The Retirement Efficient Frontier, Part I: Personalized Spending Floors and Shortfall Risk

April 2026

Martin-D. Lacasse

Corresponding author
8905 Ashford Gables Ct
Tampa, FL 33626
martin.d.lacasse@gmail.com

Martin-D. Lacasse, Ph.D., is a former Associate Director in the Department of Mathematics at Johns Hopkins University and a retired Distinguished Research Scientist from ExxonMobil Research & Engineering Company.

John D. Leonard II

2121 Stuart Ave
Richmond, VA 23220
Professor Emeritus, Virginia Commonwealth University

John D. Leonard, Ph.D., is Professor Emeritus in the Department of Computer Science in the College of Engineering at Virginia Commonwealth University. He previously served as Executive Dean of the College. Prior to VCU, Dr. Leonard served as Associate Dean for the College of Engineering and held several tenured faculty positions at the Georgia Institute of Technology.

The Retirement Efficient Frontier, Part I: Personalized Spending Floors and Shortfall Risk

April 2026

Abstract

We introduce two concepts that together make a rigorous efficient frontier framework for retirement spending possible. The *Historical Spending Floor* (HSF) is the highest real spending level that, under a given retiree's accounts, taxes, Social Security, Medicare, and longevity assumptions, survives every historical return sequence — a personalized, history-grounded anchor requiring no distributional assumptions. The *Synthetic Spending Floor* (SSF) provides the same anchor calibrated to prospective conditions, computed from a forward-calibrated Monte Carlo scenario set. Each floor defines the conservative endpoint of a spending/shortfall efficient frontier — formally analogous to the Markowitz mean-variance frontier — that maps the full trade-off between committed spending and shortfall risk as measured by the Conditional Value-at-Risk (CVaR), letting retirees choose how far above the floor to commit based on their risk tolerance. We illustrate both with case studies using Owl, an open-source retirement optimizer.

Key Takeaways

- We introduce the Historical Spending Floor (HSF) and the Synthetic Spending Floor (SSF) as two complementary concepts: the HSF grounded in the historical record, the SSF calibrated to prospective conditions via Monte Carlo simulation, each personalized to a specific retiree's accounts, taxes, Social Security, and longevity.
- Together they anchor a spending/shortfall efficient frontier — analogous to the Markowitz mean-variance frontier — that maps the full trade-off between committed spending and shortfall risk, with each floor as its conservative endpoint.
- The frontier makes the spending decision explicit and quantifiable: retirees can see exactly what spending they gain and what shortfall risk they accept at each target success rate, expressed in the same dollar units as the spending commitment itself.

Keywords: retirement income planning, safe withdrawal rate, spending floor, efficient frontier, probability of success, shortfall risk, mixed-integer linear programming

JEL Codes: C61 (Optimization Techniques; Programming Models), D91 (Intertemporal Household Choice; Life Cycle Models and Saving), G11 (Portfolio Choice; Investment Decisions)

The lasting impact of Bengen (1994) on retirement income planning rests on a methodological insight as much as on the number itself: by grounding the analysis in the historical record, it avoided speculative return assumptions and anchored retirement planning in observed reality. The 4.15% safe withdrawal rate was not derived from a model of expected equity premia or bond yields; it emerged from asking a simple empirical question — what is the highest initial spending, as a constant real fraction of portfolio value, that would have survived every 30-year historical sequence in the US data? That question requires no forecast and no distributional assumption. It needs only the data.

Building on this methodological idea, we introduce the concept of the *Historical Spending Floor* (HSF): the highest real spending level that survives every historical scenario for a given retiree's specific financial situation. Unlike Bengen's fixed fractional rule applied to a simplified portfolio, the HSF is found by solving a mixed-integer linear program (MILP) that maximizes spending subject to the retiree's full financial picture — account types, tax brackets, Social Security benefits, Medicare surcharges, Roth conversion opportunities, required minimum distributions, longevity, and lifestyle through the spending profile. The HSF is therefore case-specific: it is not a single universal number, but one tailored to a specific individual or couple, computed from the same historical data that motivated Bengen's original work.

This allows us to reframe the spending decision as a choice on a *spending/shortfall efficient frontier* that maps the entire trade-off between committed spending and mean shortfall across historical scenarios. The HSF is the frontier's most conservative special case: the committed spending level at which every historical scenario succeeds and mean shortfall is zero. Moving along the frontier toward higher committed spending implies accepting shortfalls in some scenarios, but these shortfalls are explicit, quantified, and visible — unlike the binary pass/fail framing of probability of success (PoS), which records only whether a plan fails, not by how much.

When historical data may not reflect prospective conditions, the same framework stress-tests the plan via a *Synthetic Spending Floor* (SSF) computed from Monte Carlo scenarios under explicitly stated return assumptions — complementing the HSF while clearly distinguishing what the future may hold from what the past has shown.

BACKGROUND

The Safe Withdrawal Rate

The SWR has the appeal of simplicity and a worst-case guarantee: by construction, no historical sequence defeats it, and its single actionable number made it easy for practitioners to understand and adopt. The Trinity Study (Cooley et al., 1998) extended this analysis across multiple allocation mixes and time horizons. The rule has attracted substantial criticism over the years — for its reliance on a limited historical sample, its insensitivity to current valuations and yields, and its assumption of a rigid spending path (Scott et al., 2009; Pfau, 2012b; Webb, 2021) — and yet it remains widely used in practice as a conservative planning benchmark (Kitces, 2008). Its persistence in the face of these criticisms is itself telling: despite the standard caveat that past performance does not guarantee future results, anchoring to historical evidence is a natural human response to deep uncertainty. Faced with an unknowable future, practitioners and retirees instinctively turn to what has actually happened (Tversky and Kahneman, 1974; Greenwood and Shleifer, 2014). Bengen's historical framing made the rule not just empirically defensible but psychologically compelling. Its deeper limitation, for our purposes, is discussed next.

The SWR is more akin to a "retirement index" — a long-term market indicator gauging roughly how much a savings portfolio can sustain — than an actionable rule for individual retirees: it reflects aggregate historical market performance but carries no information about the household it applies to. For example, Social Security alone accounts for at least half of retirement income for approximately three in five beneficiaries aged 65 and older (Dushi et al., 2017; Social Security Administration, 2025b) — an income source practitioners typically add as a separate offset when applying the SWR to portfolio assets. But no portfolio-level rule can provide the joint optimization of account withdrawals, Roth conversions, and Medicare bracket management that maximizes after-tax spending for a given account structure, nor any mechanism to express a bequest constraint: a retiree wishing to leave a specific estate cannot use a fixed withdrawal rule to find the highest spending level consistent with that goal.

The HSF addresses both limitations: it computes the case-specific safe spending level using the same historical data, with bequest targets embedded as MILP constraints, while the efficient frontier reveals what spending is achievable above that floor and how much shortfall risk each choice entails. The SSF extends this to forward-calibrated scenarios, grounding the result in the retiree's specific financial structure, historical performance, and explicit assumptions about the future.

Probability of Success and Shortfall Risk

Any method anchored to the historical record shares its implicit distributional assumption: the available historical sequences are treated as representative of the future. Monte Carlo simulation relaxes this by learning return statistics from historical data while generating synthetic scenarios under explicitly specified return assumptions, and reporting the fraction in which the portfolio survives the full retirement horizon — the probability of success (PoS). In most planning tools, PoS is computed by year-by-year forward simulation: a spending rule — fixed or with dynamic guardrails — is applied sequentially, and any scenario in which the balance reaches zero before the horizon ends is counted as a failure. PoS has become the dominant planning metric. Its limitation is that it is binary: each scenario is either a success or a failure, with no account taken of the severity of failure. A scenario in which the portfolio depletes in the final month of a 30-year retirement counts identically to one in which it runs out 15 years early.

Blanchett (2007) observed that PoS should be supplemented with a measure of shortfall magnitude, a theme further developed by Blanchett et al. (2012) and Gardner and Pittman (2013), who formalized expected shortfall as a mortality-adjusted metric. By using a MILP optimizer to maximize spending subject to all financial constraints for each scenario, the shortfall measure derived here is exact within the model's assumptions and applies to any set of return scenarios, whether historical or synthetic.

OWL: OPTIMAL WEALTH LAB

The case studies presented below use Owl (Lacasse, 2026), an open-source retirement optimizer available at `https://github.com/mdlacasse/Owl`. It builds on a line of LP-based retirement optimization research (Ragsdale et al., 1994; Coopersmith and Sumutka, 2011; Welch, 2015, 2016, 2017), extending it with a comprehensive federal tax model: income taxes, long-term capital gains brackets, Medicare IRMAA (Income-Related Monthly Adjustment Amount) surcharges with the two-year income lag, ACA (Affordable Care Act) premium credits, Social Security taxability, required minimum distributions (RMDs), and Roth conversion limits. State income taxes are not modeled, as their variation across jurisdictions precludes a general treatment. This level of detail may appear to over-model an uncertain future, but it is essential where the framework delivers its most actionable value: the first several years of retirement, when Roth conversion windows, Medicare bracket management, and RMD planning decisions are immediately consequential and still reversible.

A parallel line of research has addressed decumulation through stochastic programming and optimal control (Konicz and Mulvey, 2013; Forsyth, 2020; Forsyth et al., 2021); the MILP approach adopted here is complementary, enabling exact tax and benefit modeling in place of stylized assumptions. For each scenario, the optimizer solves a MILP across the retiree's accounts — taxable, tax-deferred, Roth, and health savings — to find the maximum sustainable spending scale; a typical solve takes a few seconds on a laptop. The framework supports historical return sequences (Damodaran, 2026) and a variety of adjustable stochastic return models; in spending optimization mode it collects the per-scenario spending scales and solves a fast commitment linear program (LP) to trace the efficient frontier.

THE EFFICIENT FRONTIER FRAMEWORK

Definition

Let $g_s \geq 0$ denote the *spending scale* for scenario s: the largest scalar multiplier of a normalized spending profile that remains sustainable under that scenario's return and inflation sequence while satisfying all financial and tax constraints. For a single retiree with a flat profile, g_s equals the constant annual spending amount. Because the MILP has full knowledge of each scenario's return sequence, g_s is a clairvoyant upper bound — the theoretical maximum sustainable spending achievable with perfect foresight under that scenario.

The spending scale g_s depends on several plan-level inputs beyond the return sequence:

Spending profile — sets the normalized shape of real spending over time: flat, declining in later years, geared to active- and slow-retirement phases, reduced at a spouse's death, or a dynamic guardrail rule that adjusts annually to portfolio performance (see the Dynamic Spending section below).

Longevity — the assumed planning horizon — directly bounds how long the portfolio must sustain spending; a longer horizon generally implies a lower g_s and is itself an uncertain input the planner must choose.

Constraints — bequest targets, the proceeds from a planned home sale, or scheduled debt service obligations enter as hard bounds in the MILP, directly shaping the feasible spending range.

Asset allocation — the equity/bond mix, any other asset classes held, and whether the allocation is fixed or follows a glide path — determines the return and risk profile of each scenario; a more conservative allocation generally lowers both g_s and the upside potential of the frontier.

Future income streams — Social Security, pensions, annuities, part-time earnings, or rental income, whether fixed or variable — reduce portfolio withdrawal requirements in every scenario, raising g_s because each dollar of guaranteed income is a dollar the portfolio need not supply.

Future environment — tax law, scheduled regulatory changes, and the full range of applicable retirement income rules — is embedded in each scenario solve.

Different choices on any of these dimensions yield different values of g_s and hence a different HSF, revealing what universal rules like the 4% SWR cannot.

> *The Historical Spending Floor (HSF) is the maximum real spending level that, under a fixed set of plan parameters and constraints, survives every historical scenario.*

Since $\text{HSF} = \min_s g_s$, any input that shifts the per-scenario spending scales shifts the floor in the same direction — so $g_s \geq \text{HSF}$ for every scenario s, with equality at the worst-case scenario.

The HSF extends Bengen's idea to any retirement case: it is computed from the same historical data, but for the retiree's actual financial structure rather than a simplified portfolio. It is the most conservative point on a family of spending choices: it eliminates shortfall entirely, but at the cost of committed spending that may fall well below what most historical scenarios could sustain.

The Spending/Shortfall Efficient Frontier

Suppose the retiree must commit to a single first-year spending level g^* before knowing which scenario will occur. In any scenario where g^* exceeds what that scenario can sustain, the plan faces a shortfall — a gap between the commitment and the maximum the optimizer finds for that scenario. At $g^* = \text{HSF}$, every scenario succeeds and the mean shortfall is zero. As g^* rises above the HSF, some scenarios begin to fail and the mean shortfall grows. The efficient frontier is the set of Pareto-optimal pairs of committed spending and mean shortfall.

The framework traces the frontier by sweeping a risk-aversion parameter from zero — maximize spending regardless of shortfalls, the most aggressive choice — to arbitrarily large values, which drives committed spending down to the HSF where every scenario succeeds. Intermediate values yield the full range of defensible spending options between these two extremes. The mean shortfall also has a precise statistical interpretation: divided by the fraction of scenarios that fail, it equals the Conditional Value-at-Risk (CVaR) of spending losses — a coherent risk measure widely used in financial risk management. The complete formulation is given in Appendix B. Exhibit 1 summarizes the analogy with Modern Portfolio Theory (Markowitz, 1952).

Concept	**Modern Portfolio Theory**	**This Article**
Decision variable	Portfolio weights	Committed spending g^*
Maximize	Expected return	Committed spending g^*
Risk measure	Variance	Mean shortfall $\bar{\sigma}$
Risk aversion	λ	λ
Conservative extreme	Minimum-variance portfolio	HSF (zero shortfall)
Aggressive extreme	Maximum-return portfolio	Best-case scenario spending

Exhibit 1: Analogy between Modern Portfolio Theory and the spending/shortfall efficient frontier. The HSF plays the role of the minimum-variance portfolio: it eliminates risk entirely, but at the cost of a spending level that may be far below the typical outcome.

Rather than asking the retiree to choose λ directly — a risk-aversion parameter with no intuitive scale — the framework translates the frontier into the familiar practitioner metric: a target probability of success ρ. Because PoS is already the dominant planning language, expressing the frontier in those terms makes the framework immediately accessible without requiring any change in how retirees and advisors frame the spending conversation. For any target ρ (say, 85%), the framework identifies the highest g^* achieving that PoS and reports the mean shortfall alongside the HSF or SSF and the median. The latter two are properties of the scenario ensemble and do not change with ρ; what changes is how g^* relates to them — how large the shortfall is as a fraction of the commitment, and how far the median lies above it. This context is entirely invisible in a standard PoS analysis.

Pfau (2013) proposed a broader framework for retirement income in which spending power, portfolio sustainability, and expected bequests form an efficient surface. The present article differs: the decision variable is a committed spending level rather than a product mix, shortfall magnitude is the risk measure, and the result is a precise two-dimensional frontier with a formal LP solution. The same framework applies to both historical scenarios and Monte Carlo scenario sets calibrated to prospective return assumptions; when applied to the latter, its conservative anchor defines the Synthetic Spending Floor (SSF), introduced later.

CASE STUDIES

We illustrate the efficient frontier framework with two retirement cases of increasing complexity. The first is a deliberately simplified benchmark that isolates the HSF concept and connects it directly to the familiar SWR literature. The second introduces a realistic couple with Social Security, tax-deferred and Roth accounts, and Medicare costs, demonstrating how Owl's MILP handles household complexity while producing the same frontier output.

Bill: A Stylised Benchmark

Bill is a 65-year-old single retiree with $1,000,000 in a Roth IRA (Individual Retirement Account), invested 50% in US equities (S&P 500) and 50% in 10-year Treasury notes. He has no Social Security, no pension, no Medicare costs, no taxable or tax-deferred accounts, and a flat (constant real) spending profile over a 30-year horizon. Because his savings are entirely in a Roth account, there are no income taxes on withdrawals — the MILP degenerates to a simple cash-flow problem. This case is designed to mirror Bengen's setup as closely as possible, allowing a clean comparison.

Applied to all 66 historical windows from 1928 to 1993, the minimum per-scenario spending scale falls on the 1966 start year. Bill's HSF is $36,666/yr — equivalent to a 3.67% withdrawal rate on his $1,000,000 portfolio. The 1966 window is the same worst case identified by Bengen, for the same reason: the severe real return sequence of 1966–1981, with simultaneous equity losses and rising bond yields, is uniquely damaging for a retiree drawing down assets.

The 0.48% difference between Bill's HSF (3.67%) and Bengen's 4.15% is entirely methodological and fully accounted for by two factors: beginning-of-year withdrawals and the use of 10-year rather than 5-year Treasury notes. The details and reconciliation are given in Appendix A. The key point is that the HSF framework, applied to Bill's stylized case, reproduces Bengen's result exactly once the underlying assumptions are aligned.

Bill's case is instructive as a benchmark, but it is not representative of real retirees. It isolates sequence-of-returns risk by removing tax complications, Social Security, Medicare, any partner, or bequest targets. The HSF for a real case may differ substantially from 3.67%. Social Security income, being independent of market performance, reduces the portfolio's exposure to sequence-of-returns risk and can support higher total spending per dollar of portfolio — a favorable effect. A bequest goal works in the opposite direction, as preserving wealth for heirs constrains how aggressively the portfolio can be drawn down.

Chris and Pat

Chris and Pat illustrate the HSF framework applied to a realistic retired couple. Chris, an engineer, retires at age 64; Pat, a marketing professional, at 62. They file jointly, plan to age 89 and 92 respectively, and own their home outright. Their nest egg is concentrated in Chris's tax-deferred accounts, reflecting decades of 401(k) contributions. Account balances are chosen so that after-tax savings match Bill's $1,000,000: tax-deferred assets of $937,000 at a 20% effective tax rate yield $750,000 after tax, plus $100,000 in Roth and $150,000 in taxable, totaling approximately $1,000,000. This design isolates the effect of account structure, Social Security, and taxes from the effect of total wealth: any difference in HSF relative to Bill's reflects the complexity of the real case, not a difference in starting resources.

Pat claims Social Security at 62 for early income; Chris delays to 70 to maximize the household survivor benefit. Both claiming ages are fixed for the frontier analysis. Their financial profile appears in Exhibit 2.

The optimizer solves the MILP for each of the 66 historical windows from 1928 to 1993. The range ends in 1993 because later start years do not provide a complete sequence sufficient to

Parameter	Value
Ages at plan start (2026)	63 (Chris), 60 (Pat); retire end of 2027 and 2028
Planning horizon	Age 89 (Chris), 92 (Pat)
Filing status	Married filing jointly
Taxable accounts	\$100,000 (Chris), \$50,000 (Pat)
Tax-deferred (IRA/401(k))	\$700,000 (Chris), \$237,000 (Pat)
Roth IRA	\$70,000 (Chris), \$30,000 (Pat)
Primary residence	\$450,000 (held to death; left as bequest)
401(k) contributions	Employee max + 4% employer match (through retirement)
Social Security PIA	\$2,000/mo (Chris, claiming age 70)†, \$1,400/mo (Pat, claiming age 62)
Spending profile	Flat; 40% reduction at passing of first-dying spouse
Asset allocation	60% equity / 40% bonds (fixed)
Heirs' tax rate on IRA	30%
Bequest target	\$400,000 (excluding primary residence; also serves as a reserve against longevity and long-term care risk)

Exhibit 2: Chris and Pat's financial profile. †Chris's claiming age confirmed by Owl's SS age optimizer in a preliminary run and fixed for the frontier analysis. PIA = Primary Insurance Amount (monthly benefit at full retirement age). The fixed 60/40 allocation isolates sequence-of-returns risk from allocation drift. The case file is available in the Owl repository (Lacasse, 2026).

cover the couple's 33-year planning horizon — chosen to reflect a realistic longevity target for a two-person household, in contrast to Bill's stylized 30-year horizon. Each solve jointly optimizes Roth conversions and account withdrawals, subject to US tax rules, RMDs, and Medicare costs, producing the spending scale g_s for that scenario.

The resulting HSF for Chris and Pat — the minimum g_s across all 66 scenarios — is anchored by the 1966 start year, consistent with Bill's case. Their per-scenario spending scales are higher than Bill's primarily because Social Security provides income independent of portfolio performance, reducing the withdrawal burden in every scenario — and despite the slightly longer planning horizon and the \$400,000 bequest constraint that Bill does not face.

HISTORICAL RESULTS

Exhibit 3 shows the efficient frontier for Bill's stylized case at an 85% target success rate (produced using Owl (Lacasse, 2026)). Exhibits 4 and 5 show the three-panel output for Chris and Pat at 85% and 70% target success rates. Each exhibit displays the PoS curve (left), the efficient frontier (center), and a bar chart of committed spending and shortfall by historical start year (right).

The 1966 scenario anchors the HSF for both cases; for Bill at an 85% target, committed spending is \$42,660/yr (4.27% of portfolio), with a mean shortfall of \$423/yr (1.0% of committed spending), a CVaR of \$3,098/yr (7.3%), and a median scenario spending of \$57,065/yr. Dropping to a 70% target raises committed spending to \$47,611/yr but increases the mean shortfall to \$1,444/yr (3.0%) and the CVaR to \$5,296/yr (11.1%).

```
Committed spending (today's $):  $42,660/yr
Spending-to-savings ratio:       4.27%  (ETR ratio 20%)
Target success rate:             85%  (actual: 86%)
Median scenario spending:        $57,065/yr
Historical spending floor:       $36,666/yr  (14.1% shortfall)
Mean shortfall:                  $423/yr  (1.0% of committed)
CVaR (avg loss | failure):       $3,098/yr  (7.3% of committed)
Scenarios:                       66
```

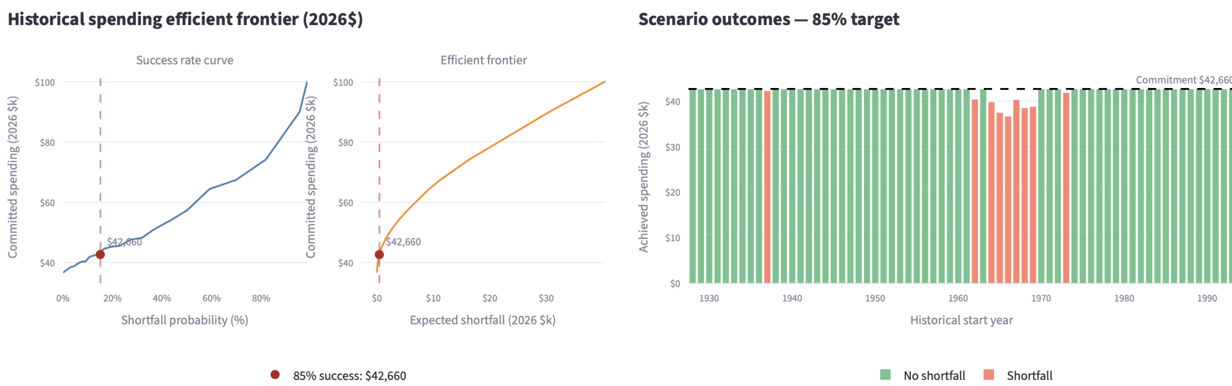

Exhibit 3: Bill's historical efficient frontier (66 scenarios, 1928–1993). *Left:* Committed spending versus shortfall probability. *Center:* Committed spending versus mean shortfall; the HSF ($36,666/yr) anchors the conservative endpoint where mean shortfall is zero. *Right:* Committed spending and shortfall by historical start year.

At an 85% target, Chris and Pat can commit to $99,003/yr in today's dollars, with a mean shortfall of only $786/yr (0.8%), a CVaR of $5,762/yr (5.8%), and a median scenario spending of $123,145/yr — well above the commitment. The 1966 scenario produces the HSF of $88,908/yr, with committed spending exceeding the floor by only 10.2% (compared to 14.1% for Bill). Knowing that the most damaging sequence in nearly a century of US market history implies a roughly 10% spending reduction — not a catastrophic failure — fundamentally reframes what "risk" means in this plan.

Accepting a 70% target raises committed spending to $106,948/yr (+8%), but the mean shortfall grows to $2,600/yr (2.4%), the CVaR to $9,535/yr (8.9%), and the gap between committed spending and the HSF to 16.9%. The median and HSF spending are unchanged between the two targets — they reflect properties of the scenario set, not of the chosen commitment level.

Comparing Chris and Pat to Bill illuminates two structural advantages of the realistic case. First, Social Security income acts as a risk buffer: at 85% PoS, Bill's worst-case shortfall is 14.1% versus 10.2% for Chris and Pat, and Bill's CVaR of 7.3% exceeds theirs (5.8%) despite the same after-tax wealth. SS payments continue regardless of market performance, compressing the distribution of per-scenario spending scales: for Bill, the range from the HSF to median is $20,399 (48%), while for Chris and Pat the absolute range is wider at $34,237 but represents only 35%. The signed percentages in Exhibit 6 make this compression visible at a glance: Bill's outcomes span -14.1% to $+33.8\%$ of committed spending, while Social Security narrows Chris and Pat's range to -10.2% to $+24.4\%$. SS effectively places a floor under each scenario's g_s, reducing the variance

```
Committed spending (today's $):  $99,003/yr
Spending-to-savings ratio:       9.90%  (ETR ratio 20%)
Spending-to-savings note:        understated due to bequest of $400k
Target success rate:             85%  (actual: 86%)
Median scenario spending:        $123,145/yr
Historical spending floor:       $88,908/yr  (10.2% shortfall)
Mean shortfall:                  $786/yr  (0.8% of committed)
CVaR (avg loss | failure):       $5,762/yr  (5.8% of committed)
Scenarios:                       66
```

Exhibit 4: Chris and Pat — historical scenarios (1928–1993), 85% target success rate. *Left:* Committed spending versus shortfall probability. *Center:* Committed spending versus mean shortfall; the HSF anchors the left endpoint at zero shortfall. *Right:* Committed spending and shortfall by historical start year; the 1966 sequence is the HSF scenario.

across scenarios and with it the exposure to sequence-of-returns risk. Second, Chris and Pat's $400,000 bequest target is embedded directly as a constraint in Owl's MILP — the committed spending $99,003/yr is the highest achievable while satisfying that constraint across at least 85% of scenarios. A fixed fractional rule has no mechanism to express this trade-off: the SWR concept simply has no meaning when a bequest constraint is present.

The bar charts (right panels) make the cost of each choice concrete and scenario-specific: advisors and retirees can examine exactly which historical periods are responsible for shortfalls and by how much, rather than absorbing a single PoS figure.

Relying on historical data carries an implicit assumption shared by all such approaches: that the 66 return sequences in the dataset are representative of the distribution of future returns. If prospective conditions differ systematically — lower expected returns, higher starting valuations, or a different inflation regime — the frontier will overstate sustainable spending. This limitation motivates the Synthetic Spending Floor introduced next, which replaces the historical record with a forward-calibrated Monte Carlo scenario set to stress-test the same framework under explicitly specified return assumptions.

```
Committed spending (today's $):  $106,948/yr
Spending-to-savings ratio:       10.70%  (ETR ratio 20%)
Spending-to-savings note:        understated due to bequest of $400k
Target success rate:             70%  (actual: 73%)
Median scenario spending:        $123,145/yr
Historical spending floor:       $88,908/yr  (16.9% shortfall)
Mean shortfall:                  $2,600/yr  (2.4% of committed)
CVaR (avg loss | failure):       $9,535/yr  (8.9% of committed)
Scenarios:                       66
```

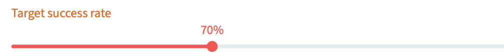

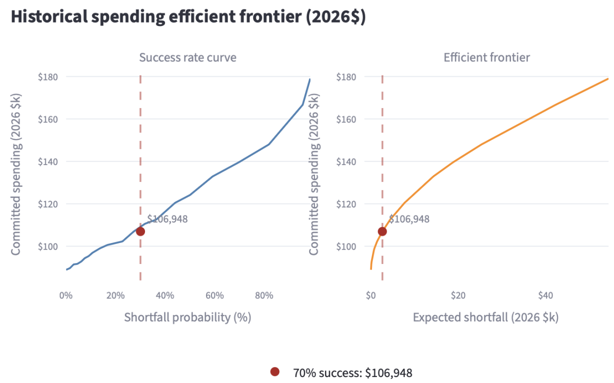

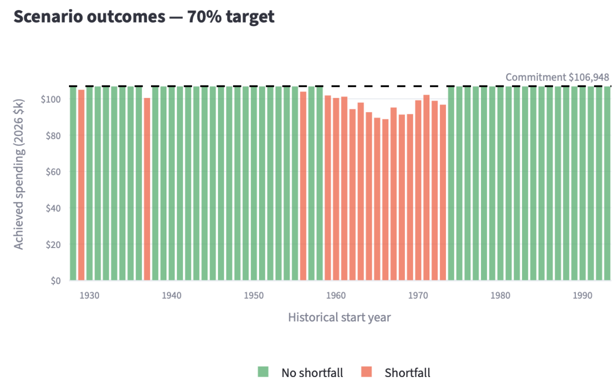

Exhibit 5: Chris and Pat — historical scenarios (1928–1993), 70% target success rate. Compared to Exhibit 4, the bar chart reveals a substantially broader and deeper shortfall distribution, illustrating the concrete cost of targeting a lower success rate.

Case	Target PoS	Committed	Mean shortfall	CVaR	HSF	Median
Bill	85%	$42,660	$423 (1.0%)	$3,098 (7.3%)	$36,666 (−14.1%)	$57,065 (+33.8%)
Bill	70%	$47,611	$1,444 (3.0%)	$5,296 (11.1%)	$36,666 (−23.0%)	$57,065 (+19.9%)
Chris & Pat	85%	$99,003	$786 (0.8%)	$5,762 (5.8%)	$88,908 (−10.2%)	$123,145 (+24.4%)
Chris & Pat	70%	$106,948	$2,600 (2.4%)	$9,535 (8.9%)	$88,908 (−16.9%)	$123,145 (+15.1%)
Bill HSF: $36,666/yr (3.67%)	Chris & Pat HSF: $88,908/yr					

Exhibit 6: Summary of historical frontier results for Bill and Chris and Pat (66 scenarios, 1928–1993). Parenthetical percentages are fractions of committed spending; signs indicate deviation from committed (− below, + above). All dollar amounts are in today's dollars per year.

SYNTHETIC SPENDING FLOOR

Motivation and Definition

The decades that drove the historical record — including the post-war economic boom, the secular equity rally of the 1980s and 1990s, and the long decline in interest rates — may not be representative of the returns available to retirees who began their decumulation in 2020 or later. Low starting bond yields, elevated equity valuations, and the potential reversal of the secular bond trend suggest that prospective returns may differ from historical averages, a concern that motivated Finke et al. (2013) to adjust the safe withdrawal rate downward under prevailing yield conditions.

To address this, the same MILP framework can be run not on historical windows but on a large set of synthetic return sequences generated from a model calibrated to current or prospective conditions. The minimum spending that survives a target fraction of those scenarios plays the same role as the HSF, but under explicitly stated prospective assumptions. We define:

> *The Synthetic Spending Floor (SSF) is the maximum real spending level that, under a fixed set of plan parameters and constraints, survives at least 95% of a Monte Carlo scenario set calibrated to current or prospective return conditions.*

Unlike the HSF — where the finite historical ensemble makes the minimum g_s a stable quantity — the worst-case draw in N independent simulations decreases without bound as N grows, so the raw minimum is not a reproducible statistic. Instead, the SSF is defined as the 5th percentile of the scenario distribution, chosen to provide a stable, reproducible quantile anchor that converges as N grows and reflects the stated return assumptions rather than the historical record. Unlike the HSF, the SSF is not a zero-shortfall anchor — it marks the 95% success-rate endpoint of the frontier, from which the frontier is explored at whatever target success rate the retiree chooses by sweeping λ exactly as in the historical case. The SSF can be higher or lower than the HSF depending on whether prospective conditions are more or less favorable than historical ones.

Monte Carlo Setup

To demonstrate the SSF in practice, we use Owl's *histolognormal* return generation model, in which log-returns for equities, bonds, and inflation are assumed jointly normally distributed, with the mean vector and covariance matrix estimated by maximum likelihood from the full 1928–2025 historical record. Each year's return vector is drawn independently from that distribution, and a scenario is the resulting sequence over the planning horizon. This parametric baseline follows Finke et al. (2013). Collins et al. (2015) recommend the log-normal over the normal distribution for this application, since portfolio values are products of gross-return factors and cannot turn negative. More sophisticated models — block bootstrap, stationary bootstrap (Politis and Romano, 1994), or GARCH (Generalized Autoregressive Conditional Heteroskedasticity) — that better capture mean reversion and autocorrelation are also available in Owl when behavior at extreme success-rate targets is of primary concern.

Five hundred simulations were run for Chris and Pat, completing in approximately two minutes on a standard laptop (Apple M5, 10 cores).

Synthetic Frontier and Comparison with Historical

Exhibit 7 shows the Monte Carlo efficient frontier for Chris and Pat at an 85% target success rate, and Exhibit 8 summarizes both targets alongside the historical results.

The Monte Carlo results are broadly consistent with the historical results — as expected, since the lognormal model is calibrated to the same historical data. For Chris and Pat at 85% PoS, committed spending is \$97,727/yr versus the historical \$99,003, and the median is \$123,279 versus

```
Committed spending (today's $):  $97,727/yr
Spending-to-savings ratio:       9.78%  (ETR ratio 20%)
Spending-to-savings note:        understated due to bequest of $400k
Target success rate:             85%  (actual: 85%)
Median scenario spending:        $123,279/yr
Synthetic spending floor:        $82,727/yr  (15.3% shortfall)
Mean shortfall:                  $1,783/yr  (1.8% of committed)
CVaR (avg loss | failure):       $12,215/yr  (12.5% of committed)
Rate method:                     histolognormal
Scenarios:                       500
```

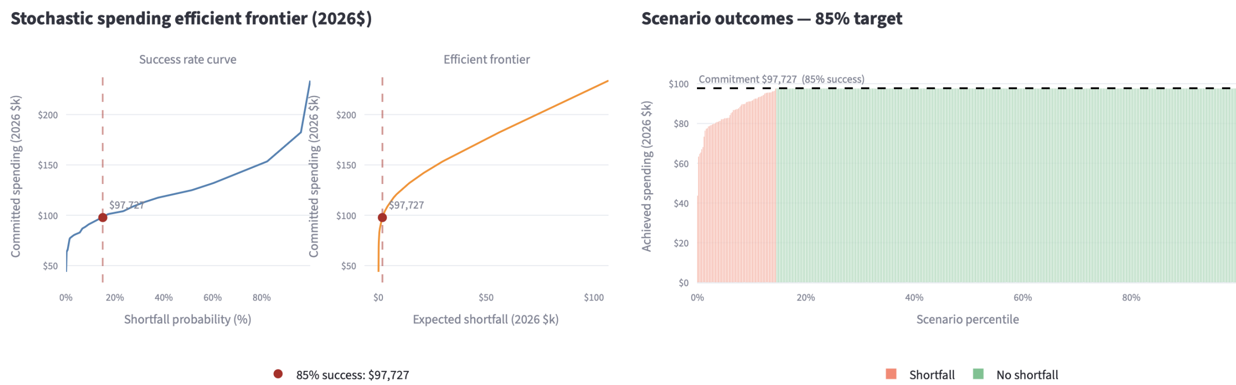

Exhibit 7: Chris and Pat — Monte Carlo scenarios (500 simulations, histolognormal model), 85% target success rate. *Left:* Committed spending versus shortfall probability. *Center:* Committed spending versus mean shortfall; the SSF anchors the conservative endpoint. *Right:* Achieved spending by scenario percentile, sorted from worst to best; red bars indicate shortfall scenarios.

$123,145 — close agreement that confirms the calibration is internally coherent. For Bill, the synthetic frontier yields $43,577/yr at 85% PoS versus the historical $42,660, again consistent. The SSF values — $35,749/yr for Bill (3.57%) and $82,727/yr for Chris and Pat — are close to but slightly below their historical counterparts ($36,666 and $88,908), reflecting that independent lognormal sampling occasionally generates sustained adverse sequences that the bounded historical record does not contain.

The differences between the two scenario sets are informative. The Monte Carlo mean shortfall at 85% PoS ($910/yr for Bill, $1,783/yr for Chris and Pat) exceeds the historical values ($423/yr and $786/yr), and the corresponding CVaRs ($6,234/yr and $12,215/yr) exceed the historical values ($3,098/yr and $5,762/yr), because independent sampling produces a wider tail than autocorrelated historical sequences. The SSF is likewise below the HSF for the same reason — independent random draws can generate more consecutive poor years than the autocorrelated historical record admits. This is a feature of the return model, not the retiree's actual risk landscape (Tharp, 2017; Campbell and Viceira, 2002), and it suggests using correlated return models when tail behavior at extreme success-rate targets matters most.

The historical PoS curve is approximately linear: each percentage-point reduction in the target buys a roughly proportional increase in committed spending. The Monte Carlo PoS curve shares this linearity in its middle range but shows nonlinearity near the extremes — a steep drop in committed spending as PoS approaches 100% (requiring defense against increasingly extreme

Case	Target PoS	Committed	Mean shortfall	CVaR	SSF	Median
Bill	85%	\$43,577	\$910 (2.1%)	\$6,234 (14.3%)	\$35,749 (−18.0%)	\$58,898 (+35.2%)
Bill	70%	\$49,550	\$2,149 (4.3%)	\$7,844 (15.8%)	\$35,749 (−27.9%)	\$58,898 (+18.9%)
Chris & Pat	85%	\$97,727	\$1,783 (1.8%)	\$12,215 (12.5%)	\$82,727 (−15.3%)	\$123,279 (+26.1%)
Chris & Pat	70%	\$108,780	\$4,140 (3.8%)	\$15,109 (13.9%)	\$82,727 (−23.9%)	\$123,279 (+13.3%)
Bill SSF: \$35,749/yr (3.57%)	Chris & Pat SSF: \$82,727/yr					

Exhibit 8: Monte Carlo frontier results for Bill and Chris and Pat (500 simulations, histolognormal model calibrated to 1928–2025). The SSF converges to a stable quantile as simulation count grows and is the appropriate tail indicator for Monte Carlo ensembles. Parenthetical percentages are fractions of committed spending; signs indicate deviation from committed (− below, + above). All dollar amounts are in today's dollars per year.

tail draws) and rapid acceleration in mean shortfall as PoS approaches 0%. The historical curve's linearity arises because its 66 scenarios are bounded by the actual range of US market history. The Monte Carlo curve shows greater curvature for two reasons: the independent sampling assumption may produce tails that are too wide (Collins et al., 2015), and the parametric model generates tail scenarios that simply did not occur in the historical record — events the statistical model considers plausible but that history never realized.

On simulation size. A natural instinct is to run more Monte Carlo simulations for greater accuracy. This holds for the mean shortfall and committed spending, whose estimation error scales as $1/\sqrt{N}$; a few hundred scenarios already yield tight estimates (Fitzpatrick and Tharp, 2022). The worst-case spending, by contrast, is unbounded: unlike the historical record, where the worst scenario is fixed by actual market history, the parametric model can always generate a more extreme draw, so the Monte Carlo minimum worsens with each additional simulation. The SSF avoids this instability, since it is estimated from $0.05N$ observations and the standard error of a sample quantile is of order $1/\sqrt{N}$. Shifting the conversation from worst-case spending to mean shortfall and a stable quantile is therefore not merely a presentational choice — it is the statistically appropriate summary of Monte Carlo tail risk.

PRACTICAL IMPLICATIONS

HSF and SSF as Complementary Anchors

The HSF and SSF serve different but complementary roles. The HSF is backward-looking: it is computed from the historical record and provides a guarantee of the same kind Bengen offered — no historical sequence would have defeated this spending level. It is therefore appropriate as a lower bound that retirees and their advisors can present with historical authority. The SSF is forward-looking: it is computed from a return model calibrated to current or prospective conditions and captures the possibility that the future differs systematically from the past. When the model suggests lower returns (higher cyclically adjusted price-to-earnings (CAPE) ratios, lower starting yields), the SSF falls below the HSF, serving as a warning. When both floors agree, it signals that prospective conditions are expected to mirror the historical record.

Together, the HSF and SSF anchor two versions of the spending frontier — one grounded in the historical record, one calibrated to prospective conditions. When the two floors are close, the frontiers align and the retiree can select a committed spending level anywhere above the floor with confidence from both perspectives. When the SSF is materially below the HSF, the frontiers diverge, and the gap reflects genuine uncertainty about future returns; the retiree can use both frontiers to understand the range of outcomes at any given target success rate.

Reframing the Spending Decision

The efficient frontier changes the retirement spending conversation in a meaningful way. Under the conventional PoS paradigm, the retiree negotiates a success rate threshold without knowing what happens in the scenarios below it. The frontier makes these scenarios visible: a mean shortfall of $x/yr, with tail spending of $y/yr at the HSF or the SSF. For many retirees, this information will reveal that a lower PoS target is entirely reasonable: 70% sounds alarming in isolation, but the frontier shows that the mean shortfall across all scenarios amounts to only 2–4% of committed spending, with even the most extreme historical sequence falling short by 16.9% (Chris and Pat at a 70% target). Committing to the median scenario would correspond to a 50% success rate — and the median for Chris and Pat is $123,279/yr, well above any reasonable target. At a 70% target, Chris and Pat commit to $106,948/yr — an 8% gain over the 85% commitment — while the floor guarantees at least $88,908/yr even in the worst sequence in nearly a century of US market history.

The choice of the word *floor* rather than *worst-case scenario* is deliberate. A floor is a foundation from which to move up; a worst case anchors the conversation to the most pessimistic outcome and can nudge retirees toward unnecessary caution. Just as a rational investor is not expected to hold the minimum-variance portfolio, a rational retiree is not expected to live at the floor. The floor defines the lower boundary of a sensible range — everything to the right of it on the frontier involves some risk, but that risk is explicit, quantified, and expressed in the same dollar units as the spending commitment. Moving slightly away from the floor typically yields a substantially higher committed spending level for only a modest increase in mean shortfall.

One caveat applies to the shortfall measure throughout: it assumes the retiree adheres rigidly to the committed spending level g^* with no adjustment in response to realized returns. This makes the shortfall a conservative upper bound on the risk faced by an adaptive retiree who would reduce spending following a poor sequence. Annual re-evaluation and dynamic guardrail rules, discussed below, address this directly. More generally, the MILP can be extended with per-year slack variables $\delta_n \in [\underline{\delta}_n, \bar{\delta}_n]$ representing allowable deviations from a prescribed spending level, so that the optimizer finds the best dynamic spending trajectory within the guardrail bounds.

Dynamic Spending

Scott et al. (2009) demonstrated that financing a constant spending plan from a volatile investment strategy is fundamentally inefficient: it generates unnecessary shortfalls when markets underperform and wasteful surpluses when they outperform. A family of dynamic spending strategies — Guyton-Klinger guardrails (Guyton and Klinger, 2006), the RMD method (Larson, 2022), spending-flexibility rules (Pfau, 2012a), and others — address this by adjusting spending annually in response to portfolio performance.

The efficient frontier framework accommodates all of them. Because spending in each scenario is computed by a full MILP optimizer that accepts any normalized spending profile — including dynamic ones — the framework is not limited to flat real amounts. In practice, a dynamic rule simply defines a different spending profile, and the HSF and frontier are computed from it in exactly the same way.

The most natural implementation is annual re-evaluation: each year, as new information arrives — updated portfolio values, revised return forecasts, tax-law changes, or changed personal circumstances — the framework re-solves the MILP for the remaining horizon and recomputes the frontier. The new HSF reflects the portfolio's current state, and the retiree selects a new committed spending level for the coming year. Guyton-Klinger guardrails and the RMD method are both special cases: each corresponds to a specific path through the frontier over time, selecting a new point each year rather than holding a static commitment. Annual re-evaluation also addresses an important practical concern: the HSF computed at retirement will differ from the HSF computed 5 or 10 years later, as the retiree's remaining horizon has shortened and the historical dataset has grown.

The computational structure of the framework supports this cycle efficiently. Constructing the frontier requires one MILP solve per scenario — whether historical or Monte Carlo — and once those solves are complete, exploring the full range of target success rates is instantaneous: the commitment LP recombines the existing g_s values without additional MILP computation.

One subtlety of dynamic spending rules deserves mention. As noted, g_s is a clairvoyant upper bound: the per-scenario MILP allocates spending with full knowledge of future returns. A real dynamic rule cannot do this: adjustments fire only after observed portfolio decline, meaning full committed spending continues through the most damaging early withdrawals before any reduction takes effect. Sequence-of-returns risk is most acute in early retirement, when the portfolio is largest and each dollar withdrawn forecloses the most future compounding — precisely the window where reactive rules provide no protection. Dynamic rules are therefore most valuable for upside management in favorable scenarios, less so for downside protection in the worst ones. This reinforces the role of the HSF and SSF as the primary tools against sequence-of-returns risk: both address the problem at its source by setting the committed level conservatively enough that adverse sequences are either fully absorbed (HSF) or statistically contained at the chosen confidence level (SSF), rather than relying on mid-course corrections that arrive after the critical damage is done.

Compatibility with Existing Tools

The efficient frontier framework is also accessible to advisors using conventional planning tools. Such tools test a user-specified spending level and report how many scenarios succeed — some go further and report shortfall magnitude in failing scenarios. The key difference is that balance-forward simulators cannot find the optimal spending level for each scenario: g_s is a decision variable, not a given. In principle, a greedy binary search over spending levels can approximate g_s for simple portfolio-only cases; for households with tax-deferred accounts, Roth conversion opportunities, and Medicare costs, however, the spending maximum depends on jointly optimizing all financial decisions simultaneously — which balance-forward simulators do not do. Because Owl's MILP maximizes spending subject to all constraints for each scenario, the resulting g_s is the true maximum sustainable spending and the shortfall measure is exact within the model's assumptions. Advisors using conventional simulators can in principle approximate the frontier by extracting per-scenario spending outcomes and applying the commitment LP as a post-processing step.

Extensions

The observation that Social Security compresses the distribution of per-scenario spending scales generalizes to any guaranteed income stream. Fixed annuities, defined-benefit pensions, and deferred income annuities (DIAs) all reduce portfolio withdrawal dependence by the same mechanism: payments continue regardless of market performance, placing a floor under each scenario's g_s. The HSF framework accommodates them directly — guaranteed income flows enter the cash-flow constraints exactly as Social Security does. A retiree weighing the purchase of an annuity against retaining full portfolio flexibility can compare HSFs and efficient frontiers under both scenarios, making the spending/risk trade-off of the annuity decision explicit and quantifiable.

The planning horizons in this paper are fixed deterministic inputs. The framework extends naturally to stochastic longevity: each scenario is assigned an independently drawn lifespan τ_s from a Social Security Administration (SSA) period life table (Social Security Administration, 2025a), coupling market risk and longevity risk within the same optimization. For a couple, the last-survivor horizon $\tau_s = \max(\tau_{1,s}, \tau_{2,s})$ is used. The MILP optimizer is solved with horizon τ_s for each scenario, producing a g_s that encodes both return and longevity risk. The commitment LP is entirely unchanged — it operates on whatever $\{g_s\}$ values the scenario solves produce. The resulting frontier spans the joint distribution of market and longevity risk, making longevity uncertainty explicit within the same framework rather than treating it as a fixed assumption.

The optionality of mid-course correction can in principle be captured more completely through multi-stage stochastic programming, which models spending as a sequence of adaptive decisions on a scenario tree rather than a single pre-committed level. Each year's spending responds to observed returns, fully capturing the value of adjustment and providing the theoretically complete treatment of adaptive risk. This approach is substantially more computationally demanding than the single-stage MILP used here, but represents a natural extension for cases where the value of adaptation is itself a quantity of interest.

CONCLUSION

This article introduces two concepts that together make a rigorous efficient frontier framework for retirement spending possible: the Historical Spending Floor (HSF), grounded in the historical record, and the Synthetic Spending Floor (SSF), calibrated to prospective conditions.

The HSF is the highest real spending level that survives every historical return sequence, computed via MILP for a given retirement case. For a stylized single retiree with a pure Roth account and no tax complications, it reproduces the established 4% benchmark exactly once methodological assumptions are aligned — a useful validation. For a real couple with Social Security, tax-deferred accounts, Medicare costs, and Roth conversion opportunities, the HSF is higher primarily because Social Security reduces the portfolio withdrawal burden in every scenario. The MILP handles the associated tax complexity optimally; bequest constraints shape the feasible spending range, and the joint optimization of Roth conversions, Medicare costs, and account withdrawals determines the true after-tax maximum in every scenario, within the model's assumptions.

The efficient frontier places the HSF in a richer context. Rather than identifying a single safe spending level, it maps the full trade-off between committed spending and mean shortfall across historical scenarios — a formally Pareto-optimal curve analogous to the Markowitz mean-variance frontier. Retirees can choose any point on this frontier based on their risk tolerance, and they can see exactly how much spending they gain and how much shortfall risk they accept at each choice. The HSF is the conservative anchor; everything to the right of it involves some risk, but that risk is explicit, quantified, and expressed in the same dollar units as the spending commitment itself.

The SSF extends the same framework to Monte Carlo scenarios calibrated to prospective return assumptions. When HSF and SSF agree, prospective conditions are expected to mirror the historical record; when they diverge, the gap is itself informative. Despite the complexity of the underlying optimization — one MILP solve per scenario, across hundreds of scenarios — the full computation completes in minutes on a standard laptop.

Together, the HSF, the SSF, and the efficient frontier give retirement planning a rigorous, personalized alternative to universal rules — grounded in both the past and the present, and capable of expressing the full complexity of a real household's financial situation.

DECLARATION OF SUBMISSION

The authors confirm that this manuscript is original work that has not been published elsewhere. It is not currently under review at any other publication in the Portfolio Management Research family or any other journal. The authors have read the PMR publication agreement and agree to its terms. The authors' institutions do not require that this paper be posted on any public-facing website, including SSRN, ResearchGate, or SciHub, prior to publication.

References

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance 9*(3), 203–228.

Bengen, W. P. (1994, October). Determining withdrawal rates using historical data. *Journal of Financial Planning*, 171–180.

Blanchett, D. M. (2007). Dynamic allocation strategies for distribution portfolios: Determining the optimal distribution glide path. *Journal of Financial Planning 20*(12), 68–81.

Blanchett, D. M., M. Kowara, and P. Chen (2012). Optimal withdrawal strategy for retirement portfolios. *Retirement Management Journal 2*(3), 7–20.

Campbell, J. Y. and L. M. Viceira (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford: Oxford University Press.

Collins, P. J., H. D. Lam, and J. Stampfli (2015). How risky is your retirement income risk model? *Financial Services Review 24*, 193–216.

Cooley, P. L., C. M. Hubbard, and D. T. Walz (1998, February). Retirement savings: Choosing a sustainable withdrawal rate. *Journal of the American Association of Individual Investors*, 16–21. Also known as the Trinity Study.

Coopersmith, L. W. and A. R. Sumutka (2011, September). Tax-efficient retirement withdrawal planning using a linear programming model. *Journal of Financial Planning*.

Damodaran, A. (2026). Annual returns on stock, t.bills, t.bonds and inflation—united states (1928–present). NYU Stern School of Business, available at `https://pages.stern.nyu.edu/~adamodar/`. Historical annual returns dataset, updated annually.

Dushi, I., H. M. Iams, and B. Trenkamp (2017). The importance of social security benefits to the income of the aged population. *Social Security Bulletin 77*(2), 1–12. Available at `https://www.ssa.gov/policy/docs/ssb/v77n2/v77n2p1.html`.

Finke, M. S., W. D. Pfau, and D. M. Blanchett (2013, June). The 4 percent rule is not safe in a low-yield world. *Journal of Financial Planning 26*(6), 46–55.

Fitzpatrick, J. and D. Tharp (2022, May). Evaluating retirement spending risk: Monte carlo vs historical simulations. Kitces.com. Available at `https://www.kitces.com/blog/monte-carlo-simulation-historical-returns-sequence-risk-calculate-sustainable-spending-levels/`.

Forsyth, P. A. (2020). A stochastic control approach to defined contribution plan decumulation: "the nastiest, hardest problem in finance". *ASTIN Bulletin: The Journal of the IAA*. arXiv:2008.06598.

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2021). Optimal control of the decumulation of a retirement portfolio with variable spending and dynamic asset allocation. *ASTIN Bulletin: The Journal of the IAA 51*(3), 905–938.

Gardner, G. and S. Pittman (2013). Measuring the risk of running out of money in retirement. *Journal of Financial Planning 26*(12), 38–44.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies 27*(3), 714–746.

Guyton, J. T. and W. J. Klinger (2006, March). Decision rules and maximum initial withdrawal rates. *Journal of Financial Planning*.

Kitces, M. E. (2008, May). Resolving the paradox – is the safe withdrawal rate sometimes too safe? *The Kitces Report*.

Konicz, A. K. and J. M. Mulvey (2013). Applying a stochastic financial planning system for an individual: Immediate or deferred life annuities? *Journal of Retirement 1*(1), 90–103.

Lacasse, M.-D. (2026). Owl: Optimal wealth lab – a mixed-integer linear programming retirement financial planner. `https://github.com/mdlacasse/Owl`. Version 2026.04.21. Additional contributors are listed in the `AUTHORS` file at `https://github.com/mdlacasse/Owl/blob/main/AUTHORS`. The case file `Case_chris+pat.toml` is included in the repository for reproducibility.

Lacasse, M.-D. and J. D. Leonard (2026). The retirement efficient frontier, Part II: The retirement efficiency score and optimal probability of success. *Journal of Retirement*. Forthcoming.

Larson, S. J. (2022, January). Required minimum distributions as a retirement strategy: The tradeoff between RMD volatility and the expected number of dollars paid out. *Journal of Financial Planning 35*(1), 60–67.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance 7*(1), 77–91.

Pfau, W. D. (2012a, March). Spending flexibility and safe withdrawal rates. *Journal of Financial Planning*.

Pfau, W. D. (2012b, January). Withdrawal rates, savings rates, and valuation-based asset allocation. *Journal of Financial Planning 25*(1), 36–48.

Pfau, W. D. (2013, February). A broader framework for determining an efficient frontier for retirement income. *Journal of Financial Planning*. SSRN 2151259.

Politis, D. N. and J. P. Romano (1994). The stationary bootstrap. *Journal of the American Statistical Association 89*(428), 1303–1313.

Ragsdale, C. T., A. F. Seila, and P. L. Little (1994). An optimization model for scheduling withdrawals from tax-deferred retirement accounts. *Financial Services Review 3*(2), 93–109.

Rockafellar, R. T. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of Risk 2*(3), 21–41.

Rockafellar, R. T. and S. Uryasev (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance 26*(7), 1443–1471.

Scott, J. S., W. F. Sharpe, and J. G. Watson (2009). The 4% rule—at what price? *Journal of Investment Management 7*(3), 31–48.

Social Security Administration (2025a). Actuarial life table, period life table 2021. `https://www.ssa.gov/oact/STATS/table4c6.html`. Accessed April 2026.

Social Security Administration (2025b). Fast facts & figures about social security, 2025. Technical report, Social Security Administration. Available at `https://www.ssa.gov/policy/docs/chartbooks/fast_facts/2025/fast_facts25.html`.

Tharp, D. (2017, July). Fat tails in Monte Carlo analysis vs. safe withdrawal rates. Kitces.com. Available at `https://www.kitces.com/blog/monte-carlo-analysis-risk-fat-tails-vs-safe-withdrawal-rates-rolling-historical-returns/`.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science 185*(4157), 1124–1131.

Webb, A. (2021, April). What can scholarly research tell us about the merits of annuitization vs. drawing down unannuitized wealth? Technical Report Literature Review #004-2021, Retirement Income Institute.

Welch, J. S. (2015). Mitigating the impact of personal income taxes on retirement savings distributions. *Journal of Personal Finance*.

Welch, J. S. (2016). Measuring the financial consequences of IRA to Roth IRA conversions. *Journal of Personal Finance*.

Welch, J. S. (2017, August). A 3-step procedure for computing sustainable retirement savings withdrawals. *Journal of Financial Planning*.

APPENDIX

A. Reconciliation of Bill's HSF with Bengen's SWR

Bill's HSF of $36,666/yr — equivalent to a 3.67% withdrawal rate — differs from Bengen's 4.15% for three methodological reasons, though the third is negligible. Beginning-of-year withdrawals reduce the effective rate because funds are drawn before returns are applied; Bengen assumed end-of-year withdrawals (+0.32%). The implementation adopts beginning-of-year withdrawals as the conservative and practically realistic choice: retirees cover living expenses at the start of each year, not after waiting for a year of investment returns. The implementation uses 10-year Treasury notes, which suffered larger price losses than Bengen's 5-year notes during the rising-rate environment of 1966–1981 (+0.16%). Both effects are independent and additive. The third is Bengen's use of long-run average returns for cohorts whose 30-year window extended beyond his 1926–1992 data; for the binding 1966 cohort only the final three years (1993–1995) are affected, when the portfolio is nearly depleted. Owl's data extends to 2025, covering all 66 windows with actual returns.

The 0.48% total gap also illustrates the sensitivity of the SWR to seemingly minor methodological choices (Pfau, 2012b; Cooley et al., 1998): two assumptions — withdrawal timing and bond maturity — that a practitioner might consider inconsequential together shift the rate by nearly half a percentage point, a material difference when applied to a seven-figure portfolio.

Factor	Bengen	Bill (Owl)
Withdrawal timing	End of year	Beginning of year
Bond type	5-year T-Notes	10-year T-Notes
Tax treatment	Simplified (not modeled)	None (pure Roth)
Planning horizon	30 years	30 years
Historical range	1926–1992	1928–1993
Safe withdrawal rate	4.15%	3.67%
+ beginning-of-year adjustment		+0.32%
+ 5-year vs. 10-year bond adjustment		+0.16%
Reconciled rate	4.15%	4.15% ✓

Appendix Exhibit A: Reconciliation of Bill's HSF with Bengen's SWR.

B. Commitment LP and CVaR Connection

Given N per-scenario spending scales $g_s \geq 0$ and a risk-aversion parameter $\lambda \geq 0$, the commitment LP is:

$$\begin{aligned}
\text{maximize} \quad & g^* - \frac{\lambda}{N}\sum_{s=1}^{N} \sigma_s \\
\text{subject to} \quad & \sigma_s \geq g^* - g_s, \qquad s = 1,\ldots,N \qquad (1) \\
& \sigma_s \geq 0, \qquad s = 1,\ldots,N \qquad (2) \\
& 0 \leq g^* \leq \max_s g_s. \qquad (3)
\end{aligned}$$

The $N+1$ decision variables are the scalar g^* and the N shortfall variables σ_s. Constraints (1)–(2) force $\sigma_s = \max(0, g^* - g_s)$ at optimality — the per-scenario spending deficit. The *mean shortfall* is then

$$\bar{\sigma}(g^*) \;=\; \frac{1}{N}\sum_{s=1}^{N} \sigma_s \;=\; \frac{1}{N}\sum_{s=1}^{N} \max\big(0,\, g^* - g_s\big), \qquad (4)$$

the average per-scenario spending deficit at commitment level g^*. Constraint (3) prevents g^* from exceeding the best achievable outcome. At $\lambda = 0$ the LP simply maximizes g^* (best-case, aggressive extreme). As $\lambda \to \infty$, g^* is driven to $\min_s g_s$ = HSF (zero shortfall, conservative extreme). Solving the LP for a grid of λ values traces the full Pareto curve; in practice, the grid is sampled on a logarithmic scale to obtain uniform visual density across the full curve.

If a scenario is infeasible — no plan satisfying all constraints exists under that return sequence — its spending scale is set to $g_s = 0$, contributing a full shortfall $\sigma_s = g^*$ for any positive commitment. The normalization constant N always equals the total number of scenarios requested, ensuring that infeasible scenarios are not silently excluded from the risk calculation.

The mean shortfall $\bar{\sigma}$ scaled by the failure rate, $\bar{\sigma}/(1-\rho)$ where ρ is the target success rate, equals the Conditional Value-at-Risk (CVaR) of spending losses at confidence level ρ (Rockafellar and Uryasev, 2000, 2002) — a coherent risk measure in the sense of Artzner et al. (1999). The commitment LP is therefore structurally identical to their CVaR minimization LP, with g^* playing the role of the Value-at-Risk threshold. Note that the SSF characterizes shortfall severity only when $\rho < 95\%$, so that more than 5% of scenarios fall below committed spending; when $\rho \geq 95\%$, the SSF lies in the success zone and the CVaR is the more informative tail statistic.

If spending surpluses — amounts by which scenario outcomes exceed the commitment — are treated as equally costly as shortfalls, the penalized term becomes the mean absolute deviation $\frac{1}{N}\sum_s |g_s - g^*|$, and the commitment LP generalizes to

$$\text{maximize} \quad g^* - \frac{\lambda}{N}\sum_{s=1}^{N} |g_s - g^*|. \tag{5}$$

The first-order condition yields $g^* = F^{-1}\big(\frac{1}{2} + \frac{1}{2\lambda}\big)$, where $F(g^*)$ is the fraction of scenarios with $g_s \leq g^*$. As $\lambda \to \infty$, g^* converges to the *median* of $\{g_s\}$ — the most conservative choice under symmetric loss. This gives the median column in the results tables a precise interpretation: it is the symmetric-loss solution at $\lambda \to \infty$, the committed spending level that minimizes expected absolute deviation from actual outcomes.

The CVaR and median together suggest a plan-quality scalar analogous to the Sharpe ratio. Part II of this series (Lacasse and Leonard, 2026) introduces the *Retirement Efficiency Score* (RES), defined per target success rate ρ as the ratio of upside spending above the HSF to tail risk, and shows that its maximum identifies a principled optimal target rate ρ^* — the tangency point on the spending frontier.