DISCIPLINE FUNDS

RESEARCH NOTE | ASSET LIABILITY MATCHING

The Price of (Un)Certainty

All investing is a temporal conundrum, a trade off between near term certainty and long run uncertainty. This research paper discusses the price of (un)certainty within asset liability matching strategies across different time horizons and the necessary trade offs an investor must consider in the context of a world with high equity valuations, 4% overnight interest rates and 3.2% real TIPS yields.

Discipline Funds · October 2026 · For financial professionals

KEY DATA POINTS

CASH
4.17%
One month T-bill yield. Very attractive today, a problem in the long run. Over 30 years cash lost purchasing power 40% of the time.

BONDS
4.91%
Five year Treasury yield on a modified duration of 4.4 years, so a year's income covers a one point rate rise. At 30 years it covers 0.4.

EQUITIES
39
US cyclically adjusted P/E, September 2026. Only the 1999 to 2000 bubble ran higher on Shiller's 150 year record. The price of equity uncertainty is near its highest ever.

TIPS
3.18%
Thirty year TIPS real yield, 2.6% to 3.2% across the curve. Attractive everywhere, more probabilistic further out. The 30 year deserves a serious look.

THE ARGUMENT IN THREE POINTS

- **Risk is not a property of an asset alone, it is a property of the asset's fragility across a specific time horizon.** Price risk tends to decay with time while inflation risk compounds with time. In this paper we price that risk and the temporal (un)certainty that corresponds to it. While T-bills buy certainty at the front end they give it up on the back end. TIPS buy it further out, but the certainty comes with a potentially significant opportunity cost in the long run. Equities offer significant long-run upside, but that upside comes with fatter tails along the way. This research explores them all in real-time and quantifies the trade-offs in optimizing temporal certainty of different portfolio structures.
- **In 1952 two seminal papers were written: *Portfolio Selection* by Harry Markowitz and *Review of the Principles of Life-Office Valuations* by Frank Redington.** Households adopted Markowitz, mean variance optimization and the efficient frontier while institutions adopted Redington's asset liability matching (ALM) framework. This research argues that Markowitz has failed too many retail investors and that Redington is better suited for financial planning processes.
- **There are three ways to match assets and liabilities.** You can lock it in with a contract, like a bond held to maturity. You can statistically match it with stocks and multi-asset blends. Or you can use a mix of both approaches. This research argues that households and institutions could be better suited using a mix of both approaches where certainty is cheap in the front end and more expensive on the long end based on the investment portfolio's risk capacity, not vague risk tolerance questionnaires matched with vague factor, style and mean variance optimized approaches.

All return figures in this paper are inflation adjusted and derived from a GDP weighted index of 16 developed markets covering 1872 through 2025. Sources and methodology are set out in the appendix. **All results shown are hypothetical.** See Important disclosures at the end of this paper.

1 The problem asset liability matching runs into

Both sides of the equation go blurry at the same time

Asset liability matching is easy when the liability is close and known. A payment due in 18 months gets matched with an 18 month Treasury and you are done. No forecast, no probability, no tracking error against the thing that actually matters.

That precision fades fast. A liability 25 years out is a guess about spending, longevity, taxes, healthcare and family. The asset side is no better, since nobody knows what equities, bonds or inflation will do over 25 years. What was arithmetic at 18 months becomes a bet at 25 years, and the advisor has to decide what kind of bet to make.

The textbook answer is TIPS. Held to maturity, a TIPS removes both unknowns at once. It pays a stated real return, so the purchasing power of the payment is fixed no matter what inflation does. For a real liability that is the theoretically correct choice, and it is why institutional liability driven investing leans so heavily on inflation linked bonds.

The catch is what that certainty costs, and the answer changes with the horizon. On the historical record, certainty is cheap against a payment due in two years and expensive against one due in thirty. At today's prices the far end is a lot cheaper than history says, and section 2 shows by how much. Either way, no single asset is the right answer across that whole span.

That is not an argument against any asset. Everything available here is right somewhere on the curve and wrong somewhere else on it.

- **Cash and short government paper** give up a lot in the long run and sometimes in the short run too. Against a near dated payment that is worth it, because nothing else matches a known amount on a known date.
- **Nominal bonds and TIPS** are strong for near and intermediate certainty, where a bond held to maturity is an exact match and the cost of that match is small. They are poor choices for long dated certainty, for different reasons in each case.
- **Equities** are expensive to hold in the short run, where the range of outcomes is wide and the floor is deep. Over long periods that cost falls, and the wait is what buys the improvement.

None of that is a forecast. It describes how two risks behave over time, and what we want to do here is price that trade at each horizon, not pick a winner.

Institutions got immunization. Retail got the efficient frontier.

Two papers published in 1952 split the investment world in half, and almost nobody notices they came out the same year.

The first is Frank Redington's *Review of the Principles of Life-Office Valuations*, in the Journal of the Institute of Actuaries. Redington gave the world immunization: the idea that an institution with known future obligations should hold assets whose cash flows and interest rate sensitivity are arranged to meet those obligations, so that a move in rates changes the value of the assets and the liabilities together and the obligation stays funded. The liability came first. The portfolio was a consequence of it.

The second is Harry Markowitz's *Portfolio Selection*, in the Journal of Finance. Markowitz gave the world mean variance optimization, and from it the efficient frontier, the risk return trade off and the idea that a portfolio should be chosen by maximizing return for a given level of volatility. There is no liability anywhere in the construction. There is only the portfolio and its own statistical properties.

Insurers and pension funds took Redington. Liability driven investing, as it came to be called, is now ordinary practice for any institution with dated obligations, and the whole apparatus of discounting a liability schedule at a risk free curve, matching duration against it and reporting a funded ratio comes straight from that 1952 paper.

Retail investors took Markowitz, and not entirely by choice. Mean variance optimization is what the retail industry was built to deliver, because you can deliver it without knowing anything about the client. All it needs is a risk tolerance score. A questionnaire, a model portfolio, a glide path keyed to a birth year and a target date fund with a year on the label are all the same idea in different packaging, and none of them requires anyone to work out what the money is actually for. The frontier scales. A liability schedule has to be built one household at a time.

The interesting thing is that the ALM process is far more appropriate for financial planning than the frontier. After all, numerous studies have shown us that optimizing the asset side of the equation is a fool's errand. Active investors do not outperform passive investors, on average (see SPIVA), and John Bogle largely solved the asset optimization problem with the democratization of low cost index funds. So, if Bogle and indexing can be paired with financial planning, then the better question is why households were never offered the technique institutions use for this exact problem?

Household liabilities are more knowable than the industry assumes

The usual objection is that a pension fund has an actuary and a household does not. True, but that matters less than it sounds.

What a household lacks is the actuarial machinery: mortality tables across thousands of lives, a valuation signed off by a qualified actuary, a regulator requiring the numbers be reported. What it has instead is a set of liabilities that are, one by one, a lot more identifiable than a pension's aggregate obligation.

A car replacement is a few years out and the amount is a known range. A wedding has a date once it has a date, and families usually know the order of magnitude years ahead. A home purchase has a down payment the household itself decides on. Tuition arrives on a calendar the registrar publishes, in an amount the institution publishes, for a child whose age is not in doubt. Monthly living expenses are the single most observable number in any financial plan, because the client has been living on them. Retirement is the one genuinely blurry item, and even there the start date is a decision and the spending level is anchored to the expenses the household already has.

Compare that with what a pension actuary actually works with. They do not know when any individual member will die, retire, leave or claim. They know the statistical shape of those events across a population, and they are routinely wrong about it, which is why mortality assumptions get revised and funded ratios move for reasons that have nothing to do with markets. A household does not have a distribution over ten thousand lives. It has one life, and it can see its own circumstances directly.

So a household's liability schedule is less precise in the far years and more precise in the near ones. The far years are a planning estimate and should be treated as one. The near and intermediate years are often known within a few thousand dollars and a few months, which is right where matching is worth the most. That is enough to run the method, though not enough to pretend it is exact.

One important advantage that households have over institutions here is their ability to use a broader framework that embeds equities. While institutions are often constrained by regulations and a requirement for greater certainty, households have the ability to hold equities or blended instruments in an ALM structure because they are not constrained by regulators and investment policy mandates. This is a significant advantage in being able to implement a more diversified form of ALM that we believe has the potential to enhance the way ALM can be applied.

Two kinds of matching, and one kind of insurance

The most common misreading of a framework like this one is that it pits guaranteed assets against risky ones and picks a winner. It does not. What matters is **what kind of answer each one gives**. Corey Hoffstein draws the line between explicit and statistical liability driven investing, and we borrow his framework here, calling the first kind **contractual** matching and the second **statistical**. Buying a six month bill to fund a car purchase six months out is contractual. The match is mechanical and the outcome is certain. Matching an equity allocation to a twenty year retirement goal is statistical. The match is probabilistic and the outcome arrives as a distribution.

- **T-bills and TIPS held to maturity deliver certainty.** Contractual matching. A known amount on a known date, nominal in the case of bills and real in the case of TIPS. There is no distribution to discuss. The outcome is written into the security.
- **Blends and equities deliver a distribution.** Statistical matching. No guarantee at any horizon, but an upside that widens as the horizon lengthens and as the temporal allocation is refined. What you get is a range of outcomes whose middle sits well above the guarantee and whose bottom, at long horizons, has historically sat close to it.
- **Alternatives deliver insurance.** Structured well, an alternative allocation hedges exposures that bills, TIPS, bonds and equities cannot reach on their own, and it earns its place by paying off when the rest of the portfolio does not. That is a third job, separate from certainty and distribution, and for some balance sheets it is a valuable one. We do not try to measure it here. **Alternatives are outside the scope of this paper**, which deals in cash flow generating assets whose defined duration can be measured from an observable yield or earnings yield. When an expected real return rests on manager skill, an illiquid mark or a convex payoff instead of a stream of cash flows, the defined duration calculation later in this section has nothing to work with. That is a limit of the measurement, not of the exposure.

The first two are complements, not substitutes. A matching strategy uses both, and the only real question is how much of each belongs against a liability at a given distance. The third sits alongside them doing something neither does. Certainty is worth paying for when the payment is close and the plan cannot absorb a shortfall. A distribution is worth accepting when the horizon is long enough for the range to work in the investor's favor and the plan can sit through it. Every table in this paper tries to price that exchange rate at a particular horizon, not to retire one side of it.

In practice the first two often sit in the same plan and even the same horizon. A 3-7Y sleeve can be a TIPS ladder, a blend, or some of each. Nothing here argues for removing guaranteed assets from a portfolio. The argument is about which maturities they should cover.

Most of our work sits in the gap between the two. Contractual matching is easy and nobody disputes it. What has been missing is a way to make the statistical side rigorous instead of hopeful, which means putting a number on how long an equity allocation actually needs. The defined duration measure, set out later in this section, does that.

Bucketing is not asset liability matching

The obvious objection comes up early, so we will deal with it early. This paper splits a portfolio by time, which looks a lot like what advisors call bucketing, and the criticism of bucketing is well established. Time segmented pools have been around since Harold Evensky popularized the cash flow reserve in 1985, and Michael Kitces rightly took the generalized version apart in 2014, showing that the rules deciding which bucket gets spent were invented, not derived. We agree. Bucketing without a liability schedule behind it is a behavioral illusion. What was missing was never discipline. It was the liability schedule.

The difference is precision. Bucketing sorts money by time loosely and never quantifies whether each segment can do its job. **Asset liability matching derives every allocation from the liability schedule and measures it.** Cutting a portfolio into segments is not the method. Where the segments come from is.

Evensky's construction shows the difference clearly. A one to two year cash reserve attached to a total return portfolio has two segments, but only one of them came from a liability. The other ninety percent is still a single pool optimized against a benchmark and matched to nothing in particular. That is a total return strategy with a cash reserve bolted on, and the reserve is doing behavioral work, not matching work.

Kitces's objection comes from the same gap. If the segments were not derived from anything, nothing determines what gets spent from where, and the rules have to be made up. Made up rules dressed as structure are worse than a disciplined total return portfolio, which is what he found. The flaw was never that the portfolio had segments. It was that the segments had no origin.

A matching process derives every allocation from the liability schedule, including the long one. The boundaries are where the plan says the money is needed. The weights are the present value of each segment of the spending schedule, discounted at the risk free real curve. The overall stock and bond mix is whatever falls out of that math, not a target anyone picked. Institutions often refer to this methodology as "liability driven investing" precisely because the investments are driven by the liability needs, not the asset optimization.

Where the time horizons come from

The horizons we use throughout are the Defined Duration horizons: **0-3Y, 3-7Y, 7-15Y and 15Y+**. Advisors would call them buckets. We avoid that word on purpose, because a bucket is usually a starting assumption and these horizons are an output of the liability schedule. They are guideposts for a financial planning process, not a model portfolio and not a fixed rule.

The planning comes first. Working out the liability schedule produces the payment dates, and the horizons group those dates so a portfolio can actually be built and rebalanced against them. Inside any horizon the matching can get as fine as the plan justifies, down to individual dated payments where the account allows it. A client with a known tuition bill in year six and a property purchase in year eleven does not need two extra horizons. They hold two dated matches inside the horizons those years fall in.

The boundaries fall where they do because **how predictable a household's spending is fades along a fairly reliable curve**, and the horizons line up with the big common outlays at each stage of it in order to help organize and structure the liability matching process.

- **0-3Y expenses are very predictable.** Property taxes, insurance, a car, a wedding, next year's living costs. The amounts and the dates are mostly known and a plan that gets them wrong is a plan that was not built.
- **3-7Y is less so.** A home purchase, tuition beginning, a business transition, a planned relocation. The events are usually identifiable but the amounts and timing move.
- **7-15Y is less predictable again.** Later education, care for an aging parent, the shape of a retirement that has not started. The category is visible from here and the specifics are not.
- **15Y+ is almost entirely unpredictable.** Longevity, healthcare, tax law, family circumstance, what a comfortable life costs three decades out. Almost nothing about it can be stated to the dollar.

Two things move together along that curve, and that is why the framework works. As the liability gets less predictable, the sequence risk attached to it grows, because there is more path between now and the payment. And as the horizon lengthens, the allocations funding it can carry more dispersion, because there is more time for outcomes to resolve and more upside for accepting them. Precision is worth the most where you can actually get it, and accepting a distribution pays the most where the liability was never precise anyway.

The four horizons give the process a common vocabulary and a place to start. They sit where those trade offs change character. What goes inside them is measured, not asserted, and the next subsection shows how.

Defined duration: how long an allocation needs

A bond's duration is the horizon at which price risk and reinvestment risk offset, so the investor earns the stated yield no matter what rates do in between. It marks the point where the bond becomes reliable for its job. Equities have no equivalent, which is why institutions matching long liabilities so often will not use them at all. An equity market behaves like a perpetual claim, not a dated one, and "stocks for the long run" is a hope, not a horizon.

So we measure it. **Defined Duration** is the compound time horizon needed to get back to breakeven in real terms after a maximum drawdown.

$$\textit{Defined Duration} = \frac{\ln\left(1 / (1 - D)\right)}{\ln\left(1 + r\right)}$$

- **D** is the maximum anticipated real drawdown, for example 50%.
- **r** is the expected real return implied by current valuations, for example the earnings yield of a cyclically adjusted price to earnings ratio of 30, which is 3.33%.

A portfolio that falls 50% needs a 100% real gain to get back to where it started. Compounding at 3.33% real, that takes **21.1 years**. That is sequence risk measured instead of described, and it is why the 15Y+ horizon sits where it does.

The same logic sizes every horizon. A blend of stocks and short government paper falls less in a bad market, so it recovers faster, and its defined duration is shorter. That is how each sleeve gets matched to its horizon: it holds the mix that can recover from a bad drawdown before the earliest payment it funds comes due. Section 5 shows how the number moves with today's prices and runs it for every sleeve.

A spectrum, not a switch

Here is the claim this paper is built to defend, before any of the evidence.

Contractual and statistical matching are not two camps. They are the two ends of one spectrum, and every horizon in this framework is a particular blend of them. The 0-3Y horizon is purely contractual. The 15Y+ horizon is purely statistical. The two in between are mixtures, and the mixture shifts toward the statistical end as the horizon lengthens.

Read the sleeves that way and the construction stops looking arbitrary.

- **A 0% equity allocation is entirely contractual.** The 0-3Y horizon holds no stocks at all. Bills, notes and short TIPS held to maturity deliver a known amount on a known date, which is exactly what a liability inside three years calls for, and inflation has almost no time to compound against it. This is the one place on the spectrum where nothing is left to chance.

- **A 15/85 blend is mostly contractual with a statistical element.** Eighty five cents of every dollar sits in short government paper whose outcome is close to written down. The remaining fifteen cents is the first statistical dose in the ladder, small enough that it cannot push the sleeve's recovery horizon past the liability it funds and large enough to collect some part of the equity premium over every period the sleeve is held.
- **A 35/65 blend is fairly balanced.** Neither kind of matching dominates, which fits, because a liability eight to fifteen years out is neither precisely known nor completely unknowable.
- **A 100% equity allocation is entirely statistical.** There is no contract anywhere in it. What justifies it is that the liability it funds was never contractually specifiable in the first place.

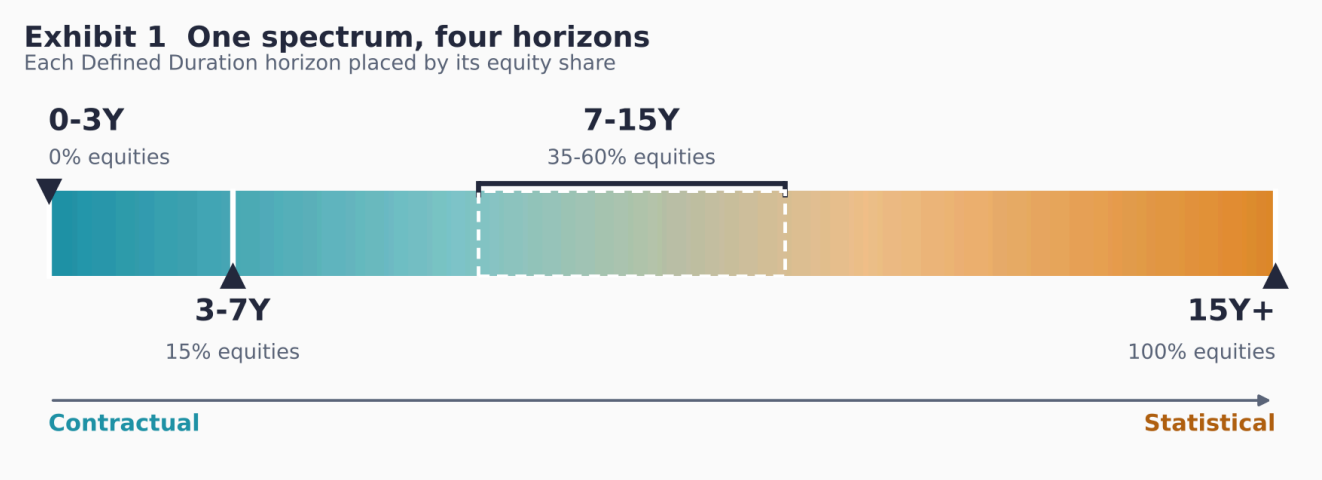


Exhibit 1. Each Defined Duration horizon placed by the equity share of the allocation that fills it. The 7-15Y range runs from a 35/65 blend to a 60/40. The further out the horizon, the further the fill sits toward the statistical end. Source: Discipline Funds.

Here is the part most readers find surprising. **Equities can play a bigger and bigger role further out the spectrum without the investor giving up much certainty in exchange.** Equities do not become safe. What shrinks is the amount you give up, because an allocation's floor rises with the time it gets and the liability it is matched against was less precise to begin with. Adding equity to a two year sleeve buys dispersion at full price. Adding the same equity to a twenty year sleeve buys it at a discount, and section 6 prices it at every horizon in between.

Where a household belongs on that spectrum comes down to risk capacity, not temperament. The statistical investor can take more risk because they have more capacity, usually because the balance sheet is large relative to what it has to fund, and surplus dollars have no due date to match. The contractual investor needs more certainty because they have less capacity, and a shortfall is something the plan cannot recover from. Section 7 puts that on a single number, the funded ratio, and shows the allocation moving on its own as that number moves. The investor's appetite for volatility never enters the calculation.

2 What long dated certainty actually costs

The give up is real, and so is the protection

Start with where we are today. In mid October 2026 the 30 year TIPS real yield was 3.18%. The 30 year nominal Treasury yielded 5.43%, implying breakeven inflation of about 2.18%. The federal funds target range was 3.75% to 4.00%, so cash was paying close to 4% nominal. Headline CPI was running at 3.7% year over year in September, with core at 2.6%. And equity valuations were high, though very unevenly. At the end of September, the most recent date we

could read every market on one consistent basis, the US cyclically adjusted price to earnings ratio stood at 39 against a long run median near 16, while developed markets outside the US were at about 22. Cap weighted, the world comes to roughly 32. That spread matters for a paper built on a global index, and we come back to it in section 5.

That mix is what makes this interesting. Everything on the table is attractive and expensive at the same time, just in different ways. Cash pays a real yield for the first time in a generation and still loses purchasing power over long horizons. A guaranteed 3.18% real for thirty years is a serious return, and on the historical record it is also the most expensive certainty on the curve. Equities carry the widest distribution, and US equities start from the highest valuation in almost thirty years. That lowers the expected return and lengthens the defined duration, but it does not change the shape of the horizon effect we measure. Later in this section we show what it does to the level. Nothing in this enviroment is cheap on every dimension, so the allocation has to be made against liability dates and not against a market view.

A guarantee caps the outcome as well as floors it. A 3.18% real yield compounded for 30 years turns $10,000 of purchasing power into $25,578, and that is the whole outcome. There is no upside case and no downside case. That figure is actually a little generous to the TIPS. It assumes coupons get reinvested at the same real rate and it ignores the annual tax on inflation accretion in a taxable account.

One condition runs through everything that follows. A TIPS only delivers its contractual real return if you hold the individual bond to maturity. That is where the certainty comes from. A TIPS fund holds the same bonds but marks them to market every day and never matures, so you get the inflation indexing without the cash flow certainty. Throughout this paper, TIPS means individual bonds held to their date. That is a constraint on the measurement, not on implementation. A short maturity TIPS fund gives up the pull to par and keeps almost everything else, and at short maturities that costs very little.

> **WHAT A TIPS DOES NOT GUARANTEE**
>
> A TIPS guarantees a real return measured by the government's own price index, and over long horizons that distinction matters more than most investors assume. The index is a national average, so a household whose costs run hot, in healthcare, housing or a particular city, can fall behind while the bond does exactly what it promised. The government also measures the inflation it owes on, and the methodology can change. Argentina is the extreme case. Starting in 2007 its statistics agency understated inflation for years, and holders of its inflation linked bonds were paid on the official number, not the real one. In 2013 the IMF formally censured Argentina over its inflation data, the first such censure of any country. Payments also track the index with a lag of about three months, which costs little at 3% inflation and a lot at 30%. And in a taxable account the inflation accrual is taxed every year, so in a high inflation the after tax real return can turn negative. None of this undoes the case for TIPS against a dated liability. It means the guarantee is against the CPI, not against the investor's actual cost of living, and that gap can widen with time and with the inflation rate.

Now compare that with equities. Across the 125 overlapping 30 year windows in our GDP weighted index of sixteen developed markets from 1872 to 2025, the worst real annualized return was 1.55%, in the window starting in 1891. That turned $10,000 into $15,875, about 38% less than the TIPS locks in today. The median window returned 5.28% real, or $46,841.

So the guarantee sits inside the distribution, above the floor and well short of the middle. The rest of this paper builds on that, and it is a more modest claim than the usual case for equities.

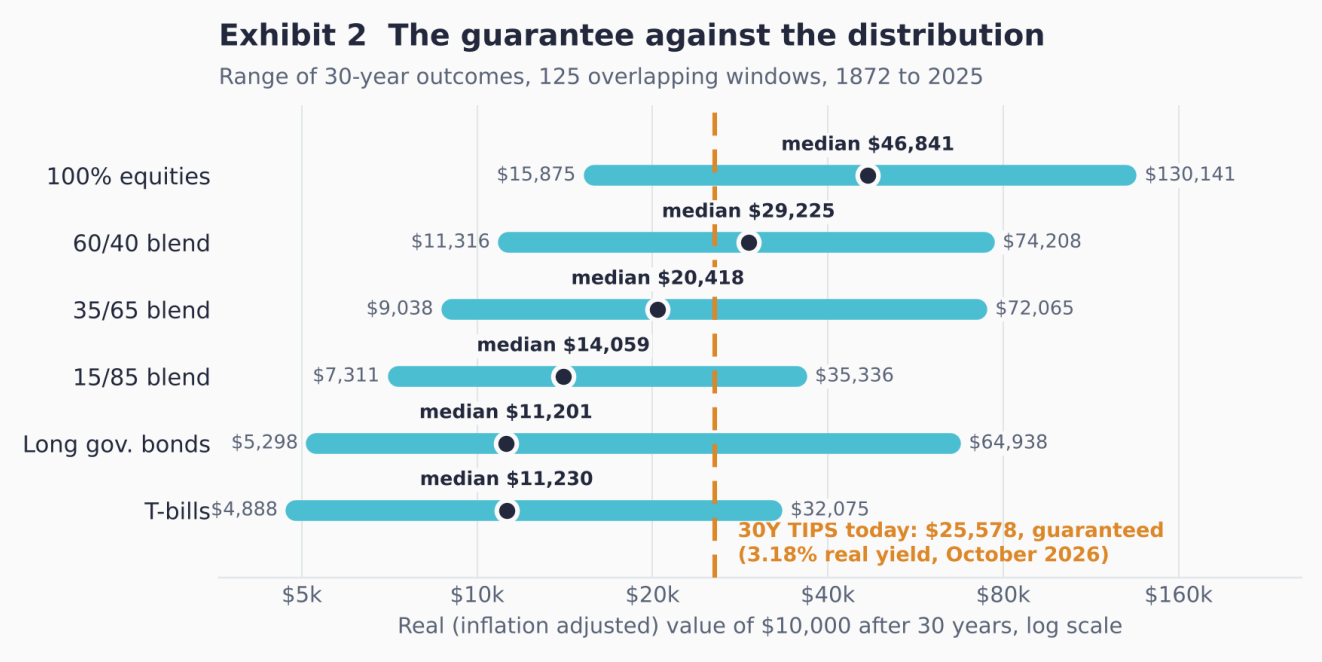


Exhibit 2. Real terminal value of $10,000 over 30 years, all 125 overlapping windows 1872 to 2025, GDP weighted across 16 developed markets. Bars span the worst and best historical outcome, the dot marks the median. Sources: Jorda Schularick Taylor Macrohistory Database, US Treasury daily real yield curve.

> **THE TRADE AT THIRTY YEARS**
>
> Buying the 30 year TIPS today trades a range of $15,875 to $130,141, with a median of $46,841, for a certain $25,578. The protection is real, since the bottom of that range sits below the guarantee. On the historical record it costs about 2.1 points a year of median return to rule out an outcome that showed up in about 6% of thirty year windows. At today's valuations the cost is closer to 0.4 to 1.1 points.

Whether that is worth paying is a question about the investor, not the security, and section 7 takes it up.

Certainty is priced along a curve

Thirty years is the extreme case, and quoting it alone would be misleading. Certainty has a price at every maturity, and that price changes a lot along the curve.

To set that price you have to name the risky alternative at each horizon, because it changes. Each horizon gets the mix whose defined duration does not exceed the nearest liability in it, meaning the mix that recovers from a bad drawdown before the earliest payment in that horizon comes due. Section 1 defines the measurement and Table 6 in section 5 runs it. The intuition is simple: more equity takes longer to recover, so more equity belongs against a later liability. That is why the front horizon holds bills and short nominal paper, the 3-7Y horizon a 15% equity blend, the 7-15Y horizon a 35% blend, and only the last horizon holds equities outright.

Table 1 extends the Exhibit 2 comparison to all four horizons.

Table 1. What certainty costs at each horizon

HORIZON	MATCHING ALLOCATION	CERTAIN OPTION	WORST	MEDIAN	BEST	ANNUAL PRICE
0-3Y, 3 yr	**T-bills and short nominals**	$10,766	$6,514	$10,491	$13,460	-0.88 pp
3-7Y, 5 yr	**15/85 equity and bills**	$11,358	$5,132	$11,359	$14,165	0.00 pp
7-15Y, 10 yr	**35/65 equity and bonds**	$13,155	$4,321	$15,205	$25,187	1.50 pp
15Y+, 30 yr	**100% equity**	**$25,578**	**$15,875**	**$46,841**	**$130,141**	**2.10 pp**

Real value of $10,000 held over each horizon's representative holding period. The certain option is a Treasury inflation protected security of that maturity held to maturity, at real yields from the US Treasury daily par real yield curve for mid October 2026: 2.49% at three years, 2.58% at five, 2.78% at ten and 3.18% at thirty. The three year figure is extrapolated linearly from the five and seven year points, since the published real curve begins at five years. Worst, median and best are the full range of outcomes across every overlapping window of that length, 1872 to 2025. Annual price is the matching allocation's median annualized real return minus the certain real yield. The 35/65 blend pairs equities with the long government bond, which is the conservative construction. Built with a short fixed income leg instead its median is 3.36% and the annual price is 0.58 points rather than 1.50. Sources: US Treasury, Jorda Schularick Taylor.

The certain option is a number and the alternative is a range. At three years the range is narrow enough that the two are almost the same thing. At thirty years the alternative runs from $15,875 to $130,141 against a guarantee of $25,578, and comparing the two takes some care.

The worst case sits below the guarantee at every horizon, including the last one. History never rules out falling short anywhere on this curve. What changes is how far short, how often, and how much of the middle you give up to avoid it.

The price rises with the horizon, and at the front it is currently negative. A three year TIPS at 2.49% real beats the 1.61% T-bills have historically delivered over three years. So at today's yields an investor funding a near dated liability is not paying for certainty at all, they are being paid to take it. At five years the two are level, with the 15/85 blend's median of 2.58% matching the five year real yield. Then the price climbs to 1.50 points at ten years and 2.10 at thirty.

The price does not rise forever, though. Measured against what each horizon actually holds, it peaks at about 2.4 points between fifteen and twenty years and eases back to 2.10 at thirty, because today's real yield curve keeps rising at the long end while the equity median does not. The ladder does not need a steady slope. It needs the level: on the historical record certainty is free or close to it at the front, expensive through the middle and long end, and the equity floor only becomes dependable at the far end. The next subsection shows how much of that expense today's prices take away.

Three forces shape that curve.

- **Nominal certainty degrades as the horizon lengthens.** A T-bill fixes the dollar and says nothing about what the dollar buys. Over three years that gap is bounded, and the worst it ever did was cost 35% of purchasing power. Over thirty years it is not. A 3% inflation gap costs 3% over a year and 59% over thirty.
- **Real certainty gets more expensive as the horizon lengthens**, because what you give up to buy it is the compounding equity premium. Today's curve blunts that past twenty years, and today's valuations blunt it a lot more.
- **The equity floor rises as the horizon lengthens**, because mean reversion and diversification have time to work. That is Exhibit 3, with Exhibit 4 showing the same thing in loss frequency, and it held up in every robustness test we ran.

The horizon boundaries sit roughly where those forces change character, rounded to something you can actually build and rebalance. They are planning guideposts, not points the data pins to the year.

What today's prices do to the curve

Table 1 has a blind spot, and it is a big one. It sets today's yields against the whole historical distribution, and most of that history started from ordinary valuations. The median starting CAPE in the US record is around 16. The global composite this paper uses sits at about 27.8 today. Section 5 argues that a high starting price lowers the expected return, and if that is true it has to show up in the price of certainty too.

So we measured it two ways. First, on Shiller's US data back to 1881, we estimated how much a starting CAPE of 27.8 has historically lowered forward real returns at each horizon, and took that much off the global equity median. Second, we used this paper's own defined duration input, the 3.60% earnings yield, as the expected real return on equities at every horizon. In both cases the fixed income legs earn today's forward real yields instead of their historical medians. Table 1B shows the result.

Table 1B. The price of certainty on history and at today's prices

HORIZON	MATCHING ALLOCATION	ON HISTORY	VALUATION ADJUSTED	EARNINGS YIELD
3 yr	T-bills and short nominals	-0.88 pp	-0.56 pp	-0.56 pp
5 yr	15/85 equity and bills	0.00 pp	-0.69 pp	-0.37 pp
10 yr	35/65 equity and bonds	1.50 pp	-0.11 pp	0.29 pp
15 yr	100% equity	2.37 pp	-0.64 pp	0.69 pp
20 yr	100% equity	2.37 pp	0.11 pp	0.57 pp
30 yr	**100% equity**	**2.10 pp**	**1.12 pp**	**0.42 pp**

Annual price of certainty is the expected real return of the matching allocation minus the TIPS real yield of that maturity in mid October 2026. On history is the median across all overlapping windows 1872 to 2025, as in Table 1. Valuation adjusted takes the global equity median at each horizon and subtracts the amount by which a starting CAPE of 27.8 lowered US forward real returns relative to the US median, estimated by regressing January start forward real total returns on the earnings yield, Shiller data 1881 to 2023. Earnings yield uses 1/27.8, or 3.60%, as the expected real return on equities at every horizon. In both of those columns the bill leg earns today's 3 month bill yield less breakeven inflation at that horizon and the bond leg earns the TIPS real yield of matching maturity. The 15 year real and nominal yields, 2.91% and 5.29%, are midpoints of the 10 and 20 year points. Sources: US Treasury, Robert Shiller, Jorda Schularick Taylor, Discipline Funds calculations.

The answer moves a lot. On history, certainty costs 1.5 points a year at ten years and more than 2 points from fifteen years out. At today's prices it is roughly zero through twenty years on the valuation adjusted method and well under a point on the earnings yield method, and the thirty year price falls to somewhere between about 0.4 and 1.1 points. The adjusted column narrows toward history at the long end because valuations have historically reverted over long periods, so a high starting price matters a lot over ten years and much less over thirty.

Neither method is precise. The valuation effect is measured on US data and applied to a global index. Today's earnings yield sits at the edge of the historical range, so the US estimate is an extrapolation. And there is exactly one thirty year US window that started from a CAPE above 25, so the thirty year adjustment leans on the statistical fit more than on experience. The earnings yield column assumes valuations never revert, and the adjusted column assumes they revert about as fast as they have before. The truth is probably somewhere in between.

The shape survives. Certainty is cheap at the front under every method. What changes is the level further out, and the plain version is this: **today equities are expensive and TIPS are not.** The price of uncertainty in equities is high and the price of certainty is low. On history, the case for TIPS ladders weakens past fifteen years. At today's prices it holds out to

roughly twenty, and the thirty year TIPS becomes a much closer call than Table 1 makes it look. We think this is the most important qualification in the paper.

That does not mean an investor with high risk capacity should give up equities, and it is worth being clear about why. Cheap certainty is a reason to buy it against dated liabilities, where a shortfall would actually hurt. It is not a reason to buy it with surplus dollars that have no due date. Those are the dollars that can wait out a bad decade, and a TIPS caps their outcome as firmly as it floors it. Even at today's prices the earnings yield on global equities, 3.60%, sits above the thirty year real yield, and outside the US it is about 4.48%. Lower expected returns are a reason to lengthen the horizon equities are asked to cover and to diversify the starting price paid, which is what section 5 argues. They are not a reason to leave. So at today's prices a plan should lean harder on certainty for the liabilities it has to meet, and keep equities for the capacity it has to spare.

> **WHAT TODAY'S PRICES MEAN**
>
> Equities are expensive and TIPS are not. The price of uncertainty in equities is high and the price of certainty is low, so dated liabilities out to roughly twenty years should lean hard on certainty. That is not an argument for an investor with high risk capacity to give up equities. Surplus has no due date, a guarantee caps it as surely as it floors it, and the right response to an expensive market is a longer horizon and a cheaper mix of markets, not an exit.

Paying twice for safety you already had

Table 1 puts the price of certainty at the front of the curve at negative 0.88 points a year, but that assumes the near dated money is actually held in Treasuries, and a lot of plans miss this. Move the same dollars into the accounts most households use for cash and the math flips, far enough that the front of the ladder can carry the *highest* price of certainty anywhere on the curve.

Two things do it. The first is the spread the bank keeps, which is not compensation for any risk, since both options are default free to the household. The second is state income tax, and in a high tax state it is the bigger of the two. Treasury interest is exempt from state and local income tax. Interest on savings accounts, CDs and prime money market funds is not. So you cannot compare them at their stated yields. You have to put them on the same after tax footing first, which for a top bracket California filer means grossing the Treasury yield up at 13.3%.

Table 2. What the same safety costs depending on where it is held

WHERE THE 0-3Y MONEY SITS	STATED YIELD	STATE TAX TREATMENT	COMPARABLE YIELD, TOP BRACKET CALIFORNIA
Treasury bill, 3 month	4.28%	**Exempt**	4.94%
Treasury bill, 1 year	4.49%	**Exempt**	5.18%
Treasury note, 3 year	4.85%	**Exempt**	5.59%
Treasury money market fund	3.69%	**Substantially exempt**	4.26%
Online savings, large provider	3.50%	**Fully taxable**	3.50%
Online savings, largest provider	**3.20%**	**Fully taxable**	**3.20%**

Treasury yields from the US Treasury daily par yield curve for mid October 2026, the date used throughout this paper. Savings rates are the posted annual percentage yields at two of the largest online savings providers in the country as of October 23 and 24, 2026. The money market row is the most recent seven day SEC yield we could source directly from the provider, at the end of August 2026, against an expense ratio of 0.07%. Treasury money market yields have risen with bill yields since, so that row understates the comparison it is used to make. The final column grosses the state exempt yields up at California's top marginal rate of 13.3%, dividing each by one minus that rate, which is the 12.3% top bracket plus the 1% mental health services surcharge on income above $1,000,000, and assumes the state tax is not federally deductible at that income. Taxable yields are unchanged, since they are already the comparable figure. Sources: US Treasury, Ally Bank, Marcus by Goldman Sachs, Vanguard, California Franchise Tax Board.

Look at the bottom row against the three year note. To that household the note is worth 5.59%. The largest online savings account in the country pays 3.20%. That is a **2.39 point gap every year** on the one part of the portfolio whose whole job is safety, and it buys nothing. It is bigger than most of the allocation decisions further out the ladder, and you can close it for free by changing where the money sits.

CDs add a third cost, which is the exit. One large provider charges 60 days of interest on terms of three to twenty four months and 90 days on terms of twenty five to thirty six. A Treasury of the same maturity has no penalty and can be sold at market if the plan changes. So a CD gives up the state tax exemption and the exit at the same time, usually for a rate below the Treasury it replaces.

With money market funds, what matters is what the fund holds. A Treasury fund passes most of its income through as state exempt and earns close to the bills it owns. A prime fund does not, and California, New York and Connecticut apply an extra asset test under which a fund that fails passes through no exemption at all. Same three words on the statement, very different after tax result.

This does not make deposits useless. Deposit insurance covers a real credit exposure, and a household that remembers a bank failure may reasonably want to pay for it. A savings account settles same day and needs no brokerage. The tax argument shrinks toward nothing in a state with no income tax and disappears inside an IRA or workplace plan. So the claim is narrow. **In a taxable account in a high tax state, where the 0-3Y money is held matters more than most of the decisions made further out the ladder**, and it is the one place in the plan where a better answer costs nothing in risk, liquidity or complexity.

How dependable is the alternative?

Every price in Table 1 compares a contract with a historical distribution, so the whole column depends on how much you trust that distribution. Probably less than most people do.

Edward McQuarrie's work is why. Extending US records back to 1792 and adding two centuries across twenty markets, he finds the equity premium much less dependable than the post 1926 US sample suggests. Across the 19th century the US equity premium averaged negative 0.29%, and a dollar in bonds grew to $594 in real terms against $322 in stocks. Same market, a century earlier, no premium at all. That finding is why this paper uses a sixteen market global record from 1872 instead of a US one. The appendix covers the construction.

The premium is obviously not guaranteed. The useful question is how much of it each part of the ladder needs. The global record puts the equity premium over bills at 4.65 points a year, measured geometrically, so we shrank it and reran the medians.

- **The 0-3Y sleeve needs none of it.** It holds no equity, its matching is contractual, and it would work identically if the premium were zero.
- **The 3-7Y sleeve needs all of it, and only draws level.** At a 2.58% five year real yield the 15/85 blend's median of 2.58% just matches the guarantee, so it uses up the entire historical premium getting there and any shrinkage puts it below the contract. At these yields the right answer for a dated liability in this horizon is a ladder.
- **The 7-15Y sleeve now needs a little of it.** A 35/65 blend's median cleared the ten year TIPS with no equity premium at all while that yield was below 2.65%. At 2.78% it needs about 8% of the historical premium. One caveat: that swaps a claim about equities for a claim about bonds earning more than today's real yield, and neither is contractual.
- **The 15Y+ sleeve needs about half.** Pure equities keep their median above the thirty year TIPS down to a premium of about 2.6 points a year, or 55% of the record. Halve the historical premium and the guarantee edges ahead. That makes it the sleeve most exposed to this question, which makes sense given it has the longest horizon and the highest price of certainty.

So the ladder does depend on the premium past the front horizon. The point is that the doses it needs are modest relative to the record and sized on purpose, not maximized.

Is the record thin, or is the mechanism sound?

There is a sharper version of this objection that does not depend on the historical record at all, and Pastor and Stambaugh make it best. They do not say the measured long horizon variance is wrong. They say an investor who does not know the true parameters faces more uncertainty than any sample can show, because uncertainty about today's expected return, and about how it will move, builds with the horizon. In their working paper's estimate the annualized thirty year variance of equities is nearly one and a half times the one year variance, the opposite of what Exhibits 3 and 4 seem to show.

That one lands on this paper and more history will not answer it. Every figure here is a realized, in sample frequency. None of it carries parameter uncertainty, and adding years to the record cannot fix a problem about not knowing the parameters behind the record.

What we can say is that an equity return is not a random draw from an unknown urn. Buffett made the case in 1977 and nobody has improved on it. A company funds itself at or near the risk free rate, puts the capital to work and tries to earn a spread on it. Earning that spread is the point of the business. He found returns on equity clustering near 12% across three decades of Dow data and two of Fortune 500 data and called it an equity coupon, a stock behaving like a bond whose coupon the holder cannot renegotiate. So the expected return on equities is anchored to something real.

But that only goes so far. The mechanism tells you roughly **where the center of the distribution sits**. It tells you almost nothing about **how wide the distribution is over any particular thirty years**, and the width is what this paper is about.

Buffett's own work makes the point. His 1977 essay concluded that the equity coupon does not rise with inflation, so shareholders end up with more dollars and no more purchasing power. In 1999 he pointed out that the Dow closed 1964 at 874.12 and closed 1981 at 875.00 while US GDP rose 370%, because long government rates went from just over 4% to more than 15%. The mechanism worked as described for a generation and shareholders got nothing. McQuarrie's nineteenth century is the same story over a longer stretch.

So the structural argument narrows where the center can be and leaves the width alone. That actually argues for this framework. If the center were all that mattered, one allocation would serve every horizon and none of this would be necessary. It is the width, and the fact that a dated liability is exposed to it, that makes matching worth doing.

The two structural results under the ladder are different, and neither leans on the premium the way the sleeves do. Risk being horizon dependent comes from how price risk and inflation risk behave over time. Blends having higher floors than their parts is a diversification result. Section 4 puts a number on the first: the falling shortfall frequency holds as long as the premium is above roughly two points a year, against 4.65 in this record, and the rising floor needs no premium at all. Both would survive a much smaller premium, which is why the argument is about where each allocation belongs and not about how much equities will return.

McQuarrie's own conclusion runs close to ours. Without certainty that stocks beat bonds, and with near certainty that a balanced portfolio is less volatile and draws down less, he argues a balanced portfolio works for any investor, not just a timid one. Table 7 and Exhibit 3 show the same thing. Where we differ is how far out the curve a guaranteed real return stays worth its price, and section 6 gets into that.

3 Nominal bonds at the wrong horizon

A promise to pay dollars is not a promise to pay for anything

Nominal government bonds are the best match there is for a dated payment through the middle of the curve. Held to maturity they pull to par, they cost almost nothing against a near liability, and the front of this framework is built from them. Nothing below disputes that.

The trouble starts when the same logic gets stretched to thirty years. Advisors who balk at the TIPS give up often reach for long nominal Treasuries instead, figuring 5.43% nominal beats 3.18% real if inflation behaves. At that maturity it swaps one problem for a worse one.

A long nominal bond guarantees a nominal cash flow and nothing about what that cash flow buys. In the global record, government bonds lost purchasing power in 44% of all rolling 30 year windows, and cash in 40%. The worst bond window, starting in 1897, compounded at negative 2.10% real and left $5,298 of purchasing power for every $10,000 put in. Over 30 years that is a slow loss dressed up as prudence. Those are sobering odds for an asset most investors think of as the safe half of the portfolio.

McQuarrie's longer record makes the point even more clearly. Across the whole 20th century, a dollar in US nominal bonds grew to six dollars in real terms while a dollar in stocks grew to $837. That is what a century of inflation does to a nominal promise.

His 19th century result is the counterweight we take most seriously. In that century bonds beat stocks, $594 against $322, and inflation averaged negative 0.27% a year. Nominal bonds are superb under deflation and ruinous under inflation, so they are really a bet on the price level. TIPS remove that bet. Blends dilute it. Long nominal bonds just take it.

The most recent example is the sharpest. In 2022 the GDP weighted global government bond aggregate lost about 21% in real terms, against about 18% for global equities. That is the second worst year for government bonds in the whole 154 year record, a hair behind 1920. What people held for safety lost more purchasing power that year than what they held for growth. Cash was not spared either, losing 5.4% in real terms in 2022 after 3.1% in 2021.

If the liability is real, meaning it grows with the cost of living, then nominal certainty is the wrong kind of certainty. It looks safe on a statement and fails the only test that counts.

Bond index funds compound the problem

In practice the nominal allocation is almost never a bond. It is a bond fund, usually one tracking the Bloomberg US Aggregate, and that makes the matching problem worse in two ways.

First, the index has no fixed duration. It has whatever duration the outstanding stock of investment grade debt happens to imply that month. The Aggregate's duration has averaged about 5 years since 1978, but it has drifted well above and below that average, and it was 5.70 years at the end of September 2026. An investor who bought the index to match a five year liability can end up holding a meaningfully longer asset years later without placing a single trade. An allocation whose duration drifts on its own is not matching anything.

Second, a bond index is weighted by amount outstanding, so the more debt an issuer sells, the bigger its weight gets. Over the past two decades the government share of the Aggregate has grown sharply as Treasury issuance surged, while agency and mortgage exposure shrank. None of that is a judgment about value. It is just what the Treasury chose to issue, and a lot of that has been low yielding and long duration. The index buys it because it has no choice.

An equity index works differently. Cap weighting looks like a similar mechanical rule, but it sits on top of competition. A company that destroys capital shrinks, its weight falls, and eventually it drops out. The index keeps moving away from failure without anyone deciding to, and that quiet selection is a big part of why passive equity investing works as well as it does.

A bond aggregate has nothing like that. Issuance weighting is close to the opposite of a quality screen. Nothing prunes the worst positions, because to the index there is no such thing as a bad position, only a large one. You end up owning the debt of whoever borrowed the most, at whatever duration suited them.

Constant duration is a flaw against a date and a feature against a horizon

The obvious fix for drift is a constant duration fund. Hold a defined slice of the curve, roll the bonds before they mature, and keep the duration where you want it. Whether that helps depends on what you are matching.

Against a single dated liability it is a step backward, because you give up the pull to par. A bond held to maturity is an exact match because its price converges on a known value on a known date. A fund that never matures never converges on anything. That is why the 0-3Y horizon holds actual bills and notes to actual maturity instead of a short bond fund. Where you can get matching precision, take it.

Against a horizon, constant duration is the whole point. A horizon is a span that rolls forward with the calendar, and a maturing bond will not stay matched to it. A seven year note becomes a three year note within four years and drifts out of the horizon it was bought for, so you either re-match it or quietly accept that the sleeve got shorter. A holding built to keep its duration stays where it was set and never needs re-matching.

The multi asset version of that holding is worth spelling out, because it does most of the work in the middle of this ladder and it is easy to mistake for a plain balanced fund. It holds a fixed mix of equities and government paper, rebalanced back to weight, and sized so its defined duration sits at or inside the near edge of the horizon it fills. Because the weights get reset instead of left to run, that recovery horizon stays put for as long as the sleeve is held. A maturing bond gets shorter every year. This does not.

That gets you two things. The sleeve is built to hold its real value over the span it is matched to, and to keep doing so as the horizon rolls forward, which is what the floor comparisons later in this section measure. And because part of it is always in equities, it collects some of the equity premium over every rolling period it is held, without that exposure ever getting big enough to push its recovery horizon past the liability it funds. Section 2 showed how much premium each sleeve actually needs, and it is a modest fraction of the historical figure.

That is what we mean by temporal diversification. A ladder of individual bonds is a set of dated points with gaps between them. Constant duration sleeves fill the gaps, so the portfolio has exposure across the whole liability curve and not just at the dates that happen to have a bond attached.

So there are really three choices, not two. A maturing bond is right for a dated payment. A targeted constant duration allocation is right for a horizon. An issuance weighted index like the Aggregate is a perpetual fund whose duration is set by the Treasury's borrowing calendar, so it has the drawback of never maturing with none of the benefit of being aimed at anything.

The drift that is a bug, and the drift that is a feature

There is a second kind of duration drift in this framework, and it runs the other way.

The Aggregate's drift is a bug because nobody chose it. The index lengthened over the past two decades because the Treasury issued long, and every holder was carried along without placing a trade. The duration moved and the liability it was bought to match did not.

A good ALM strategy has natural drift, but on purpose. That is, the overall portfolio is set to an initial targeted set of liabilities and if the liabilities do not change then the portfolio will naturally drift to something more aggressive as short duration assets are withdrawn and long duration assets remain at work earning their temporal risk premia. A static blend gets rebalanced back to target every year, which resets its duration to the same place forever. A liability matched ladder does not. The front sleeves get spent as their payments come due, and nothing tops them back up from the long end until a distant liability has become a near one. So the portfolio's overall duration rises between refills, because the short assets get used up and the long ones are left alone.

That is a rising equity glide path, built into the structure instead of imposed by rule. Pfau and Kitces argued in 2014 that a rising equity path through retirement often serves investors better than the declining one convention prescribes, and a liability matched ladder produces it without anyone deciding to. The allocation drifts up because the near term liabilities have been paid off, which is when the plan can afford it. Here duration drift is a feature of the funding mechanism, not an accident of somebody else's issuance calendar.

It also changes what a comparison against a static blend is measuring. A static 60/40 earns a rebalancing benefit the ladder gives up, and that benefit is real. What the ladder gets instead is the drift. It does not sell its longest assets to top up its shortest on a calendar, only when the liability they were bought for gets close. Until then they compound untouched. Section 6 runs both portfolios end to end and separates the two effects, and the answer is not entirely comfortable for us: with nominal reserves, the drift accounts for essentially all of the ladder's advantage.

The multi asset version dominates the bond fund version

If a horizon is going to be filled with a perpetual fund anyway, the real question is what goes in it. A constant duration bond fund holds nominal Treasuries and nothing else. A constant duration multi asset allocation holds an equity slice alongside them, sized so the blend's defined duration lands on the horizon it is meant to match. Table 3 runs both against the record.

Table 3. Constant duration blends against the bond funds they replace

HORIZON	ALLOCATION	WORST OUTCOME	MEDIAN OUTCOME	WINDOWS WITH A REAL LOSS
5 yr	**Intermediate government fund**	$4,602	$10,939	28.0%
5 yr	**15/85 constant duration blend**	$5,132	$11,359	22.0%
7 yr	**Intermediate government fund**	$4,446	$11,373	27.7%
7 yr	**15/85 constant duration blend**	$5,078	$11,813	23.0%
10 yr	**10 year government fund**	$3,883	$12,294	32.4%
10 yr	**35/65 constant duration blend**	$4,321	$15,205	20.0%
15 yr	**10 year government fund**	$4,123	$13,402	32.1%
15 yr	**35/65 constant duration blend**	$4,821	$17,392	22.9%

Real value of $10,000 held for the stated period, across all overlapping windows 1872 to 2025, GDP weighted across 16 developed markets. The intermediate government fund is proxied by a 50/50 mix of bills and the long government bond series and serves both the five and seven year rows. The 10 year government fund is the long bond series and serves both the ten and fifteen year rows. The 15/85 blend pairs equities with bills and the 35/65 blend pairs them with the long bond series, which is the conservative construction. Section 5 shows the short leg version has better floors still. Sources: Jorda Schularick Taylor.

The blend has the higher floor and the higher median at every horizon tested, and it cuts how often you lose purchasing power by 21% at five years and 29% at fifteen. You are not trading risk for return here. The multi asset version is better on the downside too, for the reason section 5 develops: once the horizon is long enough for mean reversion to work, a modest equity sleeve raises a portfolio's floor instead of lowering it.

4 Risk is a function of the horizon, not the asset

Two risks with opposite time signatures

Every matching decision runs into two risks that move in opposite directions as the horizon gets longer.

- **Price risk** is the chance the asset is worth less than you need on the date you need it. It is severe over one year and decays as the horizon extends, because mean reversion and cash flow growth have time to work.
- **Inflation risk** is the chance the asset's return fails to keep up with the cost of the liability. It is trivial over one year and compounds relentlessly. A 3% inflation gap costs 3% over a year and 59% over 30.

Cash and short bonds carry almost no price risk and all the inflation risk. Equities are the reverse. Since the two risks trade off as the horizon moves, there has to be a point where the ranking flips. In the global record, pure equities have held a higher worst case real floor than either cash or the long government bond from year nine on. The 60/40 blend starts from a much lower floor and clears both from year twelve.

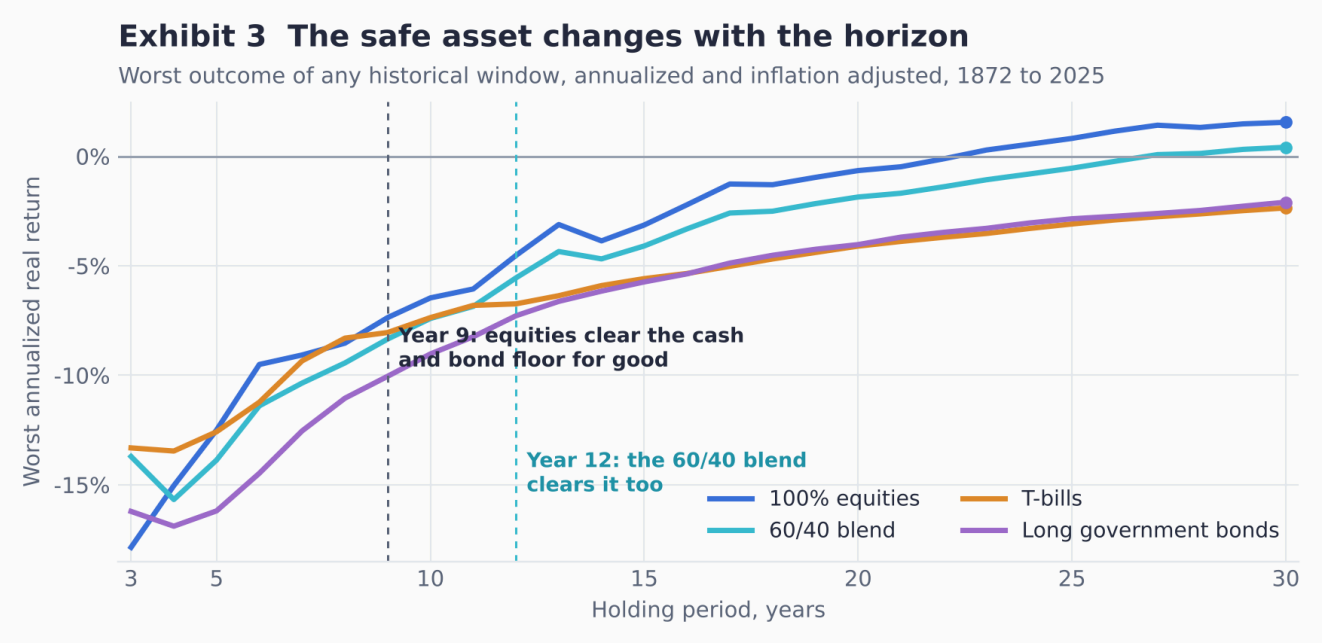


Exhibit 3. Worst annualized real return of any historical holding period of a given length, 1872 to 2025. Pure equities sit above both cash and the long government bond from year nine onward, having first drawn level at year five and briefly crossed at six and seven before slipping back at year eight. The 60/40 blend clears both from year twelve. Equities hold a higher floor than the 60/40 blend from year four onward. Sources: Jorda Schularick Taylor.

At five years the worst thing equities ever did was compound at negative 12.50% real, against negative 12.58% for T-bills. That is a tie, not a win for cash. By ten years the order has separated, with equities at negative 6.47% against negative 7.36% for bills and negative 9.03% for the long government bond. At thirty years equities never lost purchasing power in any window, while bills lost it in 40% of them and long bonds in 44%. The five year result is the surprise here. Most people assume cash is clearly safer than equities over five years. In this record it was not, because the episodes that wrecked equities also wrecked the currency they were priced in.

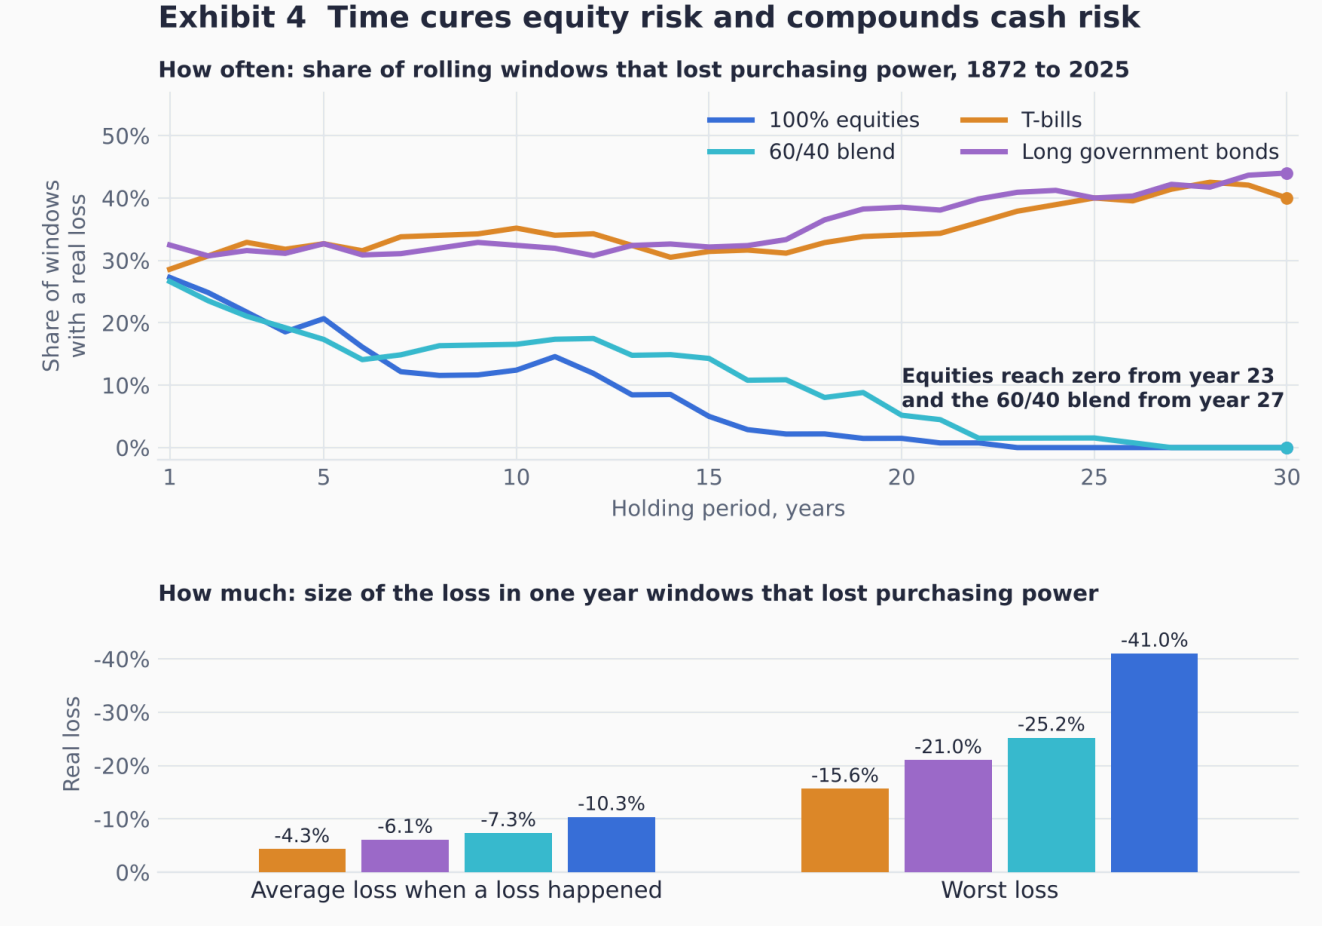


Exhibit 4. The top panel shows the share of rolling windows that failed to preserve purchasing power. Equity shortfall risk declines, unevenly, from 27.3% at one year and reaches zero from twenty three years out, with the 60/40 blend following at twenty seven. Cash and bond shortfall risk never falls below roughly 29% and 31% at any horizon, and both are higher at thirty years than at one. The bottom panel shows the average and worst real loss in the one year windows that lost purchasing power. The 60/40 blend pairs equities with the long government bond. Sources: Jorda Schularick Taylor.

How often is only half the story. Over one year cash lost purchasing power in 28.6% of windows and equities in 27.3%, so on a count they look about the same. They are nowhere near the same in size. When cash lost over a year, the average loss was 4.3% and the worst was 15.6%. When equities lost, the average was 10.3% and the worst was 41.0%. That is the volatility people actually feel, and it is why cash belongs in front of near term spending. The gap does not last, though. By five years the average losing window cost cash 14.8% and equities 17.2%, and by ten years cash was losing more, 22.7% against 19.0%. Over long periods inflation makes cash losses both more common and bigger.

The 60/40 line raises an obvious question. If the blend holds less equity, why does it take longer than equities to stop losing purchasing power, twenty seven years against twenty three? Because it carries two kinds of drawdown instead of one. It still takes most of every equity bear market, and its bond leg takes the inflation shocks that equities eventually outrun. Every 60/40 window that lost money over 23 to 26 years ended in 1920 or 1921, after the First World War inflation. Equities were up in every one of those windows. The long government bond lost between 42% and 53% in real terms. Swap the long bond for T-bills and the blend reaches zero at twenty five years instead of twenty seven, so duration accounts for about half of the extra time and holding fixed income through an inflation accounts for the rest.

None of that makes the blend riskier than equities where most investors actually feel risk. Over one year it lost purchasing power about as often as equities, 26.6% of windows against 27.3%, but its average loss was 7.3% against 10.3% and its worst was 25.2% against 41.0%. Its average loss stays below the equity figure through five years and its

worst loss through three. Past about seven years the order flips and equities lose purchasing power less often than the blend. That is the trade a 60/40 makes. It cushions the near term and pays for it with a longer tail, which is why this framework uses it in the 7-15Y horizon rather than as a whole portfolio.

> **THE REFRAMING**
>
> There is no such thing as a safe asset, only an asset that is safe for a given horizon. A T-bill is about as close as it gets to a match for a two year liability. Against a thirty year one it cleared today's guaranteed real yield in 9% of historical windows. The guarantee is always a guarantee. What changes with the horizon is what it costs to hold one.

The strongest objection

Everything above rests on the idea that an equity allocation gets more reliable as the horizon lengthens. That idea has been attacked in print for thirty years, and the attack is better than its popular summaries suggest.

Zvi Bodie made the case in the Financial Analysts Journal in 1995. If stocks really got safer with time, he argued, the cost of insuring against a shortfall relative to the risk free rate should fall as the horizon lengthens. He showed it does the opposite. The price of that insurance rises with the horizon, so whatever is happening to equities over long holding periods, they are not becoming safe.

The tempting reply is that his demonstration assumes returns are independent from year to year, and our record shows they are not. That reply is wrong, and we will admit we reached for it first. Bodie dealt with mean reversion in the original paper, and Feinstein later proved the result holds under any binomial price process, mean reverting ones included. The shortfall insurance gets more expensive with time as a matter of structure, and nothing in this paper or anywhere else overturns that.

The way out is that his number and ours are measured under different assumptions, and both are true at once.

An option price is a risk neutral quantity. To price options, every asset is assumed to drift at the risk free rate, which means the equity premium is defined away before the calculation starts. What the option prices is the shortfall against a strike that compounds at the risk free rate, and that strike climbs with every year you add. That is what drives Feinstein's result. Even when mean reversion limits how spread out the outcomes get, the strike keeps rising, so the option moves further into the money and has to cost more. The frequencies in Exhibits 3 and 4 are real world numbers. They are what actually happened to investors who were paid a premium for holding equities over those windows.

So the insurance gets more expensive while the realized shortfall frequency falls, and that is not a contradiction. The market charges more to remove a thirty year shortfall for the same reason it expects to pay more for bearing one.

It does mean the falling shortfall frequency in Exhibit 4 depends on the premium, which is worth keeping in mind when reading section 2. A realized shortfall frequency only falls toward zero with horizon while the premium is bigger than roughly half the variance, which on this record's equity volatility is about two points a year. The record supplies 4.65. So the result holds with well under half the historical premium and fails below roughly 43% of it. That is real headroom, but not unlimited. The rising floor in Exhibit 3 is sturdier, since narrowing dispersion does not need a premium at all.

What does not follow from Bodie's result is the allocation conclusion people usually draw from it. Samuelson and Merton showed in 1969 that with independent returns the horizon does not change the optimal equity weight at all, in either direction. Campbell and Viceira, coming from the other side, found that once mean reversion is allowed the optimal

equity weight is *materially higher* at long horizons than at one year, though not in a straight line. Their weight peaks near ten years and declines after it. Bodie's result does not settle that question either way, because a rising insurance price and an optimal allocation are different things.

Now the part that is less comfortable for us, and the reason this is a concession and not a rebuttal.

For a liability that has to be met to the dollar on a known date, the risk neutral price is arguably the right number to look at, because it is literally what defeasance costs. We already work that way. Section 7 discounts the client's spending schedule at the risk free real curve and refuses to discount it at expected returns, for Bodie's reason: a liability is worth what it costs to retire, not what you hope to earn against it. On the funded part of a plan, Bodie and we are making the same argument.

So the disagreement is narrower than the literature makes it sound. The question is not whether to buy certainty. It is how far out the curve to buy it, and what to do with the surplus and with liabilities so distant nobody can state them to the dollar. Bodie's own recommendation, that most savers should hold inflation linked bonds as the core of a retirement portfolio, looked poor through a decade of real yields at or below zero and looks a lot better at a 3.18% thirty year real yield. That is our point about the right fill depending on the curve, arriving from the other direction.

5 Defined duration at today's prices

How valuation and the fill change the horizon

Section 1 set out the measurement: the compound time an allocation needs to get back to breakeven in real terms after a maximum drawdown, given the real return today's prices imply. This section puts it to work. It shows how the answer moves with valuation, how it applies to blends, and how it lines up with what the historical record says on its own.

Defined duration is not one number, it moves with price

Because r comes from valuation, an equity allocation's defined duration changes as the market reprices. That makes it a live input to a plan, not a constant.

Table 4. Defined duration in years, by drawdown assumption and starting valuation

MAXIMUM REAL DRAWDOWN	CAPE 20	CAPE 25	CAPE 30	CAPE 35	CAPE 40
30%	7.3	9.1	10.9	12.7	14.4
40%	10.5	13.0	15.6	18.1	20.7
50%	14.2	17.7	21.1	24.6	28.1
60%	18.8	23.4	27.9	32.5	37.1

Expected real return taken as the inverse of the cyclically adjusted price to earnings ratio. Highlighted row is the 50% drawdown case used throughout this paper. These are arithmetic consequences of the two inputs, not historical outcomes.

Read across the highlighted row in Table 4. **An expensive equity market has a longer defined duration than a cheap one**, because the same drawdown takes longer to recover from when the starting expected return is lower. The market has not become more volatile. It has become a match for a later liability.

That matters for what goes inside the equity sleeve. This is where temporal thinking stops being about stocks versus bonds and starts being about which stocks. At the end of September 2026 the spread across developed equity markets was unusually wide.

Table 5. What each equity market's price implies about its horizon

EQUITY ALLOCATION	CAPE	IMPLIED REAL RETURN	DEFINED DURATION
Developed markets excluding the US	22.3	4.48%	15.8 yrs
GDP weighted developed, as used in this paper	27.8	3.60%	19.6 yrs
Cap weighted world	32.3	3.10%	22.7 yrs
United States	**39.4**	**2.54%**	**27.6 yrs**

Cyclically adjusted price to earnings ratios at September 30, 2026, from the Barclays historic CAPE series: United States 39.35 and developed Europe 22.34. The excluding the US figure is developed Europe, used as the proxy for developed markets outside the United States. Defined duration at a 50% maximum real drawdown. The GDP weighted figure is derived from those two readings at the GDP weights this paper's index uses, where the US carries 45.5% of the sixteen market total against about 71% of world market capitalization. That difference in weight is what makes the composite cheaper. The cap weighted figure holds the US share of capitalization at its July 2026 level, the last date we had a published world reading. Using developed Europe as the ex US proxy understates Japan, which carries 10.1% of the GDP total and has traded above Europe. Replacing Japan's contribution with a ratio of 28 moves the composite from 27.8 to 28.5 and the defined duration from 19.6 to 20.1 years, so the approximation does not carry any conclusion here. These readings are two weeks older than the October 15 yields used elsewhere in this paper, the closest consistent cross market date available.

The gap is nearly twelve years. A US equity allocation at these prices matches a liability about twenty eight years out. The same allocation outside the US matches one about sixteen years out. The drawdown assumption did not change and neither did the investor. Only the price paid did.

An advisor matching a fifteen or twenty year liability should read that as an argument for diversifying the equity sleeve, not as a market call. We are not forecasting that US equities will underperform. We are saying that at today's price the uncertainty in US equities, and in the concentrated technology exposure driving that valuation, is more than the matching need calls for. If a cheaper allocation gives a shorter defined duration for the same drawdown assumption, it is the better match for that liability, whatever either one goes on to return.

It is also the best practical argument we have for the GDP weighted construction in the appendix. Weighting by the size of the economy instead of the price the market puts on it gives a composite at 27.8 instead of about 32, which is a defined duration three years shorter for the same assumed drawdown. We are not claiming one will return more, only that one is a better match for a dated liability.

Why valuation belongs in a horizon framework and nowhere else

The formula treats starting valuation as an input to a matching decision. That is deliberate, because the two usual takes on valuation are both wrong for this purpose. One says valuation tells you when to get out. The other says it tells you nothing. We looked at this in a companion note, *Valuations Matter. Just Not the Way Wall Street Says.*, and the finding that matters here is narrow and useful.

On Shiller's data back to 1881, how much of future real returns a starting CAPE explains depends almost entirely on how long you hold. At one year the R-squared is about 0.04, which is nothing. At ten years it is 0.26. At twenty years it is 0.38. A relationship that is useless for timing is meaningful for duration, and that is how this section uses it.

The second finding matters more for matching. A high valuation does not show up as higher year to year volatility. In that record the richest CAPE quintile was actually slightly less volatile at one year than the cheapest. It shows up in the cumulative outcome. From the richest fifth of starting valuations, the next decade produced a negative real return 27% of the time, and never once from the cheapest fifth, and the median investor in that group sat through roughly a 49% real drawdown somewhere in the ten years. A high price means a high expectation, and that leaves a thinner cushion. That is where sequence risk actually bites, in the cumulative outcome over the horizon the money was needed for.

Both findings point the same way. Valuation does not tell an advisor whether to hold equities. It tells them which liability those equities match, and at a high starting price that is a later one. The right response to an expensive market is to lengthen the horizon the allocation covers and diversify the starting price paid, not to get out and wait for a cheaper entry that may be a decade away.

Defined duration for a blend, not just for equities

The formula takes its two inputs from whatever allocation you are measuring, and that turns it from a fact about equities into a tool for building a ladder. For a blend, both inputs change. The expected real return is the weighted mix of the equity earnings yield and the real yield on the fixed income leg. The drawdown is the equity drawdown scaled by its weight, assuming the fixed income leg holds its value while equities fall.

Table 6. Defined duration of each matching allocation

MATCHING ALLOCATION	HORIZON	DRAWDOWN D	EXPECTED REAL R	DEFINED DURATION	AT US VALUATIONS
15/85 equity and bills	**3-7Y**	7.5%	2.73%	2.9 yrs	3.1 yrs
35/65 equity and bills	**7-15Y**	17.5%	2.94%	6.6 yrs	7.6 yrs
60/40 equity and bills	**7-15Y**	30.0%	3.19%	11.4 yrs	14.1 yrs
100% equity	**15Y+**	**50.0%**	**3.60%**	**19.6 yrs**	**27.6 yrs**

Drawdown is a 50% maximum real equity drawdown scaled by the equity weight, on the assumption that the fixed income leg holds its value while equities fall. Expected real return blends the earnings yield of the GDP weighted developed ratio of 27.8 with the 2.58% real yield on short TIPS. The final column repeats the calculation using the US ratio of 39.4 in place of 27.8. These are arithmetic consequences of the stated inputs, not historical outcomes.

The ladder falls out of that column instead of being imposed on it. A 15/85 blend recovers in about three years, which is where the 3-7Y horizon starts. A 35/65 blend takes about six and two thirds and a 60/40 about eleven and a half, which bracket the 7-15Y horizon from its near edge to its middle. Pure equities take twenty, inside 15Y+. The first two sleeves land just short of their horizon's near edge on purpose. Sizing a sleeve to the middle of its horizon would leave the earliest liability in that horizon exposed to a recovery longer than it has, so each sleeve is sized against the nearest payment it has to fund.

One assumption in that table matters a lot. Scaling the drawdown by the equity weight assumes the fixed income leg holds its value while equities fall. That is a statement about what the fixed income leg is *for*, and it is only true if the leg is short. In a growth shock, government paper rallies while equities fall and the assumption is conservative. In an inflation shock both fall together, which is why every blend in that table pairs equities with bills and not with duration.

Getting that wrong is expensive. Fill the 7-15Y sleeve with long government bonds instead and apply a two point real rate shock to a seven year duration leg, and the 35/65 blend's drawdown rises from 17.5% to 26.6%, which pushes its defined duration from 6.6 years to **10.2**. The 60/40 version moves from 11.4 years to 13.7. Those two figures use the ten

year real yield of 2.78% for the fixed income leg instead of the 2.58% short TIPS yield in Table 6, since the leg here is a long one. The sleeve kept its equity weight and its label. It just became a match for a different liability, by nearly four years in the 35/65 case and more than two in the 60/40.

So this lands in the same place as the floor comparisons later in this section, from another direction: **a blend's defined duration is only short if its fixed income leg is short**. Pair equities with duration and the recovery horizon gets longer. Long bonds are the conservative choice for a return comparison and the optimistic one for a drawdown calculation, and it helps to keep that straight.

The last column applies the valuation point to the whole ladder. At US valuations every sleeve stretches, and the 60/40 blend goes from about eleven and a half years to just over fourteen, which carries it from the middle of the 7-15Y horizon past its edge. A ladder built on expensive assets is matching later liabilities than the same ladder built on cheap ones, whatever the labels say.

The 0-3Y horizon is left out of that table on purpose. It holds no equity, its matching is contractual, and recovery time does not matter for a payment funded by a maturing bill. Nothing has to recover if nothing falls.

What the historical record says about the same question

The formula looks forward and depends on today's price. The 154 year record looks backward and depends on nothing. They are different tools aimed at the same question, and they land in roughly the same place.

Start with the record. For every blend of equities, long government bonds and bills, we computed the full distribution of real returns across every overlapping holding period from one to thirty years, then found the blend that maximized the fifth percentile outcome at each horizon. The result is not a smooth curve, because 154 years of overlapping windows is still a noisy sample and the optimum reverses direction at several horizons. But the general direction is clear. Equity weight is low at one and two years, rises quickly, and hits 100% at seventeen years and stays there. In between it bounces around. Across the four to seven year horizons the raw optimum averaged 75%, and across the eight to fifteen year horizons it averaged 73%. We read that as noise, since even at ten years the sample holds only about fifteen non overlapping observations, and at fifteen years about ten.

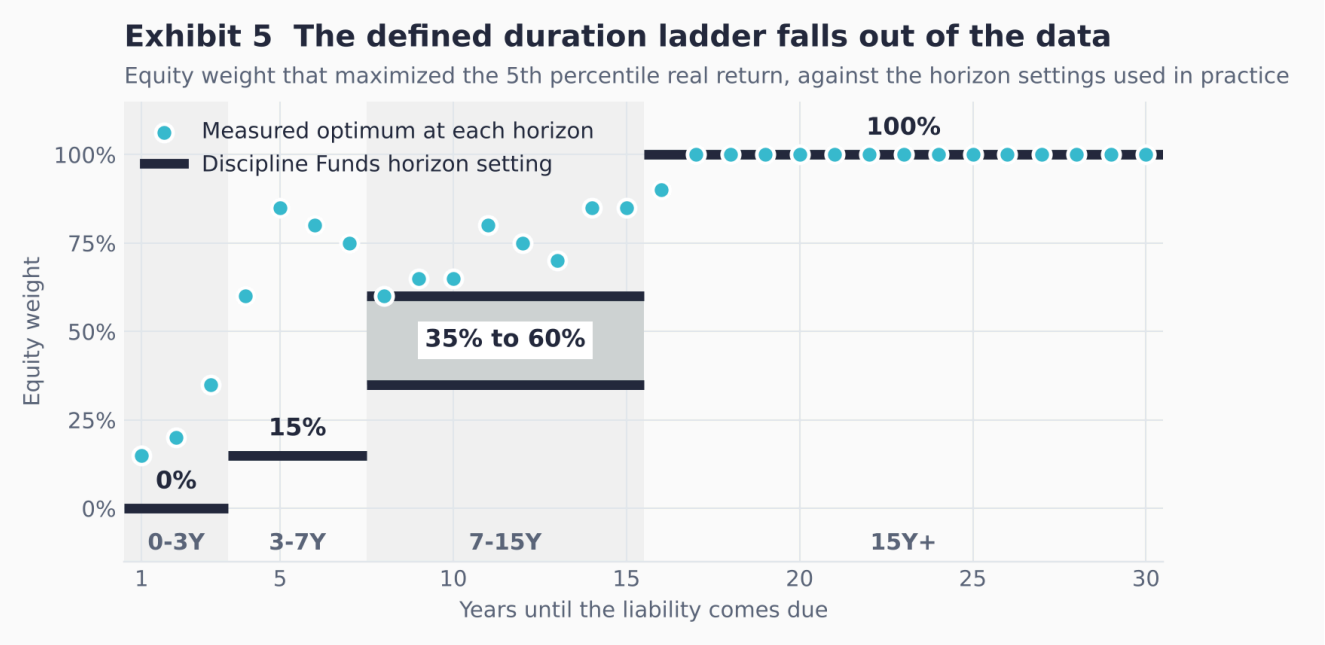


Exhibit 5. Dots are the raw per horizon optimum. Bars are the horizon settings used in practice. In the first three horizons they sit well below the measured optimum, by a wide margin, for the reason given below. In the 15Y+ horizon they converge, the ladder sitting slightly above the optimum at sixteen years and matching it exactly from seventeen. At three years the 35% optimum edges out 30% by two basis points, so read the shape rather than any single dot. Sources: Jorda Schularick Taylor.

So three readings point at the same stretch of the curve. The formula, at this index's 27.8 ratio and a 50% drawdown, gives 19.6 years. Twenty three years is the shortest holding period with no real loss in any window. The optimizer reaches 100% equity at year 17 and stays there. Together they span 17 to 23 years.

Two caveats. The three readings are not fully independent, because the zero loss horizon and the optimizer come from the same 154 year record and only the formula is separate. And the formula's 19.6 years follows from assuming a 50% drawdown. It is not a measurement. Assume 30% and the same ratio gives about ten years.

Even so, the convergence says something, and it is a little awkward for the ladder. None of the three readings is fifteen. The closest sits two years past the horizon boundary, and today's US valuation puts the formula at nearly 28. Our read is that the 15Y+ horizon's contents become dependable somewhere between seventeen and twenty three years on a global valuation, and later on a US one. A liability at exactly fifteen years is at the edge of what pure equities should be asked to cover, which is why the boundary is a guidepost and why a plan with a liability there should look at the 7-15Y allocations too.

The horizon settings in Exhibit 5 sit well below the measured optimum in the first three horizons. Averaged across the 0-3Y horizon the optimizer wants 23% equity, and across the 3-7Y horizon it wants 75%. The ladder holds 0% and 15%.

That gap is deliberate, and it needs a defense. The optimizer maximizes a fifth percentile real return across a pooled distribution. It does not know that a liability due in two years has to be met to the dollar on a specific date. And it does not know that the global fixed income record it is measuring against is dominated by two world wars and a hyperinflation, when every asset failed and cash just failed less. When a payment is close and known, matching precision matters more than an optimized percentile. From the 3-7Y horizon outward that precision is not available anyway, so the ladder starts taking the premium, just more slowly than the optimizer would.

A blend can beat both of its own parts

At ten years the worst real outcome for a 35/65 blend of equities and T-bills was negative 6.06% a year. Pure equities did negative 6.47% and bills alone negative 7.36%. The blend has a higher floor than either of its own parts, and the same is true at five years. At fifteen it draws level with equities, so the effect is strongest through the near and middle horizons. That is diversification at work, and for matching it means a blend of stocks and short government paper can beat either piece on its own.

There is a catch, and it cuts against the conventional bond allocation. The floor only improves when the fixed income leg is *short*. Pair the same equity weight with the long government bond and the blend's floor sits below pure equities from five years out, because the long bond's worst stretches line up with the equity ones. Inflation is the common cause. It is the one thing that hurts long nominal bonds and equities at the same time, so adding duration to an equity sleeve adds correlated tail risk instead of diversifying it. Blends beat their components, but only the short duration ones.

The mechanism underneath is mean reversion, and we are not the only ones to find it. Backtesting rolling thirty year periods of Shiller's US data back to 1871, Sharkansky found that an all equity portfolio supported better retirement spending than a balanced one, which is hard to square with returns that are independent from year to year. Two different datasets and two very different methods pointing the same way is worth more than either result alone.

The duration matched Treasury loses on both counts

Table 7 sets each horizon's allocation against the long government bond an investor would otherwise reach for. The matching allocation has the higher floor at every horizon in the ladder. On the median it wins from the 3-7Y horizon out, by 63 basis points there and by 219 to 490 basis points further out. Only in the 0-3Y horizon does the bond win on the median, by 65 basis points, and that is the horizon where return is not the goal.

Table 7. Defined duration matching allocations against the nominal alternative

HORIZON	MATCHING ALLOCATION	WINDOW	WORST REAL	5TH PCTILE	MEDIAN REAL	BOND WORST	BOND MEDIAN
0-3Y	**T-bills and short government bonds**	3 yr	-13.32%	-7.55%	1.61%	-16.22%	2.26%
3-7Y	**15/85 equity and bills**	5 yr	-12.49%	-5.75%	2.58%	-16.19%	1.95%
7-15Y	**35/65 equity and bonds**	10 yr	-8.05%	-4.01%	4.28%	-9.03%	2.09%
7-15Y	**60/40 equity and bonds**	10 yr	-7.41%	-2.16%	4.78%	-9.03%	2.09%
15Y+	**100% equity**	20 yr	-0.65%	1.11%	5.40%	-4.03%	1.64%
15Y+	**100% equity**	**30 yr**	**1.55%**	**3.01%**	**5.28%**	**-2.10%**	**0.38%**

Annualized real returns across all overlapping windows of the stated length, 1872 to 2025. The bond columns show the long government bond series held over the same horizon. The 0-3Y row is proxied with bills, since the underlying dataset carries no separate short government bond series. The 7-15Y rows pair equities with the long bond series, which is the conservative reading. As shown above, the same blends have better floors when the fixed income leg is short, which is how they are built in practice. Sources: Jorda Schularick Taylor.

Short nominals stay short nominals

The 0-3Y horizon is the one place where nominal certainty is worth paying for, and the only horizon with no equity. A liability due inside three years can be matched precisely, and inflation has almost no time to compound. So the 0-3Y horizon is T-bills, custom bill ladders and short government bonds, held to maturity. Its worst three year real outcome in

the record was negative 13.32% a year against negative 16.22% for the long government bond, with no give up on principal certainty. The spectrum starts from full certainty and layers in probability as the horizon gets longer.

6 Where each answer stops working

The TIPS hurdle, matched to every horizon

The fairest test is simple. Take today's TIPS real yield as the hurdle a matching strategy has to clear, using the TIPS that matures at the same horizon, then ask how often each mix cleared it historically. No single answer wins. Each one has a stretch of the curve where it is right and a stretch where it is not.

Keep section 1's distinction in mind when reading Exhibit 6. A 94% clearing rate and a contractual 3.18% are not the same kind of thing on one scale. One is a probability and the other is a promise. The exhibit shows the exchange rate between them at each horizon, which is the only fair way to compare them at all.

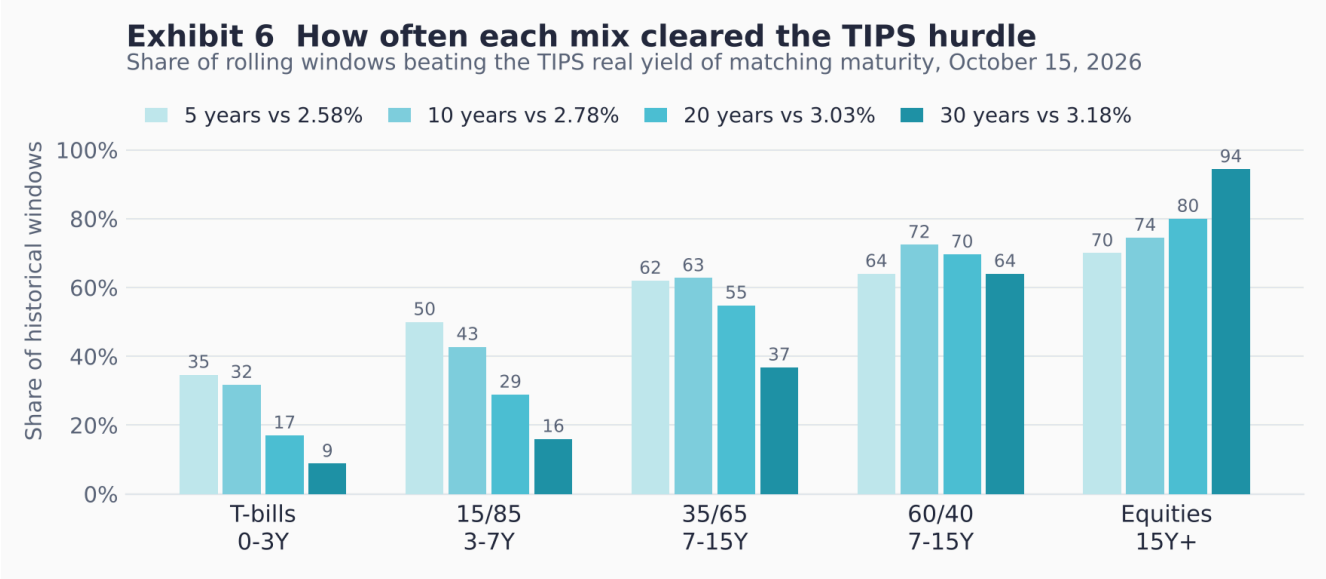


Exhibit 6. Share of rolling windows in which each mix delivered at least the TIPS real yield of matching maturity on October 15, 2026: 2.58% at five years, 2.78% at ten, 3.03% at twenty and 3.18% at thirty. Sources: Jorda Schularick Taylor Macrohistory Database, US Treasury.

This is a narrower result than the usual case for equities. Over 30 year windows, 100% equities cleared the 3.18% hurdle 94% of the time and the 60/40 blend 64% of the time. Neither cleared it every time. Cash cleared it in 9% of 30 year windows and the 15/85 blend in 16%, which is why neither is used past the 3-7Y horizon.

The direction matters more than any one number. For equities the clearing rate rises steadily with horizon, from 70% at five years to 74% at ten, 80% at twenty and 94% at thirty. For T-bills it holds in the low to mid thirties out to ten years and then falls away, to 17% at twenty and 9% at thirty. Time works in opposite directions for the two, and that is the point. But 94% is not 100%. An advisor putting a thirty year liability into equities is accepting roughly a one in eighteen historical chance of falling short of what a TIPS would have locked. That is a real probability against a real liability, and it is why the front horizons exist.

What a locked outcome gives up

A 30 year TIPS bought today cannot be revisited without a mark to market loss. A temporal ladder is a set of positions that can be rebalanced every year for the next three decades. If real yields back up to 4%, the ladder buys them. If equities run hard, the ladder can harvest the gain into the front horizons and deliberately shorten the plan's duration. That is a choice, and section 3 covers what happens when it is not made. If the liability changes, and it will, the ladder adjusts. That flexibility is worth something, and it never shows up in a yield to maturity comparison.

There is a quieter benefit too. A matching strategy works best when its long duration assets are left alone. Because the front horizons have already funded the near term liabilities, the equity sleeve is never a forced seller. Funding the front buys the back the one thing equities need most, which is uninterrupted time.

And that flexibility grows as the plan runs, because the long duration assets have been working the whole time while the front horizons did the spending. Equities that have compounded for fifteen years, and longer bonds that stayed invested instead of being sold to meet a payment, turn into surplus. Surplus can be spent, gifted, moved into the front horizons, or just left alone. A 30 year TIPS delivers the same $25,578 whether the investor ends up needing $20,000 or $40,000. The ladder delivers a distribution, and historically that distribution widens in the investor's favor right when the near term needs are covered and the only thing left to do is wait.

Why not simply hold 60/40?

The obvious objection is that a single balanced portfolio already solves the problem. Why bother splitting the portfolio by horizon when a 60/40 or a 40/60 has worked fine for decades?

It deserves a real answer, because 60/40 is a good portfolio, and it is good for the reason we have been describing. It does not carry the duration of stocks or the duration of bonds. It carries its own, somewhere in between. At ten years its worst real outcome was negative 7.41% a year against negative 9.03% for the long government bond it replaces, and its median was more than double. It is already a synthetic matching allocation. Across the eight to fifteen year horizons the optimizer averaged 73% equity, which is above a 60/40, not below it. The ladder is the conservative reading of its own evidence.

The problem is that it has one defined duration. An investor does not have one liability. They have a schedule of them running from next quarter to three decades out, and a single blend gives every dollar the same intermediate duration no matter when it is needed. The result is a portfolio that is too aggressive at the front and too conservative at the back at the same time. Table 12 shows the same 60% equity exposure spread across several durations instead.

At the front, the worst two year stretch for a 60/40 blend cost 32.3% of purchasing power against 28.2% for T-bills. The gap is smaller than most advisors expect, because the record includes inflations that punished cash badly, but it still runs the right way, and the certainty of a matured bill is worth more than a 4 point edge in a bad tail. Against a payment due in eighteen months, a blend is a hope, not a match. At the back, $10,000 in a 60/40 blend grew to a median of about $29,000 in real terms over 30 year windows, while the same $10,000 in equities grew to about $47,000. Dollars with three decades to run paid for protection they had time to do without. In a 40/60 blend the median was about $22,000, under half of what equities produced.

So 60/40 is not wrong, and defined duration does not reject it. It uses the idea behind it in every rung. Table 6 puts a 60/40 blend's defined duration at about eleven and a half years, inside the 7-15Y horizon. Take the idea that a blend can be built to a duration neither of its parts has, and apply it across the whole liability schedule instead of at one point. The

60/40 stops being the portfolio and becomes the way each rung is built: a 15/85 blend for 3-7Y, a 35/65 blend for 7-15Y, pure equities beyond that, each sized so its defined duration sits where its liability does.

Shorter TIPS, where the case is strongest

The same objection comes in another form. If the thirty year TIPS is the expensive one, why not just buy shorter TIPS and skip the argument?

Inside about fifteen years that is a defensible position on history alone, and at today's prices it holds out to roughly twenty. It is the strongest claim any competing approach makes. Pair the same equity weight with TIPS held to maturity instead of nominal notes and the floor improves dramatically. At ten years, 35% equities alongside 65% in a TIPS ladder would have turned $10,000 into no less than $10,345 in its worst historical case, against $4,321 for the nominal blend. At fifteen years it is $12,165 against $4,821.

That protection is not quite free. The TIPS version gives up $657 on the ten year median, and that is the only horizon where it gives up anything. At fifteen years it adds $179 to the median. Nearer the front, where the 3-7Y horizon holds a 15/85 blend, the ladder version wins on both counts, adding $246 to the five year median and $561 to the seven year median while lifting the floor from $5,132 to $10,424 and from $5,078 to $10,993. Whether that is the right trade depends on how binding the liability is. For a dated obligation that has to be met, a floor that roughly preserves starting purchasing power is worth more than a few percent of median.

Put each horizon's fixed income leg in a ladder and the worst case on record clears $10,000 at five, seven, ten and fifteen years. At ten years the nominal blend's worst case was $4,321, so that is the difference between keeping your purchasing power in the worst case on record and losing more than half of it. Where a ladder held to maturity is practical, today's yields make it the better fill through roughly the middle of the curve.

The form matters more than the label

Two things temper that, and neither is really an argument against TIPS. They are cautions about how the case usually gets made.

First, the math on today's curve. The 10 year nominal Treasury yields 5.10% against 2.78% real, so breakeven inflation is 2.26%. If inflation lands on that breakeven the two deliver the same real return, and the market sets the breakeven where that is a fair bet. What the TIPS buys is protection against inflation running above it. The $6,024 floor advantage above is an inflation surprise premium, concentrated in the war and postwar windows and the 1970s, which is the scenario the breakeven now prices. So the nominal blend is not reaching for return. It is passing on a specific tail hedge whose expected cost is close to zero, in exchange for liquidity, no customization and daily rebalancing. Whether that is right depends on the client, and at today's real yields it is a genuine case by case call.

Second, the whole TIPS ladder case rests on holding individual bonds to maturity, and that does not survive being put in a fund. Sharkansky, who makes the strongest published version of this argument, says so explicitly: the ladder delivers predictable, essentially risk free real income only on the condition that the individual TIPS are held to maturity. He draws the same line we drew in section 3, against TIPS funds as well as conventional bond funds. A TIPS fund is a perpetual constant duration allocation whose holders trade at market value. It offers inflation indexing and none of the cash flow certainty the argument depends on.

That leaves a clean division of labor. Where the liability is dated, the account allows it and the scale justifies the work, buy actual bonds and hold them, and at today's real yields give TIPS a serious look out to roughly twenty years. Where matching precision is not available anyway, because the money is pooled, the horizon is long or the liability schedule is not set yet, a constant duration blend is the better allocation, and Table 3 shows why.

There are really two questions here. Using TIPS as the fixed income leg inside a blend improves the floor a lot against nominal notes at today's yields out to roughly fifteen years, wherever a ladder can actually be built. Using TIPS instead of the risky blend altogether is a narrower claim, and it weakens as the horizon extends. The median give up is 1.50 points a year at ten years against the 35/65 blend that fills that horizon. It peaks at about 2.4 points between fifteen and twenty years, where the twenty year real yield is 3.03% against an equity median of 5.40%, and eases to 2.10 at thirty as the real yield curve rises faster than the equity median falls. Those are the historical prices. At today's valuations Table 1B puts them at well under a point through twenty years and between about 0.4 and 1.1 points at thirty. And a ten year TIPS does not match a thirty year liability anyway. Rolling it three times swaps return risk for two reinvestment decisions at real yields nobody can know today, which leaves no locked outcome and a lower expected one. The only holding that really fixes a thirty year result is the thirty year TIPS, and the question is whether its price, more than two points a year on history and closer to 0.4 to 1.1 points at today's prices, is worth paying.

The temporal allocation is doing the work, not the aggregate mix

This is easy to miss and it is the heart of the design. The sample portfolio in Table 12 has exactly 60% aggregate equity. Held as a plain static 60/40, it would have compounded at a median of 3.64% real over 30 year windows, turning $10,000 into about $29,000, with a worst case of 0.41% and about $11,300.

Structured by time, the same 60% behaves very differently. The 50% earmarked for the 15Y+ liability sits in equities and earns the equity distribution instead of the blended one, compounding at a median 5.28% real into about $47,000. For the dollars that needed protection least, that is 60% more purchasing power at the median.

On floors, the worst 30 year equity outcome turned $10,000 into $15,875 against $11,316 for a static 60/40, so the equity sleeve had the better floor too. At the other end, the 10% in the 0-3Y horizon sits in bills and short notes whose worst two year real loss was 28.2%, against 32.3% for a 60/40.

Same headline allocation, better outcome at both ends, because each dollar is matched to when it is actually needed instead of averaged across everything.

Running the whole thing against a control

So far we have compared sleeves one horizon at a time. The obvious next step is to run a whole plan end to end and see whether the structure actually helps. That sounds simple, but a fair test has to separate three different questions, and each one needs its own comparison.

WHAT THIS TEST IS ASKING

1. Does a liability matched ladder beat a plain portfolio with the same stock exposure? We compare each ladder with a static 60/40 that funds the same spending from the same $2 million.

2. If it does, is that the horizon structure, or just the fact that the ladder ends up owning more stocks? We compare each ladder with a glide matched control: one rebalanced pool that holds exactly the ladder's own stock weight every year, with no separate horizons.

3. Does it matter what sits in the middle of the ladder? We run four ladders that are identical except for how the intermediate reserves are held.

The test household has $2,000,000 and a real spending schedule: $60,000 for a car in year one, $75,000 for a wedding in year two, $50,000 for a remodel in year three, $100,000 toward a downsizing in year seven, and then $60,000 a year, three percent of the starting portfolio, in real terms from age 65 on. Both portfolios carry 60% aggregate equity. The ladder spends front first and refills its front sleeves every three years, drawing from whichever sleeve behind it has done best, and it will not sell equities into a three year drawdown while anything else is left. The first control is a single 60/40 pool rebalanced every year, which pays each bill by selling a slice of everything. Every portfolio runs across all 125 overlapping thirty year windows in the record.

The second control needs a word of explanation, because it is the one that matters most. A ladder does not stay at 60% stocks. As it spends its bills and short notes and only tops them up every few years, the stocks become a bigger share of what is left, and by year thirty the median ladder is 91% stocks. So if a ladder beats a static 60/40, it might simply be because it ended up owning more stocks, which is the rising glide path Pfau and Kitces described in 2014, and any portfolio could do that without a horizon structure. The glide matched control strips that out. It is a single rebalanced pool that copies the ladder's stock weight year by year, with no separate horizons and no spending order. Whatever difference is left between the ladder and that control is the structure itself.

Table 8. What each portfolio holds at the outset

SLEEVE	LIABILITIES IT FUNDS	BASELINE	BALLAST OUT	REAL BALLAST	CONSTANT DURATION
Bills and short notes	**Car, wedding, remodel, years 1 to 3**	$340,000	$200,000	$200,000	$200,000
Intermediate government, nominal	**Years 4 to 7, including the downsizing**	$320,000	$460,000	n/a	n/a
Intermediate, held in real terms	**Years 4 to 7, including the downsizing**	n/a	n/a	$460,000	n/a
60/40 multi asset bridge	**Years 8 to 12 of retirement spending**	$350,000	$350,000	$350,000	$1,500,000
Global equity	**Everything beyond year 12**	$990,000	$990,000	$990,000	$300,000
Aggregate equity content		**60%**	**60%**	**60%**	**60%**

Each portfolio totals $2,000,000 and carries identical aggregate equity, so the only thing varying across the four is where the non equity 40% sits and whether it is nominal or real. The bridge is 60% global equity and 40% intermediate government bonds. Sleeves are specified as asset classes and weights, not products, and each can be built from broad low cost funds. This is a different household from the one in section 7, which works a spending schedule through to horizon weights rather than testing an existing ladder.

The four versions differ only slightly, on purpose. The baseline is the case study ladder. Pushing the ballast out moves $140,000 from bills into the intermediate sleeve. Holding it in real terms keeps the weights and only swaps the intermediate sleeve's return series. The constant duration version collapses the middle into one large bridge. In every case the overall stock and bond mix is the same 60/40 the control holds.

Table 9. The same 60/40, structured four ways, against both controls

PORTFOLIO	WORST	5TH PCTILE	MEDIAN	BEATS STATIC 60/40	BEATS GLIDE MATCHED	GLIDE MATCHED WORST
Static 60/40, annually rebalanced	$1,128,034	$1,363,145	$3,291,046	n/a	n/a	n/a
Ladder, nominal ballast	$1,206,627	$1,637,915	$4,528,768	100.0%	54.4%	$1,244,804
Ladder, ballast pushed further out	$1,181,997	$1,542,338	$4,415,107	100.0%	52.8%	$1,245,018
Ladder, ballast held in real terms	$1,650,838	$2,303,258	$4,645,425	96.0%	73.6%	$1,185,894
Ladder, large constant duration sleeve	$1,140,876	$1,459,188	$3,548,684	81.6%	54.4%	$1,204,297

Terminal real wealth after thirty years of the stated withdrawals, across 125 overlapping windows, 1872 to 2025. Beats static 60/40 is the share of windows in which the ladder finished ahead of the static control. Beats glide matched is the share in which it finished ahead of a rebalanced pool holding the ladder's own year by year stock weight, and the last column is that pool's worst outcome. A win rate near 50% against the glide matched control means the structure added little beyond the glide path. All five portfolios start at $2,000,000 with 60% aggregate equity. The real ballast row replaces the nominal intermediate sleeve with one earning a certain 2.78% real, today's ten year TIPS yield. That is a forward looking scenario, not a historical outcome, since no investor in 1905 could buy it, and it is the single most important caveat on this table. Sources: Jorda Schularick Taylor, US Treasury.

Question 1: yes, the ladder beats the static 60/40 comfortably, and the result does not hinge on the refill rule. The median runs about 37% ahead. We picked a three year refill window and a rule against selling equities into a drawdown, and since we chose those after looking at the data, we checked whether the choice mattered. It mostly does not. Refilling every year, every two years and so on out to every five, with or without the no selling rule, the nominal ladder's worst case runs from about $1.02 million to $1.47 million against $1.13 million for the static control, and it beats the control in 94% to 100% of windows. Longer refill windows help a little because they leave the long end alone for longer. The rule against selling equities into a drawdown barely moves anything.

Question 2: with nominal ballast, most of that edge is the glide path, not the structure. Against the glide matched control the nominal ladder's median is about the same, $4.53 million against $4.56 million, and it wins 54% of windows. A coin flip. The glide matched pool even has the better floor, $1,244,804 against the ladder's $1,206,627. With nominal ballast, a temporal ladder is basically a rising glide path in disguise.

Question 3: the middle matters, and real ballast is what makes the structure pay on its own. Hold the intermediate sleeve in real terms instead and the worst case rises from $1,206,627 to $1,650,838, a 37% better floor, and at today's real yield the median rises too, from about $4.53 million to $4.65 million. The only cost is four points on the win rate against the static control, 96% against 100%, because real ballast gives up ground in the windows where inflation never showed up. And it holds up against its own glide matched control, with a worst case of $1,650,838 against $1,185,894, a 39% better floor with the equity path held constant, winning 74% of windows.

The two fixes that failed tell you as much as the one that worked. Pushing the same nominal ballast further out the curve does nothing, and a large constant duration sleeve is worse at both ends. What repairs the tail is not more ballast or longer ballast. It is ballast that inflation cannot reach. The worst windows in this test start between 1905 and 1913 and run into the inflation of the following decade, and nominal reserves get wrecked by the same force wrecking everything else. That is section 3's argument showing up in a portfolio simulation.

So our takeaway is narrower than an advisor might hope, and more useful. A temporal ladder does beat a static blend with the same exposure, reliably and by a wide margin on the median. But if its reserves are nominal, most of that is a glide path any portfolio could adopt without the structure. The structure earns its keep when the reserves are real, and it earns it in the tail, which is where a matching strategy is supposed to earn anything.

What those two year figures are actually measuring

Those two year comparisons have a name. A liability due in year two is a liquidation event, and the dollar funding it does not care what the thirty year average return was. It cares what the portfolio is worth on that date. That is sequence risk, and it shows up in any plan with dated liabilities, not just one drawing an income.

A static blend exposes every dollar to the entire path no matter when that dollar is needed. A temporal ladder exposes each dollar only to the path over its own horizon. The 0-3Y sleeve meets a two year liability after two years in bills. The equity sleeve meets a thirty year liability after thirty years in equities, never having been sold into a drawdown to pay for something else. Every horizon boundary is a place where that difference pays.

The effect flips between phases. For an investor drawing income, early losses are the dangerous ones, since withdrawals compound against a shrinking base. For an investor still saving, early losses help and the danger is at the end, because a drawdown in the final years hits the largest balance the plan will ever have. Only a single lump sum with no cash flows and one liquidation date is truly indifferent to the order of returns, and almost nobody is that investor. Matching by horizon addresses both versions of the problem.

The strongest objection: a ladder plus equities already does this

A reader who has followed along this far can throw a sharp version of the argument back at us. If short dated certainty is worth paying for and long dated certainty is not, then buy a TIPS ladder for the liabilities inside the next decade, put everything else in equities, and stop. Two holdings, no horizon boundaries, no blends, no refill rule. And if that turns out too conservative or too aggressive, do not add a third or fourth sleeve, just move the split.

That is a better objection than most, and the first thing to say is that it is not an alternative to the framework. It is the framework with two rungs. A TIPS ladder for near liabilities and equities for far ones is a temporal allocation. It sorts dollars by when they are needed and gives each a different risk profile, which is the whole proposition of section 1. Anyone making this argument has already conceded the main point. What is left is a narrower and more answerable question: how many rungs, and what goes in them.

So the disagreement is about resolution, not principle, and our answer cuts both ways.

In one direction it goes toward the objection, further than an advocate of this framework might like. Nothing in the defined duration approach rejects TIPS. Our own Table 9 is the best evidence. Of the four ladder variants we ran, the best one held its intermediate reserves in real terms, and it was the only one that kept its edge against a glide matched control. Its worst case was 37% better than the identical ladder with nominal reserves, and at today's real yield its median was higher too. Section 7 gets to the same place from the other end, where at today's real yields a ladder fill for the 3-7Y

sleeve is a free improvement and the 7-15Y fill is a real trade. If there is one practical takeaway here about what to hold in the middle of the curve, it is that a dated real allocation belongs there wherever an account can carry one. That is the objection's central claim, and we agree with it.

In the other direction, two rungs are not enough, and the reason is the middle of the curve. In a two rung structure every dollar is either fully certain or fully at risk. No dollar sits at 15% or 35% equity, because no such rung exists. That is a real cost, because the middle of the curve is where neither extreme is right. The blend that beat both of its parts on the floor earlier in this section is exactly what a two rung portfolio cannot hold. At ten years that 35/65 mix of equities and bills bottomed at negative 6.06% a year, against negative 6.47% for pure equities and negative 7.36% for bills, so it beats both rungs a two rung portfolio has.

That is also where the second half of the objection breaks down. Adding equity sounds like a free adjustment, but in a two rung structure there is only one dial and it moves every horizon at once. The only way to add equity is to shorten the ladder, and that does not just raise the risk on the twelve year dollar. It removes the funded runway behind the six year one. You cannot keep year six certain while year ten takes a third of the equity premium, because that setting does not exist in a two rung portfolio. It exists in a four rung one, and being able to set risk separately at each horizon is the whole reason for the extra rungs. Separate horizons are how risk becomes a function of time instead of one number applied across the plan.

A third point, which we hold less firmly. A ladder is matched to the liability schedule as it stood the day it was built, and it adapts as fast as its rungs mature. It is not frozen. But moving faster than the maturity schedule means selling bonds at market, which is the cost holding to maturity was supposed to avoid. Since the liability schedule will change and real yields will move, some of that cost will get paid eventually. That is not a defect in the ladder. It is the price of the precision that makes it valuable, the same trade in miniature that section 2 prices across the curve. But it is a cost, and it never shows up in a yield to maturity comparison.

So the objection is right about the direction and wrong about where it stops. Once you accept the argument for two rungs, the argument for four is the same argument applied to the part of the curve two rungs have to round off.

> **WHAT A TEMPORAL ALLOCATION ACTUALLY DOES**
>
> A static blend gives every dollar the same risk profile no matter when it gets spent. A single ladder gives every dollar certainty whether it needs it or not. A temporal allocation gives each dollar the risk profile that fits its own due date, which is what asset liability matching was supposed to do in the first place.

7 A sample temporal allocation

Where the weights come from, and what they buy

Every plan here starts with the liability, not the portfolio. The horizon weights are the present value of each segment of the client's spending schedule, and the overall allocation is whatever falls out. Table 10 runs the math for a 65 year old with $1 million who needs $40,000 a year in today's dollars for 30 years.

Table 10. Turning a spending schedule into horizon weights

HORIZON	SPENDING YEARS	DISCOUNT RATE	PRESENT VALUE	SHARE OF PORTFOLIO
0-3Y	**1 through 3**	2.49%	$114,263	11.4%
3-7Y	**4 through 7**	2.58%	$139,140	13.9%
7-15Y	**8 through 15**	2.78%	$233,915	23.4%
15Y+	**16 through 30**	3.03%	$304,503	30.5%
Funded need			**$791,821**	**79.2%**

Each segment is discounted at the published real yield nearest the start of that segment, from the US Treasury daily par real yield curve for mid October 2026. Using the near end rather than the duration weighted center of each segment raises the present value slightly, so the funded need shown is the conservative figure. Illustrative and pre tax.

The liability uses up 79% of the portfolio and leaves a 21% surplus. Put that surplus in the longest horizon, where it has the most time to work and the least chance of being needed early, and the weights become 11.4, 13.9, 23.4 and 51.3. Their combined equity content is 61.5%.

Nobody chose 60/40. It is just what a $40,000 spending schedule and today's real yield curve produce when each dollar goes to the horizon that will spend it.

Risk capacity is just the funded ratio

Table 10 is one client at one funding level, and the framework generalizes from it with no new machinery. The liability has a due date. The surplus does not. A dollar with no date attached has no horizon to match, so it belongs at the long end by default. Run the same spending schedule against different balance sheets, as Table 11 does, and the allocation moves on its own.

Table 11. The same liability against different balance sheets

PORTFOLIO	FUNDED RATIO	0-3Y	3-7Y	7-15Y	15Y+	EQUITY CONTENT
$2,000,000	253%	5.7%	7.0%	11.7%	75.6%	80.8%
$1,500,000	189%	7.6%	9.3%	15.6%	67.5%	74.4%
$1,000,000	126%	11.4%	13.9%	23.4%	51.3%	61.5%
$900,000	114%	12.7%	15.5%	26.0%	45.9%	57.3%
$791,821	**100%**	**14.4%**	**17.6%**	**29.5%**	**38.5%**	**51.4%**

The same $40,000 real spending schedule over 30 years, valued at $791,821 per Table 10, funded from balance sheets of different sizes. Horizon weights are each segment's present value as a share of the portfolio, with any surplus assigned to the 15Y+ horizon. The highlighted row is the client illustrated in Table 12.

Same liability, same spending, same yield curve. Equity content moves from 51% to 81% on nothing but the size of the balance sheet relative to what it has to fund. Nobody chose those weights either.

Income pulls the same lever from the other side. Social Security, a pension and human capital are not portfolio assets. They reduce the liability. A client whose fixed income covers half the spending has half as much liability to fund, which roughly doubles the funded ratio and moves the allocation with it. The framework does not need a second variable for

that.

The relationship is not linear, though, and the bottom of the table is where advisors should be most careful. There are three different situations.

- **Large surplus.** The undated dollars carry the equity weight, aggregate exposure rises, and a guaranteed real return has little left to offer. This client can absorb a long stretch of disappointing equity returns because the plan does not depend on them.
- **Thin surplus, roughly 100% to 115% funded.** The liability is affordable and there is no margin for error. This is where certainty earns its price, because a shortfall here cannot be recovered from, and it is the client for whom TIPS at today's real yields deserve the hardest look.
- **Underfunded.** No allocation fixes it. Buying a guaranteed real return at this point does not manage the shortfall, it locks it in by contract, which is the pension industry's mistake at household scale. Reaching for return is the same mistake in the other direction. The answer is on the liability side: spend less, work longer, or transfer the longevity risk instead of carrying it.

That bottom row is also where our own argument comes back around. Section 2 conceded that the equity premium cannot be relied on. Surplus is what lets a plan not need it, and the lack of surplus is what makes certainty worth buying.

The whole framework in a few sentences

Those three situations, the defined duration measurement from sections 1 and 5 and the fills compared in section 6 all add up to one argument. Here it is in one place.

> **THE PRACTICAL CONCLUSION**
>
> Risk capacity determines the dispersion an allocation can withstand. Defined duration determines which allocations qualify at each horizon. Where capacity is thin, real ballast should dominate through the middle of the curve. Where capacity is ample, inflation linked bonds can be confined to roughly the three to twenty year range and blends and equities carry the rest. At today's prices that means more certainty against dated liabilities, not fewer equities where capacity allows. Where the plan is underfunded, no allocation solves it and the answer sits on the liability side.

One part of that is easy to get backwards. **The demand for a guaranteed real return does not keep rising as capacity falls. It peaks in the middle and drops off at both ends.** A client with a large surplus has little use for it, because the plan does not depend on the outcome. A client at roughly 100% to 115% funded has the most use for it, because the liability is affordable and there is no margin for error. And an underfunded client has the least use for it of all, because a contract at that point does not manage the shortfall, it guarantees it.

Watch that last case, because intuition says the opposite. Low capacity feels like it should call for more certainty, and the client who most looks like they need a guarantee is often the one a guarantee will hurt. Demand for certainty is shaped like a hump, the same way its price turned out to be.

Nobody funds year one out of the tenth rung

Section 1 put off the sharpest criticism of segmented portfolios, which is Kitces's, and this is where we answer it. The objection is that the rules for which segment gets spent are arbitrary. The one usually quoted first is that if equities are up, spending comes from equities.

On a liability matched ladder that rule is a category error. It takes an allocation matched to a twenty year liability and spends it on a one year one, just because its price went up. Nobody holding a ten rung bond ladder funds year one out of the tenth rung because ten year bonds rallied. The tenth rung is there to be worth something in year ten, and selling it early defeats the reason it was bought.

So **a properly matched ladder needs no decision tree at all**. What gets spent is whatever came due. When the front sleeve's bills mature they fund the liability they were bought for, and the question of which asset to sell never comes up because nothing is being sold. Rebalancing runs on the same schedule as the liabilities, and its size is set by the payment, not by a rule about relative performance. What Kitces correctly found missing from generalized bucketing is the liability schedule itself.

That is also why the drift from section 3 is not something anyone has to defend each year. The long end is left alone because nothing has come due against it yet.

What that displaces

Look at what did not change across Table 11. All five rows are the same 65 year old, spending the same $40,000 a year, with the same stomach for volatility. An age based glide path would put all five in the same allocation. So would a target date fund with the same retirement year on the label. So would a risk questionnaire. Only the balance sheet moved, and the right answer moved with it, from 51% equities to 81%.

That is a real break from how allocations usually get set. **Age is a proxy for horizon, and a poor one**, because it says nothing about how much has to be funded or when. A 65 year old with two years of spending in the portfolio and a 65 year old with forty are not the same client, and their birth year will not tell you that. **Risk tolerance is a real constraint, but it is the wrong input**, because it measures what a client is willing to hold, not what their circumstances can absorb. Capacity tells you the right answer. Tolerance tells you whether the client can stick with it, which is an implementation and communication problem. It is still a real constraint: a 100% equity sleeve only works as a 15Y+ matching allocation if the client can actually hold it through the kind of drawdown this record contains, which at its worst cost 41% of purchasing power in a single year.

The safe withdrawal rate literature runs the same thing backwards. It takes the allocation as given, usually a static 60/40, and asks how much can be drawn from it. This framework starts from what has to be drawn and when, and lets the allocation fall out. The order matters, because a rule that fixes the portfolio first has no way of knowing which dollar is needed in year two and which in year twenty five, so it cannot put them in different places.

None of this needs a new input. The liability schedule is already the first thing any planning process collects. We are only arguing that the portfolio should be built from it, instead of checked against it afterward.

A liability is worth what it costs to retire

The discount rate in Table 10 is the one choice in the whole exercise that really changes the answer, so we treat it as a rule. Discount the same $40,000 schedule at each horizon's median historical return instead and the liability shrinks from about $792,000 to about $642,000. Same client, same spending, same market. The plan now looks 156% funded instead

of 126%, and a $208,000 surplus has become $358,000 without a dollar changing hands.

That version feels better and it is wrong. It values the liability by assuming the portfolio earns the equity premium, which is the assumption this whole paper says you cannot count on, and it is the same math that left a generation of pension plans underfunded. A liability is worth what it costs to defease, and that cost is set by the risk free real curve. The expected return on the assets belongs on the asset side of the page, where it can be wrong without quietly shrinking the obligation.

Table 12 takes the 126% funded case and rounds it to weights a portfolio can actually be built and rebalanced in.

Table 12. Illustrative defined duration ladder

SLEEVE	HORIZON	WEIGHT	ROLE	EQUITY CONTENT
Bill and short note ladder	**0-3Y**	10%	**Held to maturity. Nominal or, at real yields like today's, short TIPS**	0%
15/85 blend	**3-7Y**	20%	**Constant duration blend, bond like with an equity premium tilt**	3%
35/65 blend	**7-15Y**	20%	**Constant duration blend sized to the near edge of the horizon**	7%
Global equity	**15Y+**	50%	**Terminal horizon, no fixed income drag**	50%
Total		**100%**	**Effective aggregate exposure**	**60%**

Illustrative only. Weights start from Table 10's present values for a 126% funded client and are then simplified to practical increments a portfolio can be built and rebalanced in, with the surplus assigned to the longest horizon. The simplification moves individual sleeves by several points in both directions and lands the aggregate equity content at 60% against the 61.5% Table 10 computes. Horizon boundaries and weights are set by the client's liability schedule, not by a model portfolio. The three constant duration sleeves are pooled allocations. Where the account permits individual bonds held to maturity and the liability dates are fixed, the fixed income leg of the 3-7Y and 7-15Y sleeves can instead be built from a TIPS ladder, for the reasons section 6 sets out.

Four sleeves and seven to ten underlying holdings carry the whole structure, which keeps cost and complexity down. Only the front sleeve is held to maturity. The other three are liquid, rebalanceable and priced daily, and nothing is locked to a date the client cannot revisit.

Where a TIPS ladder fits inside this

A TIPS ladder is not a rival to this framework. It is the best fill available for part of it, and the purest version of the contractual matching from section 1. Where a liability is dated and the account allows actual bonds held to their dates, you get certainty no blend can match, and the framework wants it there.

Think about what each sleeve is being asked to do. The 0-3Y and 3-7Y sleeves are being asked for precision, and where an account allows individual bonds held to maturity, a ladder supplies precision no pooled allocation can. At today's real yields the ladder version of the 3-7Y sleeve roughly doubles the worst case, from $5,132 to $10,424 over five years, and it does it with a median gain of $246 instead of a cost. On today's curve that is a free improvement for a dated liability. The 7-15Y sleeve is more interesting. The floor gain is even bigger, but so is the median cost, so it is a real trade. Past fifteen years the historical case for a ladder weakens, though it does not disappear. On the record, the median give up grows past two points a year and equities cleared the guarantee in 94% of windows. At today's prices, though, Table 1B puts the give up at well under a point out to about twenty years, so the ladder reaches further than history alone suggests. Past that it is a judgment call, not arithmetic.

So the split is simple. Use a ladder wherever the account, the scale and the liability dates allow it, from the front of the plan through roughly the middle of the curve. Use constant duration blends for horizons a ladder cannot reach, for pooled money, and for schedules that are not set yet. Both do the same job, putting each dollar in an allocation whose reliable horizon matches its own. They just differ in how much precision is available.

One caution applies to every choice inside the framework. Which fill wins depends on the real yield curve the day you decide, and the curve moves. In 2021 every ladder row in this section was deeply negative and the nominal blends won easily. The horizon structure is permanent. What goes in each horizon is a relative value call to redo as the curve moves.

The ladder nobody can afford to run

Certainty is not the only thing with a price here. Complexity has one too, and it is paid in something no table in this paper measures, which is the operating capacity of the practice running it.

Everything above understates what a ladder costs to run. Not in fees, in operations. Every client's ladder is a different set of securities, because their liability dates and amounts are their own. There is no model portfolio to apply and no way to onboard a household without building a new ladder at whatever the curve is that week. It cannot be rebalanced, because held to maturity means held. Odd lot spreads on off the run issues are wide. The maturity schedule has years where no TIPS matures, and those gaps have to be bridged some other way. And the client's statement has thirty line items instead of a handful.

And the burden compounds. A hundred household ladders are a hundred seperate reinvestment calendars, each throwing off maturities that have to be reinvested at rates nobody knows in advance, each needing its own decision. Performance reporting, billing, trade supervision and error correction all scale with the number of distinct positions, not the number of clients. Every one of those is a chance for a mistake, and the odds of at least one somewhere go up with the count. Meanwhile the scarcest resource in any advisory practice is advisor hours, and hours spent on ladder mechanics are hours not spent on the liability planning that produced the ladder in the first place.

That last point is the one we would push hardest. This framework starts with financial planning, and the planning is where the value comes from. A process that eats its own planning capacity to get a marginal gain in precision has made a bad trade, whatever the historical tables say.

A certain cost for a conditional benefit

The case for simplicity is not about looks or fees. Complexity has a certain cost and a conditional benefit. The operational burden of a ladder gets paid every year, in every account, no matter what. The floor protection it buys only shows up when inflation surprises to the upside. One is a certainty and the other is a probability, and people rarely weigh them that way.

And a plan that actually gets implemented beats one that does not. A framework a practice can run across its whole book, consistently, and revisit when a client's life changes, will beat a theoretically better one that gets applied to the twelve households big enough to support it and dropped for everyone else.

That may sound strange after several pages on where ladders win, but **for a lot of practices the TIPS ladder question is moot, not because the analysis went against it but because the operating model cannot carry it.** That is a legitimate reason to skip it, and it costs less than you might think.

Defined duration in four or five holdings

The framework does not need securities held to maturity to work. It needs each dollar in an allocation whose reliable horizon matches when that dollar is needed, and that can be built entirely from pooled funds. At its simplest that is four sleeves and four or five holdings: a T-bill fund for 0-3Y, a bond heavy blend for 3-7Y, a multi asset blend for 7-15Y, and a global equity allocation for 15Y+. Low cost, daily liquid, rebalanceable without touching a client account, and runnable as one model across a whole book. Most builds end up with seven to ten holdings, like Table 12, for diversification inside the sleeves, not because the structure requires it. Four is the minimum, and the point is that the minimum is low.

That version is not contractual, not statistical, but blended. It is purely liability driven and its weightings should change depending on the risk capacity of the investor. The investor who needs more certainty needs more contractual holdings and the investor who has a higher risk capacity can afford more of the statistical blends.

We tend to defer towards simple is better because complexity has its own costs, but the reality is that the plan needs to be as complex as the financial planning requires and that is something only the advisor and investor can agree on.

8 Conclusion

> **THE CONCLUSION IN THREE POINTS**
>
> - **Inside about five years, match with contracts.** Bills and TIPS held to maturity cost nothing there, and today they pay you.
> - **From about five to twenty years, risk capacity sets the mix.** The more surplus, the more equity in the blend. The thinner the cushion, the more TIPS, and the further out they should go.
> - **Past about twenty years, statistical matching has historically been the better tool.** The equity floor sits above every alternative, and stocks never lost purchasing power over any window of 23 years or more. At today's prices the 30 year TIPS is a closer call than usual, and for a thin cushion it deserves a serious look.

The long horizon matching problem is real. The liability is unpredictable and so is the asset return, and no allocation makes both knowable at once. Nobody can forecast the future, so the useful question is not which asset wins. It is which uncertainty to accept, and what a reasonable person should pay to get rid of it at each point on the curve.

The distinction from section 1 organizes the answer, and it is worth coming back to now that the numbers are in. Contractual matching puts a known amount in the client's hands on a known date. A bill, a note or a TIPS held to maturity does that, and nothing else does. Statistical matching puts a distribution in their hands instead, with a middle well above the contract and a bottom, at long horizons, close to it. Neither is the better method. They answer different questions, and which question you are answering depends on when the money is needed and whether the plan could survive being wrong.

What we tried to do here is price the exchange rate between the two at every point on the curve, and the shape of that price is the main result. At three years the contract is free, and on today's yields better than free. At five years the two are level. On the historical record, by ten years the contract costs about a point and a half a year of median return, and past the middle of the curve the price peaks at about 2.4 points before easing at the very long end. At today's valuations most of that disappears out to about twenty years, and the thirty year price shrinks to roughly 0.4 to 1.1 points. Certainty is expensive at some maturities and cheap at others, and knowing which is which is most of the work.

So the allocation moves along the spectrum instead of picking a side. The front of the plan is contractual because precision is available there and costs nothing. The middle is a mix, leaning toward the contract where the liability is dated and the plan cannot absorb a miss, and toward the distribution where it can. The long end is statistical because the liability it funds was never contractually specifiable, because the one holding that can promise a thirty year real outcome charges for it, more than two points a year on history and less at today's prices, and because the floor on a distribution held that long has historically sat close to the guarantee it replaces. Where a household sits on that spectrum is about capacity, not courage. A large surplus has no due dates to match and can live at the statistical end. A plan funded within a few points of its liability belongs closer to the contractual end, and an underfunded plan is not an allocation problem at all.

Read the record that way and it comes down to three sacrifices, each smart in its place and expensive outside it.

- **Cash gives up return, permanently.** Over thirty year windows it lost purchasing power 40% of the time and cleared today's real yield in 9%. Against a payment due inside three years it is still the best money in the portfolio, because nothing else delivers a known amount on a known date.
- **Bonds held to maturity give up upside for precision.** Through the near and intermediate horizons that is a great trade, and at today's real yields the TIPS version deserves a serious look out to roughly twenty years. Stretched to thirty, it turns a distribution with a floor of $15,875 and a median of $46,841 into a certain $25,578. The problem is the maturity it is asked to cover, not the security.
- **Equities give up near term stability for long run reliability.** Inside a few years the range is punishing and the floor is deep. Past twenty the floor sits above every alternative in this record, and from year twenty three they never lost purchasing power in any window. The wait is the price, and only dollars with time can pay it.

A temporal allocation is not a bet against certainty or a bet on equities. It recognizes that certainty has a term structure and that each of those sacrifices has a horizon where it is worth making. Buy certainty outright at the front, where it is cheap and the match is precise. Buy it through dated bonds in the middle, where the price is still small. Layer in probability as the horizon lengthens, using allocations you can measure against the record, and keep the whole thing liquid so it can respond to a liability schedule that will change.

That is why nothing here argues for removing guaranteed assets from a plan. Bills and TIPS held to maturity make an amount and a date certain. Blends and equities offer a range whose middle sits well above that guarantee. A plan wants both, and should not ask either one to cover a maturity it is poorly suited to. The industry's mistake is not picking the wrong one. It is picking one for the entire portfolio.

Today's prices sharpen that instead of overturning it. Equities are expensive and TIPS are not, so certainty is cheap and the price of uncertainty in equities is high. A plan should buy more certainty against the liabilities it has to meet. But an investor with high risk capacity should not walk away from equities because of it. Their surplus has no due date, it can sit through a poor decade, and a guarantee would cap it as surely as it floors it. The answer to expensive equities is a longer horizon and a more diversified starting price, not an exit.

I am convinced that asset liability matching will become the dominant way retail financial planning and investment management get done, and nothing in this record changed my mind. What the record did was show what the approach costs at each horizon, and where an advisor should stop paying for certainty. That point is a good deal further out the curve than I expected when we started, and at today's prices it is further out still.

Nobody gets to know the future, but time is the one problem every investor has to confront at every stage of their financial plan. We can't know in advance what uncertainty will cost us, but we can know what certainty costs today. A disciplined asset liability matching process lets an investor make that trade deliberately, one dollar and one due date at a time, which is exactly what institutions have been doing for more than seventy years and what households deserve too.

Methodology and Sources

Everything in this paper is reproducible from public data

Data

The 1872 to 2020 portion of the record is built from the Jorda Schularick Taylor Macrohistory Database, release 6, which carries annual total returns for equities, long government bonds and short government bills, together with consumer prices, real GDP per capita, population and dollar exchange rates, for sixteen developed markets. The markets used are Australia, Belgium, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom and the United States. Canada and Ireland are dropped because their return histories begin too late to contribute to the early sample.

Release 6 ends in 2020, so the years 2021 through 2025 are built separately and spliced on. That extension is described below and is a genuine departure in construction, disclosed here because the period it covers, an inflation shock and the sharpest rise in policy rates in four decades, is exactly the kind of episode a note about the price of certainty should not be allowed to skip.

Current yields are the US Treasury daily par yield curve and daily par real yield curve for mid October 2026, and current inflation is the Bureau of Labor Statistics CPI release of October 2026 covering September 2026. Bloomberg US Aggregate duration figures are taken from the Hartford Funds publication cited below.

How the index is weighted, in one paragraph

The series is **GDP weighted for its entire length, 1872 through 2025**. There is no capitalization weighted segment anywhere in it. The series is built in two pieces, both weighted the same way: the Jorda Schularick Taylor record to 2020, and a separately sourced extension for 2021 through 2025. To be precise about what is GDP weighted: it is the **weighting across countries**, in every year of the record. Each country's own equity return is whatever its national index delivers, and for the 2021 to 2025 years those national indices are MSCI country indices, which are capitalization weighted inside each market as any index is. What never happens is a switch to capitalization weighting across countries. Cap weighting enters the cross country construction at exactly one point, which is the validation described below, where the GDP weighted series is checked against MSCI World ex USA over their 1970 to 2020 overlap. That comparison is a test of the construction, not a switch in it, and the only thing carried forward from it is the equity haircut.

Construction of the global index

Each country's nominal total return is converted to a local real return year by year using (1 + nominal) divided by (1 + local inflation), minus 1. The three asset series are then aggregated across countries each year as a weighted average, using the prior year's total real GDP as the weight, where total real GDP is Maddison real GDP per capita multiplied by population. Prior year weights are used so that no weight embeds information from the return year itself. A year enters the

index only if at least nine markets report, which produces a continuous annual series from 1872 to 2020, 149 observations, with coverage rising from nine markets in 1872 to all sixteen. With the extension the full record runs to 2025 and 154 observations.

GDP weighting rather than capitalization weighting is often assumed to be a departure from convention. For long run data it is the convention. No capitalization weighted global index exists back to the 1870s, because reliable country level capitalizations do not exist that far back, and equal weighting would give Portugal the same influence as the US.

The most authoritative long run dataset in the field makes the same choice for the same stated reason. Dimson, Marsh and Staunton weight each country in their world equity index by its starting equity market capitalization or, in years before capitalizations were available, by its GDP. That means the early decades of their record, including the Great War, the Depression and the Second World War, are GDP weighted. Their world bond index is built the same way but weighted by GDP throughout, which is the same choice made for the bond leg here. They publish no world bill index and use US Treasury bills as the world risk free rate, where we build a GDP weighted world bill series instead.

One clarification is worth making because it is easy to get backwards. The global market portfolio, meaning the aggregate of financial assets actually held, is capitalization weighted by definition. GDP weighting is not a proxy for it. It is a deliberate choice to weight by the size of the underlying economy rather than by the price the market happens to be putting on it, which keeps the composite from being dominated by whichever market is in a bubble. Japan was roughly 45% of world equity capitalization in 1989 on a far smaller share of world output. Readers should treat the choice as a modeling decision, and the section below reports what happens when it is varied.

To test that construction, the index is validated against MSCI World ex USA over the 1970 to 2020 overlap, converting the Jorda Schularick Taylor local returns to US dollars using the database's own exchange rates and excluding the United States to match the benchmark. The GDP weighted series is not replaced by the benchmark at any point. The comparison exists to size the gap between the two and to charge that gap against the equity leg. Measured on nominal US dollar returns, annual returns correlate at 0.934 with a mean absolute annual difference of 5.7 points, and the volatilities are close, 20.5% against 20.9%. On real returns the correlation is 0.938. The Jorda Schularick Taylor construction runs 0.95 points a year higher over the overlap on an arithmetic mean basis, which is the basis the haircut is applied on, and 0.91 points higher geometrically. That 0.95 point gap is subtracted from every equity return year in the Jorda Schularick Taylor portion of the record before we compute any statistic. The haircut is applied only to equities, is applied uniformly across those 149 years, and is a conservative adjustment, not a measured one. It is meant to stand in for index construction differences, implementation cost and the impossibility of actually holding a sixteen market portfolio for most of the period. Bond and bill series are used unadjusted. Over the Jorda Schularick Taylor years the index compounds at 6.32% real for equities before the haircut and 5.36% after, the small difference from a flat 0.95 being the effect of compounding. Across the full spliced record, 1872 to 2025, the figures are 5.50% for equities, 1.58% for long government bonds and 0.81% for bills. Measured geometrically, as (1 + equity return) divided by (1 + bill return) minus 1, the equity premium over bills is 4.65 points a year, which is the figure section 2 uses.

Does the weighting drive the result?

Two checks, both reported in full including where they go against us.

The first asks whether the haircut closes the gap to real capitalization weighted data. Over 1970 to 2020, the fifty one years where a capitalization weighted developed market index exists, we compared the GDP weighted equity series against MSCI in US dollar real terms. The raw GDP weighted construction runs hot, as the haircut assumes. After the haircut it lands almost exactly on the capitalization weighted series.

Table 13. GDP weighting against capitalization weighting, 1970 to 2020

EQUITY SERIES, REAL USD GEOMETRIC MEAN	EXCLUDING THE US	ALL 16 MARKETS
Capitalization weighted (MSCI)	5.39%	5.88%
GDP weighted, before the haircut	6.30%	6.61%
GDP weighted, after the haircut	5.34%	5.65%
Correlation of annual returns	**0.938**	**0.954**

Both panels are US dollar returns deflated by US CPI, 1970 to 2020. The excluding the US panel is the well sourced one and is the comparison the haircut was calibrated on. The all 16 panel is the out of sample check, and there the haircut overshoots by 23 basis points, charging equities slightly more than capitalization weighting would. Sources: MSCI, Jorda Schularick Taylor, Bureau of Labor Statistics.

The series used in this paper is therefore within **5 basis points a year of a capitalization weighted index** excluding the US at a correlation of 0.938, and within 23 basis points across all sixteen markets at 0.954. On the structural findings the two weightings agree closely but not perfectly, and the exceptions matter.

Whether a short fixed income blend beats both of its components comes out the same under both weightings in all eight cells tested, at three, five, ten and fifteen years in both panels. The gap between the equity floor and the cash floor agrees in sign at three, ten and fifteen years, and differs only where that gap is within a rounding of zero, at five years excluding the US and at ten years across all sixteen. The price of certainty traces the same shape under both, near zero at the front of the curve and roughly three points at ten to fifteen years, but at three and five years the two weightings straddle zero instead of agreeing on sign. None of those disagreements is large in magnitude. All of them sit at horizons where the quantity being measured is close to nothing.

One finding also reverses inside this fifty one year window relative to the full record, and it reverses under capitalization and GDP weighting alike, which makes it a property of the period, not of the weighting. Cash's real loss frequency falls with horizon here instead of rising, because 1970 to 2020 was an unusually good half century for real cash rates. Which is itself an argument against building a long horizon paper on a fifty year window.

The second check varies the weighting across the full record and is less comfortable. Holding the sixteen countries and the entire construction fixed and changing only the weight vector, equal weighting materially changes three results.

Table 14. The same record under three treatments, 1872 to 2020

FINDING	GDP WEIGHTED	EQUAL WEIGHTED	US ONLY
Worst 30 year real equity return	1.55%	1.53%	2.46%
Median 30 year real equity return	5.28%	4.73%	5.75%
Share of 30 year windows clearing 3.18%	94.2%	76.7%	95.8%
Zero real loss windows from year	23	23	21
Bonds negative over 30 years	45.8%	40.8%	20.0%
Cash negative over 30 years	**40.8%**	**40.0%**	**13.3%**

Identical construction and identical country set in the first two columns, with only the cross country weight changed. Both weightings run on the same 149 Jorda Schularick Taylor years and therefore on the same 120 thirty year windows. The third column is not a weighting but the United States series alone, drawn from the same database. Jorda Schularick Taylor years only, so the figures are the pre extension versions of those quoted in the body. Sources: Jorda Schularick Taylor.

Three things in that table and the run behind it deserve stating instead of burying. Under equal weighting the share of 30 year windows clearing the hurdle falls from 94.2% to 76.7%, which is the difference between a one in seventeen and a one in four shortfall. The finding that a short fixed income blend beats both of its components fails at every horizon tested. And the price of certainty stops rising: measured on the Jorda Schularick Taylor years alone and against the 15/85 blend at three years rather than Table 1's bill row, it runs 0.67, 1.52 and 2.10 points at three, ten and thirty years under GDP weighting, but 0.93, 1.41 and 1.55 under equal weighting, which is a curve that flattens after ten years, not one that keeps climbing. That last one bears on the historical price curve in Table 1, though at today's prices Table 1B already shows that curve flattening for a different reason.

Those are real vulnerabilities and an allocator who reruns this will find them. What holds across all three treatments is the worst 30 year equity outcome, which sits at 1.53% to 1.55% in both global versions, the arrival of zero real loss windows at 21 to 23 years everywhere, and the direction that the equity floor eventually clears both cash and bonds, though the horizon at which it does so moves a great deal, from year five under equal weighting to year nine under GDP weighting to year fifteen in the United States alone.

The US only column matters for a different reason. The US column is the most flattering treatment in the table on the three headline equity measures, with the best floor, the best median and the highest clearing rate. It is also the one where cash and bonds look dramatically safer than anywhere else, and the one where the equity floor takes longest to clear them, at year fifteen. Building on it would have produced a stronger sounding paper on the numbers most often quoted and a less honest one overall.

The 2021 to 2025 extension

The extension keeps the same country set, the same weighting rule and the same real return arithmetic. Only the underlying series change. Equity returns are the MSCI country indices, gross total return in local currency, for each of the sixteen markets. Bill returns are the OECD three month interbank rate, annual average, published through the Federal Reserve Bank of St Louis. Inflation is the annual average change in each country's consumer price index, from the World Bank, with United Kingdom headline CPI taken from the Office for National Statistics and the United States 2025 figure from the Bureau of Labor Statistics. Weights are prior year GDP at purchasing power parity from the World Bank, which replaces the Maddison series the earlier years use.

Government bond total returns are the weakest link and are disclosed as such. Local currency government bond index returns were obtained for ten of the sixteen markets, from index levels published on fund factsheets that name the benchmark: gilts for the United Kingdom, ICE US Treasury 7-10 Year for the United States, Bloomberg Euro Aggregate Treasury sub indices for France, Germany, Italy and Spain, FTSE Japan World Government Bond for Japan, Bloomberg AusBond Treasury for Australia, SBI Domestic Government 7-15 for Switzerland, and the NBP Norwegian Government Bond index for Norway. For the remaining six, Belgium, Denmark, Finland, Netherlands, Portugal and Sweden, no local currency government bond index return was obtainable. Each is proxied by the market whose ten year yield path it tracks most closely, taking the average of France and Germany for Belgium, the Netherlands and Finland, the average of Italy and Spain for Portugal, and Germany for Denmark and Sweden. Those six markets carry 7.0% of the index weight, and dropping them from the bond aggregate entirely moves it by no more than 0.29 percentage points in any of the five years, so the proxy choice does not drive any of our results.

The equity haircut described above is **not** applied to the extension years, because the extension is already built from MSCI data and the haircut exists to close the gap between the Jorda Schularick Taylor construction and MSCI. Applying it twice would double count.

What the extension changes is less than a reader might expect. No horizon in the note acquired a new worst case: every multi year floor quoted here is still set by a window beginning before 1997, the single worst calendar year for equities excepted. The thirty year statistics barely moved, because five new years add only five windows to the hundred and twenty that were already there. What did move is the front of the curve. Cash lost purchasing power in three consecutive years, and the share of rolling windows in which bills lost real value rose at every horizon inside twenty years. The 2022 calendar year is the sharpest single illustration in the whole record of the point section 3 makes: in real terms the global government bond aggregate lost 20.91% that year, more than the 17.8% lost by global equities, and the only worse year for government bonds in the whole record is 1920.

Construction of the blends

Blends are annually rebalanced mixes of the global equity series and a fixed income sleeve that is itself a mix of the bill and long bond series. The intermediate Treasury proxy in Table 3 is a 50/50 mix of those two. The one exception is the TIPS ladder comparison in sections 6 and 7. There the TIPS leg is held to maturity at the stated real yield and the equity leg beside it compounds untouched, since a ladder held to maturity cannot be rebalanced without breaking the property that makes it work. The nominal blends it is compared against are annually rebalanced, which is how a pooled allocation is actually run. The comparison is therefore between two allocations as each would really be held, not between two identically constructed ones. The database carries no separate short government bond series, so the bill series stands in for the 0-3Y row of Table 7, which understates that horizon slightly. Rolling statistics use every overlapping window of the stated length, giving 154 one year, 145 ten year, 140 fifteen year and 125 thirty year windows.

What the sample cannot tell you

Four properties of this record limit what any figure drawn from it can carry, and they are stated here rather than left to the reader to find.

Overlapping windows are not independent observations. There are 125 rolling thirty year windows in the sample and only five non overlapping ones, and all seven thirty year windows that fell short of today's 3.18% real yield start between 1889 and 1903, which is really one episode. To put a rough size on that uncertainty we resampled the 154 year record in blocks of ten and twenty years and recomputed the headline figures 1,500 times. Table 15 shows the results. Block resampling breaks up some of the long horizon mean reversion this paper relies on, so these ranges are if anything too wide, but they are the right order of magnitude. This is the largest single caveat on everything above.

Table 15. How much the headline figures could move on a different draw of history

STATISTIC	AS REPORTED	90% RANGE, 10 YEAR BLOCKS	90% RANGE, 20 YEAR BLOCKS
30 year equity median, real	5.28%	3.47% to 7.50%	3.78% to 7.13%
30 year equity worst case, real	1.55%	-1.94% to 3.58%	-1.29% to 3.53%
30 year windows clearing 3.18%	94.4%	55% to 100%	60% to 100%
10 year 35/65 median, real	4.28%	1.95% to 5.09%	2.22% to 5.24%
5 year 15/85 median, real	2.58%	1.06% to 3.34%	1.06% to 3.42%
Equities never lose in real terms from year	23	14 to beyond 30	14 to beyond 30

Circular block bootstrap of the 1872 to 2025 annual record, 1,500 resamples at each block length, with equities, bonds and bills resampled together so their cross correlations are preserved. Each range is the 5th to 95th percentile of the statistic across resamples. "Beyond 30" means a meaningful share of resamples never reach a horizon with zero real loss windows inside thirty years. Sources: Jorda Schularick Taylor, Discipline Funds calculations.

An aggregate hides the worst national outcomes, and the sample is of survivors. A GDP weighted composite is not any one investor's experience. Run the same windows country by country and the tails are much worse, including multi decade real equity losses in markets that were invaded, defeated or hyperinflated, and the composite assumes an investor who could hold a global portfolio continuously, rebalance it across borders and never have it expropriated. That investor did not exist for most of the sample. The country set is also developed markets only. Russia in 1917 and China in 1949 are not in it, and neither are the markets that never developed enough to be measured.

The 0.95 point equity haircut is a judgment, not a measurement. The haircut is our estimate of the gap between an unimplementable index and a real portfolio net of costs and frictions. Remove it and every equity figure in this paper improves by roughly that amount a year.

A guarantee is a contract and history is not. The 1.55% worst thirty year case describes what happened, not a floor anyone underwrites. Nothing in the record rules out a worse thirty years, and individual markets inside this sample delivered exactly that. The figures on the statistical side of every comparison in this paper are historical frequencies. The figures on the contractual side are terms.

The valuation adjustment in Table 1B

Table 1B conditions the price of certainty on today's starting valuation in two ways. The valuation adjusted column uses Robert Shiller's monthly US series, January observations from 1881, to regress forward annualized real total returns at each horizon on the starting earnings yield, the inverse of the cyclically adjusted ratio. The fitted return at a ratio of 27.8, minus the unconditional US median at that horizon, gives a valuation effect of about negative 3.1 points at ten years, negative 2.3 at twenty and negative 1.0 at thirty. That effect is subtracted from the global equity median at the same horizon. The earnings yield column simply sets the expected real equity return at 3.60% at every horizon. Fixed income legs in both columns earn forward real yields from the October 2026 curves. The US forward return series ends in 2023, so the most recent windows are not in the regression, and only one thirty year window started from a ratio above 25. The adjustment is an estimate of direction and rough size, not a measurement.

The fifth percentile optimization

For each horizon from one to thirty years we searched a grid of equity weights from 0% to 100% in 5% steps, crossed with a fixed income sleeve ranging from 100% bills to 100% long bonds in 25% steps. The reported optimum maximized the fifth percentile of the real annualized return distribution across all overlapping windows of that length. It is an in sample result on one historical path and indicates the shape of the relationship, not a precise prescription. The raw optimum is noisy at short horizons where the fifth percentile is estimated from few effectively independent observations. The horizon settings in Exhibit 5 and Table 7 are the allocations used in practice. They sit below the raw optimum at every horizon inside sixteen years, sit above it at sixteen, where the optimum is 90%, and match it exactly from seventeen years out.

Corporate credit, real estate and commodities are excluded, as are fees, taxes and trading costs beyond the equity haircut described above, which would reduce every figure shown.

Sources

- Oscar Jorda, Moritz Schularick and Alan M. Taylor. The Macrohistory Database, Release 6. Federal Reserve Bank of San Francisco Working Paper series and macrohistory.net/database. Underlying paper: Jorda, Knoll, Kuvshinov, Schularick and Taylor, The Rate of Return on Everything, 1870 to 2015, Quarterly Journal of Economics 134 (3), 2019.
- MSCI Inc. MSCI World ex USA Index, annual gross returns in US dollars, 1970 to present, and the MSCI single country indices, gross total return in local currency, 2021 to 2025. msci.com/indexes
- Federal Reserve Bank of St Louis (FRED). OECD Interest Rates, 3-Month or 90-Day Rates and Yields, Interbank Rates, and 10-Year Government Bond Yields, annual averages, series IR3TIB01 and IRLTLT01 by country. fred.stlouisfed.org
- World Bank. Inflation, consumer prices, annual percent (FP.CPI.TOTL.ZG), and GDP, PPP, current international dollars (NY.GDP.MKTP.PP.CD), 2020 to 2025. data.worldbank.org
- UK Office for National Statistics. Consumer Price Index annual rate, series D7G7. ons.gov.uk
- Government bond index total returns 2021 to 2025 are taken from published fund factsheets that name their benchmark index: iShares Core UK Gilts (FTSE Actuaries UK Conventional Gilts All Stocks), iShares Treasury Bond 7-10yr (ICE US Treasury 7-10 Year), iShares France, Germany, Italy and Spain Government Bond UCITS ETFs (Bloomberg Euro Aggregate Treasury and country Treasury indices), iShares Core Japan Government Bond ETF (FTSE Japan WGBI), iShares Treasury ETF Australia (Bloomberg AusBond Treasury 0+ Yr), iShares Swiss Domestic Government Bond 7-15 (SBI Domestic Government 7-15), and KLP Statsobligasjon annual reports (NBP Norwegian Government Bond Index).
- US Department of the Treasury. Daily Treasury Par Yield Curve Rates and Daily Treasury Par Real Yield Curve Rates, October 2026. home.treasury.gov/resource-center/data-chart-center/interest-rates
- US Bureau of Labor Statistics. Consumer Price Index Summary, September 2026, released in October 2026. bls.gov/news.release/cpi.nr0.htm
- Board of Governors of the Federal Reserve System. FOMC statement of September 2026, and Federal Funds Target Range, series DFEDTARU and DFEDTARL. federalreserve.gov and fred.stlouisfed.org
- Barclays. Historic CAPE ratios by country and region, September 30, 2026. United States 39.35, developed Europe 22.34. indices.cib.barclays/IM/21/en/indices/static/historic-cape.app
- S&P Dow Jones Indices. SPIVA U.S. Scorecard, Year-End 2025. In 2025, 79% of actively managed large cap US equity funds underperformed the S&P 500, and 93% did so over the 20 years ending in 2025. spglobal.com/spdji/en/spiva/article/spiva-us-year-end-2025
- Robert Shiller. US Stock Markets 1871 to Present and CAPE Ratio, ie_data.xls, data through early October 2026. shillerdata.com. The forward return regressions behind Table 1B use the same monthly series through September 2023, from the public mirror at github.com/datasets/s-and-p-500, checked against published readings.
- Zvi Bodie. On the Risk of Stocks in the Long Run. Financial Analysts Journal 51, no. 3, May 1995, pages 18 to 22. See also Mike Dempsey, Robert Hudson, Kevin Littler and Kevin Keasey. On the Risk of Stocks in the Long Run: A Resolution to the Debate? Financial Analysts Journal 52, no. 5, 1996, pages 57 to 62.
- Steven P. Feinstein. Measuring Risk with the Bodie Put When Stocks Exhibit Mean Reversion. Journal of Risk 1, no. 3, Spring 1999, pages 25 to 36. Establishes that the Bodie result holds under any binomial price process.
- Lubos Pastor and Robert F. Stambaugh. Are Stocks Really Less Volatile in the Long Run? Journal of Finance 67, no. 2, 2012. Earlier as NBER Working Paper 14757, 2009, the source of the one and a half times figure in section 2.
- Paul A. Samuelson. Lifetime Portfolio Selection by Dynamic Stochastic Programming, and Robert C. Merton, Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case. Both Review of Economics and Statistics 51, no. 3, 1969.

- John Y. Campbell and Luis M. Viceira. The Term Structure of the Risk Return Tradeoff. Financial Analysts Journal 61, no. 1, January 2005.
- Warren E. Buffett. How Inflation Swindles the Equity Investor. Fortune, May 1977. And Mr. Buffett on the Stock Market, Fortune, November 1999, the source of the 1964 to 1981 Dow and gross domestic product comparison.
- Cullen Roche, Discipline Funds. *Valuations Matter. Just Not the Way Wall Street Says.* Research note, October 2026. The CAPE regressions, quintile statistics and timing tests cited in section 5 are set out in full there.
- Frank M. Redington. Review of the Principles of Life-Office Valuations. Journal of the Institute of Actuaries 78, 1952, pages 286 to 340. Harry Markowitz. Portfolio Selection. Journal of Finance 7 (1), 1952, pages 77 to 91.
- Cap weighted world cyclically adjusted ratio of 33 at the end of July 2026, which implied a 71.3% United States share of developed market capitalization. Holding that share at the September 30 readings gives a cap weighted ratio of about 32.3. The GDP weighted ratio of 27.8 used throughout section 5 is computed from the September 30 readings at this paper's own GDP weights.
- Hartford Funds. Are You Prepared for Duration Risk? Duration of the Bloomberg US Aggregate Bond Index, 1978 to September 30, 2026. hartfordfunds.com/dam/en/docs/pub/whitepapers/CCWP013.pdf
- Edward F. McQuarrie. Stocks for the Long Run? Sometimes Yes, Sometimes No. Financial Analysts Journal 80 (1), 2024, pages 12 to 28, and Stocks for the Long Run, Setting the Record Straight, CFA Institute Enterprising Investor, 2024.
- Elroy Dimson, Paul Marsh and Mike Staunton. Global Evidence on the Equity Risk Premium. Journal of Applied Corporate Finance 15 (4), 2003, and The Worldwide Equity Premium, A Smaller Puzzle, 2006. Index weighting methodology cited from these. The Global Investment Returns Yearbook dataset itself is licensed and is not used here.
- Corey Hoffstein (@choffstein). Post on X, May 17, 2026, distinguishing explicit from statistical liability driven investing. x.com/choffstein/status/2056106136017736079. The dual framework in this paper is taken from him. We use contractual in place of his explicit, because the guarantee it describes is a contract.
- Michael Kitces. Managing Sequence Of Return Risk With Bucket Strategies Vs A Total Return Rebalancing Approach. Nerd's Eye View, November 12, 2014.
- Wade D. Pfau and Michael E. Kitces. Reducing Retirement Risk with a Rising Equity Glide Path. Journal of Financial Planning 27 (1), January 2014, pages 38 to 45.
- Harold Evensky. The cash flow reserve strategy, 1985, as described in Evensky, Horan and Robinson, The New Wealth Management.
- Stefan Sharkansky. The Only Other Spending Rule Article You Will Ever Need. Financial Analysts Journal 81 (4), 2025, pages 59 to 83.
- Edward F. McQuarrie and William J. Bernstein. Long TIPS Yield 3%. Time to Buy? Advisor Perspectives, August 13, 2026.
- M. Barton Waring and Laurence B. Siegel. The Only Spending Rule Article You Will Ever Need. Financial Analysts Journal 71 (1), 2015, pages 91 to 107.

Important disclosures

General. This paper is for informational and educational purposes for financial professionals. It is not investment, tax or legal advice, and it is not an offer or a recommendation to buy or sell any security or to adopt any strategy. Any allocation decision should be made in light of an individual investor's full financial circumstances, liability schedule and risk capacity.

Conflict of interest. Discipline Funds is an SEC registered investment adviser. Registration does not imply a certain level of skill or training. Discipline Funds manages investment strategies built on the asset liability matching framework described here and therefore has a commercial interest in its adoption. No Discipline Funds product, fund, account or strategy is named, described or shown in this paper. The blends used throughout, such as 15/85 and 35/65, are generic illustrative mixes built from broad index return series and are not intended to represent any fund, account or strategy managed by Discipline Funds.

Disclosure of interest. Discipline Funds manages investment strategies built on the framework described here and therefore has a commercial interest in its adoption. No Discipline Funds product is named or recommended anywhere in this paper, and every allocation shown is constructed from broad index return series rather than from any managed strategy. The sleeve structures discussed are presented as a method, not as a product. Readers should weigh the argument accordingly.

Hypothetical results. Every return, outcome, backtest and allocation result in this paper is hypothetical. None of it reflects actual trading or the performance of any client account, and no investor achieved these results. The figures were constructed with the benefit of hindsight, and some modeling choices, including the refill rule in the Table 9 backtest, were made after reviewing results. They do not reflect advisory fees, fund expenses, transaction costs or taxes, which would lower every figure shown, and they do not reflect the effect of real economic and market conditions on actual decision making. Hypothetical results have inherent limitations and actual results would differ, possibly materially.

Estimates and projections. Defined duration figures, the expected returns behind them and the valuation adjusted figures in Table 1B are estimates based on stated assumptions. They are not forecasts or guarantees, and different reasonable assumptions would produce different figures.

Historical data. Past performance and historical return distributions do not predict or guarantee future results. The historical analysis relies on overlapping windows drawn from a sixteen market composite, which overstates the statistical confidence a reader should place in any single figure, as Table 15 shows. Indices and composites are unmanaged and cannot be invested in directly. Third party data is believed to be reliable but is not guaranteed.

Securities and accounts. Treasury Inflation Protected Securities held to maturity pay a contractual real return backed by the US government. Sold before maturity they can lose value, and inflation accruals may be taxable before they are received. The blends discussed provide no guarantee of any outcome. Diversification does not ensure a profit or protect against loss. Deposit and money market yields shown are as of the dates stated and change often. Bank deposits are insured by the FDIC up to applicable limits. Money market funds are not insured and can lose value. Tax treatment depends on individual circumstances and state law, and readers should consult a tax professional.

All investing involves risk, including the possible loss of principal.