Upper Left Wealth Management, LLC

August 2026

Completely Re-wiring The Model Portfolio

This is a longer letter than I usually send, and I want to tell you up front why it exists.

When I do the semi-annual rebalance, I am pulling a lot of data and running a lot of numbers. And tinkering, and experimenting, and going down rabbit holes. It is very nerdy work, but I enjoy it. I am obsessed with one question: "how can you make the best portfolio, keeping both risk and return in mind?" The process always spans days. The most recent rebalance took over a month. This was a BIG haul.

I spend too much time with AI. Recently I have spent more time with Claude's Fable 5 model than I have with actual humans. Don't judge. I wanted to see if Claude would understand the data-driven investment process, and then act as an investment committee comprised of Nobel-prize winners and thought leaders. This was not a simple task.

Over the past several weeks I spent countless hours to **put our own investment process through the most demanding review it has ever received**: an AI-assisted examination that reviewed and distilled the entire process, roughly 100,000 formulas spanning about 50 tabs across 9 interconnected Excel workbooks. We* examined every number in the portfolio and process and asked it to justify itself. Some numbers passed with flying colors. Some needed better documentation. Some were resized. A few were retired entirely. Then we rebuilt the machinery that runs the process, tested it, and ran the semi-annual rebalance through it live. **Our process and portfolio have been completely rebuilt from the ground up.**

This letter is the complete account: what we found, what we changed, what it cost, and what we are *not* claiming. A process that only reports its wins is a marketing document. This is meant to be the other thing. I was never meant to be an English major, so I have long hired a writing professional to proofread my letters to make them more readable, and that editing process has recently been replaced with AI. This letter was largely written *by* AI and edited by me (the normal process inverted). I did this because it was better to ask Claude "keep a detailed record of everything we are doing, with the goal of explaining it to an outsider when we are done." This letter is largely that output.

*Throughout this letter I will use "we" to refer to either AI and myself, or you and me as "we" collectively.

1. What we believe, which has not changed

Everything in this letter rests on the same philosophy we have always run:

1. **Own the world.** The portfolio starts from the global market portfolio (everything meaningfully investable, weighted the way the world actually holds it) and deviates from that starting point only where we have documented evidence.
2. **Choose risk first.** We decide how much risk the portfolio carries; return is the output.
3. **Diversify for the times that matter.** An asset earns its place by behaving differently in the moments you need it to, for reasons we can explain, at a price that makes sense. An asset's true contribution to portfolio risk is its volatility multiplied by its correlation with everything else you own, which is why a volatile asset that zigs when the rest zags can *lower* the portfolio's risk.
4. **Discipline over reaction.** Rebalancing runs on a schedule and on pre-set tolerance bands, or on rapid price changes (as in March 2020). Never on headlines or nerves.

Nothing above is new. What is new is the machinery underneath it, and the honesty with which every number now has to carry its own papers.

2. I audited the entire process. Here is what that audit uncovered.

I asked a hard question of my own process: if a skeptical outside committee examined every cell, assumption, and formula, what would survive?

Quite a lot survived, and we will say so plainly. The overall structure (global, diversified, systematically rebalanced, valuation-aware, tax-aware) held up well. In several places the review vindicated the old judgment: our hand-set rebalancing tolerances turned out to rhyme closely with what the new formulas derive from first principles. And the single best trade in the process's history, moving money from bonds into stocks on March 17, 2020, two days before the bottom of the Covid crash, because the model said to… that trade *was* this process, executed by hand before the machinery existed.

But we also found things we would not accept from anyone else, and we will not hide them from you:

1. Some position sizes existed because they had always existed. Asked why gold was 11 percent, the honest answer was history and our best guess on sizing, not math-driven analysis.
2. Our reliance on Shiller P/E to size the country-specific equity weightings leaned too hard on country-specific data and correlation, and not enough on global-market data. When we tested it, part of it worked and part of it did not.
3. Some of the data feeding the process had quietly gone stale, and nothing in the old spreadsheets would ever have said so.

Every one of these has been fixed, and the process now contains permanent machinery designed to catch this entire class of problem automatically.

3. Every number now has papers

The central reform is simple to state. Every parameter in the portfolio now carries exactly one label (anchor, constraint, view, or policy), one recorded source, one confidence grade where that applies, and one written condition that would change it. Nothing is inherited, nothing is unexplained, nothing wears two hats.

An *anchor* is the neutral starting point: the world portfolio itself. A *constraint* is a non-negotiable, like never holding assets we consider uninvestable: Russian securities (untradeable in the US), or Chinese local-currency government bonds (largely untradeable in the US). A *view* is an opinion we hold with evidence, a stated confidence, and a written trigger for changing our mind. A *policy* is an implementation choice, like how often we rebalance. The labels matter because they set the burden of proof: views must earn their size, constraints must name their reason, and policies must be applied consistently.

All of it lives in a single document and executable program. Adjusting the model as the world changes can now be done in minutes with AI, versus over a few days, laboriously, by me.

4. The old portfolio and the new one, side by side

The easiest way to understand this work is to put the old model and the new one next to each other. Some things moved a great deal, and each of those moves has a documented reason. And a surprising amount barely moved at all, which is worth pausing on: when a hostile, evidence-first review leaves a position standing, that is the strongest endorsement the old model could have received.

Asset class	Old	New	Change
US stocks	13.5%	30.9%	+17.4
Real estate	10.0%	3.5%	−6.5
International developed stocks	17.4%	22.4%	+5.0
Broad commodities	5.0%	0.0%	−5.0
Gold	11.0%	7.0%	−4.0
Bonds	21.3%	17.5%	−3.8
Emerging market stocks	9.3%	7.3%	−2.1
Bitcoin	5.0%	4.0%	−1.0
Managed futures	7.0%	7.3%	+0.3

One thing the table does not show: the portfolio's expected risk did not jump with the stock weight. The new construction still targets roughly 10–13 percent annualized volatility (the same range as before) but is built to deliver more zig when other things zag: **more expected return per unit of risk taken, not more risk**.

5. The big changes, decision by decision

The risk target itself. The old process targeted a 75/25 stock/bond split. The new process targets *risk*: the portfolio is engineered to carry roughly 10–13 percent annualized volatility, and

the stock/bond split becomes whatever delivers that. The number came from two points, one data-driven and one behavioral. First, the historical data suggests that risk beyond 10–13% is uncompensated via expected returns; you merely get a larger range of outcomes. Second, you, collectively: in the 2022 bear market and the March 2020 Covid crash, almost no clients abandoned the plan. A plan that gets abandoned at exactly the wrong time is a plan set up to fail. This has long been my goal: a more robust and formulaic way to achieve the desired portfolio risk level.

More stocks, and more of them American (stocks 40% to 60%; US stocks 13.5% to 30.9%). This is the headline, and it has two causes. The first is where the starting point comes from. The old model built its country weights from both a GDP weighting and an equity market-cap weighting, and this punished the US market. We removed the GDP weighting due to lack of evidence. The second cause is how the old model responded to expensive markets, which deserves its own paragraph. The new level is not an opinion that markets are cheap. It is the level that matches the risk target we are actively choosing, with a measured valuation trim on top.

Valuation: adjusted, not abandoned. We have always paid attention to how expensive markets are, using a measure called CAPE (the Shiller P/E), and the old model's valuation formula scaled stocks down without a limit. When we tested that formula family against a century of data, the verdict was mixed. The caution is real at true extremes, but applied too aggressively it costs real return in long bull markets, and comparing *countries* against each other on valuation turned out to be a weak and unstable signal; cheap countries were often cheap for good reason. So we rebuilt the machinery to match the evidence. Country comparisons still exist but with strictly limited authority: capped tilts. The heavy lifting moved to a new overlay we call the valuation thermostat: when the world as a whole becomes historically extreme, the portfolio's risk target steps down along a pre-committed schedule, and steps back up as valuations normalize. We do not try to time markets. We lower the temperature, pay a known cost in expected return, and let the schedule restore risk later.

Gold, resized (11% to 7%). The position was rebuilt on two documented legs: gold's tendency to hold or gain in extended bear markets, and a considered view about governments and debt (central banks themselves have long been heavy owners). If you don't own gold you have to ask yourself, "what does virtually every central bank in the world know that I don't?" The trim of the excess, after gold roughly doubled, was harvesting a win, not abandoning a belief.

Real estate, right-sized after an autopsy (10% to 3.5%). The old weighting traced to an institutional survey about *private* real estate that had been applied to publicly traded REITs, a different animal. Our audit showed REITs fell harder than stocks in every modern bear market. Your stock ETFs already contain real estate at market weight (including international property, which is why the separate international real estate fund was retired), so total exposure now matches the global market's roughly 5 percent rather than a category overweight.

Broad commodities, 5% to zero (but the job didn't disappear; it moved). The long-run return evidence for a static commodity basket did not survive scrutiny, and five years of actual tax returns showed the commodity fund was among the most tax-expensive positions in the book. The inflation-protection job moved into managed futures, conditionally: trend funds in 2021–22

were long commodities and short bonds *because the trends said so*, which is why DBMF-type funds printed +20% in 2022. The honest caveat: managed futures' inflation protection is conditional on the episode being *trended* rather than a spike. A 1973-style overnight oil shock is the scenario where static commodities beat trend.

Managed futures, kept on evidence (7.0% to 7.3%, same funds). This strategy (systematic trend-following) was the portfolio's quiet hero in 2022, gaining 20%+ in the year stocks and bonds fell together. We rebuilt its statistics from actual fund data rather than assumptions, and the evidence supported the position. It is the portfolio's primary defense against inflation-driven bear markets, the bear scenario where ordinary bonds fail too.

Bitcoin, kept smaller (5% to 4%). Treated as what it is: a small, survival-weighted call option on the possibility that digital gold becomes a permanent monetary asset. It is sized by that math rather than by round number, with its volatility assumption now updating mechanically as history accumulates. Still valuable for its volatility alone: rebalancing loves buying low-ish and selling high-ish.

Emerging markets, trimmed and simplified (9.3% to 7.3%). The dedicated China fund was retired; China now enters at its investable weight through the broad funds. Emerging markets are now owned the same way everything else is: at anchor weight, tilted only by documented rules.

Bonds, rebuilt around different jobs (21.3% to 17.5%, different insides). Every bond in the portfolio now names its job. Inflation-protected Treasuries more than doubled (3% to 7%) because they match your future spending directly. Intermediate Treasuries provide ballast. A long-Treasury position was *added* (new, ~2%) as capital-efficient insurance for the one crisis nothing else in the portfolio covers: a deflationary crisis, the scenario where growth collapses and interest rates collapse with it. Gold, TIPS, and trend-following carry the inflationary crises; long Treasuries show up in the deflationary one, and their long duration means a thin 2 percent sliver buys protection that would otherwise cost roughly triple the capital. Hedged international bonds diversify interest-rate cycles, and a measured slice of unhedged foreign bonds matches the imported share of what you actually buy.

Three exits from the bond side. The short-term bond fund was cut for having no job our other funds don't already do better. Emerging market local currency bonds went to zero: they fail precisely in the crashes they were supposed to diversify. In equity bear markets they behave like risk assets, falling alongside stocks, which is the moment we want bonds to perform. And the unhedged foreign bond position was cut roughly in half, resized to match the imported share of what we actually spend rather than a larger currency / dollar devaluation bet.

6. What barely changed, and why that should give you confidence

A review this aggressive could have rebuilt everything. It did not, because much of the old model turned out to be right, and in several places the new machinery independently derived nearly the same answer the old judgment had reached.

The building blocks themselves. Look at the two portfolios and you will see mostly the same tickers, unchanged or merely resized, not replaced: the same core international stock funds in the same internal proportions; the same TIPS, intermediate Treasury, and international bond funds; the same gold vehicles; the same two managed futures funds; the same bitcoin fund; the same US REIT fund; the same small-cap completion approach beside the S&P 500 fund (now set to the market's own proportions); and cash where it was. The old model's fund selection survived the audit nearly intact. What changed is that every *size* now traces to a stated tolerance, a piece of evidence, or a written rule.

The discipline. Semiannual rebalancing, daily tolerance-band monitoring, and tax-aware trading all predate this review and all continue. The new formula-driven tolerance bands, derived from each asset's measured volatility, landed remarkably close to the bands the old model had set by hand over the years. And the single best decision in the old model's history (the March 2020 rebalance into stocks, made by hand, because intuition said to) is precisely the behavior the new machinery is built to repeat.

The honest summary: the old model's instincts were substantially vindicated; its bookkeeping was not. The new portfolio keeps the instincts, sizes them with evidence, writes every reason down, and adds machinery that catches errors. Same philosophy, with a more robust process.

7. Your taxes: we read the actual forms, and tax-efficient asset location is now policy

Rather than theorize about tax efficiency, we sat down with five consecutive years of actual tax forms from real client accounts and traced every dollar of tax-relevant income to the specific fund that generated it, validating each year to the penny.

What the analysis showed, in plain English:

1. **The US stock funds were exonerated — on income.** They generated almost none of the ordinary-income drag (about 1 percent of it). The income drag came, year after year, from the same categories: bond funds, commodity funds, real estate funds, managed futures funds. Capital gains from *selling* appreciated funds are a separate tax, and stocks are part of those when we rebalance (see point 4); those are managed by how we trade, not by where funds live.
2. **The drag was material.** Across five years, one large account produced roughly $400,000 of ordinary dividends, only a third of it tax-favored. Most of the rest could have compounded untaxed simply by living in a different account type. Same investments, same risk, same expected return; different location going forward.
3. **The fix is a policy, not a product.** Funds that throw off ordinary income belong in retirement accounts. Stocks belong in taxable accounts (their dividends are tax-favored, and foreign stock funds carry a foreign tax credit that is *forfeited* inside an IRA).
4. **Even the process's best moment had a tax bill, and we measured it.** The biggest tax year in the study was the bill for the best trade in the study's history (the March 2020 rebalance into stocks, harvested in 2021). The trade was right and we would make it again. The lesson was different: keep the process, stop paying full freight on its income.

So we are re-organizing how we manage households with multiple account types. If you have different account types (an IRA and a brokerage account, a Roth and an IRA, or all three) we will treat those accounts as one large bucket and allocate asset classes between them in the most tax-efficient manner. This is a big change in how I manage the portfolios and how the reporting works, and it should lower your taxes over time. A side benefit for smaller accounts: because the household portfolio is now built from asset-class sleeves, a small account can participate properly in the household allocation. A $15,000 Roth cannot reasonably hold 18 different securities, but it can hold the one security that benefits most from being in a Roth, while the household's total allocation stays exactly on target. So expect your small IRA or Roth to look odd, but trust that it fits into your total diversified portfolio.

This will be a measured rollout, executed with taxes in the foreground. We are in August, and I don't want to generate taxes at this point in the year if I don't have to. Two things are happening, one now and one more gradually. First, positions we are exiting completely, like emerging market bonds, will be sold; those likely have some gains, and we will realize them. Throughout, sales are routed through the positions with losses or the smallest gains first, and trades may spread across tax years where that helps. Second, the location moves (such as "managed futures belong in tax-deferred accounts") will not be immediate for everyone. Appreciated positions that remain in the model are *not* sold merely to move them between accounts. Instead, I have created a program that analyzes your cost basis and your account sizes (how large your IRA is versus your brokerage account) and produces two numbers before anything moves: what a move costs in realized tax this year, and what it saves every year afterward. Based on that ratio and your current realized-gains position, I will decide when, and if, to make the move. For many of you with large gains, I will not sell anything at all relating to location-aware investing; I will simply direct future purchases of certain funds to the IRA instead of the brokerage account. Nothing about this requires action from you; it is explained here so the changes in your statements have their reasons attached.

8. Numerous portfolios, one household

Some of you hold one portfolio with us; some more. They are specialized tools, and they are measured together, at the household level. All three of these have been rebuilt from the ground up.

Core is the subject of most of this letter: the best balance of risk and return we know how to build, engineered to our desired tolerance.

Conservative is the spending shield. It has one job: protect one to three years of your planned spending so that a bear market never forces you to sell Core at depressed prices. It holds Treasury bills for the first year of spending, short-maturity inflation-protected Treasuries for years two and three, and a measured position in gold (sized so that even a bad gold outcome should not cost more than 5 percent of the spending target). In most years it will lag Core, sometimes by a lot. That gap is the insurance premium, and the account is working precisely when it looks boring. It is optional, and whether it fits you is a conversation, not a formula.

Aggressive is the higher-octane satellite, for the specific situations where it fits. It holds no ordinary bonds at all, with one deliberate exception: a small position in long-term Treasuries, because the aggressive portfolio keeps exactly one bond, the one with the most crash protection per dollar. Everything else in it obeys the same rules, caps, and machinery as Core; it is a second column in the same constitution, not a second philosophy. One honest sentence belongs in every conversation about it: it is built to fall further in a crisis. That is not a flaw; it is the price of the higher expected return, and the household's ballast lives in Core.

9. Closing

If we had to compress this whole letter into two sentences, they would be these. Every parameter in your portfolio now carries exactly one label, one provenance, one confidence, and one revision condition; nothing is inherited, nothing is unexplained, nothing wears two hats. And the AI investment committee's closing note, which we are keeping: "*the manager showed up with a spreadsheet process and is leaving with a constitution — and he won several of the arguments that mattered, and left with a sound process where he didn't*."

There is nothing you need to do. The August rebalance will roll out over the next few weeks (it is much more complicated than usual, with client-by-client tax analysis that slows the process), and the changes described here will be reflected in your accounts. If you would like to walk through any part of this, from the philosophy to a single number, we would genuinely enjoy that conversation.

As always, thank you for your trust and stay diversified.

Randy Kurtz, CFP®

8/15/2026

Disclaimer:

Risk is defined as standard deviation for the purposes of this letter. *Diversified portfolio is defined as equally split between U.S. stocks, international stocks, U.S. bonds, international bonds, real estate, gold, and commodities, rebalanced annually and excluding fees, taxes, or other costs. This is not a template for how I manage money, but it is a template for how I show the benefits of diversification. The views expressed are the views of Randy Kurtz as of the date on this document and are subject to change at any time based on market and other conditions. This is not an offer or solicitation for the purchase or sale of any security and should not be construed as such. Past performance is not indicative of future performance. Performance quoted is taken from generally representative accounts, is after fees but before any taxes, and individual results will vary, sometimes substantially, from any results discussed here. Any references to historical data refer to the period 1970–present unless otherwise specified. This letter was edited by AI for clarity. Any information, statement, or opinion set forth herein is general in nature, is not directed to or based on the financial situation or needs of any particular investor, and does not constitute, and should not be construed as, investment advice, a forecast of future events, a guarantee of future results, or a recommendation with respect to any particular security or investment strategy.