

A SPENDING-FIRST STEERING FRAMEWORK FOR SUSTAINABLE RETIREMENT MODELLING

Author: Nick May

OKINAWA ZEST LIMITED [Company address]

Abstract

Retirement planning involves uncertainty, multiple income sources, and financial resources that evolve over time. The sustainability assessments of traditional portfolio-centric models can vary sharply because all adjustments are driven by the liquid portfolio rather than the household's broader financial position. This paper presents a holistic, spending-first modelling framework that incorporates all relevant components of net worth, including the net present value of future guaranteed income, and adjusts discretionary spending in response to deviations from an expected path.

We introduce a 'steering' mechanism, used in this framework solely as a modelling device rather than a behavioural rule. Its purpose is to produce spending trajectories that respond coherently to changes in long-term financial capacity and that remain relatively stable in the presence of market noise. By adjusting the variable that households naturally monitor, namely discretionary spending, the model generates projections that are easier to interpret and that align more closely with the way advisers and households discuss long-term plans.

The framework is intended to support structured exploration of uncertainty, to illustrate the scale and direction of adjustments that might arise from annual re-assessment, and to provide a clearer basis for communication. It does not predict behaviour or prescribe specific actions, but offers a coherent representation of how lifestyle might adapt under uncertainty.

1. Introduction

Retirement planning has long centred on the challenge of sustaining a desired standard of living over an uncertain future. Households must make decisions about spending, saving, and drawing down assets in the face of unpredictable investment returns, inflation, longevity, and the timing of income streams. Financial planners and researchers have developed a wide range of frameworks to address these uncertainties, yet most share a common underlying structure: they begin with the investment portfolio and work outward from there.

The dominant question in both academic research and financial-planning practice has been: *How much can a retiree safely withdraw from their portfolio each year?* This portfolio-first perspective has shaped decades of thinking. The well-known "4% rule" exemplifies this approach by proposing a fixed, inflation-adjusted withdrawal rate intended to last for a typical retirement horizon. Later research introduced dynamic withdrawal strategies, such as the guardrails framework proposed by Guyton and Klinger, which adjusts withdrawals in response to portfolio performance. These innovations improved flexibility and reduced the rigidity of fixed withdrawal rules, but they retained the same foundational assumption: that retirement sustainability is fundamentally a problem of managing withdrawals from a portfolio.

In practice, however, retirees do not experience their financial lives through the lens of withdrawal rates. They experience them through spending. Spending is the objective; withdrawals are merely one of the mechanisms used to support it. Moreover, a portfolio is only one component of a retiree's financial position. Most households rely on a combination of pensions, guaranteed income streams, tax-advantaged accounts, property, and cash reserves. A framework that focuses solely on portfolio withdrawals risks overlooking the broader financial capacity available to support retirement spending.

While the professional community has increasingly embraced dynamic guardrails to bridge this gap, these rules are almost universally applied in a portfolio-centric silo. By ignoring the broader financial ecosystem, traditional models overstate the volatility of a client's true wealth, leading to overly conservative initial spending recommendations. An embedded steering framework corrects this distortion by anchoring systematic adjustments to holistic net worth rather than the liquid portfolio alone. Because the model accounts for the stabilizing effect of lifetime guaranteed income and other non-portfolio assets, it can safely unlock a higher, more optimal initial discretionary spending target at the outset of retirement, without altering the underlying risk profile.

In this paper, we propose a different organising principle for retirement sustainability. Instead of beginning with the portfolio and determining how much can be withdrawn, we begin with spending and determine how it can be supported holistically. We introduce a steering model that applies flexibility not to withdrawals but to discretionary spending, and we base these adjustments not on portfolio value alone but on total net worth, including the net present value of future fixed income streams. This approach aligns more closely with lifecycle economics, which emphasises consumption smoothing, and it reflects the real-world complexity of household finances.

Our contribution is twofold. First, we frame the steering concept in a way that places spending at the centre of the model, consistent with how households naturally make decisions. Second, we formalise a net-worth-based steering framework that provides a more stable and comprehensive measure of financial capacity than portfolio value alone. By integrating these ideas into a coherent mathematical and simulation-based structure, we offer a new paradigm for sustainable retirement planning – one that is both theoretically grounded and practically applicable.

This framework is not purely theoretical; it has been implemented in practice within the EvolveMyRetirement® planning tool, where it supports long-term modelling across a wide range of household circumstances.

2. Background and Literature Review

Retirement-planning research has evolved over several decades, shaped by attempts to balance simplicity, sustainability, and behavioural realism. In this section, we review the major strands of the literature that inform our work: fixed real withdrawal rules, dynamic withdrawal strategies, and lifecycle-economics models of consumption. We also highlight the gap between academic theory and financial-planning practice, which motivates the framework we propose.

2.1. Fixed Real Withdrawal Rules: The 4% Rule

Modern withdrawal-rate research is often traced to Bengen [1], who analysed historical US market data to identify a withdrawal rate that would have survived the worst observed sequence of returns. His conclusion – that a retiree could withdraw 4% of their initial portfolio value, adjusted annually for inflation – became widely known as the “4% rule”.

The appeal of the rule lay in its simplicity. It provided a single number that appeared to offer safety across a wide range of historical scenarios. However, its limitations are well documented:

- It assumes a fixed, inflation-linked spending path.
- It is sensitive to the specific historical period analysed.
- It ignores the presence of other income sources.
- It treats the portfolio as the sole determinant of sustainability.
- It does not adapt to changing circumstances.

Despite these limitations, the 4% rule remains influential in both public discourse and adviser practice, largely because of its simplicity and ease of communication.

2.2. Dynamic Withdrawal Strategies: Guardrails and Beyond

Recognising the rigidity of fixed real withdrawals, researchers developed dynamic withdrawal strategies that adjust spending in response to portfolio performance. Among these, the guardrails framework proposed by Guyton and Klinger [2] is one of the most widely cited.

Their approach introduced rules for adjusting withdrawals:

- Increase withdrawals when the portfolio grows sufficiently above its initial value.
- Decrease withdrawals when the portfolio falls sufficiently below its initial value.
- Constrain the frequency and magnitude of adjustments.
- Incorporate decision rules for inflation adjustments and portfolio composition.

This framework improved on the 4% rule by allowing withdrawals to flex in response to market conditions. It reduced the risk of premature portfolio depletion while still permitting lifestyle improvements in favourable markets.

Guyton and Klinger originally presented their guardrails as a real-world decision framework for retirees to apply annually. In practice, however, advisers rarely implement the rules literally for their clients. Most either re-baseline the plan each year using fresh projections or use the guardrails purely as a modelling tool. Within Monte Carlo simulations, guardrails can be implemented as heuristic rules applied along each projection path, functioning as a proxy for annual re-baselining.

Importantly, a retiree cannot simultaneously follow the guardrails mechanically and also re-baseline annually, as the two approaches are logically incompatible. This distinction highlights the need for a framework that clearly separates modelling heuristics from real-world decision processes.

2.3. Lifecycle Economics and Consumption Smoothing

While withdrawal-based research focuses on determining how much can safely be taken from a portfolio, the lifecycle-economics literature begins from a different premise: households seek to maintain a stable standard of living over time. This idea, rooted in the life-cycle hypothesis of Modigliani and Brumberg [3] and the permanent-income hypothesis of Friedman [4], asserts that individuals base their consumption decisions not on current income or portfolio value alone, but on their expectations of lifetime resources.

Lifecycle-economics models begin from the normative assumption that households aim to smooth their consumption over time, maintaining a relatively stable standard of living despite fluctuations in income. In practice, real-world behaviour varies widely: some households adjust spending almost fully in response to income changes, while others smooth consumption more

strongly. The lifecycle framework therefore represents an idealised benchmark rather than a descriptive model of typical behaviour.

Modern lifecycle models, including those developed by Kotlikoff et al. [5], formalise consumption smoothing using stochastic dynamic programming. These models incorporate uncertainty in returns, longevity and income, but the optimisation itself is deterministic. In contrast, some contemporary tools, including our own, perform strategy generation directly over simulated stochastic paths, allowing the model to capture a wider range of real-world financial complexities. Both dynamic-programming and simulation-based approaches face significant computational challenges: dynamic programming suffers from the curse of dimensionality, while simulation-based optimisation can become expensive when Monte Carlo evaluation is embedded inside a search procedure. In practice, both approaches rely on approximation and algorithmic techniques to remain tractable, particularly when modelling high-dimensional financial decisions.

2.4. Gaps Between Theory and Practice

The literature reveals a clear divide between how retirement spending is modelled and how it is managed in practice. Withdrawal-based approaches remain centred on the portfolio, using its value or derived withdrawal rates as the primary determinant of spending. This structure is straightforward to model and easy to communicate, but it does not reflect the broader set of resources that typically support retirement or the way households think about their standard of living.

Lifecycle-economics models take the opposite perspective by placing spending at the centre of the analysis and defining financial capacity in terms of lifetime resources. Although this framework aligns more closely with how households might rationally organise their finances over time, the underlying optimisation structure is difficult to communicate clearly in adviser-client settings and often impractical to implement at scale. As a result, these models have had limited influence on mainstream financial planning.

In practice, advisers and planning tools tend to rely on simple heuristics such as fixed withdrawal rates, portfolio-based guardrails or probability-of-success thresholds. These methods are operationally convenient but remain fundamentally portfolio-centric and do not integrate the full balance sheet or the behavioural considerations that shape real-world retirement decisions. They also provide no coherent way to incorporate periodic re-baselining, even though re-baselining is essential as new information replaces old uncertainty.

These gaps point to the need for a framework that is both consumption-centric and practical: one that reflects the household's total financial capacity, supports periodic re-baselining, and remains intuitive for advisers and consumers. The next section introduces such a framework.

3. Conceptual Framework

3.1. Spending as the Primary Objective

Spending is not the only lifetime objective; many households also value leaving a legacy. But spending is the only variable that can be chosen directly, whereas legacy is an outcome shaped by investment returns, longevity, and other uncertainties. For this reason, spending serves as

the primary objective in the modelling framework, with legacy treated as a consequence of those choices rather than as an independently controllable target.


3.1.1. The causal inversion: spending → drawdown

Traditional models assume:

Portfolio → Withdrawal Rate → Spending

But in reality:

Desired Lifestyle → Spending → Required Withdrawals → Portfolio Trajectory



This inversion matters because:

- Withdrawals are an output, not an input.
- Optimising the wrong variable leads to distorted decisions (e.g., underspending due to fear of running out).

3.1.2. Essential vs discretionary spending

A robust framework must distinguish:

- **Essential spending:** Non-negotiable baseline (housing, food, utilities, healthcare).
- **Discretionary spending:** Lifestyle-enhancing but flexible (travel, hobbies, gifts).

This distinction is crucial because:

- Essential spending must be protected with high confidence.
- Discretionary spending can flex with changes in financial capacity.
- Treating all spending as equally flexible produces unrealistic or overly conservative plans.

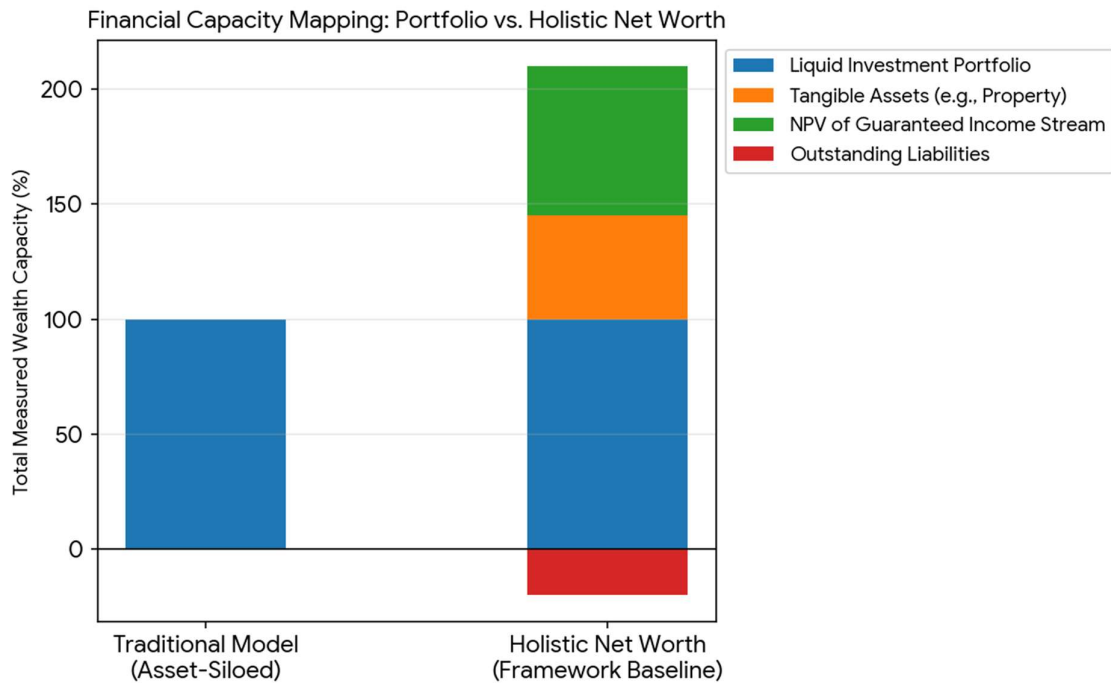
3.2. Holistic Net Worth

3.2.1. Definition

Portfolio value alone is an incomplete and often misleading measure of retirement readiness. A more accurate measure is holistic net worth, which is a broader measure of financial capacity. It takes into account:

- Financial assets.
- Tangible assets such as property or valuable possessions.

- The net present value (NPV) of guaranteed income streams, including pensions, annuities and state benefits.
- Outstanding liabilities.



3.2.2. Inclusion of NPV of guaranteed income

Guaranteed income streams such as defined-benefit pensions, annuities and state benefits are included in holistic net worth through their net present value. This reflects the fact that these income streams can be replicated, at least in principle, by purchasing an annuity with capital. Including their NPV prevents the model from interpreting annuitisation as a reduction in financial capacity and ensures that steering assessments remain consistent before and after the conversion of capital into guaranteed income. Non-guaranteed income is not included in this measure, as it is time-limited and already modelled directly within the simulation.

Guaranteed income can be valued using standard present-value methods or market-based annuity pricing. The choice of method does not affect the conceptual role of guaranteed income within holistic net worth.

3.3. Steering as a Modelling Heuristic

In this framework, *steering* is used solely as a modelling heuristic within the simulation. It provides a continuous, proportional mechanism that adjusts discretionary spending in line with changes in financial capacity, keeping the spending path aligned with the evolving wealth trajectory. This allows the simulation to show how spending might evolve under different future scenarios.

3.3.1. Conceptual rationale

The steering paradigm works because:

- It recognises that future outcomes are uncertain and incorporates this uncertainty directly into the simulation.
- It provides a continuous, proportional adjustment mechanism rather than relying on thresholds or discrete triggers.
- It prevents modelled projections from drifting into unrealistic extremes when financial capacity diverges from expectations.
- It is designed purely as a modelling device, not a behavioural rule, so adjustments occur mechanically whenever the simulated wealth path departs from the mid-case.

In essence, the paradigm maintains coherence in the model by continuously steering spending towards the mid-case wealth trajectory as uncertainty unfolds.

3.3.2. Why spending, not withdrawals, should flex

Traditional guardrail systems adjust withdrawals, but this places the focus on the funding mechanism rather than the objective. In a spending-first framework, spending is the primary quantity of interest, and withdrawals simply finance it.

Flexing spending instead of withdrawals:

- aligns the model with the underlying objective of sustaining a chosen lifestyle;
- allows essential spending to remain stable while discretionary spending adjusts; and
- avoids the false impression that a fixed withdrawal rate can remain appropriate as financial capacity changes.

A spending-based steering system therefore models how spending evolves as financial capacity diverges from expectations, producing projections that are more coherent and economically grounded.

While steering functions strictly as a modelling device within the simulation architecture, it fundamentally mirrors the dynamic 'guardrails' conversations that progressive financial planners already conduct with clients. Rather than forcing clients into a rigid, artificial spending rule, this approach provides a robust framework for managing real-world expectations. It translates abstract stochastic risk into tangible, manageable course corrections – a paradigm shift that transforms retirement planning from a static projection into an active, responsive navigation system.

4. Mathematical Formulation

4.1. Notation

This section introduces the notation used throughout the mathematical formulation. The aim is to define a consistent vocabulary for assets, liabilities, income streams and holistic net worth, without imposing any specific modelling assumptions beyond those required for clarity.

4.1.1. Income and spending

We define the following, all at time t , and all in real terms:

- Y_t = total income received

- G_t = lifetime guaranteed income received (see 4.1.3 for definition)
- E_t = essential spending
- \bar{C}_t = average-case discretionary spending
- C_t = realised discretionary spending

4.1.2. Assets and liabilities

We define the following, all at time t , and all in real terms:

- A_t = the value of financial assets
- T_t = the value of tangible assets
- L_t = outstanding liabilities
- Γ_t = NPV of future guaranteed income
- W_t = the holistic net worth (see 4.1.4 for definition)

4.1.3. Lifetime guaranteed income

Lifetime guaranteed income, denoted by G_t , represents income from sources that are contractually secure, not subject to market or employment risk, and guaranteed to continue to be received for the rest of the life of the recipient. Examples are state pensions, defined benefit pensions and lifetime annuities.

4.1.4. Holistic net worth

Holistic net worth, denoted by W_t , represents the combined value of financial assets, tangible assets, the NPV of guaranteed income streams, less outstanding liabilities. The formula is:

$$W_t = A_t + T_t + \Gamma_t - L_t$$

For the rationale of including the NPV of guaranteed income streams, see section 3.2.2.

4.1.5. Expected net worth

Let $\mathbb{E}[W_t]$ denote the expected value of holistic net worth at time t under the model's average-case assumptions.

This expected trajectory provides a deterministic reference path used later when defining the steering reference path in Section 4.3.

4.2. NPV of Future Income

4.2.1. Definition

The NPV of a guaranteed income stream G , evaluated at time t , is written as:

$$\Gamma_t = \sum_{s=t}^{\infty} \frac{G_s}{(1+r)^{s-t}}$$

where r is the real discount rate.

This expression represents the economic value, at time t , of a sequence of future real payments.

4.2.2. Discounting assumptions

The discount rate r reflects the real rate of time preference used for valuation. No specific value is assumed here. The formulation is compatible with:

- Constant real discounting.
- Term-structure-based discounting.
- Any internally consistent choice of real discount rate.

The framework does not require a particular calibration method.

4.2.3. Treatment of inflation

All quantities in the formulation are expressed in real terms. If nominal income amounts G_s^{nom} are used, they can be converted to real terms via:

$$G_s = \frac{G_s^{\text{nom}}}{\Pi_s}$$

where Π_s is the cumulative inflation index between time t and time s .

This ensures that the NPV calculation is performed on a consistent real basis.

4.3. Steering Reference Path

In the modelling framework, the expected net worth path $\{\mathbb{E}[W_t]\}$ serves as the reference against which realised outcomes are compared. The purpose of this reference path is to provide a consistent baseline for determining how discretionary spending is adjusted (steered) in the model when financial conditions differ from the model's average case. The baseline is established only once, and reused across all iterations of a Monte Carlo simulation.

Adjustments to discretionary spending are proportional to the deviation between realised holistic net worth W_t and the reference path $\mathbb{E}[W_t]$. The steering paradigm uses a single reference trajectory rather than thresholds or bands.

The magnitude and direction of adjustments are determined by the spending adjustment rules described in Section 4.4.

4.3.1. Deterministic baseline

The sequence $\{\mathbb{E}[W_t]\}_{t \geq 0}$ forms the deterministic baseline path. This path reflects:

- Expected investment returns.
- Expected inflation.
- Expected income receipts.
- Planned spending levels.
- The amortisation of liabilities.

The expected path is defined over a planning horizon that extends to the maximum possible lifetime of the last surviving member of the plan, with average case mortality assumptions being used for the other member, if any.

4.3.2. Role of average case assumptions

The steering path of expected net worth depends on the model's chosen average case assumptions. These may include:

- Long term real return assumptions for asset classes.
- Long term inflation expectations.
- Expected changes in income or liabilities.
- Any other structural assumptions required for internal consistency.

The framework does not prescribe specific values for these assumptions. Their role is simply to generate a coherent baseline against which the stochastic simulation can measure deviations.

4.4. Spending Adjustment Rules

Spending adjustment rules specify how discretionary spending changes when realised holistic net worth differs from the steering reference path. Recalling that this steering mechanism serves strictly as a modelling device rather than a behavioural prescription (see Section 3.3), the mathematical mechanics operate as follows.

4.4.1. Average discretionary spending

\bar{C}_t is a deterministic real-terms spending path defined by the model. The simplest form is constant real spending, but other patterns may be used, such as tapering after a specified date.

4.4.2. Adjustment as a function of deviation

The adjustment mechanism maps the deviation between realised holistic net worth and the expected path to a deviation from the average-case spending level:

$$C_t = \bar{C}_t + f(W_t - \mathbb{E}[W_t])$$

The function f determines both the direction and magnitude of the adjustment. A positive deviation increases discretionary spending relative to the average case, while a negative deviation reduces it.

4.4.3. Proportional adjustment

A common formulation is a proportional rule:

$$C_t = \bar{C}_t \left(1 + k \cdot \frac{W_t - \mathbb{E}[W_t]}{\mathbb{E}[W_t]} \right)$$

where k is a non-negative parameter controlling sensitivity. This formulation scales naturally with wealth. However, in sufficiently adverse scenarios the adjustment factor may fall below zero, so the rule does not guarantee non-negative discretionary spending in all cases. The modelling framework therefore applies an explicit non-negativity constraint to ensure that discretionary spending cannot fall below zero, regardless of the adjustment factor.

4.4.4. Tempering upside adjustments

If the household has a legacy motive, it would be appropriate to temper upside adjustments. One approach is to scale the sensitivity to positive deviations by a factor that depends on the legacy motive parameter $L \in [0,1]$:

$$k_{\text{up}} = (1 - L) k, k_{\text{down}} = k$$

By dampening upside spending adjustments during market expansions ($k_{\text{up}} < k_{\text{down}}$), the model directly accommodates a client's legacy motives. Rather than mechanically inflating lifestyle consumption during market booms, which often yields diminishing marginal utility for retirees, the asymmetric sensitivity rule automatically captures and preserves some of these windfall gains. This builds a structural bridge between lifetime spending optimization and estate preservation, allowing planners to mathematically demonstrate how market upside is efficiently captured to fulfil a client's intergenerational wealth objectives.

4.4.5. Absolute adjustment

An alternative formulation expresses adjustments in absolute terms:

$$C_t = \bar{C}_t + h (W_t - \mathbb{E}[W_t])$$

where h maps deviations in net worth to absolute changes in spending.

Absolute rules are less likely to be used, as they can produce negative discretionary spending and do not scale naturally with wealth. They are included here only for completeness.

4.5. Drawdown as a Residual Funding Mechanism

Drawdown is treated as the mechanism that supplies whatever funding is required to meet the spending level implied by the adjustment rules. Our framework does not attempt to optimise withdrawals or prescribe a specific sequencing strategy. Instead, withdrawals arise mechanically from the interaction between spending, income, and the structure of the household's assets.

4.5.1. How withdrawals are derived

Y_t denotes the total income received at time t , including both guaranteed and non-guaranteed components. The total spending requirement at time t is:

$$S_t = E_t + C_t$$

Withdrawals are then defined residually as:

$$D_t = S_t - Y_t$$

If $Y_t \geq S_t$, no withdrawal is required and any surplus income is added to the household's assets. If $Y_t < S_t$, the shortfall is funded through withdrawals from the available asset accounts.

This formulation ensures that drawdown is a consequence of the spending rule rather than an independent decision variable.

4.5.2. Tax-aware sequencing

When a withdrawal is required, the model allocates it across the household's asset accounts in a tax-aware manner. The purpose is not to prescribe a specific tax strategy, but to ensure that the simulation reflects the fact that different accounts have different tax implications.

The sequencing mechanism respects the following principles:

- Withdrawals from tax-advantaged accounts may give rise to taxable income.
- Withdrawals from tax-free accounts do not.
- Withdrawals from taxable investment accounts may realise capital gains.
- The ordering of withdrawals affects the household's tax position in the current and future periods.

The model applies a sequencing rule that is internally consistent and sensitive to these tax characteristics. The specific ordering is not prescribed by the framework, and different implementations may adopt different approaches. The essential requirement is that the sequencing reflects the tax structure of the accounts and that withdrawals remain a residual mechanism driven by the spending requirement.

The framework does not prescribe a specific tax regime; implementations should reflect the relevant jurisdiction while maintaining the residual nature of withdrawals.

5. Simulation Methodology

5.1. Stochastic Return Modelling

Stochastic return modelling specifies how the liquid investment portfolio evolves over time within the simulation. Tangible assets such as property follow separate growth assumptions and are not included in the portfolio return distribution. The liquid portfolio is represented by a single risk level $r \in [0,1]$, where higher values correspond to higher expected returns and higher volatility.

5.1.1. Return distributions

Let R_t denote the real return on the liquid portfolio in year t . The simulation draws:

$$R_t \sim \mathcal{D}(r)$$

where “ \sim ” means “is drawn from”, and $\mathcal{D}(r)$ is the lognormal distribution associated with risk level r .

The distribution is defined by specifying the mean μ_0, μ_1 and standard deviation σ_0, σ_1 for risk levels 0 and 1. For intermediate values of r , the mean is obtained by linear interpolation:

$$\mu(r) = (1 - r)\mu_0 + r\mu_1$$

The distribution is stationary across time, meaning that its parameters do not change from year to year.

5.1.2. Volatility assumptions

Volatility refers to the standard deviation of annual real returns. For intermediate risk levels, the standard deviation is derived by treating the portfolio as a mix of two uncorrelated assets:

$$\sigma(r) = \sqrt{(1-r)^2\sigma_0^2 + r^2\sigma_1^2}$$

This formulation ensures that volatility increases smoothly with risk and reflects the assumption of zero correlation between the underlying components.

5.1.3. Correlations

The framework does not model correlations between asset classes, because the portfolio is represented by a single stochastic return series. As a result, no covariance matrix is required, and no correlation assumptions are needed.

5.2. Inflation Modelling

Inflation modelling specifies how nominal and real quantities evolve over time within the simulation. The framework distinguishes clearly between real and nominal modelling, and requires that inflation be represented in a way that is internally consistent with the treatment of returns, spending, and asset values.

5.2.1. Stochastic inflation

Inflation is modelled as a stochastic process. Let π_t denote the inflation rate in year t . The simulation draws:

$$\pi_t \sim \mathcal{J}$$

where “ \sim ” means “is drawn from”, and \mathcal{J} is the distribution chosen for inflation.

The distribution may be normal, lognormal, or any other form that is consistent with the modeller’s assumptions. The framework does not prescribe a specific inflation model, but it requires that:

- Inflation is stochastic rather than fixed.
- Inflation is independent of nominal portfolio returns.
- The inflation process is stationary across time (its parameters do not change year to year).

This ensures that inflation uncertainty is reflected in the evolution of nominal spending and nominal asset values.

5.2.2. Real vs nominal modelling

The framework is expressed primarily in real terms, meaning that:

- Spending paths \bar{C}_t and C_t are defined in real terms.
- Portfolio returns R_t are real returns.
- The expected net worth path $\mathbb{E}[W_t]$ is a real trajectory.

However, sometimes nominal quantities are required, for example:

- Tax calculations.
- Nominal account balances.
- Sequencing rules involving taxable gains.
- Nominal withdrawal amounts.

In such cases they are obtained by applying inflation:

$$X_t^{\text{nominal}} = X_t^{\text{real}} \cdot \prod_{s=1}^t (1 + \pi_s)$$

This separation ensures that:

- The conceptual model remains cleanly real-terms.
- Implementation can convert to nominal terms whenever needed.
- Inflation uncertainty is correctly propagated into nominal outcomes.

The framework does not mandate a specific inflation model, only that the treatment of real and nominal quantities be internally consistent.

5.3. Sequencing Risk

Sequencing risk refers to the sensitivity of long-term outcomes to the order in which investment returns occur. Even when the average return over the projection horizon is the same, different sequences of returns can lead to materially different net worth paths. This arises because withdrawals occur throughout retirement, and negative returns early in the projection reduce the asset base from which subsequent returns compound.

5.3.1. How sequencing affects net worth

Let R_t denote the real return in year t , and let D_t denote the withdrawal in that year. The evolution of the liquid portfolio is:

$$A_{t+1} = (A_t - D_t)(1 + R_t)$$

This expression shows that:

- A negative return in a year with a large withdrawal reduces the portfolio more than the same return occurring later.
- Early-period losses reduce the base on which future returns compound.
- Identical average returns can produce different outcomes depending on their order.

Thus, sequencing risk is an inherent feature of any drawdown-based retirement plan.

5.3.2. Why steering stabilises modelled spending

Steering adjusts discretionary spending mechanically and proportionally in response to deviations from the expected net worth path. Because the expected path is fixed and not

re-baselined within the simulation, the steering mechanism provides a stable reference trajectory against which sequencing effects are measured.

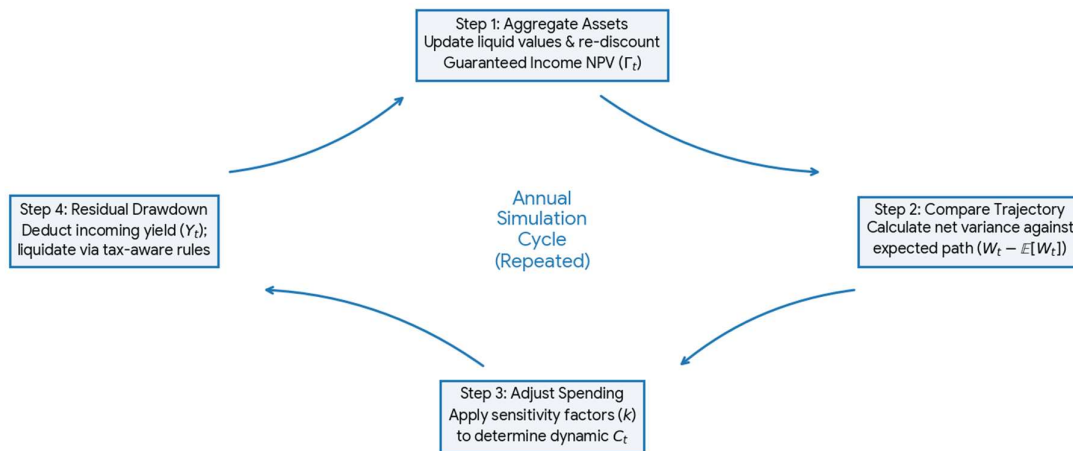
This stabilises the modelled spending path in two ways:

- **Adjustments depend on deviations from the fixed baseline, not on short-term returns.** Sequencing risk affects realised net worth, but the spending response is governed by the deviation from the expected trajectory rather than by the return sequence itself.
- **The mechanism smooths the modelled response.** Because spending changes are continuous and proportional, the model avoids abrupt shifts in spending that would otherwise arise from unfavourable return sequences.

The purpose of the steering mechanism is to produce a coherent and stable modelled spending path, not to represent real-world behaviour or decision-making.

5.4. Annual Projection Cycle

The annual projection cycle follows the steering approach described in Section 3.3. Each cycle updates the household's financial position by applying returns, inflation, spending, income, and withdrawals in a consistent sequence. The expected net worth path remains fixed throughout the simulation and is not re-baselined.



5.4.1. Recalculating net worth

At the start of each year, the model updates the value of all components of holistic net worth:

- Liquid assets grow according to the stochastic return R_t .
- Tangible assets follow their specified growth assumptions.
- Liabilities evolve according to their contractual terms.
- The NPV of future guaranteed income Γ_t is recalculated using the updated income schedule and discounting rules.

These updates produce the realised holistic net worth W_t for the year. While calculating these path-dependent spending rules does not require nested optimization loops, it does demand a data architecture that dynamically aggregates all components of holistic net worth at every simulation interval. Legacy planning software is typically hardcoded around isolated, portfolio-centric projections. Consequently, executing this integrated approach requires a modern, institutional-grade calculation engine specifically engineered to evaluate complex, multi-asset dependencies across a unified lifetime timeline.

5.4.2. Comparing to the expected trajectory

The realised net worth W_t is compared with the expected net worth $\mathbb{E}[W_t]$. The expected trajectory is fixed at the start of the simulation and does not change in response to realised outcomes. This ensures that deviations are measured consistently across all Monte Carlo trials and all time points.

5.4.3. Adjusting spending

Discretionary spending for the year is determined by applying the spending adjustment rules described in Section 4.4 to the deviation:

$$W_t - \mathbb{E}[W_t]$$

This produces the realised discretionary spending C_t , which may be higher or lower than the average-case level \bar{C}_t . Essential spending E_t is added to obtain total spending:

$$S_t = E_t + C_t$$

5.4.4. Recomputing drawdown

Income Y_t is applied, and any shortfall relative to total spending is funded through residual drawdown:

$$D_t = S_t - Y_t$$

Withdrawals are allocated across accounts according to the tax-aware sequencing rule adopted by the implementation. Surpluses, if any, are added to assets.

This completes the annual cycle, and the process repeats for the next year.

6. Comparative Analysis

6.1. Portfolio-Based Guardrails vs Spending-Based Steering

Portfolio-based guardrails and spending-based steering represent two different ways of introducing conditional adjustments into a retirement projection model. Both approaches respond to deviations from an expected trajectory, but they differ in what they adjust and how those adjustments propagate through the projection.

6.1.1. Conceptual comparison

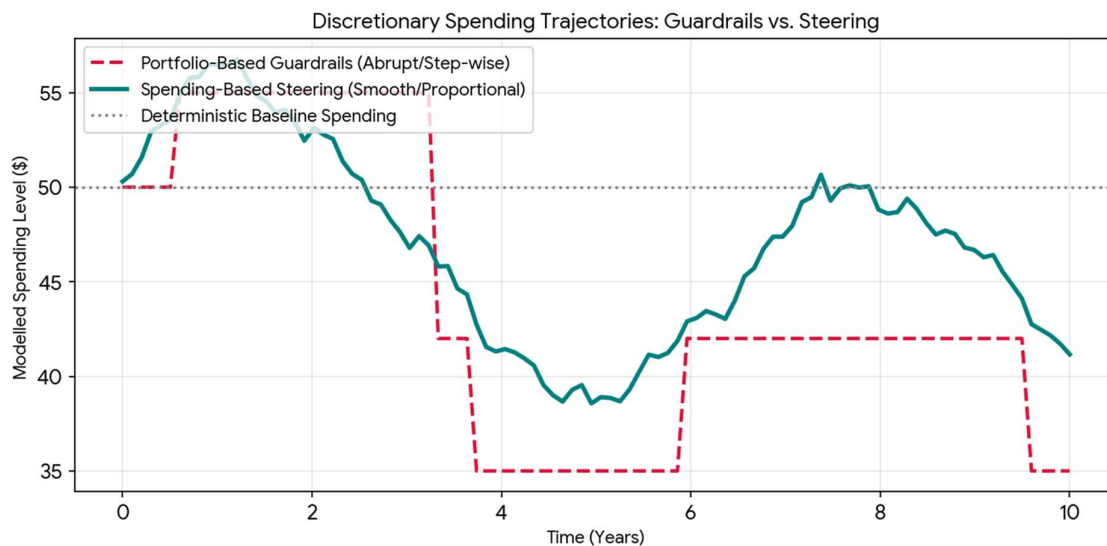
Portfolio-based guardrails adjust **withdrawals** when the portfolio value moves outside predefined bands. The adjustment mechanism is anchored to the liquid portfolio alone and does not incorporate other components of holistic net worth.

By contrast, spending-based steering adjusts **discretionary spending** in response to deviations in **holistic net worth**, which includes the NPV of future guaranteed income streams, as well as liquid and illiquid assets and liabilities. This aligns the adjustment mechanism with the broader financial position rather than with the portfolio in isolation.

6.1.2. Stability comparison

Portfolio-based guardrails can produce abrupt changes in withdrawals when the portfolio crosses a threshold, particularly in the presence of sequencing risk. Because withdrawals are adjusted directly, the resulting spending path may exhibit discontinuities that reflect short-term market movements rather than long-term financial capacity.

Spending-based steering adjusts discretionary spending gradually and proportionally to deviations from the expected holistic net worth path. Because holistic net worth can include components that are less volatile than the liquid portfolio, such as property or other tangible assets, the deviation measure itself is smoother. This leads to more stable modelled spending paths and avoids the threshold-driven changes that arise from portfolio-based rules.



Steering is a modelling mechanism rather than a behavioural assumption. Its purpose is to maintain coherence and stability in the modelled spending path as uncertainty unfolds.

6.2. Fixed Real Spending vs Spending-First Planning

Fixed real spending and spending-first planning represent two different ways of structuring retirement projections. The distinction lies in what is treated as the primary decision variable. Fixed real spending treats discretionary spending as constant in real terms, while spending-first planning allows discretionary spending to adjust in response to financial conditions.

6.2.1. Why fixed real spending is fragile

A fixed real spending rule specifies a constant inflation-adjusted level of discretionary spending each year. This creates fragility in two ways.

- **It ignores sequencing risk.** When spending is fixed, the model removes the same real amount from the household's resources regardless of market conditions. Negative

returns early in retirement reduce the asset base, and fixed spending accelerates depletion. The model therefore produces a wide dispersion of outcomes driven largely by the order of returns.

- **It disconnects spending from financial capacity.** Because spending is fixed, the model does not respond when the household's financial position diverges from expectations. This can lead to unrealistic projections where spending remains constant even when net worth has moved significantly above or below the expected path.

These features make fixed real spending sensitive to market noise and sequencing effects, which limits its usefulness as a modelling framework.

6.2.2. Why spending-first planning is more coherent

Spending-first planning has two advantages.

- **It links spending to overall financial capacity.** Because adjustments are based on holistic net worth, which includes future income and less volatile assets, the model produces spending paths that reflect the household's broader position rather than the liquid portfolio alone.
- **It avoids the rigidity of fixed spending.** Spending adjusts gradually when outcomes diverge from expectations, which prevents the model from producing unrealistic constant spending in the face of large financial shocks.

The purpose of spending-first planning is not to prescribe behaviour, but to provide a modelling structure that responds coherently to financial conditions and avoids the fragility inherent in fixed real spending rules.

6.3. Holistic vs Portfolio-Centric Models

Holistic models and portfolio-centric models differ in how they represent the household's financial position. Portfolio-centric models focus on the liquid investment portfolio alone, while holistic models incorporate all relevant components of net worth, including the NPV of future income streams and the value of tangible (non-portfolio) assets. This distinction affects both the interpretation of financial capacity and the stability of modelled outcomes.

6.3.1. The advantage of including future guaranteed income NPV

Portfolio-centric models treat the portfolio as the sole resource available to support spending. This can understate financial capacity when the household has significant future guaranteed income, such as state pensions, defined benefit pensions, or annuity payments. Because these income streams are not represented as assets, the model may imply a need for higher withdrawals or reduced spending even when future income is sufficient to support long term plans.

Holistic models include the NPV of future guaranteed income within net worth. This provides a more complete representation of the household's financial position and avoids overstating the reliance on the liquid portfolio. It also ensures that spending adjustments respond to the entire resource base rather than to the portfolio in isolation.

6.3.2. Avoiding overreaction to market noise

Portfolio-centric models are sensitive to short term market movements because the portfolio is the only component of net worth. A temporary decline in portfolio value can therefore trigger large changes in modelled withdrawals or spending, even when the household's broader financial position is stable.

Holistic models reduce this sensitivity by incorporating components that are less volatile than the liquid portfolio, such as property and the NPV of future guaranteed income. Tangible assets contribute to stability because they represent resources that can support spending over the long term, whether through planned changes in housing, borrowing capacity, or the option to convert assets if required. These components act as stabilisers, so the model responds to genuine changes in long term financial capacity rather than to short term fluctuations. This leads to smoother modelled spending paths and avoids overreaction to market noise.

7. Behavioural and Practical Implications

Although spending-based steering is a modelling construct rather than a behavioural rule, it produces projections that are easier to interpret and more useful in practice. This is because the model adjusts the quantity that households naturally focus on when thinking about their lives: discretionary spending. And it does so in a way that reflects the household's overall financial position rather than short-term market movements. The result is a representation of lifestyle evolution that is coherent, stable and aligned with long-term financial capacity.

The framework also supports clearer communication. Advisers often find that clients relate more readily to discussions framed around spending and affordability than to conversations centred on withdrawal rates or portfolio mechanics. Steering provides a structured way of illustrating how plans might adapt when circumstances change, without implying that clients should follow any specific rule. Because the model responds to holistic net worth rather than to the liquid portfolio alone, it anchors conversations in long-term financial capacity rather than recent performance, helping to keep planning discussions focused on the broader financial picture.

A further advantage is that the framework reflects the multi-component nature of real-world household finances. Most households rely on a combination of liquid investments, property, cash reserves, state pensions, defined benefit pensions and other income streams. Annual planning processes already involve reassessing these elements and updating assumptions. A model that adjusts spending in response to holistic net worth mirrors this structure, producing projections that are more representative of the complexity advisers and households manage in practice.

This makes the approach particularly useful in contexts where uncertainty, multiple income sources and evolving financial capacity all play a meaningful role. It is well suited to long-term planning scenarios, annual review processes and situations where discretionary spending can adjust over time. Its purpose is not to predict behaviour or prescribe actions, but to provide a stable and interpretable framework for exploring how lifestyle might evolve under different economic conditions.

8. Conclusion

Retirement planning involves uncertainty, multiple income sources, and financial resources that evolve over time. A modelling framework that focuses solely on the liquid portfolio cannot fully capture this complexity, and it can produce projections that are difficult to interpret or that overreact to short-term market movements.

The holistic, spending-first, steering-based approach presented in this paper offers a more coherent way of representing long-term financial capacity. By incorporating all relevant components of net worth, including the net present value of future income, and by adjusting discretionary spending in response to deviations from an expected path, the model produces spending trajectories that are smoother, more interpretable, and more closely aligned with the structure of real-world finances.

The purpose of this framework is not to prescribe behaviour or to predict how households will respond to changing circumstances. Instead, it provides a structured way of exploring how lifestyle might adapt under uncertainty, and it offers advisers and households a clearer basis for discussing long-term plans. It also illustrates the scale and direction of adjustments that might emerge from annual re-assessment, without implying that any specific rule should be followed.

Ultimately, this methodology transitions retirement planning away from legacy forecasting models that often paralyse clients with fear of ruin. By embedding systemic, gentle adjustments directly into the math, planners can confidently demonstrate how a portfolio actively defends itself over time. While executing these multi-layered, path-dependent optimisations is computationally beyond the reach of traditional linear planning tools, the framework represents a highly practical evolution in methodology, one that is already computationally viable within advanced, institutional-grade planning platforms.

By focusing on holistic financial capacity and on the variable that households naturally monitor, namely discretionary spending, the framework supports clearer communication, more grounded expectations, and a more realistic representation of the dynamics that shape retirement planning. The approach is already operational within the EvolveMyRetirement® planning tool, demonstrating that the framework is computationally tractable and suitable for use in real-world planning workflows.

References

1. Bengen, W. (1994). *Determining withdrawal rates using historical data*. Journal of Financial Planning, 7(4), 171–180.
2. Guyton, J. and Klinger, W. (2006). *Decision rules and portfolio management for retirees*. Journal of Financial Planning, 19(10), 48–57.
3. Modigliani, F. and Brumberg, R. (1954). *Utility analysis and the consumption function*. In: Kurihara, K. (ed.) Post-Keynesian Economics. Rutgers University Press, 388–436.
4. Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton University Press.
5. Kotlikoff, L. J., Smetters, K., and Walliser, J. (1999). *Privatizing Social Security in the United States*. Review of Economic Dynamics, 2(3), 532–574.