), q_w(i, 1, n)] &= -1, \\
A_e[J_7(n), q_x(i, n)] &= -1, \\
A_e[J_7(n), q_b(i, 0, n)] &= -\sum_{k \neq 0} \alpha_{i0kn} \tau_{kn}, \\
A_e[J_7(n), q_w(i, 0, n)] &= \sum_{k \neq 0} \alpha_{i0kn} \tau_{kn}, \\
A_e[J_7(n), q_d(i, n)] &= -\sum_{k \neq 0} \alpha_{i0kn} \tau_{kn}, \\
&\quad \forall t \in \{0, \dots, N_t - 1\}, \\
&\quad \forall i \in \{0, \dots, N_i - 1\}, \\
&\quad \forall k \in \{0, \dots, N_k - 1\}, \\
&\quad \forall n \in \{0, \dots, N_n - 1\},
\end{aligned} \tag{7.16}$$

with

$$v[J_7(n)] = \mathcal{A}_n^x + \Psi_n \sum_i \bar{\zeta}_{in} + \sum_i [\omega_{in} + v_{in} + \nu_{in} + \bar{\pi}_{in}] + 0.5 \sum_i \kappa_{i0n} \sum_{k \neq 0} \alpha_{i0kn} \tau_{kn} \tag{7.17}$$

7.3 Other considerations

Beneficiaries Tax-free and tax-deferred accounts have special tax rules that allow giving part or the entire value of tax-free accounts to a spouse who can then consider it as his/her own. These accounts typically use percentages to designate beneficiaries. Let ϕ_j be the fraction of the account j that a spouse i_d wishes to leave to his/her surviving spouse i_s in year n_d , the year following the passing. To account for that

event in year n_d , Eq. (5.21) needs to be rewritten as

$$\begin{aligned}
b_{ij(n+1)} = & [1 - \delta_{n,n_d-1}\delta_{i,i_d}] \\
& \times \left\{ [b_{ijn} + \delta_{j,0}d_{in} - w_{ijn} + (\delta_{j,2} - \delta_{j,1})x_{in}] (1 + \mathcal{T}_{ijn}) + \kappa_{ijn}(1 + \mathcal{T}_{ijn}/2) \right\} \\
& + [\phi_j \delta_{n,n_d-1} \delta_{i,i_s}] \\
& \times \left\{ [b_{i_djn} + \delta_{j,0}d_{i_dn} - w_{i_djn} + (\delta_{j,2} - \delta_{j,1})x_{i_dn}] (1 + \mathcal{T}_{i_djn}) \right. \\
& \left. + \kappa_{i_djn}(1 + \mathcal{T}_{i_djn}/2) \right\}. \tag{7.18}
\end{aligned}$$

The first multiplier $[\dots]$ on the right-hand side will always be one except for i_d in year $n_d - 1$ when it will be zero. This will result in emptying all accounts for i_d for years n_d and beyond. The second special multiplier $[\dots]$ before the second set of curly braces $\{\}$ will always be zero except for the surviving spouse i_s in year $n_d - 1$, who will then inherit a fraction ϕ_j of account j that was scheduled to go into i_d 's j account at the beginning of year n_d .

Rewriting the last equation as a constraint results in

$$\begin{aligned}
& b_{ij(n+1)} \\
& - [1 - \delta_{n,n_d-1}\delta_{i,i_d}] \\
& \times \left\{ [b_{ijn} + \delta_{j,0}d_{in} - w_{ijn} + (\delta_{j,2} - \delta_{j,1})x_{in}] (1 + \mathcal{T}_{ijn}) \right\} \\
& - [\phi_j \delta_{n,n_d-1} \delta_{i,i_s}] \\
& \times \left\{ [b_{i_djn} + \delta_{j,0}d_{i_dn} - w_{i_djn} + (\delta_{j,2} - \delta_{j,1})x_{i_dn}] (1 + \mathcal{T}_{i_djn}) \right\} \\
= & [1 - \delta_{n,n_d-1}\delta_{i,i_d}] \kappa_{ijn}(1 + \mathcal{T}_{ijn}/2) \\
& + [\phi_j \delta_{n,n_d-1} \delta_{i,i_s}] \kappa_{i_djn}(1 + \mathcal{T}_{i_djn}/2). \tag{7.19}
\end{aligned}$$

We are now ready to replace Eq. (7.8) for $A_e y = v$ by

$$\begin{aligned}
A_e[J_0(i, j, n), q_b(i, j, n+1)] &= 1, \\
A_e[J_0(i, j, n), q_b(i, j, n)] &= -[1 - \delta_{n, n_d-1} \delta_{i, i_d}](1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_d(i, j, n)] &= -[1 - \delta_{n, n_d-1} \delta_{i, i_d}] \delta_{j,0} (1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_w(i, j, n)] &= [1 - \delta_{n, n_d-1} \delta_{i, i_d}](1 + \mathcal{T}_{ijn}), \\
A_e[J_0(i, j, n), q_x(i, n)] &= -[1 - \delta_{n, n_d-1} \delta_{i, i_d}] (\delta_{j,2} - \delta_{j,1}) (1 + \mathcal{T}_{ijn}), \\
\text{when } N_i = 2 \text{ and } i = i_s, \\
A_e[J_0(i, j, n), q_b(i_d, j, n)] &= -[\phi_j \delta_{n, n_d-1} \delta_{i, i_s}](1 + \mathcal{T}_{idjn}), \\
A_e[J_0(i, j, n), q_d(i_d, n)] &= -[\phi_j \delta_{n, n_d-1} \delta_{i, i_s}] \delta_{j,0} (1 + \mathcal{T}_{idjn}), \\
A_e[J_0(i, j, n), q_w(i_d, j, n)] &= [\phi_j \delta_{n, n_d-1} \delta_{i, i_s}](1 + \mathcal{T}_{idjn}), \\
A_e[J_0(i, j, n), q_x(i_d, n)] &= -[\phi_j \delta_{n, n_d-1} \delta_{i, i_s}] (\delta_{j,2} - \delta_{j,1}) (1 + \mathcal{T}_{idjn}),
\end{aligned}$$

$$\forall i \in \{0, \dots, N_i - 1\},$$

$$\forall j \in \{0, \dots, N_j - 1\},$$

$$\forall n \in \{0, \dots, N_n - 1\},$$

where v is

$$\begin{aligned}
v[J_0(i, j, n)] &= [1 - \delta_{n, n_d-1} \delta_{i, i_d}] \kappa_{ijn} (1 + \mathcal{T}_{ijn}/2) \\
&\quad + [\phi_j \delta_{n, n_d-1} \delta_{i, i_s}] \kappa_{idjn} (1 + \mathcal{T}_{idjn}/2).
\end{aligned} \tag{7.20}$$

While the last two equations may look cumbersome, their net effect is only to include a few more terms when $n = n_d - 1$.

Asset allocation ratios When asset allocation ratios α are imposed, they should also be applied to how contribution amounts κ_{ijn} are invested, such that

$$\kappa_{ijkn} = \alpha_{ijkn} \kappa_{ijn}. \tag{7.21}$$

For other allocation schemes, just substitute $\alpha_{ijkn} = \alpha_{ikn}$ or α_{kn} depending on the scheme selected.

Asset allocation has been handled easily by assuming that the accounts are always rebalanced and only using a single multiplier \mathcal{T} , defined as

$$\mathcal{T}_{ijn} = \sum_k \alpha_{ijkn} \tau_{kn}, \tag{7.22}$$

to compute the return on the total balance of each savings account.

Spousal deposits and withdrawals While keeping the problem linear, a simple constraint can be imposed on surplus deposits in taxable savings accounts. One can specify a spousal ratio η such as

$$d_{0n} = \eta d_{1n}. \tag{7.23}$$

A similar spousal ratio can be imposed on withdrawals from tax-deferred accounts

$$w_{01n} = \eta w_{11n}, \tag{7.24}$$

but this can cause drawing from an empty account while the other spousal account is not. Only the deposit scheme has currently been implemented in Owl.

8. Objective Functions and Their Formulation

The objective function is a simple scalar defined as $c \cdot y$ that will be minimized.

8.1 Maximizing net spending

There are a few ways by which a retirement plan can be optimized. For maximizing the net spending under the constraint of a desired bequest, we introduce the following relation

$$E_n = \sum_{i,j} (1 - \nu \delta_{j,1}) b_{ijn}, \quad (8.1)$$

which is the value of the estate in nominal dollars at year n , taking into consideration the heir's marginal income tax rate ν on the ($j = 1$) tax-deferred account.

For a desired bequest ϵ_{N_n} , expressed in today's dollars, the final amount in year N_n will need to be

$$E_{N_n} = \bar{\epsilon}_{N_n} = \epsilon_{N_n} \gamma_{N_n}. \quad (8.2)$$

Fixing a bequest value amounts to adding the following constraint

$$\sum_{i,j} b_{ijn} (1 - \nu \delta_{j,1}) = \epsilon_{N_n} \gamma_{N_n}, \quad (8.3)$$

which would add one more row to $A_e y = v$ as

$$\begin{aligned} A_e[I(0), q_b(i, j, N_n)] &= (1 - \nu \delta_{j,1}) \\ v[I(0)] &= \epsilon_{N_n} \gamma_{N_n} \end{aligned} \quad (8.4)$$

$$\begin{aligned} \forall i &\in \{0, \dots, N_i - 1\}, \\ \forall j &\in \{0, \dots, N_j - 1\}, \end{aligned} \quad (8.5)$$

where $I(0)$ is used only to provide the proper row offset C_l . See Eq. (7.1).

For maximizing the net spending under the constraint of a fixed bequest, one may either maximize the first-year net spending basis g_0 (with the profile fixed by Eq. (5.24)), or maximize the sum of net spending over all years in today's dollars. In the first case, minimize the inner product $c \cdot y$ with

$$c[q_g(0)] = -1, \quad (8.6)$$

and 0 otherwise; the whole time series g_n then follows from Eq. (5.24). In the second case, when the slack variable λ is used over net spending (as in Eq. (8.10) below), the objective is the sum $\sum_{n=0}^{N_n-1} g_n/\gamma_n$, maximized by setting

$$c[q_g(n)] = -1/\gamma_n, \quad (8.7)$$

$\forall n \in \{0, \dots, N_n - 1\}$, and 0 otherwise. This sum-of-net-spending formulation is the one implemented in Owl.

Variable profile with slack

When the slack variable λ is used (Eq. (5.25) or Eq. (8.10) below), the objective function is the sum of net spending over the full duration of the plan in today's dollars, $\sum_{n=0}^{N_n-1} g_n/\gamma_n$. The coefficient vector is

$$c[q_g(n)] = -1/\gamma_n, \quad (8.8)$$

$\forall n \in \{0, \dots, N_n - 1\}$ and 0 otherwise. In that case, constraint equality Eq. (5.24) is replaced by an inequality. Instead of obeying

$$g_n \xi_0 - g_0 \bar{\xi}_n = 0, \quad (8.9)$$

we impose the inequality constraint

$$(1 - \lambda)g_0 \bar{\xi}_n / \xi_0 \leq g_n \leq (1 + \lambda)g_0 \bar{\xi}_n / \xi_0, \quad (8.10)$$

where λ is the fractional tolerance by which annual net spending is allowed to deviate from the desired profile. It is set via the `spendingSlack` option (an integer percentage, e.g. `spendingSlack=10` gives $\lambda = 0.10$); the default is $\lambda = 0$ (strict profile). It should be noticed that when $\lambda = 0$ the two last equations are equivalent.

8.2 Maximizing bequest

If, on the other hand, one would like to maximize the bequest under the constraint of a desired net spending g_0 , specified for the first year, one would add the following row to $A_e y = v$

$$\begin{aligned} A_e[I(0), q_g(0)] &= 1, \\ v[I(0)] &= g_0, \end{aligned} \tag{8.11}$$

subject to the net spending g_n obeying Eq. (5.24) over time.

The objective function would then be derived from Eq. (8.1) as minimizing the inner product $c \cdot y$, where c is

$$c[q_b(i, j, N_n)] = -(1 - \nu \delta_{j,1}), \tag{8.12}$$

$$\forall i \in \{0, \dots, N_i - 1\},$$

$$\forall j \in \{0, \dots, N_j - 1\},$$

$$\tag{8.13}$$

and 0 otherwise.

8.3 Lexicographic tie-breaking

The coefficient vector c can be augmented with a small positive weight ε to break ties without changing the primary objective in practice. In addition to the coefficients of the chosen primary objective (max spending or max bequest), the implementation adds the following contributions to c :

$$\begin{aligned} c[q_x(i, n)] &\leftarrow c[q_x(i, n)] + \varepsilon, \\ &\quad \forall i \in \{0, \dots, N_i - 1\}, \forall n \in \{0, \dots, N_n - 1\}, \\ c[q_w(1, j, n)] &\leftarrow c[q_w(1, j, n)] + \varepsilon, \\ &\quad \forall j \in \{0, \dots, N_j - 1\}, \forall n \in \{0, \dots, N_n - 1\}, \quad \text{when } N_i = 2. \end{aligned} \tag{8.14}$$

The first line penalizes Roth conversions x_{in} , so that among solutions with the same primary objective value the solver prefers fewer or smaller conversions (reducing churn). The second line penalizes withdrawals w_{1jn} from the second individual ($i = 1$), so that withdrawals from the first individual ($i = 0$) are preferred when otherwise equivalent. The weight ε is chosen small (e.g. of order 10^{-9} in the same units as y) so that the primary objective dominates; it is often enabled by default when optimizing Medicare (IRMAA) to stabilize bracket selection.

9. Rate models for return scenarios

The optimization model described in the preceding chapters requires a sequence $\{\tau_{kn}\}_{n=0}^{N_n-1}$ of annual returns for each asset class k over the plan horizon. Owl provides several models for generating this sequence, ranging from simple fixed-rate assumptions to fully stochastic simulation. This chapter describes the models available and the statistical algorithms underlying the stochastic ones.

9.1 Historical data

All built-in models draw on an annual dataset covering the period from $\mathcal{Y}_{\min} = 1928$ through the most recently completed calendar year \mathcal{Y}_{\max} , containing $N_y = \mathcal{Y}_{\max} - \mathcal{Y}_{\min} + 1$ annual observations [19]. Four asset-class return time series are included, corresponding to the index k :

- $k = 0$ S&P 500 total return, including dividends.
- $k = 1$ Moody's Baa corporate bond total return (investment-grade bonds with a moderate default premium).
- $k = 2$ 10-year U.S. Treasury note total return.
- $k = 3$ Consumer Price Index (CPI) inflation. This series serves a dual role: it drives the cumulative inflation factor γ_n (see Chapter 2), and, as the *cash asset*, it proxies a Treasury Inflation-Protected Security (TIPS) with a real yield of zero so that holding cash preserves but does not grow purchasing power.

All values are stored in percent (e.g. 7.0 for a 7% annual return). A user-selected sub-range $[\mathcal{Y}_{\text{frm}}, \mathcal{Y}_{\text{to}}] \subseteq [\mathcal{Y}_{\min}, \mathcal{Y}_{\max}]$ is used by the historical, histogaussian, bootstrap, and VAR(1) models. We denote the window width by $N_w = \mathcal{Y}_{\text{to}} - \mathcal{Y}_{\text{frm}} + 1$, and the observed return of asset class k in calendar year y by r_{ky} .

Sequence transformations: reverse and roll. Two optional transforms apply to the extracted sequence before it is used. The *reverse* option reflects the sequence in time, $\tau_{kn} \leftarrow r_{k, \mathcal{Y}_{\text{to}} - n}$, so that the earliest calendar year in the window maps to plan year $N_n - 1$ and the latest to plan year 0. The *roll* option shifts the sequence by s positions, $\tau_{kn} \leftarrow r_{k, \mathcal{Y}_{\text{frm}} + ((n+s) \bmod N_w)}$, with positive s advancing toward later years (and wrapping). Both transforms preserve the set of returns in the window while altering the order; they are useful for exploring sequence-of-returns sensitivity without changing the underlying data. Reverse and roll are supported by all methods that use a historical window (historical, histogaussian, histolognormal, bootstrap, VAR, GARCH). Constant-rate methods ignore them.

9.2 Constant rate methods

The simplest approach assigns a constant return to every year of the plan, $\tau_{kn} = \tau_k$ for all n . Four variants are supported:

trailing-30

Trailing 30-year geometric mean of annual returns, updated annually.

conservative and optimistic

Preset rates based on long-term expert forecasts.

historical average

The geometric mean over the selected historical window,

$$\tau_k = \exp\left(\frac{1}{N_w} \sum_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}} \ln(1 + r_{ky})\right) - 1. \quad (9.1)$$

user Freely specified by the user for each asset class.

9.3 Deterministic varying rates: historical

The **historical** method assigns recorded returns directly to each year of the plan:

$$\tau_{kn} = r_{k, \mathcal{Y}_{\text{frm}} + n}, \quad n = 0, \dots, N_n - 1. \quad (9.2)$$

If the plan horizon N_n exceeds the window width N_w , the sequence is repeated cyclically. The reverse and roll transforms (see § 9.1) are applied before the plan runs.

9.4 Stochastic models

Stochastic methods draw N_n annual return vectors $\boldsymbol{\tau}_n = (\tau_{0n}, \tau_{1n}, \tau_{2n}, \tau_{3n})^\top \in \mathbb{R}^{N_k}$ from a probability distribution calibrated on historical data. Because random number generation is involved, a seed can be fixed to obtain reproducible results across runs. Each fresh draw of N_n vectors constitutes one scenario; running many independent scenarios yields a Monte Carlo distribution of outcomes.

9.4.1 Multivariate normal: histogaussian and gaussian

Both methods model annual returns as draws from a multivariate normal distribution,

$$\boldsymbol{\tau}_n \stackrel{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad n = 0, \dots, N_n - 1. \quad (9.3)$$

For **histogaussian**, the mean vector $\boldsymbol{\mu}$ and the $N_k \times N_k$ covariance matrix $\boldsymbol{\Sigma}$ are estimated from the selected historical window:

$$\mu_k = \frac{1}{N_w} \sum_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}} r_{ky}, \quad (9.4)$$

$$\Sigma_{kk'} = \frac{1}{N_w - 1} \sum_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}} (r_{ky} - \mu_k)(r_{k'y} - \mu_{k'}). \quad (9.5)$$

For **gaussian**, the user supplies $\boldsymbol{\mu}$ directly, together with the standard deviations $\sigma_k = \sqrt{\Sigma_{kk}}$ and the Pearson correlation coefficients $\rho_{kk'} \in [-1, 1]$, from which the covariance matrix is assembled as $\Sigma_{kk'} = \rho_{kk'} \sigma_k \sigma_{k'}$. In both cases, samples are generated as $\boldsymbol{\tau}_n = \boldsymbol{\mu} + L \mathbf{z}_n$, where L is the lower-triangular Cholesky factor of $\boldsymbol{\Sigma}$ and $\mathbf{z}_n \sim \mathcal{N}(\mathbf{0}, I)$. By construction, the multivariate normal model treats each year's returns as independent of all other years; it captures cross-sectional correlation among asset classes but no serial (year-to-year) autocorrelation.

9.4.2 Log-normal models: lognormal and histolognormal

Geometric Brownian Motion (GBM) [20] — the foundation of most option-pricing theory — implies that the price relative $1 + \tau_{k,n}$ is log-normally distributed. The log-normal model enforces $\tau_{k,n} > -1$ by construction (no total-loss artifacts are possible), produces a right-skewed marginal distribution (large gains are more likely than a symmetric normal would suggest), and is consistent with continuous-time finance theory.

Parameter conversion Given the arithmetic mean m_k and arithmetic standard deviation s_k of asset-class k (both in decimal), the corresponding log-space parameters are [20]

$$\sigma_{Z,k}^2 = \ln \left(1 + \left(\frac{s_k}{1 + m_k} \right)^2 \right), \quad (9.6)$$

$$\mu_{Z,k} = \ln(1 + m_k) - \frac{1}{2} \sigma_{Z,k}^2. \quad (9.7)$$

One can verify that $\mathbb{E}[\exp(Z_k)] - 1 = m_k$ and $\text{Var}[\exp(Z_k) - 1] = s_k^2$ when $Z_k \sim \mathcal{N}(\mu_{Z,k}, \sigma_{Z,k}^2)$.

Sampling A log-space covariance matrix is assembled from the log-space standard deviations $\sigma_{Z,k} = \sqrt{\sigma_{Z,k}^2}$ and the Pearson correlation matrix \mathbf{C} with entries $\rho_{kk'}$,

$$(\boldsymbol{\Sigma}_Z)_{kk'} = \rho_{kk'} \sigma_{Z,k} \sigma_{Z,k'}. \quad (9.8)$$

Then N_n vectors of log-returns are drawn jointly,

$$\mathbf{Z}_n \stackrel{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{\mu}_Z, \boldsymbol{\Sigma}_Z), \quad n = 0, \dots, N_n - 1, \quad (9.9)$$

and annual returns are recovered by exponentiation,

$$\tau_{k,n} = \exp(Z_{k,n}) - 1. \quad (9.10)$$

Applying the correlation matrix directly in log-space is a standard approximation, exact for jointly log-normal variates and accurate for moderate return magnitudes.

lognormal — user-supplied parameters The user provides arithmetic means m_k and standard deviations s_k (in percent) together with Pearson correlations $\rho_{kk'}$. These are converted to log-space via Eqs. (9.6)–(9.7) before sampling.

histolognormal — history-calibrated parameters Parameters are estimated directly in log-space from the selected historical window without any intermediate conversion to arithmetic parameters. Let $\ell_{ky} = \ln(1 + r_{ky})$ be the log-return of asset class k in calendar year y . Then

$$\mu_{Z,k} = \frac{1}{N_w} \sum_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}} \ell_{ky}, \quad (9.11)$$

$$(\boldsymbol{\Sigma}_Z)_{kk'} = \frac{1}{N_w - 1} \sum_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}} (\ell_{ky} - \mu_{Z,k})(\ell_{k'y} - \mu_{Z,k'}). \quad (9.12)$$

Arithmetic statistics displayed in the UI are recovered from the log-space moments via the inverse of Eqs. (9.6)–(9.7):

$$m_k = \exp\left(\mu_{Z,k} + \frac{1}{2}(\Sigma_Z)_{kk}\right) - 1, \quad (9.13)$$

$$s_k = (1 + m_k)\sqrt{\exp((\Sigma_Z)_{kk}) - 1}. \quad (9.14)$$

9.4.3 Bootstrap sequence-of-returns: `bootstrap_sor`

The bootstrap sequence-of-returns model resamples N_n annual return vectors directly from the observed window $\{\mathbf{r}_y\}_{y=\mathcal{Y}_{\text{frm}}}^{\mathcal{Y}_{\text{to}}}$, where $\mathbf{r}_y = (r_{0,y}, \dots, r_{N_k-1,y})^\top$. Resampling from actual data preserves the empirical marginal distributions exactly, including skewness, fat tails, and any departure from normality. Four resampling strategies are supported.

IID bootstrap Each year is drawn independently and with replacement from the full window [21]:

$$\boldsymbol{\tau}_n = \mathbf{r}_{y_n}, \quad y_n \stackrel{\text{iid}}{\sim} \text{Uniform}\{\mathcal{Y}_{\text{frm}}, \dots, \mathcal{Y}_{\text{to}}\}. \quad (9.15)$$

This is the simplest and default variant; it destroys all serial correlation present in the original data.

Block bootstrap To preserve short-run serial correlation, the block bootstrap of Künsch [22] draws $\lceil N_n/b \rceil$ starting positions s_ℓ uniformly from $\{\mathcal{Y}_{\text{frm}}, \dots, \mathcal{Y}_{\text{to}} - b + 1\}$ and concatenates the corresponding blocks of b consecutive years, truncating to N_n years:

$$(\boldsymbol{\tau}_0, \dots, \boldsymbol{\tau}_{N_n-1}) = (\mathbf{r}_{s_1}, \dots, \mathbf{r}_{s_1+b-1}, \mathbf{r}_{s_2}, \dots, \mathbf{r}_{s_2+b-1}, \dots). \quad (9.16)$$

The parameter b is the block size, chosen by the user.

Circular block bootstrap A variant due to Politis and Romano [23] treats the historical window as circular, wrapping from \mathcal{Y}_{to} back to \mathcal{Y}_{frm} . Starting positions are drawn uniformly from the full window $\{\mathcal{Y}_{\text{frm}}, \dots, \mathcal{Y}_{\text{to}}\}$, and block membership wraps around modulo N_w . This removes the boundary effect present in the block bootstrap, where observations near the endpoints of the window are underrepresented.

Stationary bootstrap In the stationary bootstrap of Politis and Romano [24], block lengths are random rather than fixed. Each block length L_ℓ is drawn from a geometric distribution,

$$P(L_\ell = \ell) = p_b (1 - p_b)^{\ell-1}, \quad \ell = 1, 2, \dots, \quad (9.17)$$

with parameter $p_b = 1/b$, so that the expected block length equals b . This choice produces a strictly stationary bootstrap distribution, resolving the boundary discontinuities of the fixed-block method while still approximately preserving serial correlation at lags up to order b .

9.4.4 Vector autoregression: var

The VAR(1) model [25, 26] extends the multivariate normal approach by explicitly capturing year-to-year serial correlations in asset returns, such as momentum and mean-reversion. The generating model is

$$\mathbf{y}_t = \mathbf{c} + \mathbf{A}\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad (9.18)$$

where $\mathbf{y}_t \in \mathbb{R}^{N_k}$ is the vector of asset-class returns at time t , $\mathbf{c} \in \mathbb{R}^{N_k}$ is a constant intercept vector, $\mathbf{A} \in \mathbb{R}^{N_k \times N_k}$ is the transition matrix encoding one-period predictability, and $\boldsymbol{\Sigma}$ is the contemporaneous residual covariance matrix. For $\mathbf{A} = \mathbf{0}$, Eq. (9.18) reduces to the multivariate normal model of the **histogaussian** method.

Parameter estimation Given N_w historical observations $\mathbf{y}_1, \dots, \mathbf{y}_{N_w}$, parameters \mathbf{c} and \mathbf{A} are estimated jointly by ordinary least squares (OLS). Stacking observations into matrices

$$Y = [\mathbf{y}_2 \mid \dots \mid \mathbf{y}_{N_w}]^\top \in \mathbb{R}^{(N_w-1) \times N_k}, \quad (9.19)$$

$$X = [\mathbf{1} \mid \mathbf{y}_1 \mid \dots \mid \mathbf{y}_{N_w-1}]^\top \in \mathbb{R}^{(N_w-1) \times (N_k+1)}, \quad (9.20)$$

the OLS solution $B = (X^\top X)^{-1} X^\top Y$ yields the intercept $\mathbf{c} = B[0, :]$ and transition matrix $\mathbf{A} = B[1 :, :]^\top$. The residual covariance is estimated as

$$\boldsymbol{\Sigma} = \frac{(Y - XB)^\top (Y - XB)}{N_w - N_k - 1}. \quad (9.21)$$

Stationarity and spectral shrinkage Stationarity of the VAR(1) process requires that all eigenvalues of \mathbf{A} have modulus strictly less than one. Let $\rho(\mathbf{A}) =$

$\max_i |\lambda_i(\mathbf{A})|$ denote the spectral radius. If $\rho(\mathbf{A}) \geq 0.95$ (near or above the unit circle), the transition matrix is rescaled as

$$\mathbf{A} \leftarrow \frac{0.95}{\rho(\mathbf{A})} \mathbf{A}, \quad (9.22)$$

ensuring $\rho(\mathbf{A}) = 0.95 < 1$ while preserving the direction of every eigenmode. This spectral shrinkage step can be disabled by the user, at the cost of potentially non-stationary sequences.

Sequence generation Starting from $\mathbf{y}_0 = \boldsymbol{\mu}$ (the historical sample mean), sequences are generated recursively from Eq. (9.18) using $\boldsymbol{\varepsilon}_t = L\mathbf{z}_t$, where L is the Cholesky factor of $\boldsymbol{\Sigma}$ and $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, I)$. The VAR(1) model produces synthetic sequences of arbitrary length while capturing both the cross-sectional covariance structure and the serial autocorrelation of asset returns.

9.4.5 DCC-GARCH(1,1): garch_dcc

The DCC-GARCH(1,1) model [27, 28] extends the multivariate normal approach by allowing both the per-asset conditional variance and the cross-asset correlation matrix to vary over time. It is fitted in two steps on the selected historical window.

Step 1 — per-asset GARCH(1,1) Let $\varepsilon_{kt} = y_{k,t} - \mu_k$ be the demeaned return of asset class k at year t . The conditional variance h_{kt} evolves as

$$h_{kt} = \omega_k + \alpha_k \varepsilon_{k,t-1}^2 + \beta_k h_{k,t-1}, \quad (9.23)$$

with $h_{k,0} = \text{Var}(\varepsilon_k)$. The parameters $(\omega_k, \alpha_k, \beta_k)$ are estimated independently for each asset by maximizing the Gaussian log-likelihood

$$\ell_k = -\frac{1}{2} \sum_{t=1}^{N_w} [\ln h_{kt} + \varepsilon_{kt}^2 / h_{kt}], \quad (9.24)$$

subject to $\omega_k > 0$, $\alpha_k, \beta_k > 0$, and the covariance-stationarity constraint $\alpha_k + \beta_k < 1$. Optimization uses L-BFGS-B; if convergence fails, an EWMA fallback ($\alpha_k = 0.06$, $\beta_k = 0.94$) is used instead.

The standardized residuals $z_{kt} = \varepsilon_{kt} / \sqrt{h_{kt}}$ are passed to the second step.

Step 2 — DCC dynamics Let $\overline{\mathbf{Q}} = \frac{1}{N_w} \mathbf{Z}^\top \mathbf{Z}$ be the unconditional covariance of $\mathbf{Z} = [z_{kt}]$. The quasi-correlation matrix \mathbf{Q}_t evolves as

$$\mathbf{Q}_t = (1 - a - b) \overline{\mathbf{Q}} + a \mathbf{z}_{t-1} \mathbf{z}_{t-1}^\top + b \mathbf{Q}_{t-1}, \quad (9.25)$$

with $a, b > 0$ and $a + b < 1$. The time-varying correlation matrix is obtained by normalizing:

$$\mathbf{R}_t = \text{diag}(\mathbf{Q}_t)^{-1/2} \mathbf{Q}_t \text{diag}(\mathbf{Q}_t)^{-1/2}. \quad (9.26)$$

The DCC parameters (a, b) are estimated by maximizing the concentrated log-likelihood

$$\ell_{\text{DCC}} = -\frac{1}{2} \sum_{t=1}^{N_w} [\ln |\mathbf{R}_t| + \mathbf{z}_t^\top \mathbf{R}_t^{-1} \mathbf{z}_t]. \quad (9.27)$$

Sequence generation Warm-start states \mathbf{Q}_0 and $\mathbf{h}_0 = (h_{1,N_w}, \dots, h_{N_k,N_w})$ are obtained from a forward pass through the historical data. Each simulated year is then drawn as

$$\mathbf{y}_t = \boldsymbol{\mu} + \text{diag}(\mathbf{h}_t)^{1/2} L(\mathbf{R}_t) \mathbf{z}_t, \quad \mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, I), \quad (9.28)$$

where $L(\mathbf{R}_t)$ is the lower Cholesky factor of \mathbf{R}_t , followed by GARCH and DCC updates via Eqs. (9.23) and (9.25). This produces sequences that exhibit *volatility clustering* — large shocks in one year increase conditional variance in the next — and *time-varying cross-asset correlations* that spike during simulated market-stress episodes, consistent with empirical evidence in financial return data.

9.5 User-supplied rates: dataframe

The `dataframe` method allows the user to bypass all built-in rate models and supply the return matrix $\{\tau_{kn}\}$ directly as a pandas `DataFrame` (one column per asset class, one row per plan year). This is useful for custom scenarios, external forecasts, or deterministic stress tests not expressible through the other methods. When the `DataFrame` has fewer rows than the plan horizon N_n , the sequence is repeated cyclically, matching the behavior of the `historical` method. No statistical estimation or sampling is performed; the user bears full responsibility for the consistency and units of the supplied data. This method is accessed programmatically via the Python API (`setRates("dataframe", ...)`); it is not available through the TOML configuration file.

9.6 Monte Carlo and stress testing

In Monte Carlo mode, Owl calls one of the stochastic models (`histogaussian`, `histolognormal`, `gaussian`, `lognormal`, `bootstrap_sor`, `var`, or `garch_dcc`) independently N_s times, each time solving the full optimization problem with a different rate scenario. The distribution of resulting optimal spending or bequest values across the N_s scenarios provides a probabilistic assessment of the retirement plan.

In stress-testing mode (Historical Range), the `historical` method is applied repeatedly: the starting year \mathcal{Y}_{frm} is advanced one year at a time over a specified calendar range, producing one deterministic solve per starting year (by default). This sweep is equivalent to asking: “How would this plan have fared had the retiree begun at every historical starting year?” When the *augmented* option is enabled, the run is expanded: for each starting year, every *(reverse, roll)* pair in $\{\text{false}, \text{true}\} \times \{0, \dots, N_n - 1\}$ is executed, yielding $2N_n$ sequences per year instead of one. For a 30-year plan over 50 years of history, that produces $50 \times 60 = 3000$ outcomes rather than 50, providing a much denser sample for the result histogram while still using only observed historical returns. Reverse and roll (see § 9.1) alter the ordering of the same returns; augmented mode systematically explores all such orderings.

A. Source File Reference

The Owl Python package is located under `src/owlplanner/`. The table below summarizes the primary source files and their roles.

<code>plan.py</code>	The core optimization engine. Constructs and solves the LP/MIP using HiGHS (default, via <code>highspy</code>) or MOSEK directly. Implements the self-consistent iteration loop, all constraint-building methods, the objective functions, and post-solution analysis. This is the largest and most complex file.
<code>abcap.py</code>	A solver-neutral constraint-building API used by <code>plan.py</code> . Provides an abstract interface for adding rows to the constraint matrix, keeping the model description independent of any specific LP solver.
<code>tax2026.py</code>	Federal tax calculations: ordinary income brackets and standard deductions, LTCG tax thresholds, SS taxability formula, NIIT, ACA Premium Tax Credit (with <code>acaCosts</code> and <code>acaVals</code>), and Medicare IR-MAA premiums and bracket thresholds. Also computes RMD fractions from the IRS Uniform Lifetime Table.
<code>socialsecurity.py</code>	Social Security benefit modeling: FRA schedules, own-benefit and spousal adjustment factors, spousal and survivor benefit computation, and optional benefit trim for policy scenarios.
<code>timelists.py</code>	Reads and validates the Wages and Contributions time-horizon spreadsheet (one tab per individual) and the Fixed Assets and Debts tables. Returns year-indexed data structures consumed by <code>plan.py</code> .
<code>fixedassets.py</code>	Models fixed-income assets (real estate, annuities, collectibles, etc.),

	computing year-by-year income, proceeds, and capital gains arising from assets listed in the Fixed Assets table.
<code>debts.py</code>	Models debt obligations (loans, mortgages), computing annual payment schedules from the Debts table.
<code>rates.py</code>	Loads and pre-processes the historical market-return data from <code>data/rates.csv</code> (equity, bond, and inflation series from 1928 to present) and provides utilities for constructing rate scenarios.
<code>rate_models/</code>	A pluggable system of rate model classes inheriting from <code>BaseRateModel</code> . Built-in models (historical, fixed, gaussian, histogaussian, lognormal, histolognormal, bootstrap SOR, VAR, GARCH) are in <code>builtin.py</code> and <code>_builtin_impl.py</code> . Custom rate models are discoverable at run-time if they subclass <code>BaseRateModel</code> and are importable.
<code>config/schema.py</code>	Pydantic v2 models defining the full case configuration schema. All plan parameters (account balances, claiming ages, rate method, options, etc.) are validated here before being passed to <code>plan.py</code> .
<code>config/toml_io.py</code>	TOML file loading and saving for <code>.toml</code> case files. Also handles backward compatibility for older file formats.
<code>plotting/</code>	Visualization backends: <code>plotly_backend.py</code> for interactive plots in the Streamlit UI, and <code>matplotlib_backend.py</code> for static plots used in the CLI and Jupyter notebooks.
<code>cli/</code>	A Click-based command-line interface registered as the <code>owlcli</code> entry point. Supports <code>list</code> (enumerate available case files) and <code>run</code> (solve a case from a TOML file) commands.

References

- [1] M.-D. Lacasse, “Optimal Wealth Lab (Owl),” GitHub repository, 2025. <https://github.com/mdlacasse/Owl>.
- [2] Social Security Administration, *Program Operations Manual System (POMS)*, RS 00615.003 (Full Retirement Age), RS 00615.015 (SSA age computation), <https://secure.ssa.gov/poms.nsf/lnx/0300615003>.
- [3] Internal Revenue Service, *Publication 915: Social Security and Equivalent Railroad Retirement Benefits*, <https://www.irs.gov/pub/irs-pdf/p915.pdf>.
- [4] Internal Revenue Code § 408A, “Roth Individual Retirement Accounts,” <https://www.law.cornell.edu/uscode/text/26/408A>.
- [5] Internal Revenue Service, *Publication 590-B: Distributions from Individual Retirement Arrangements (IRAs)*, <https://www.irs.gov/pub/irs-pdf/p590b.pdf>.
- [6] Internal Revenue Code § 223, “Health Savings Accounts,” <https://www.law.cornell.edu/uscode/text/26/223>.
- [7] Internal Revenue Service, *Publication 969: Health Savings Accounts and Other Tax-Favored Health Plans*, <https://www.irs.gov/pub/irs-pdf/p969.pdf>.
- [8] Internal Revenue Service, *Form 8889: Health Savings Accounts (HSAs)*, <https://www.irs.gov/pub/irs-pdf/f8889.pdf>.
- [9] Internal Revenue Code § 36B, “Refundable credit for coverage under a qualified health plan,” <https://www.law.cornell.edu/uscode/text/26/36B>.
- [10] Internal Revenue Service, *Rev. Proc. 2024-35*, “Applicable Percentage Table for the Premium Tax Credit for Taxable Years Beginning in Calendar Year 2025,” 89 Fed. Reg. 66278 (2024).

- [11] Internal Revenue Service, *Rev. Proc. 2025-25*, “Applicable Percentage Table for the Premium Tax Credit for Taxable Years Beginning in Calendar Year 2026,” <https://www.irs.gov/pub/irs-drop/rp-25-25.pdf>.
- [12] U.S. Department of Health and Human Services, “Federal Poverty Guidelines” (annual), <https://aspe.hhs.gov/poverty-guidelines>.
- [13] Centers for Medicare & Medicaid Services, “2026 Medicare Parts A & B Premiums and Deductibles,” <https://www.cms.gov/newsroom/fact-sheets/2026-medicare-parts-b-premiums-deductibles>.
- [14] Internal Revenue Service, “Topic No. 409: Capital Gains and Losses,” <https://www.irs.gov/taxtopics/tc409>.
- [15] Internal Revenue Code § 1411, “Imposition of Tax (Net Investment Income),” <https://www.law.cornell.edu/uscode/text/26/1411>.
- [16] Internal Revenue Service, *Publication 17: Your Federal Income Tax*, <https://www.irs.gov/pub/irs-pdf/p17.pdf>.
- [17] One Big Beautiful Bill Act (OBBBA), signed July 4, 2025. Extends and modifies individual income tax provisions; adds § 1002 senior deduction.
- [18] SECURE 2.0 Act of 2022, Pub. L. No. 117-328, Division T, 136 Stat. 5966.
- [19] R. J. Shiller, *Irrational Exuberance*, Princeton University Press, Princeton, NJ, 2000. Historical return data available at <http://www.econ.yale.edu/~shiller/data.htm>.
- [20] J. C. Hull, *Options, Futures, and Other Derivatives*, 10th ed., Pearson, Hoboken, NJ, 2017. (Chapters 15–16 derive the log-normal / GBM return model and the arithmetic-to-log-space parameter conversion.)
- [21] B. Efron, “Bootstrap methods: another look at the jackknife,” *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979.
- [22] H. R. Künsch, “The jackknife and the bootstrap for general stationary observations,” *The Annals of Statistics*, vol. 17, no. 3, pp. 1217–1241, 1989.
- [23] D. N. Politis and J. P. Romano, “A circular block resampling procedure for stationary data,” in *Exploring the Limits of Bootstrap*, R. LePage and L. Billard, Eds. Wiley, New York, 1992, pp. 263–270.

- [24] D. N. Politis and J. P. Romano, “The stationary bootstrap,” *Journal of the American Statistical Association*, vol. 89, no. 428, pp. 1303–1313, 1994.
- [25] J. Y. Campbell and L. M. Viceira, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, Oxford, 2002.
- [26] H. Lütkepohl, *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin, 2005.
- [27] R. F. Engle, “Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models,” *Journal of Business & Economic Statistics*, vol. 20, no. 3, pp. 339–350, 2002.
- [28] T. Bollerslev, “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, 1986.