Glidepath Lab: Frequently Asked Questions

lazyeconomist.com · Methodology series, Paper 3 of 3 · September 2026 · engine v0.4

Engine, API, the browser workbench and the six-factor memo generator — which builds a versioned evidence document from a saved run — are live at glidepath.lazyeconomist.com.

About the tool

What is Glidepath Lab? A simulator for testing a private-markets sleeve inside a target-date glidepath. It projects participant retirement income under a range of return assumptions, simulates plan-level cash flows against the vehicle's redemption cap, computes the fee hurdle the sleeve must clear, and assembles that evidence into a versioned six-factor memo aligned with the proposed DOL rule.

What can I do today? Build a plan and run it in the browser workbench at `/app/build` — fan charts, the liquidity dashboard, the fee panel. You can also post the same inputs to the demo endpoint and read back the JSON, browse the interactive API documentation at `/docs`, and read these three papers. Run it anonymously and you get demo mode: at most 500 paths, two stress scenarios, premiums up to +3 pp, nothing stored. Sign in and a run is saved to your workbench, where it can be reopened and turned into a six-factor memo.

Who is it for? Plan committees, investment consultants and target-date designers evaluating or building a private sleeve. A public demo mode lets anyone explore the mechanics on an illustrative plan.

Is it advice? No. It is an educational and analytical tool. Its default assumptions are illustrative and cited in the technical note. It does not recommend an allocation, a vehicle or a manager for any plan, and it is not affiliated with any provider.

Does it use live fund data? No. All inputs — capital market assumptions, fees, vehicle terms — are user-supplied. The engine can estimate a vehicle's smoothing parameter from an uploaded reported-return history, but that upload path is not yet wired into the API; nothing is fetched automatically.

What is the difference between demo and workbench mode? Demo mode is public and stateless; workbench mode is access-controlled and persistent. The same engine runs behind both.

	DEMO	WORKBENCH
Sign-in	none	Cloudflare Access
Monte Carlo paths	up to 500	up to 20,000
Stress scenarios per run	2	8
Premium grid	up to +3 pp	any grid the schema accepts: up to eight points in -10 to +20 pp
Where the run executes	inline, inside the request	a background worker, polled
Saved plans and assumption sets	no	yes
What is stored	nothing	the run, its inputs, its summary, and every memo generated from it

A workbench run can be turned into a six-factor memo, versioned per run and printed from the browser. Demo runs are not stored, so they cannot produce one.

Expected returns

Why do you not use a single expected return for private equity? Because credible estimates of the premium over public equity range from roughly zero to three percentage points a year. Choosing one would make the allocation decision for the committee. The Lab runs every configuration across a grid of premiums and reports the break-even premium the sleeve needs after fees. Where your own assumption sits on that grid is the decision.

What is the break-even premium? The annual return, over the public asset the sleeve replaces, at which the median participant's retirement income with the sleeve equals income without it, after all fees. In v0.3 it is reported net of the sleeve's fee stack; a gross figure is planned.

How do you handle manager dispersion? Each simulation path draws a manager alpha from a distribution you set, centred on zero by default with a ±3 point range. A committee confident of top-quartile access can shift the centre.

Are private credit returns treated differently? Yes. Private credit has a deferred-loss mechanism: an economic loss can be applied immediately while the reported NAV recognises

the same loss, in full, some months later. That reflects the way loan modifications and payment-in-kind defer rather than remove losses. The engine reports the gap.

Does the model believe in an illiquidity premium? It has no opinion. It makes the consequences of each belief visible.

Risk, valuation and diversification

Why are private assets so much riskier in the model than their reported volatility? Appraisal-based returns are smoothed, so reported volatility and correlation understate the economic reality. The Lab keeps two series for each private sleeve — economic and reported — and uses the economic series for risk. The smoothing parameter θ is user-set (or, in the engine though not yet the API, estimable from an uploaded history) and reported alongside the results.

What is θ? The fraction of true return recognised in the current period under a first-order smoothing filter. θ = 1 means marks are fully current; θ = 0.4 (the PE default) means reported volatility is about half of economic volatility.

Do private assets diversify? Less than reported correlations suggest. The model derives economic correlation from the funding asset plus an idiosyncratic term rather than from reported correlations. Private credit's diversification is further complicated by the deferred-loss mechanism.

What is the stale-mark scenario? A participant retiring into a public drawdown may see the private sleeve's reported value lag and then catch down after withdrawals begin. Scenario S5 measures the gap between reported and economic balance at the target date.

Liquidity

If a participant leaves, is her share of the private fund sold? Not directly. Participant outflows are met at plan level, first from public sleeves and only from the private sleeve if public sleeves are exhausted; the private sleeve otherwise trades only to rebalance, within its cap. The Lab enforces this as a tested invariant: participant-level liquidity is always met.

Then why does liquidity matter? Because the sleeve can only be reduced at the rate the vehicle allows — typically 5% of net asset value per quarter — and because plan-level events (layoffs, mergers, participants switching to cash) and public drawdowns push the sleeve away from target faster than that. Every run answers the following on every path — the first four per sleeve, the last three per cohort:

THE QUESTION A COMMITTEE ASKS	WHAT THE RUN REPORTS
How often does the cap bind?	the share of paths in which it binds, year by year
How long does a gate last?	the gate-duration histogram and the sustained-gate probability
How large does the queue get?	peak backlog, for the plan and for the whole vehicle
How long does a sale wait?	the dollar-weighted months a redemption request waits
How far does the weight drift?	the peak private weight a cohort reaches, and the probability it passes your tolerance
How long is the overweight carried?	the total months a cohort's private weight spent above the top of its band — a count of months, not the length of one spell
What did all this cost?	forced-sale cost and cash drag, per head

The last three are measured on each cohort's own private weight, aggregated across its private sleeves, and reported as the worst cohort on each path. None of them is the plan's own weight, which stays nearer the tolerance because the excess is sold as soon as a cohort breaches it.

What happened in 2026 that the scenarios reproduce? Large evergreen private credit funds received redemption requests well above 5% quarterly caps, upsized once, then enforced the cap with pro-rata fills and multi-quarter backlogs. Scenario S3 reproduces that dynamic; S2 reproduces the 2022 pattern where private weights drifted up because they could not be rebalanced.

Where should the cash buffer sit? The model supports both: at the TDF level (no drag on the sleeve; redemptions met from public assets) or inside the vehicle (a return drag proportional to the cash held). You can compare them.

Does vehicle scale matter? The model takes the vehicle's assets and investor mix as inputs. It does not model other investors' behaviour in detail beyond a steady base subscription rate and the request-spike scenarios, which are the proxy for a run on the vehicle.

Cost

How are fees modelled? Per sleeve: management fee, incentive fee with hurdle and optional catch-up (paid on reported returns, as in practice), acquired-fund fees and platform fees. The fund's total expense ratio is shown at every age with and without the sleeve.

Why is the default private equity fee around 3% all-in? Because that is what surveys of semi-liquid vehicles show for the average expense ratio, often before incentive fees. It is the priciest of the four sleeves: at the shipped defaults and a 2 pp premium, private equity works out at about 3.0% all-in, and private credit, real estate and real assets at 1.7-1.8%. Institutional pricing through a large TDF may be lower; enter it.

What does the break-even chart tell a committee? Where the break-even premium sits relative to your own assumption. If it sits above, the sleeve is a bet on manager access; if below, the fee case is made and the liquidity case becomes the binding one.

Glidepath design

Taper or constant sleeve? Providers disagree. Some taper private exposure toward retirement and place different private assets at different ages; one holds a constant multi-asset sleeve to avoid forced trades as the path rolls. Run the two shapes under the same seed and they draw identical returns, so the difference in income and in forced-trade and gate cost is attributable to the shape alone — the metric that separates them. Reading the two runs against each other is still a manual comparison: nothing in the workbench puts them side by side for you.

What is a funding map? The public sleeve each private sleeve is carved from (for example PE from equity, private credit from high yield). The no-sleeve baseline returns the weight to that asset, so the comparison is like-for-like.

Can I import a provider's glidepath? CSV import is a v2 feature. In v1 you enter the rows.

Fiduciary use

Why six factors? The proposed DOL safe harbor lists expected performance, fees, liquidity, valuation, benchmarking and complexity as the factors a fiduciary must evaluate. The memo is organised so that each factor is populated from the run's own evidence and the plan's record rather than from assertion, and so that the judgement sits where it belongs: in the committee note each section ends in. The technical note shows what evidence stands behind each.

Is the memo enough for the safe harbor? No. The rule is not final, the safe harbor is process-based, and a memo is one input to that process, not the whole of it. Counsel and the

consultant own the rest.

How is a run versioned? Every run stores its engine version, an input hash and its seed. Same inputs and seed reproduce identical results. Memos layer on top of that: regenerating a memo writes the next version rather than overwriting the last, so earlier versions stay readable.

Why does the memo report a stated objective and benchmarks? Because pending litigation turns on whether a plaintiff has pleaded a meaningful benchmark, and the objective and benchmarks chosen at design time become the yardstick. The Lab records governance inputs — objective text, benchmarks, revaluation triggers — with the plan, and the memo reports what the plan record holds. It does not ask for them at generation time: where the record carries none, the memo names the gap as pending rather than inviting someone to supply it as the document is written.

Practicalities

How long does a run take? Demo runs complete in seconds. A full run on the demo plan — eight cohorts, four scenarios, five premium variants each (a no-sleeve baseline plus four premium points), 2,000 paths per simulation — takes about a minute on the deployment hardware.

Are results reproducible? Yes. Same inputs and seed give identical outputs; the engine version is recorded.

What data do you keep? Demo mode stores nothing. Workbench data is held on a private server and contains cohort aggregates, the memos generated from them and the notes a committee writes into those memos; the schema has no participant-level fields.

What is not in v0.3?

NOT BUILT	WHERE IT STANDS
A memo that names a stored assumption set	the memo compares a run's assumptions against the shipped defaults and says how many series differ; it cannot yet cite a saved set by name, so no watermark means "not the shipped defaults", not "a reviewed set"
A side-by-side comparison of two runs, or of two memos	not built; the workbench lists what you have saved and reopens one at a time
Drawdown funds with capital calls and pacing	v2; the engine assumes evergreen or notice-period vehicles
Regime-conditioned scenarios	v2
Provider glidepath import from CSV	v2
Re-enrolment events	v2
The switching-fairness transfer	v2; the mechanism is described in the technical note
Uploading a return history to estimate θ	the engine function exists; no API route accepts the upload
A gross break-even premium	planned; the net figure is reported today

Where do the assumptions come from? The technical note lists a source or basis for every default, and the introduction to this series sets out the reasoning behind the three commitments the model rests on. An in-interface source note, and a landing page linking a fuller research brief, are planned along with the interface; today the landing page links directly to these three papers.

AN INTRODUCTION TO THE METHODOLOGY · TECHNICAL NOTE · FREQUENTLY ASKED QUESTIONS