Rebalance, Don't De-Risk: Lessons from 10,000 Simulated Retirements

Under Market Regime Uncertainty

Wen Rong (wr2247) Lena Ho (yh6628) Filip Zaleski (fz2139)

Stochastic Modeling and Simulation

August 2026

The Problem

A very ordinary situation

Someone retires with \$1,000,000 and needs it to last 30 years.

Each year they take out money to live on. What is left stays invested in stocks, bonds and cash.

Two things can go wrong:

1. They run out of money before they run out of life.
2. They are forced into a drastic cut to their standard of living.

Every retiree, every pension fund, and every financial adviser faces some version of this.

The obvious approach: just use the average

Historically, a 60/40 stock/bond portfolio returned about 6.5% a year after inflation.

Withdraw 4%, earn 6.5%, and you are comfortably ahead. On average, the money should never run out.

This reasoning is wrong — and it is wrong in a way that
cannot be fixed by picking a better average.

The same returns, in a different order

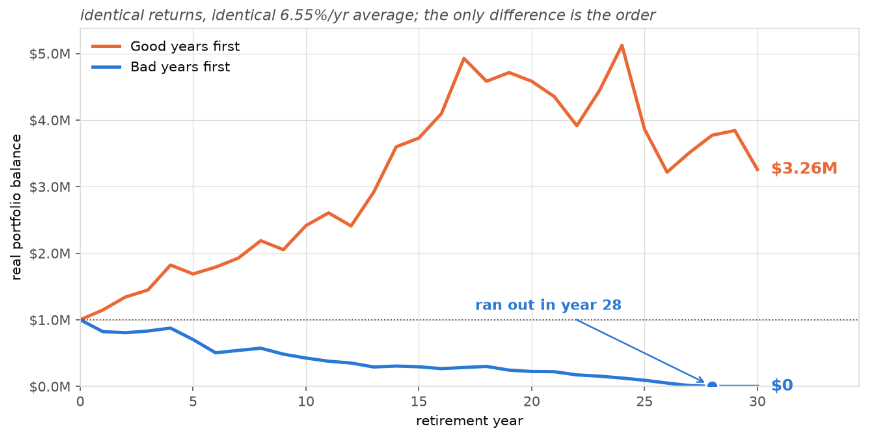

Both lines: a 60/40 portfolio drawing 4% of the initial balance, over the same 30 annual real returns.

Source: R. J. Shiller, *Online Data — U.S. Stock Markets 1871–Present*, Yale University. US real total returns, 1969–1998.

Why order matters: sequence-of-returns risk

A crash in year 2 forces you to sell a large slice of a shrunken portfolio just to pay that year's bills. The base that has to compound back up is permanently smaller.

The identical crash in year 27 barely matters.

So the outcome depends on the entire path of returns — not on their average.

Average-based formulas throw the ordering away. That is exactly the information that decides whether the retiree is fine.

So we simulate

We generate thousands of complete 30-year lifetimes for a retiree and count how many end badly.

Our decision question

For a retiree drawing down a portfolio over 30 years, which combination of **asset allocation**, **withdrawal rule** and **rebalancing policy** minimises the chance of running out of money, while still supporting a decent standard of living?

What we fed the model

Shiller dataset (Yale)

Monthly, 1871–2026

- S&P 500 total return
- 10-year Treasury bonds
- Consumer Price Index

FRED 3-month T-bills

Daily, 1981–2026

- our "cash" asset

Everything is in real terms

All returns are adjusted for inflation, so a $40,000 withdrawal buys the same amount in year 30 as in year 1.

154 years of stock and bond history.

Long enough to contain the Depression, the 1970s, 2008 and COVID.

The Model

Markets are not the same thing all the time

Market conditions vary over time. Some periods are relatively calm and rises steadily, while others—such as 1929–1933, 1974, 2008–2009, and 2020—experience much greater volatility.

We classify each year into one of **two market regimes** based on realized stock-market volatility:

Calm
Higher average returns
Lower volatility

Stressed
Lower average returns
Higher volatility

Once the regimes are classified, their year-to-year transitions are modeled as a **two-state Markov chain**.

How we classified historical years as calm or stressed

For each year in our historical sample, we measure **stock-market realised volatility** s_y: the standard deviation of its twelve monthly stock returns.

$$R_y = \begin{cases} \text{stressed}, & \text{if } s_y > q_{2/3} \\ \text{calm}, & \text{otherwise} \end{cases}$$

$q_{2/3}$ is the 66.7th percentile of historical volatility: **the most volatile one-third** of years are classified as stressed.

The years it picks out:

1873 1893 1907 1929–35 1937 1974 1987 1998 2001–02 2008–09 2011 2020 2022

Essentially every famous crisis in US financial history.

Estimating transition probabilities and parameters within each regime

Transition probabilities are estimated from how often each regime is followed by calm or stressed.

After classifying each historical year as calm or stressed, we estimate the **return characteristics** separately within each regime.

For regime $k \in \{\text{calm}, \text{stressed}\}$:

$$\boldsymbol{\mu}_k = \begin{pmatrix} \bar{R}_{\text{stock},k} \\ \bar{R}_{\text{bond},k} \\ \bar{R}_{\text{cash},k} \end{pmatrix} \qquad \boldsymbol{\sigma}_k = \begin{pmatrix} \sigma_{\text{stock},k} \\ \sigma_{\text{bond},k} \\ \sigma_{\text{cash},k} \end{pmatrix} \qquad C_k = \begin{pmatrix} 1 & \rho_{SB,k} & \rho_{SC,k} \\ \rho_{SB,k} & 1 & \rho_{BC,k} \\ \rho_{SC,k} & \rho_{BC,k} & 1 \end{pmatrix}$$

Mean returns — Volatilities — Correlations

These parameters are calculated using only the historical years classified into regime k.

Calm vs. Stressed Markets: Estimation Results

Transition probabilities P

	next year	
now	calm	stressed
calm	**72.8%**	27.2%
stressed	54.0%	**46.0%**

Calm periods last 3.7 years on average; stressed periods last 1.9 years on average.

Stationary distribution ($\pi P = \pi$)
In the long run, the model spends about **67%** of years in the calm regime and **33%** in the stressed regime.

Real returns in each state

	Calm		Stressed	
	avg	vol	avg	vol
S&P 500	+12.6%	15.4%	+0.6%	20.1%
Bonds	+2.5%	7.7%	+3.3%	10.5%
T-bills	+1.0%	2.5%	+0.6%	3.0%

In calm years stocks earn 12.6% after inflation. In stressed years they earn essentially nothing — and exhibit substantially higher volatility.

A result we did not expect

Calm regime

103 years (32 with cash)

	S&P	10y	cash
S&P 500	1.00		
10y bond	**0.23**	1.00	
3m T-bill	0.12	0.54	1.00

Stressed regime

51 years (12 with cash)

	S&P	10y	cash
S&P 500	1.00		
10y bond	**0.12**	1.00	
3m T-bill	0.19	0.67	1.00

The standard claim is that in a crisis everything falls together, so diversification stops working. Our data show the opposite: stock–bond correlation drops from $+0.23$ to $+0.12$.

A possible explanation is flight to quality — when stocks fall sharply, investors move into Treasuries. Bonds may therefore diversify best in exactly the years that threaten a retiree.

Markov-Chain Simulation of a 30-Year Retirement

Start with \$1M split according to the strategy and an initial market state according to stationary distribution. Then each year:

1. Move to the next market state using the estimated transition probabilities.
2. Draw stock, bond, and cash returns jointly from the current state's distribution.
3. Withdraw that year's spending.
4. Apply the simulated returns to the remaining portfolio.
5. Rebalance, if required by the strategy.
6. If the portfolio runs out of money, record the year and stop.

$$\mathrm{r}_t \mid S_t = k \sim \mathcal{N}(\boldsymbol{\mu}_k, \Sigma_k) \qquad B_t = (B_{t-1} - W_t)(1 + \rho_t)$$

S_t :	market state (calm or stressed)	W_t :	withdrawal in year t
r_t :	stock, bond, and cash real returns	ρ_t :	portfolio return in year t
$\boldsymbol{\mu}_k$:	mean returns in state	B_t :	portfolio balance at the end of year t
Σ_k :	covariance matrix in state k		

Monte Carlo Simulation: Repeating the 30-Year Paths

The Markov-chain procedure generates **one possible 30-year retirement path**.

We repeat this process **10,000 times**, independently generating different market states and asset returns.

One 30-year path $\longrightarrow$ **10,000 simulated paths** $\longrightarrow$ Retirement outcomes

From these 10,000 outcomes, we estimate:

- Probability of portfolio ruin
- Distribution of terminal wealth
- Downside risk (CVaR)
- Sustainable withdrawal rate

This repeated random simulation is the Monte Carlo step.

The Five Strategies We Compared

	Allocation	Withdrawal Rule	Rebalancing
S1	60% stocks / 40% bonds	\$40,000, fixed real \$	Annual
S2	60% stocks / 40% bonds	4% of current balance	Annual
S3	Glide path: 60/40/0 → 30/60/10	\$40,000, fixed real \$	Annual
S4	Glide path: 60/40/0 → 30/60/10	4% of current balance	Annual
S5	Starts at 60% stocks / 40% bonds	\$40,000, fixed real \$	None

- **Fixed real withdrawal:** Year 1 withdrawal is 4% of initial wealth. The same real dollar amount is then withdrawn each year.
- **Percentage withdrawal:** Each year, the retiree withdraws 4% of the portfolio value remaining at that time.
- **Glide path:** Over 30 years, stocks decline linearly from 60% to 30%, cash rises from 0% to 10%, and bonds rise from 40% to 60%.

Fair comparison: All five strategies are evaluated on the same 10,000 simulated market paths, so differences in outcomes come from the strategies rather than different random draws.

What We Found

The headline numbers

	Chance of running out	Typical final balance	Worst 10% end with	Sustainable withdrawal
S1 baseline	6.5%	\$1.97M	\$190k	4.34%
S2 % rule	0.0%	\$1.53M	\$666k	3.93%
S3 glide	6.9%	\$1.49M	\$141k	4.29%
S4 glide, %	0.0%	\$1.23M	\$610k	4.08%
S5 no rebal.	7.5%	\$2.15M	\$123k	4.23%

The 4% rule roughly holds up

A 6.5% failure rate and a sustainable rate of 4.34% land squarely on the published retirement literature (Bengen 1994; Trinity study 1998) — which got there by completely different means. Useful evidence we have not built something broken.

Rebalancing helps — but not for the reason you would guess

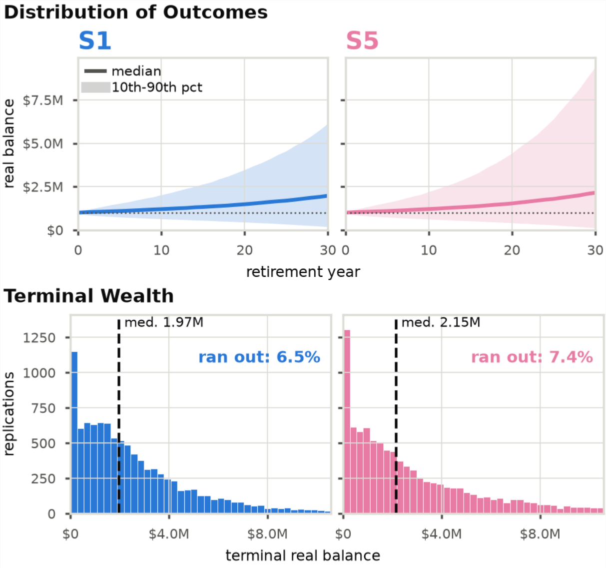

	Allocation	Withdrawal	Rebal.
S1	60/40	$40k fixed	Annual
S5	60/40, drifts	$40k fixed	None

The clean test: identical but for rebalancing. Failure falls 7.5% → 6.5%.

Yet buy-and-hold ends *richer* in the typical case, $2.15M vs $1.97M — see how much wider its band runs.

The whole difference is in the bad tail: worst-10% $123k vs $190k.

Rebalancing does not make you richer. It stops you drifting to maximum stock exposure right before a crash.

The glide path did not work

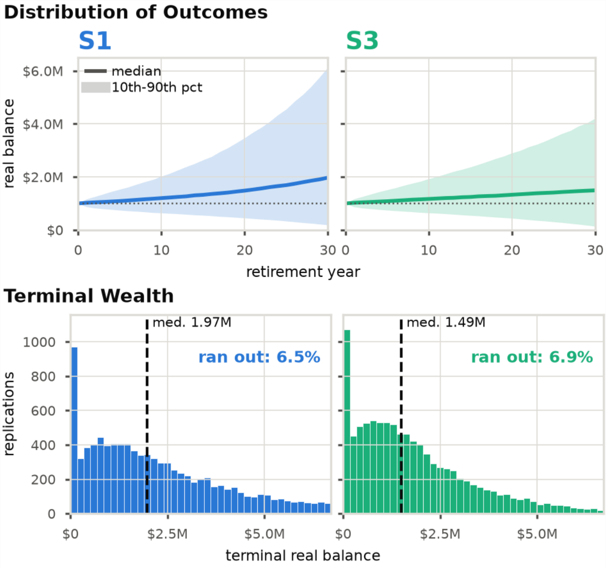

	Allocation	Withdrawal	Rebal.
S1	60/40	$40k fixed	Annual
S3	Glide 60→30%	$40k fixed	Annual

De-risking with age is standard advice. Here it makes things slightly worse: 6.9% failure against 6.5%.

It protects the wrong end of the retirement. The danger sits in the first few years — exactly where the glide path still carries its full 60% stock weight.

By the time it turns defensive, the survivors no longer need it. The whole distribution shifts left: $1.49M vs $1.97M.

All of the cost of de-risking, none of the benefit.

The % rule does not remove risk — it moves it

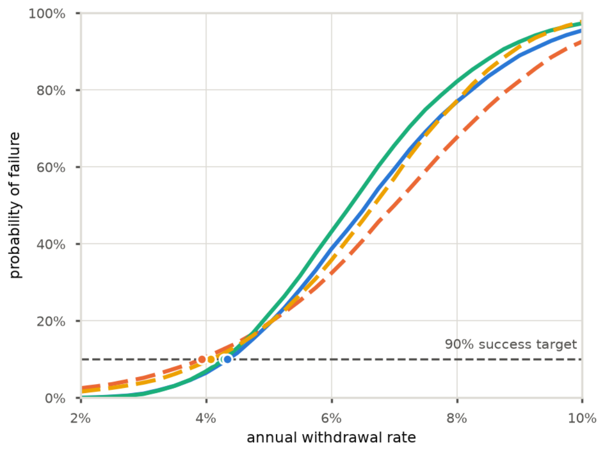

Sustainable rate, 90% success bar

S1 4.34% S2 3.93% S3 4.29% S4 4.08%

solid = runs out dashed = income halves

	Allocation	Withdrawal	Rebal.
S1	60/40	\$40k fixed	Annual
S2	60/40	4% of balance	Annual
S3	Glide 60→30%	\$40k fixed	Annual
S4	Glide 60→30%	4% of balance	Annual

S2 and S4 never run out — but that is arithmetic, not skill: 4% of a shrinking balance is always positive.

The cost lands on income, which swings by roughly a quarter of its own average, hardest when markets are worst.

Require that income never halves and S2 supports 3.93% — *below* the baseline's 4.34%.

Their strength is the other tail: worst-10% of \$666k against \$190k.

What we would actually tell a retiree

1. **Rebalance annually.** The cheapest improvement available — a full percentage point off the failure rate, and a 55% better worst-case, at no cost but discipline.

2. **Skip the conventional glide path.** On our model it costs both money and safety.

3. **Pick the withdrawal rule that matches which risk you can stomach.** Fixed dollars protect your lifestyle and accept a 6.5% chance of ruin. A % rule guarantees you never go broke but lets your income swing. Neither dominates.

4. **Treat 4% as a ceiling, not a target.** It clears a 90% success bar only before taxes and fees.

Where we would push back on ourselves

- **No taxes or fees.** Every number here is therefore optimistic.
- **We assume the past repeats.** 154 years of US data is the most successful market in history — a real survivorship problem.
- **The cash leg is thin.** Only 44 years of T-bill data, and just 12 stressed years in it.
- **Exactly 30 years.** Nobody knows their lifespan; adding mortality would change the picture.
- **Two regimes may be too few.** A separate short, violent "crisis" state would fit 1929–33 and 2008 better.

Thank you

Questions?

Backup: is 10,000 runs enough?

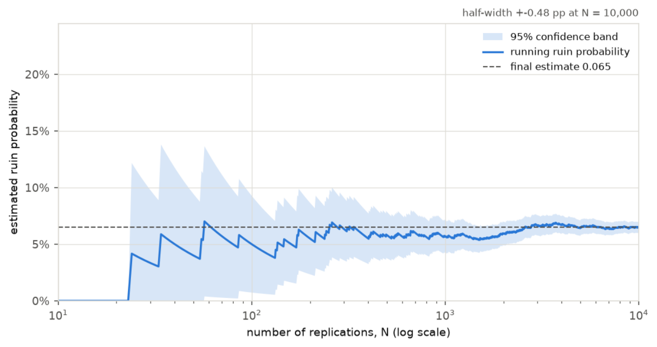

More runs give a more precise answer, but take longer.

At $N = 10{,}000$ our 95% confidence interval on the failure rate is ± 0.48 percentage points.

We had set ourselves a ± 0.5 pp target, which needs 9,312 runs. So 10,000 clears it.

Backup: the input data

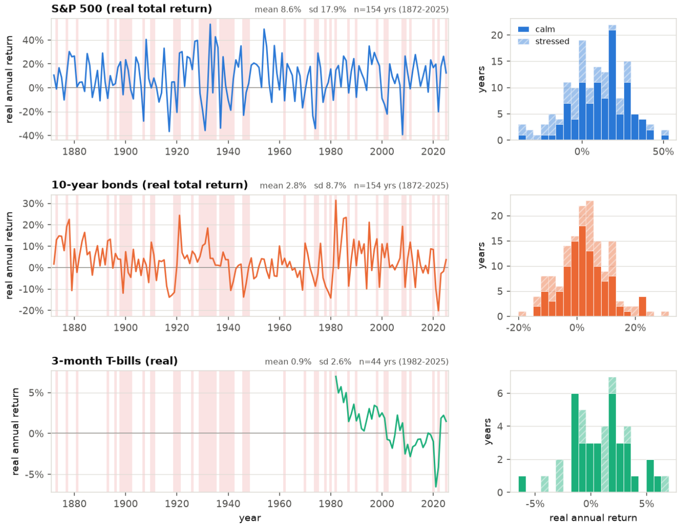

stressed-regime year (top tercile of realised equity volatility)

Backup: regime-fitting methods we rejected

Gaussian mixture on annual returns. Splits years into “good” and “bad” at roughly the average. The resulting transition probabilities came out near 50/50 regardless of the current state — no persistence, so no clustering, so no sequence-of-returns risk. Unusable.

Markov-switching model fitted by EM. Highly persistent, but it found the wrong thing: its “stressed” state had *lower* stock volatility, higher bond volatility, and lasted 11–15 years. It had detected the pre/post-1970s shift in bond markets, not market stress — and our model explicitly assumes no structural breaks.

The volatility threshold avoids both failures, and results are stable for any cutoff between the 60th and 80th percentile.

Backup: joint simulation of asset returns in practice

$$\mathrm{r}_t \mid S_t = k \sim \mathcal{N}(\boldsymbol{\mu}_k, \Sigma_k) \qquad\qquad B_t = (B_{t-1} - W_t)(1 + \rho_t)$$

Joint return draw:

$$\mathrm{z}_t \sim \mathcal{N}(0, I_3), \qquad \mathrm{r}_t = \boldsymbol{\mu}_k + L_k \mathrm{z}_t, \qquad L_k L_k^\top = \Sigma_k$$

This preserves the estimated stock–bond–cash correlations within each regime.

Why this generates tail risk: returns come from a mixture of calm and stressed distributions, rather than one single normal distribution, producing a heavier-tailed long-run return distribution.

S_t :	market state (calm or stressed)	W_t :	withdrawal in year t
r_t :	stock, bond, and cash real returns	ρ_t :	portfolio return in year t
$\boldsymbol{\mu}_k$:	mean returns in state	B_t :	portfolio balance at the end of year t
Σ_k :	covariance matrix in state k	L_k :	Cholesky factor
		I_3 :	3×3 identity matrix