Custom benchmarks for 13F portfolios: relative performance and what they cost to trade

Evidence from 13F filings: Berkshire Hathaway, Pershing Square, Tiger Global, Lone Pine and Baupost, plus 2,801 filed portfolios and 144 large wealth managers

Part 2 of a series on hierarchical risk decomposition. Part 1 benchmarked a single stock with market, sector and subsector ETFs. Because those ETF exposures add across positions, the same construction extends directly to a stock portfolio.

By Conrad Gann, Founder, RiskModels

Growth-oriented Fidelity Contrafund and the value-oriented Dodge & Cox Stock Fund pursue opposing strategies, yet both benchmark to the S&P 500. Their management discussions describe the themes behind performance, but standard reporting does not measure those themes separately from company-specific selection. A thematic bet expressed through several companies and a return unique to one company can both appear as selection. Investors get the narrative, but not the decomposition.

An alternative approach turns that thematic narrative into a measurable benchmark. Benchmark each stock with a small set of traded ETFs representing its market, sector and subsector exposures, then add those ETF weights across the portfolio. The exposures shared across holdings become the portfolio's custom benchmark; the remaining return is measured relative to those common exposures. The S&P 500 still shows performance against the broad market. The custom benchmark shows performance after accounting for the risks embedded in the holdings.

This paper applies the construction to the 13F filings of Berkshire Hathaway, Pershing Square, Tiger Global, Lone Pine and Baupost. It asks: **what benchmark is implied by the stocks the portfolio actually owns?**

RiskModels.org publishes the open ERM3 specification, research and visual dossiers for inspection in a browser. RiskModels.app provides the production risk engine through REST, Python, CLI and MCP for attribution and ETF hedge-basket construction on a reader's own holdings.

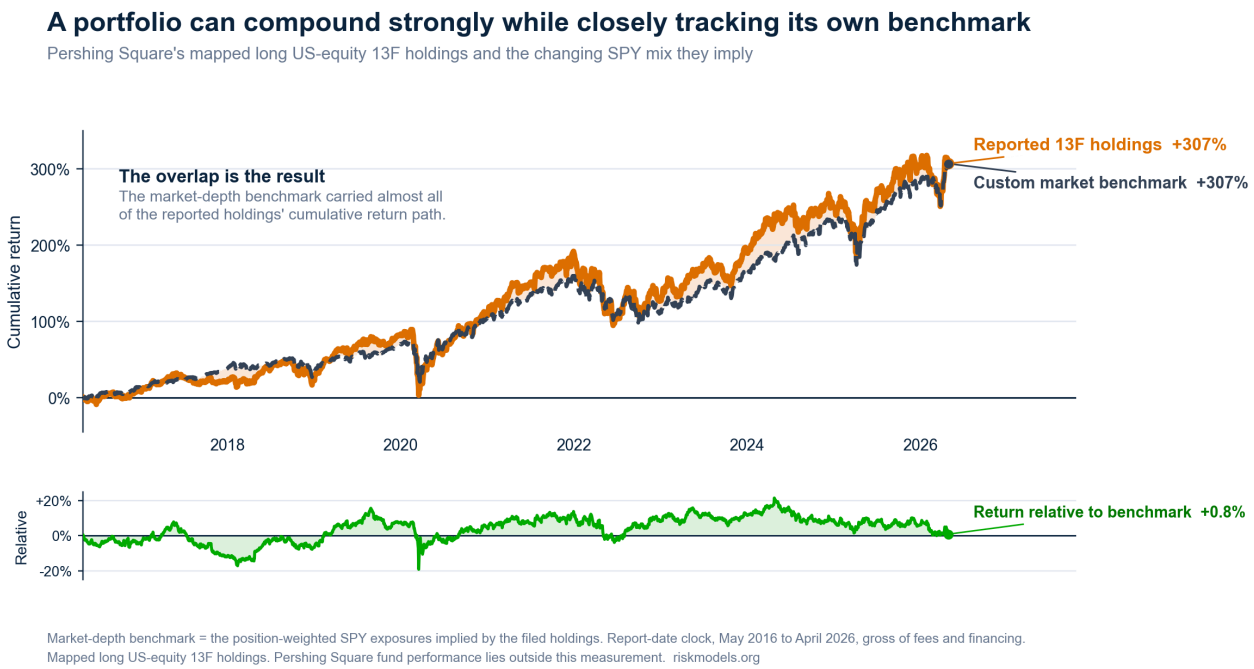


Figure 1. Cumulative return of Pershing Square's mapped long US-equity holdings reported on Form 13F, the model-implied market-depth benchmark built from their position-level SPY exposures, and the return relative to that benchmark, May 2016 to April 2026, report-date clock. Rounded endpoints are +307%, +307% and +0.8%. Returns are gross of fees and financing and represent

the reported holdings rather than fund performance.

The nearly identical upper lines are the finding. Pershing's reported 13F holdings earned a high absolute return, but so did the market exposure implied by those holdings. The lower panel shows the difference.

A portfolio benchmark starts with its stocks. For NVIDIA, the ETF benchmark can stop at any of three depths:

Benchmark depth	ETF benchmark for NVDA	Return left outside the benchmark
L1: Market	SPY	Sector, subsector, style and stock-specific return
L2: Sector	SPY + XLK	Subsector, style and stock-specific return
L3: Subsector	SPY + XLK + SOXX	Style and stock-specific return

The table identifies the ETFs at each depth. The model estimates their weights afresh using information available before the return day. Adding XLK changes the SPY weight; adding SOXX changes both.

The portfolio step is arithmetic:

1. Estimate the ETF weights for every stock.
2. Multiply each ETF weight by that stock's weight in the portfolio.
3. Add the weights of identical ETFs across positions.

NVDA and Apple can both contribute SPY and XLK legs. At portfolio level, those common exposures collapse into one SPY position and one XLK position; SOXX and RSPT remain separate subsector legs. Repeating the calculation across every holding produces one ETF benchmark for the book. Any stock portfolio with mappable constituents can be treated the same way.

The return identity is:

> **Portfolio return = custom ETF benchmark return + return relative to the benchmark.**

The second term is often called residual or benchmark-relative return. It shows where return sits in this decomposition; it is neither total return nor a direct estimate of manager skill. High absolute return can coexist with little benchmark-relative return when the benchmark exposures also perform well.

Filed holdings supply observable stock weights. That makes it possible to build the benchmark for 2,801 portfolios, the five named filers and 144 large wealth managers. The unit of analysis is the mapped long US-equity holdings reported in each filing; the manager's full fund lies outside the filing.

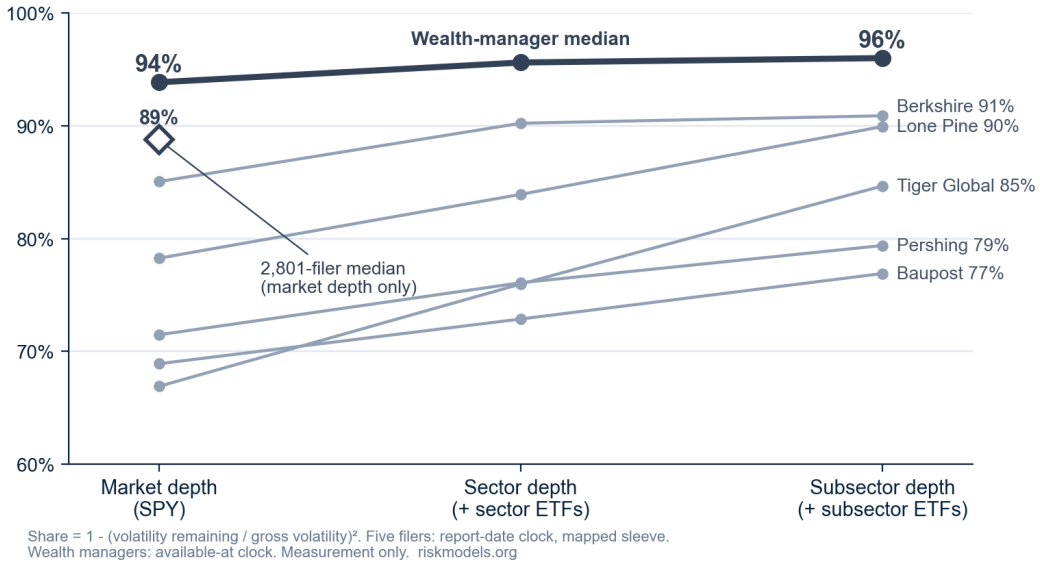


Figure 2. Share of each book's variance carried by its custom ETF benchmark at the market, sector and subsector depths: five named filers (report-date clock, May 2016 to April 2026), the median of 144 large wealth managers, and the median of 2,801 filers at market depth. The share is one minus the squared ratio of the volatility remaining after the benchmark to the gross volatility. The five filers are model-layer series, and the wealth managers are implemented series.

Tiger Global shows why benchmark depth matters. SPY alone carries 67% of its variance; sector ETFs raise the share to 76%, and subsector ETFs to 85%. The improvement comes from themes repeated across holdings: software through Microsoft, Adobe, Salesforce, ServiceNow and CrowdStrike; internet platforms through Meta, Alphabet, Sea and Reddit; retail and e-commerce through Amazon, JD.com and Carvana; and semiconductors through NVIDIA, Taiwan Semiconductor, Lam Research and Applied Materials. Berkshire begins at 85% with SPY and reaches 90% with sectors, but subsectors add only one point. Its concentrated book leaves less repeated within-sector structure for another ETF layer to capture.

RESULTS IN BRIEF

Ten years ending April 2026. Return statistics for the five named filers and the 2,801 filers are on the report-date clock. ETF positions, turnover and wealth-manager figures begin when each filing becomes public.

Most filed equity books can be closely replicated by a compact ETF benchmark built from their own holdings. A SPY-only benchmark carries about 89% of measured variance for the median of 2,801 filers. At subsector depth, custom benchmarks carry 77% to 91% for the five named filers and about 96% for the median large wealth manager. The five named benchmarks require only 1 to 29 distinct ETFs.

Benchmark-relative return is small beside gross return in much of the sample. The median filer returned 13.8% a year gross and -0.7% relative to the market-depth benchmark. Among 100 large filers with meaningful stock-specific risk, 92 to 97 realized manager-minus-benchmark returns remain within two standard errors of zero, depending on benchmark depth.

> "Factors are the language of investing that everyone should be speaking." - Andrew Ang, BlackRock

Here, small benchmark-relative returns show that market, sector and subsector exposures carried much of the measured return and most of the measured variance. The return left after those common exposures is the narrower claim on company selection, style and themes outside the traded hierarchy.

Depth has an implementation cost. Daily turnover rises as sector and subsector ETFs enter the benchmark. Monthly refresh cuts turnover to 34% to 45% of the daily rate and adds no more than 0.4 volatility points for any named filer at a fixed depth.

The filing delay has surprisingly little annual effect in most cases. Across 2,505 filers present on both clocks, market-relative returns have a rank correlation of 0.945 and a median annual change of -0.12 percentage points. Three of the five named books change by less than one point a year; Lone Pine and Tiger Global are the larger exceptions.

A 13F covers reported long US-equity holdings and omits shorts, many options, non-US positions and other assets. The results therefore describe the reported holdings rather than whole funds. The study ends before constructing an investable ETF-plus-cash return series after financing and actual trading costs.

RUN THE SAME MEASUREMENT ON A COVERED 13F FILER

Install with `pip install riskmodels-py`, then set `RISKMODELS_API_KEY` to your key from riskmodels.app/get-key.

```
from riskmodels import RiskModelsClient

client = RiskModelsClient.from_env()
filer = "BW-FILER-CIK0001336528"  # Pershing Square
snapshot = client.get_filer_snapshot(filer)
print(snapshot)
```

The snapshot supplies holdings and portfolio diagnostics; decomposition fields depend on what has been materialized. It does not return the complete ten-year series or promise ETF hedge ratios. Section 8 and the API reference describe the additional calls.

Two clocks appear throughout. The report-date clock treats the quarter-end portfolio as known on that date; it is a research measure. The publicly available clock begins on the first session after filing, when a follower could trade. It is called `available_at` in the data. All returns are gross of fees, trading costs and financing. A filing has no daily net asset value, leaving the filer's realized fund return unavailable for comparison. Persistence lies outside this study.

1. From one stock benchmark to a portfolio benchmark

RiskModels estimates each stock's benchmark hierarchically: SPY for market exposure, a sector ETF net of market, and a subsector ETF net of both broader layers. Every estimate uses information available before the return day. Reversing the hedge signs produces the stock's long ETF benchmark. For a portfolio, multiply those ETF weights by each stock's portfolio weight and add identical ETFs across positions; exposures shared by several stocks collapse into single book-level positions. The 13F analysis uses daily-drifted portfolio weights, but the construction applies to any mappable stock portfolio.

The paper reports three fixed depths: market (SPY), sector (SPY plus sector ETFs) and subsector (SPY plus sector and subsector ETFs). Ratios are re-estimated at every depth rather than appended mechanically. The model-selected specification, L*, assigns each position its selected depth. Benchmark-relative return is book return minus benchmark return. The volatility left after the benchmark is often called tracking error; the benchmark's variance share is one minus the squared ratio of that remaining volatility to gross volatility.

The hierarchy stops at the tradability frontier. Market, sector and subsector risks can be represented with liquid ETFs. Style cannot be added as a clean fourth layer: once value, growth, quality or momentum is conditioned within a subsector, the baskets become too sparse to support stable tradable benchmarks. Style is therefore split by construction. Some travels with the sector and subsector ETFs; the rest remains in the residual. This paper chooses an observable, reproducible benchmark over a more complete but synthetic style attribution.

ETF weights need not sum to one. The benchmark consists of the ETF mix plus a cash remainder and is levered when ETF weights exceed one; reported returns exclude financing. At market depth, the SPY weight is the book's SPY beta. For Pershing Square it averages 0.95, falling to 0.44 at sector depth and 0.33 at subsector depth as more specific ETFs absorb exposure. The hedge-ratio construction removes the broader exposure already embedded in each raw ETF.

Tables 1 and 2 apply position-level ratios on the report-date clock. Table 4 aggregates the same ratios into book-level ETF positions after each filing becomes public. Differences between the model components and aggregated ETF positions are implementation diagnostics and are reported outside the main argument.

The benchmark locates the return and risk carried by the filed book. The residual after subsector depth includes company-specific return, style and any themes the traded hierarchy does not capture; it is not, by itself, an estimate of manager skill.

2. For the median filer, the market-depth benchmark carries about 89% of variance

The analysis begins with daily holdings-based returns for 15,333 filers. After applying the history, coverage, breadth and sanity requirements, 2,801 have the ten-year histories reported in Table 1.

Item	Value
Filers with a daily return series in the store	15,333
With at least 80% of business days usable, May 2016 to April 2026	3,112
After mapped coverage of at least 80%, at least five names, and a sanity filter on return and volatility	2,801
Filers meeting the same rules for May 2021 to April 2026	4,705
Reported 13F value of the 2,801 (latest filing; includes overlaps between parents and sub-advisers)	$53.7 trillion; median $0.96 billion (10th to 90th percentile $0.19 to $14.5 billion)
Annualized gross return, median	13.8%
Annualized return from the book's SPY beta, median	14.9%
Return relative to the market-depth benchmark (SPY-hedged), median; 5th / 25th / 75th / 95th percentile	-0.7%; -6.5% / -2.2% / +0.4% / +3.0%
Median relative to the sector-depth benchmark; relative to the subsector-depth benchmark (stock-specific layer)	0.0%; +0.5%
Filers with a positive return relative to the market-depth benchmark	34%
Median share of layer variance: market / sector / subsector / stock-specific	90% / 2% / 1% / 6%
Median volatility, gross and relative to the market-depth benchmark	19.0% and 6.3%

Table 1. The 2,801 filers with ten-year histories, May 2016 to April 2026, report-date clock, model-layer series. Medians across filers unless stated; returns annualized. The SPY-hedged return is the return relative to the market-depth benchmark.

The median filer returned 13.8% a year gross. The return associated with the SPY betas of its positions was 14.9% a year, and the return relative to that market-depth benchmark was -0.7%. As separate cross-sectional medians, their sum has no interpretation. The positive relative-return share was 34%. Market exposure accounted for 90% of layer variance for the median filer, and benchmarking that exposure reduced volatility from 19.0% to 6.3%.

Requiring ten years creates survivorship bias by excluding filers that closed after 2016. Among filers observed in 2016-2021, those absent in 2021-2026 had a median first-window relative return of -0.61% a year, versus -0.37% for survivors. Their positive shares were 41% and 44%. The reported -0.7% median and 34% positive share therefore modestly overstate the full population. Appendix I compares gross and relative-return ranks.

3. Realized manager-minus-benchmark returns cluster near zero across 100 large filers with material stock-specific risk

The five named filers are recognizable examples, but they are too small a sample to establish the distribution. Figure 3 expands the test to the 100 largest ten-year filers whose stock-specific layer carries at least 10% of variance. The screen uses risk, not realized benchmark-relative return, and is an operational definition of active risk rather than a legal classification of the manager.

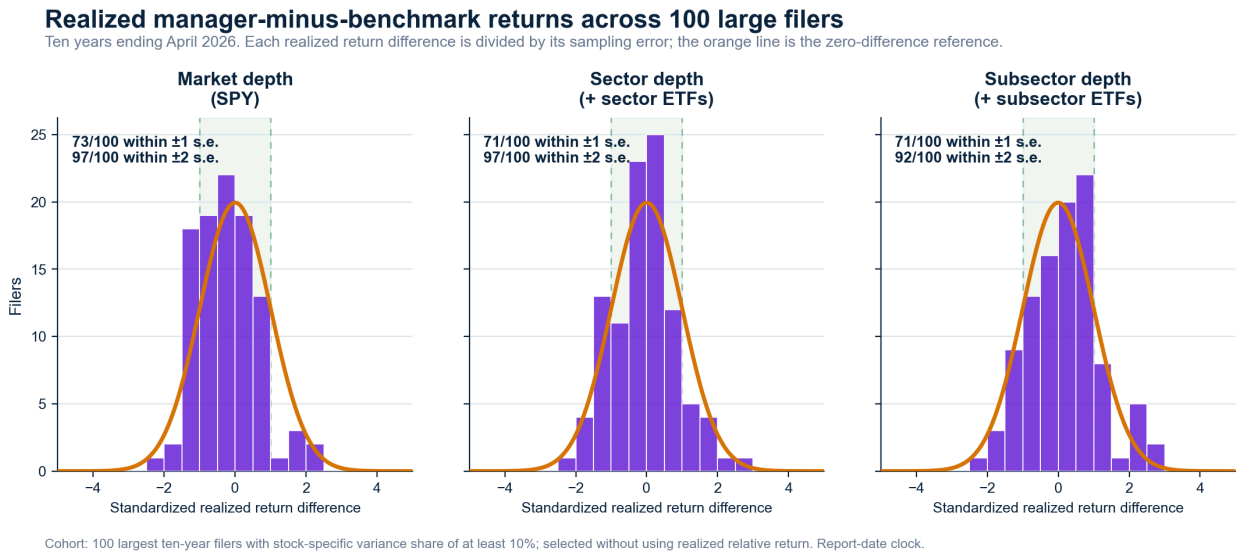


Figure 3. Realized ten-year benchmark-relative returns divided by their sampling errors for 100 large filers with material stock-specific risk, May 2016 to April 2026, report-date clock. The filers are the 100 largest by reported 13F assets among the ten-year sample with at least 10% stock-specific variance. The orange line is the marginal distribution expected if the true benchmark-relative return is zero; the three depths for a filer are not independent.

At market depth, 73 of 100 realized differences lie within one standard error of zero and 97 lie within two. The corresponding counts are 71 and 97 at sector depth, and 71 and 92 at subsector depth. The distribution shifts slightly as the benchmark becomes more specific, but most realized return differences remain small relative to their sampling error at every depth.

The five named filers show how individual books sit inside that broader result.

Pershing's return relative to its subsector-depth benchmark is +1.9% a year, with a standard error of 3.0 points

Filer	Gross	Relative to market-depth benchmark (SE)	Relative to sector-depth benchmark (SE)	Relative to sub-sector-depth benchmark, stock-specific (SE)	Volatility, gross to relative (market depth)	Max drawdown, gross / relative (market depth)	Rank, gross / relative (market depth)
Berkshire Hathaway	+13.5%	–1.3% (2.5)	–1.1% (2.0)	–1.4% (1.9)	20.4% to 7.9%	–39.3% / –20.7%	1,576 / 1,771
Lone Pine Capital	+16.9%	–0.2% (3.7)	–0.9% (3.2)	–0.3% (2.5)	25.0% to 11.7%	–52.9% / –43.0%	368 / 1,095
Tiger Global	+18.9%	+1.6% (5.1)	+0.6% (4.4)	+2.2% (3.5)	28.2% to 16.2%	–65.0% / –59.7%	175 / 310
Pershing Square	+15.1%	+0.1% (3.6)	–0.1% (3.3)	+1.9% (3.0)	21.0% to 11.2%	–45.5% / –30.1%	782 / 912
Baupost Group	+5.1%	–9.9% (4.0)	–9.1% (3.8)	–7.4% (3.5)	22.8% to 12.7%	–47.7% / –68.6%	2,743 / 2,741

Table 2. Five filers, annualized, May 2016 to April 2026, report-date clock, model-layer series, mapped long 13F holdings only. Each return column is the return relative to the benchmark at that depth. Mean mapped weight coverage is 98% for Berkshire, 97% for Lone Pine, 96% for Pershing Square, 91% for Tiger Global and 84% for Baupost (Appendix F). Standard error is the volatility of each relative-return series divided by the square root of ten years. Ranks are among the 2,801 filers.

Depth determines which exposures enter the ETF mix and which remain. Pershing Square realized +0.1% a year relative to the market benchmark, -0.1% relative to sector and +1.9% relative to subsector. With annual relative volatility near 10%, the ten-year standard errors are about three points. None of Pershing's realized differences is distinguishable from zero or from the others. The same is true for Tiger Global (+1.6%, +0.6% and +2.2%) and Berkshire Hathaway (-1.3%, -1.1% and -1.4%).

Figure 1's 307% cumulative gross return and 0.8% cumulative market-relative return answer different questions. The first reconstructs the reported long-equity holdings. The second shows that a changing SPY mix, sized from those holdings' measured market exposure, earned almost the same daily return. Neither series is the fund's net asset value: the filing omits shorts, many options, non-US positions and other assets.

Baupost is the only named filer more than two standard errors from zero: -9.9% a year relative to market, with a 4.0-point standard error, and -7.4% relative to subsector, with a 3.5-point standard error. The mapped holdings cover 84% of filed value; the rest of Baupost's portfolio lies outside the analysis.

For an allocator, the comparison separates common thematic exposures from the return left for company selection and unmodeled risks. Managers who own different stocks may still share those exposures. Section 6 then tests how much return survives the wait for a public filing.

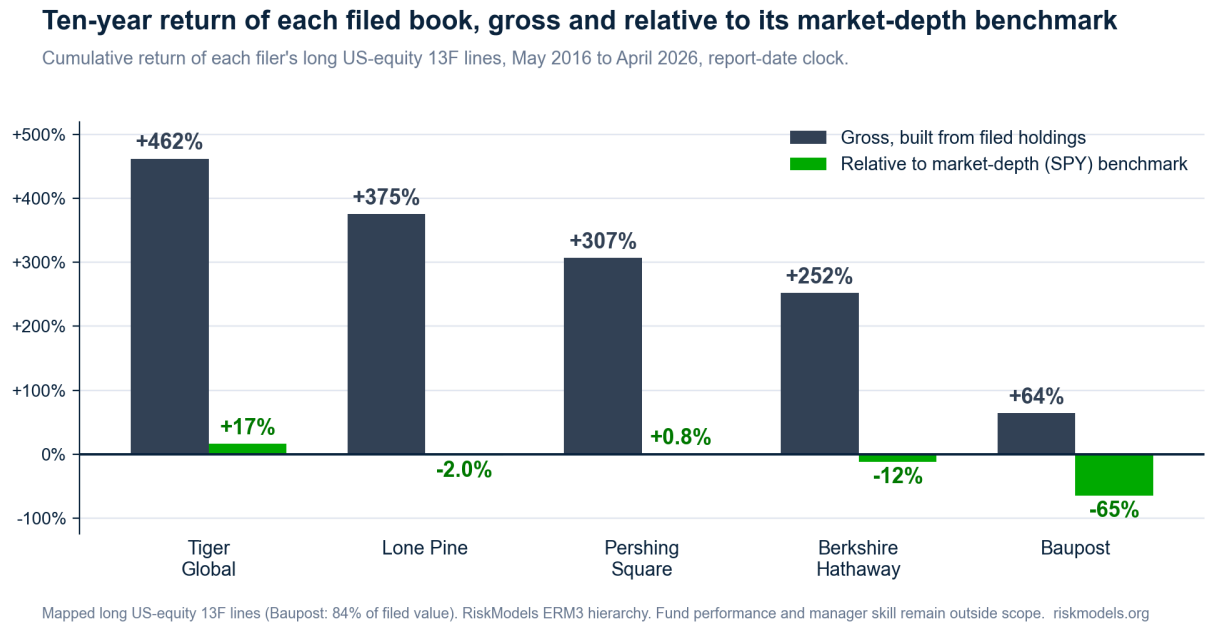


Figure 4. Cumulative return of each filer's long US-equity 13F lines, gross and relative to its market-depth benchmark, May 2016 to April 2026, report-date clock. These returns come from filed holdings rather than fund net asset values: the 13F excludes non-equity positions, and Baupost's mapped coverage is 84% of filed value.

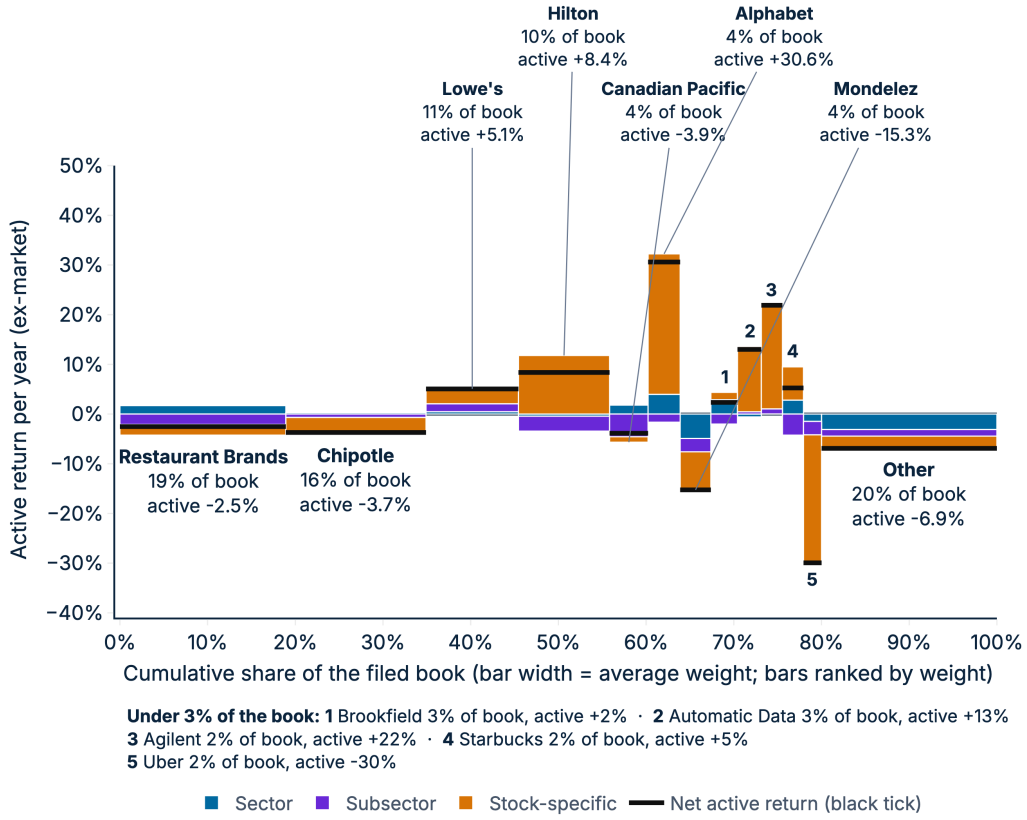


Figure 5. Pershing Square by stock, market layer removed, May 2016 to April 2026. Bar width is the stock's average weight in the whole mapped book, height is its annualized return while held split into sector, subsector and stock-specific layers, and the black tick is the net. Bars are ordered by weight; issuers under 1.5% of the book are pooled as Other holdings.

The stock cascade locates Pershing's remaining return. Over ten years, Hilton Worldwide contributes +28.0 points, Chipotle +20.5 and Alphabet +28.2 across its two listings. Uber contributes -15.7, Howard Hughes -22.0 across two entities and Netflix -9.3. The subsector layer contributes -40.0 points, including -19.5 from the leisure and entertainment cohort containing Hilton, Restaurant Brands and Chipotle.

Pershing's realized difference changes sign across depths, yet all three values remain within one standard error of zero. Appendix D reports Berkshire Hathaway's sector cascade: technology contributes about 5 points a year of ex-market return, while consumer staples and financials each contribute about -4 points.

4. For 144 large wealth managers the subsector-depth benchmark carries about 96% of variance

The wealth-manager books are broader and more diversified. The median adviser reports 683 13F lines. Directly held model stocks cover 39% of filed value at the median and 62% after looking through funds and ETFs. Results refer only to those covered long US-equity holdings. Of the 144 advisers with at least 500 return days, 74 have the full May 2016 to April 2026 history; the shortest history contains 553 days. Appendix G defines the Form ADV cohort and reports look-through, concentration and style.

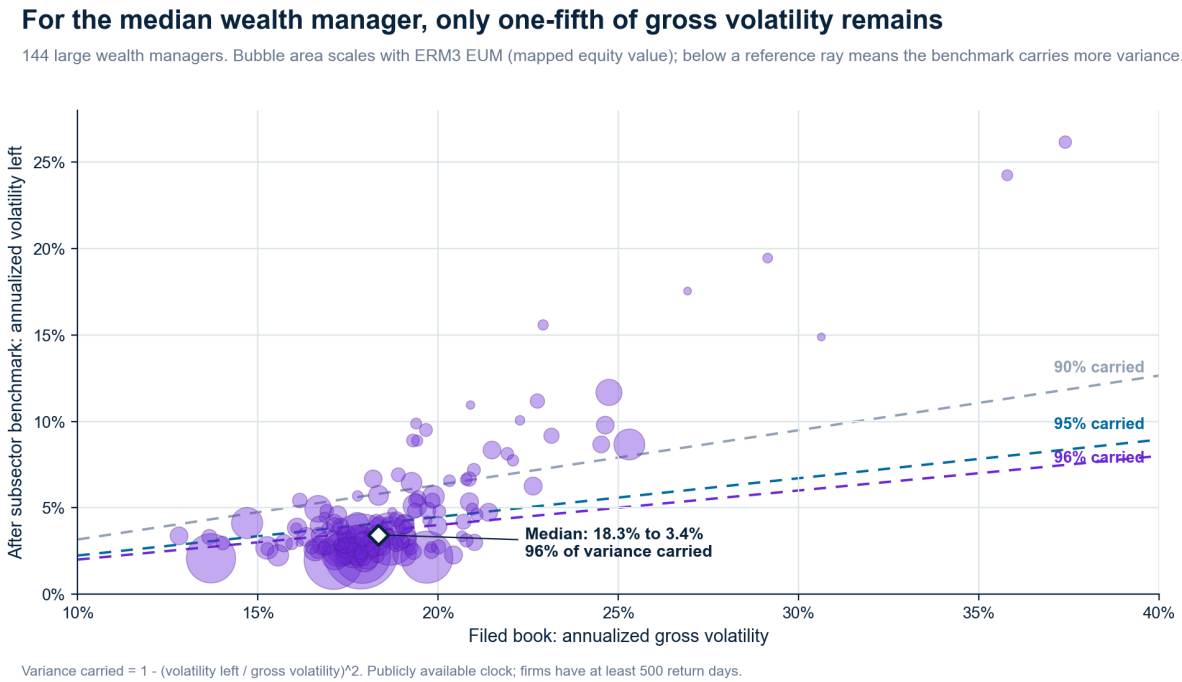


Figure 6. Annualized gross volatility against volatility left after the subsector-depth benchmark, daily refresh, publicly available clock, 144 advisers with at least 500 return days. Bubble area scales with ERM3 EUM (mapped equity value). The reference rays show where the benchmark carries 90%, 95% and 96% of variance; points below a ray have a higher benchmark share.

Median annualized gross volatility is 18.3%, with quartiles of 17.3% and 19.8%. Market depth leaves 4.4% volatility and 6.1% of variance at the median of the per-firm ratios, so the benchmark carries 94%. Sector depth leaves 3.7% volatility and 4.4% of variance. Subsector depth leaves 3.4% volatility and 4.0% of variance: a 96% benchmark share. The scatter shows that this is not only a median result: 114 of 144 managers have at least 90% of variance carried by the benchmark, and 84 have at least 95%. The remainder includes company-specific return, style and themes outside the traded hierarchy.

The median adviser holds ten of eleven sectors at weights of at least 1%, eight at weights of at least 3%, and 36% in the largest sector; the quartiles are 31% and 41%. Sector depth reduces median volatility by 0.8 points beyond SPY. Subsector depth removes another 0.2 points, based on unrounded medians of 4.45%, 3.65% and 3.43%. The causal link between sector breadth and the smaller gain at the final depth remains untested.

Benchmark depth	ETFs	Volatility left, % (daily)	Relative return, %	Turnover D / W / M	Cost, bp of book: daily / monthly
Filed book	-	18.3% (17.3, 19.8)	13.5% (12.3, 15.1)	-	-
Market depth (SPY)	1 (1, 1)	4.4% (3.2, 6.9)	-0.6% (-1.6, 0.4)	0.9/0.6/0.3	0.9/0.3
Sector depth	12 (11, 12)	3.7% (2.6, 5.6)	0.5% (-0.4, 1.2)	3.4/2.1/1.2	3.4/1.2
Subsector depth	33 (29, 35)	3.4% (2.7, 5.4)	1.0% (0.2, 1.6)	6.6/4.6/2.4	6.6/2.4
Model-selected level (L*)	19 (16, 21)	4.0% (3.1, 5.8)	0.8% (-0.1, 1.4)	7.3/6.1/2.6	7.3/2.6

Table 3. Cohort medians with 25th and 75th percentiles in parentheses, 144 advisers with at least 500 return days, publicly available clock, before financing and differences between ETF and index returns. Turnover in multiples of book per year; cost at 1 basis point per unit traded.

The median annualized gross return is 13.5%. Median return relative to the benchmark is -0.6% at market depth, +0.5% at sector depth, +1.0% at subsector depth and +0.8% at the model-selected depth. The interquartile range is 1.4 to 2.0 points wide at each depth. The 1.6-point span between the lowest and highest depth medians is therefore similar to the cross-sectional spread between advisers at a single depth.

For an RIA, the practical result is a benchmark of about 33 ETFs for a covered client book with hundreds of lines. It provides a common basis for comparing accounts and explaining their exposures within an existing reporting system. The results cover each adviser's available history and precede costs, financing and differences between ETF and index returns.

5. The benchmark is compact, but it is not static

At portfolio level, the same ETFs recur across many holdings. Combining those exposures turns the stock-by-stock benchmark into a short ETF list. Market depth requires only SPY. Sector depth requires 7 to 11 ETFs for the five named filers and 12 for the median wealth manager. Subsector depth requires 16 to 29 for the named filers and 33 for the median wealth manager. A portfolio with many stocks can therefore be benchmarked with a few dozen traded instruments.

Table 4 reports the ETF count, volatility left after the benchmark, relative return, turnover and illustrative cost for the five named filers on the publicly available clock. Table 3 gives the corresponding wealth-manager results. An ETF counts when its mean absolute weight over the measured history exceeds 0.5%.

Filer	Benchmark depth	ETFs	Volatility left (D)	Relative return	Turnover (D / M)	Cost, bp (daily)
Pershing Square	Filed book	-	20.9%	+15.0%	-	-
	Market depth (SPY)	1	11.2%	0.0%	1.2 / 0.5	1.2
	Sector depth	9	10.3%	+0.5%	4.7 / 2.1	4.7
	Subsector depth	16	9.5%	+2.3%	7.9 / 3.4	7.9
	Model-selected level (L*)	14	9.7%	+1.3%	8.8 / 3.3	8.8
Berkshire Hathaway	Filed book	-	20.5%	+13.1%	-	-
	Market depth (SPY)	1	7.9%	−1.6%	1.0 / 0.4	1.0
	Sector depth	9	6.4%	−1.1%	4.6 / 1.6	4.6
	Subsector depth	21	6.2%	−1.4%	8.8 / 3.3	8.8
	Model-selected level (L*)	11	7.0%	−1.2%	9.8 / 4.0	9.8
Tiger Global	Filed book	-	28.0%	+16.0%	-	-
	Market depth (SPY)	1	16.4%	−0.8%	1.9 / 0.8	1.9
	Sector depth	7	14.3%	−0.7%	6.3 / 2.2	6.3
	Subsector depth	21	11.3%	+1.1%	10.7 / 4.1	10.7
	Model-selected level (L*)	15	12.8%	+0.3%	12.5 / 4.3	12.5
Lone Pine	Filed book	-	24.8%	+15.1%	-	-
	Market depth (SPY)	1	11.9%	−1.5%	1.5 / 0.5	1.5
	Sector depth	11	10.4%	−0.9%	5.7 / 2.1	5.7
	Subsector depth	29	8.1%	+0.6%	10.7 / 4.4	10.7
	Model-selected level (L*)	20	9.1%	0.0%	12.1 / 4.3	12.1
Baupost	Filed book	-	21.7%	+4.8%	-	-
	Market depth (SPY)	1	12.3%	−9.0%	1.6 / 0.6	1.6
	Sector depth	11	12.2%	−5.8%	5.2 / 1.9	5.2
	Subsector depth	27	12.0%	−4.2%	13.5 / 5.3	13.5
	Model-selected level (L*)	17	12.2%	−5.1%	7.8 / 2.7	7.8

Table 4. ETF positions that implement each benchmark depth for five filers, aggregated from per-stock ratios, filings credited from the first session after the filing date (publicly available clock), May 2016 to April 2026. Return is annualized and gross of financing and differences between ETF and index returns. Turnover is in multiples of book per year; cost is turnover at 1 basis point per unit traded.

More detail improves the match, but it also increases trading. With daily refresh, the subsector benchmark turns over 7.9 to 13.5 times the book each year for the five named filers and 6.6 times for the median wealth manager. At the paper's illustrative cost of one basis point per unit traded, those turnover figures translate to 7.9 to 13.5 basis points a year for the named filers and 6.6 basis

points for the median wealth manager. These estimates exclude bid-ask spreads, financing and differences between ETF and index returns.

Pershing Square shows what the extra detail buys. Volatility left after the benchmark falls from 11.2% with SPY alone to 9.5% with the subsector benchmark. Its annual return relative to the benchmark changes from 0.0% at market depth to +0.5% at sector depth and +2.3% at subsector depth. Each result remains inside its roughly three-point standard error. The deeper benchmark changes the explanation of Pershing's return, but not the statistical conclusion.

L* makes a different choice for each stock: it stops at market, sector or subsector depth according to the model's selected level. That reduces the median wealth manager's ETF list from 33 to 19 and the named filers' lists from 16-29 to 11-20. It does not reduce trading. At the median, L* turns over 7.3 times the book with daily refresh, compared with 6.6 times for the fixed subsector benchmark, and leaves 4.0% rather than 3.4% volatility unexplained. L* is therefore a smaller but less stable ETF mix, not a uniformly cheaper one.

Monthly refresh preserves the match while cutting ETF turnover by about two-thirds

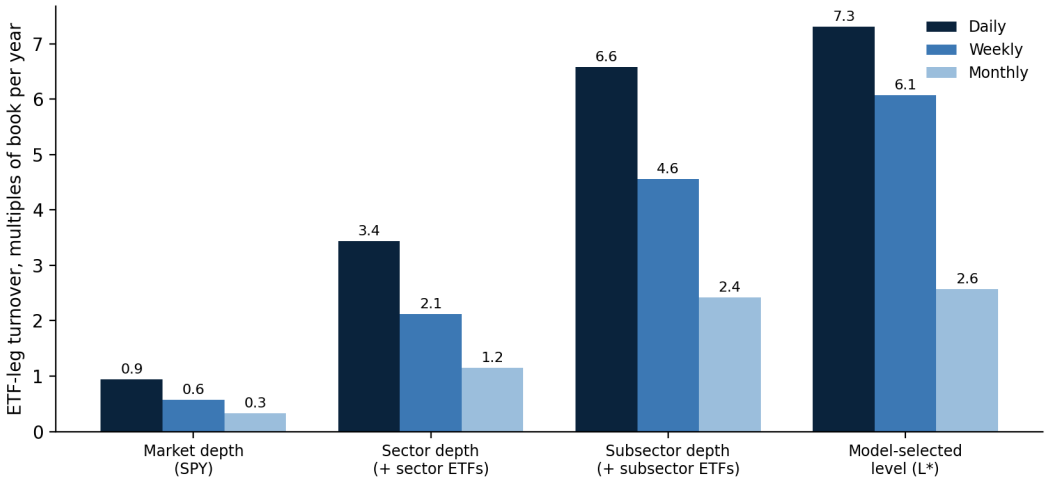


Figure 7. Median turnover of the ETF benchmark, in multiples of book per year, at daily, weekly and monthly refresh with weights frozen between refreshes. The median ETF counts are 1 at market depth, 12 at sector depth, 33 at fixed subsector depth and 19 for L. Fewer ETFs does not imply lower turnover. 144 wealth managers.*

Monthly refresh reduces median subsector turnover from 6.6 to 2.4 times the book and L* turnover from 7.3 to 2.6. The median increase in unexplained volatility is only 0.05 to 0.14 percentage points across depths; for every named filer it is no more than 0.4 points. Monthly refresh captures most of the benchmark's explanatory value with roughly one-third of the daily trading.

The practical point is simple. Shared ETF exposures combine across holdings into a short benchmark portfolio. The return left after that benchmark belongs to the individual stocks, style and themes outside the hierarchy rather than to the market, sector and subsector exposures they share.

6. Public filings preserve most of the reported book's return

A 13F becomes public as much as 45 days after quarter-end. To test whether that delay defeats a holdings-based strategy, the analysis runs each reported book on two clocks. The first applies a filing at quarter-end, before an investor could know it. The second begins on the first trading session after the filing becomes public. The difference is the return effect of waiting for information an investor could actually use.

For the 2,505 filers present on both clocks, market-relative annual returns have a rank correlation of 0.945. The median change is -0.12 percentage points a year. Waiting for the public filing therefore usually preserves both the manager's relative standing and most of the return measured from the reported long book.

The five named filers show the range of outcomes. Berkshire Hathaway changes from 13.46% to 12.54% a year, Pershing Square from 15.10% to 15.95%, and Baupost from 5.10% to 5.58%. Lone Pine falls by 1.60 points a year and Tiger Global by 1.80 points. Those two are material exceptions, but even there the public holdings retain most of the book's annual return.

This makes the public filing useful for following the disclosed portfolio. It does not reproduce the manager's fund return: 13Fs omit shorts, cash, non-US securities and most derivatives, and the calculation is gross of fees and trading costs. The narrower conclusion is still important. For most filers, an investor who waits until the holdings are public can follow the reported long US-equity book without losing its basic return profile. The stock-specific risk share moves by no more than two points for the five named filers, a secondary confirmation that the delay does not materially alter the measured structure.

7. Scope, limitations, and next tests

The sample covers long US-equity 13F lines that map to the model universe from May 2016 to April 2026. Returns are gross of fees, financing and differences between ETF and index returns. Table 2 standard errors assume independent daily relative returns. The cost calculation charges every ETF 1 basis point per unit traded, including thin subsector funds, so the cost columns simply restate turnover under that assumption.

The wealth-manager cohort relies on self-reported Form ADV fields and excludes firms with blank CIKs. Direct model stocks cover a median 39% of filed value and 62% after fund look-through. Every wealth-manager return and volatility statistic refers to that covered portion of the reported holdings.

Filer record for Pershing Square. Pershing Square Capital Management (CIK 0001336528) is flagged `stopped_filing` in the filer master because its latest filing is for 31 March 2026, one quarter before the latest quarter in the store. A second filer, Pershing Square Inc. (CIK 0002026053), has filings for 30 June 2025 through 30 June 2026, so the two filed in parallel for four quarters. The filer master lacks a succession record, this analysis leaves the relationship unresolved, and the ten-year series covers the capital-management filings only. EDGAR (checked 14 November 2026) lists Pershing Square Inc. under that name since 23 April 2026, formerly Pershing Square Holdco, L.P., and shows a 13F notice (form 13F-NT) rather than a holdings report from Pershing Square Capital Management for 30 June 2026.

The next tests are:

1. The return and volatility of the long ETF mix plus its cash remainder on its own, after each filing becomes public, at daily and monthly refresh, for the five filers and the 144 advisers.
2. Placebos for the depth results: a benchmark built with each filer's own-history OLS beta and one built with a beta of one to SPY.
3. Financing and ETF bid-ask costs in the leg-cost estimate, using the ETFs' own spreads.
4. The same table for 20 to 30 additional named filers, extending the named evidence beyond the five cases reported here.

8. How a reader runs the same measurement

RiskModels.org is the open research and methodology reference: the ERM3 specification, published papers and visual dossiers. RiskModels.app is the production API engine. Install the Python SDK with `pip install riskmodels-py`, use REST or the CLI, or connect an agent through MCP. Self-service access is billed per successful call, with $20 in free starter credits after card setup and no subscription or seat fees; current terms are on the pricing page.

An allocator can compare managers' common thematic exposures, the return left after those exposures, and the return lost while waiting for public filings. An RIA can construct a benchmark for a client book with hundreds of lines: the median wealth-manager benchmark here uses 33 ETFs at subsector depth. The same measurement can sit in an existing reporting stack. Developers can use the structured, point-in-time model output as a deterministic quantitative input to code or an agent, retaining the dates, coverage and model version returned with each response.

RiskModels provides portfolio measurement rather than order execution. For a stock, `get_hedge_basket` returns the L* ETF legs as signed ratios per $1 of stock. `get_alpha`, with `depth` set to `L1`, `L2` or `L3`, returns fixed-depth legs. For a list of positions, `analyze_portfolio` returns layer attribution and portfolio hedge ratios. Weight the stock-level legs and aggregate them by ETF to form the book-level benchmark.

Select a covered 13F filer through the RiskModels API to retrieve its disclosed holdings and the returns calculated from those holdings. `get_filer_holdings` returns the latest-quarter holdings and `get_filer_concentration` returns concentration measures. Return attribution accepts portfolios of up to 100 stocks, and longer return histories can be assembled through the API. The results can be returned directly to an agent or integrated into a research or production code stack.

Coverage includes daily holdings-based return series for 15,333 filers keyed by `bw_filer_id`, which embeds the CIK. Parent filers, sub-advisers and funds of funds can overlap. No single call returns the complete ten-year series. The API reference is at riskmodels.app/docs/api; keys are issued at riskmodels.app/get-key.

Appendix A. The decomposition

Every long-equity portfolio return decomposes into four layers in the ERM3 hierarchy, orthogonal by construction at the stock level:

- **Market beta (layer 1)**: exposure to the market factor (SPY).
- **Sector (layer 2)**: sector tilts relative to that market factor.
- **Subsector (layer 3)**: finer subsector positioning within sectors.
- **Stock-specific (layer 4)**: company-level return remaining after the three structural layers above.

For stock *i* on day *t*, the four layers add to the stock's return:

$$r_{i,t} = r_{i,t}^{\mathrm{mkt}} + r_{i,t}^{\mathrm{sec}} + r_{i,t}^{\mathrm{sub}} + r_{i,t}^{\mathrm{spec}} \tag{1}$$

> **What the stock-specific layer holds.** The hierarchy stops at the subsector layer on purpose: the market, sector and subsector layers each map to a tradeable ETF, so each can be hedged with an instrument. Style tilts (growth, value, momentum, quality) stay in the stock-specific residual unless the stock's subsector ETF carries them. A negative stock-specific contribution is therefore the return left after market, sector and subsector, combining name selection with unmodeled style exposure rather than isolating pure stock picking.

The layers group into three buckets:

Table A1. Layers and risk buckets in the ERM3 hierarchy.

Bucket	ERM3 layers	What it captures
Market risk	Layer 1	Broad market beta, excluded from active risk in this paper (Appendix C)
Thematic risk	Layers 2 + 3	Sector tilts, subsector concentration, structural macro-expression through sectors and subsectors
Stock-specific risk	Layer 4	Name-level selection — what remains after market and thematic structure are stripped out

Total portfolio risk partitions into **market risk** plus **active risk**. Active risk separates into **thematic risk** and **stock-specific risk**. Sector and subsector tilts express thematic positioning, distinct from passive drift and stock-specific skill.

Within the ERM3 hierarchy, market beta is the first layer. The remaining non-market exposures form the portfolio's **active structure**, which separates into **thematic positioning** (sector and subsector tilts) and **stock-specific** company selection.

The first three ERM3 layers are tradeable: each can be represented with an index, sector or subsector exchange-traded fund (ETF) whose estimated turnover and cost are reported, which is how the model's hedge ratios map to tradable instruments. The stock-specific layer is the return remaining after the three modeled ETF layers. It is a residual of the model and carries no verdict on skill.

Throughout the paper, **share of portfolio risk** means each layer's share of the sum of the four layers' return variances: how much of the portfolio's month-to-month volatility came from each layer of the decomposition. Appendix F describes the same object as a variance partition. The betas behind the layers are estimated on an earlier sample, applied to the period that follows, then re-estimated each cycle, so the ten-year history is a stitch of out-of-sample segments, as an investor would have seen it at the time. A whole-window fit would introduce future information; the rolling cycle keeps the tool usable in real time. For a single stock, the four layers are orthogonal in the window where its betas are fitted, so its variance splits into four shares that add to 100% there. A portfolio combines stocks with different sector and subsector factors, so some correlation can remain between its layers even then. The

four layer returns add to the gross return (identity residual below 2e-5), but they are only approximately uncorrelated: across the five managers the sum of the four layer variances is between 83% and 111% of the variance of the gross return (monthly returns, May 2016 to April 2026), above 100% for two managers and below for three.

With $R_{k,t}$ the portfolio's return in layer *k*, built from the weights that have drifted since the last filing, the layers, the share of risk and the variance identity are:

$$R_{k,t} = \sum_i w_{i,t}\, r^k_{i,t}, \qquad R_t = \sum_k R_{k,t}, \qquad k \in \{\mathrm{mkt}, \mathrm{sec}, \mathrm{sub}, \mathrm{spec}\} \tag{2}$$

$$\begin{aligned} s_k &= \frac{\mathrm{Var}(R_k)}{\sum_j \mathrm{Var}(R_j)} \\ \mathrm{Var}(R) &= \sum_k \mathrm{Var}(R_k) + 2\sum_{j<k} \mathrm{Cov}(R_j, R_k) \end{aligned} \tag{3}$$

Each layer is read two ways, and the model is symmetric across them: the same four layers partition a portfolio's return and partition its risk.

a) **Return attribution** (Appendix C, Appendix E): *how did each layer contribute to the portfolio's cumulative return?* Contributions are measured in percent of starting net asset value (NAV) and add to the compound gross return over the window. After the market layer, the focus is thematic against stock-specific contribution. A book can be market-dominated in risk and still show a large stock-specific contribution over ten years. Lone Pine is the clearest example in this cohort.

b) **Risk attribution** (Appendix B, Appendix E): *how is total portfolio risk distributed?* Each month we attribute the portfolio's return variance across the four layers. The four shares **sum to 100%** for each manager and partition **total portfolio risk**. Example: Berkshire's **80%** market risk share means roughly four-fifths of its monthly portfolio variance came from broad market exposure, and its **14%** stock-specific share means the name-level sleeve explained about a seventh of the same total.

Together the two views describe how each book is built.

To keep cross-manager comparisons clean, **every table and chart uses the same calendar window: May 2016 through April 2026** (about ten years). That is the intersection panel where all five filers have continuous mapped history with the same start and end dates. Each filer's return series begins with its first observable month.

Appendix B. Risk shares of the five books

How is portfolio risk distributed? Figure B2 partitions each book's total portfolio risk, the month-to-month volatility of the long book, across market beta, sector, subsector and stock-specific sources. Bar length is the book's annualized volatility; the segments are scaled by each layer's share of risk. The four shares sum to **100%**. Managers are ordered by stock-specific share, lowest first, and the text below follows the same order.

Market beta is at least 65% of every book's risk; stock-specific risk runs from 14% to 29%

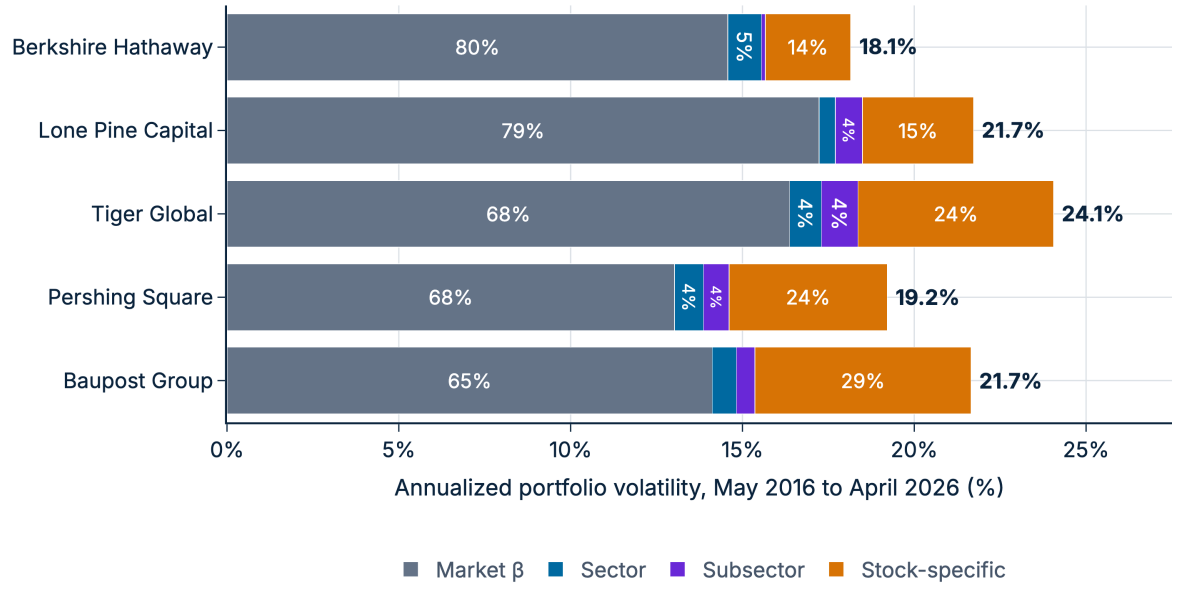


Figure B2. Risk attribution, May 2016 to April 2026 (120 months), holdings as reported, buy-and-hold weights. Bar length is annualized portfolio volatility from monthly returns built from the daily buy-and-hold series; segment length is that volatility times the layer's share of variance. Stock-specific is the residual after the subsector layer and includes style. Labels omit segments under 4% (Berkshire's subsector layer is under 1%, Lone Pine's sector layer about 2%, Baupost's sector and subsector layers 2% to 3%), so a bar's labels can sum to 98% or 99% by rounding. Shares are point estimates from 120 monthly observations and carry sampling error. ERM3 fit as of April 2026.

The bars describe portfolio risk structure. Total active risk and skill ranking lie outside their scope. Market risk dominates the partition because broad market exposure moved the book most often over the window; thematic tilts and name selection remain economically meaningful.

Reading across, in the order of Figure B2:

Berkshire carries the lowest stock-specific share (**14%**) on an annualized volatility of 18.1%. Market beta is **80%** of portfolio risk and thematic risk (sector plus subsector) is **6%**. The concentrated positions (Apple, Coca-Cola, American Express, Bank of America) are very large, market-correlated companies, so the market layer dominates and the stock-specific share sits well below Tiger, Pershing or Baupost.

Lone Pine puts **79%** of portfolio risk in market beta and **15%** in stock-specific risk, on a volatility of 21.7%. The book holds a median of about 30 mapped names at quarter-end and co-moves heavily with the broad market over the ten-year window. Thematic risk is **6%**.

Tiger Global puts **24%** in stock-specific risk, **8%** in thematic and **68%** in market beta, on the highest volatility in the group (24.1%). The stock-specific share is about 1.7 times Berkshire's, on a faster-moving book of internet and platform names.

Pershing Square puts **24%** in stock-specific risk, **68%** in market beta and **8%** in thematic risk, on a volatility of 19.2%. Its stock-specific share exceeds Tiger's by about three tenths of a point (23.9% against 23.7%); the two are tied within the precision of the estimate.

Baupost carries the highest stock-specific share (**29%**) and the lowest market share in the cohort (**65%**), on a volatility of 21.7%. This section describes how risk was structured over the window; Appendix E shows how much the stock-specific slice moves from one period to the next.

How those structures evolve, and how many names carry the book, is where stock-specific intensity, concentration, turnover and filing lag enter (Appendix E).

Appendix C. Return attribution and the arithmetic of contributions

Appendix B asked how portfolio risk is distributed across layers. This section asks how each layer contributed to return, keeping thematic positioning and stock-specific positioning separate throughout.

Market risk dominates total risk and can be trimmed actively. This paper excludes it from active risk to isolate thematic and stock-specific risk. Active structure therefore begins after the market layer. Within it, sector and subsector tilts are thematic positioning and the name-level sleeve is stock-specific positioning. Two managers can deliver similar realized returns while leaning on thematic structure, on stock-specific selection, or on different mixes of the two. The decomposition shows that split instead of collapsing active structure into one number.

How the returns are built. For the five filers, the weights drift with prices, as in a book held without trading, and the daily portfolio returns are linked geometrically. This is the holdings-based, buy-and-hold, time-weighted construction of the mutual fund literature (Grinblatt and Titman 1989; Daniel, Grinblatt, Titman and Wermers 1997), and geometric linking is what the Global Investment Performance Standards (GIPS) require of a time-weighted return. Each layer's daily return is weighted by the portfolio's net asset value at the start of that day relative to its starting value and summed over the window, so the four layers add exactly to the compounded gross return, with an identity residual below 2e-5. With no benchmark return, this weighting corresponds to Frongello's (2002) recursive linking. Cariño (1999) and Menchero (2000) reproduce the same compounded total with different period coefficients, which moves the cross-term between layers. The construction follows the arithmetic that performance-attribution systems apply to holdings; reconciliation against a particular vendor's output remains outside this study. In symbols, with NAV the net asset value of the book:

$$C_k = \sum_{t=1}^{T} \frac{\mathrm{NAV}_{t-1}}{\mathrm{NAV}_0}\, R_{k,t}, \qquad \mathrm{NAV}_t = \mathrm{NAV}_{t-1}\,(1 + R_t)$$
$$\sum_k C_k = \frac{\mathrm{NAV}_T}{\mathrm{NAV}_0} - 1 \tag{4}$$

Peak-to-trough losses of the mapped books: Berkshire 39% (February to March 2020), Pershing Square 46% (February to March 2020), Lone Pine 53% (September 2021 to June 2022), Tiger Global 65% (February 2021 to November 2022) and Baupost 48% (January 2018 to March 2020).

Holding the filed weights constant (daily rebalancing back to the filed book) instead of letting them drift raises the ten-year total for Lone Pine (+413.6% against +375.5%), Tiger Global (+531.1% against +462.2%) and Baupost (+75.0% against +64.2%) and lowers it for Berkshire (+243.9% against +252.3%) and Pershing Square (+293.4% against +306.8%). We use drift because it is the book a holder of the filing would have carried. The layers below are reported on that basis only.

Table C2. Cumulative gross return and layer contributions, May 2016 to April 2026, in % of starting NAV. Managers are ordered by stock-specific share of risk.

Manager	Cumulative gross	Return contribution by layer			
		Market beta	Thematic risk		Stock-specific
			Sector	Subsector	
Berkshire Hathaway	**+252.3%**	+260.1%	+2.3%	+9.0%	**–19.1%**
Lone Pine Capital	**+375.5%**	+452.5%	+13.6%	–26.5%	**–64.1%**
Tiger Global Management	**+462.2%**	+560.2%	+22.4%	–81.8%	**–38.6%**
Pershing Square	**+306.8%**	+324.7%	–0.4%	–40.0%	**+22.5%**
Baupost Group	**+64.2%**	+175.4%	–8.6%	–24.8%	**–77.8%**

Sector and subsector are the thematic layers; stock-specific is the name-level layer. Market beta is exposure to the market factor and sits outside active positioning in this paper.

The market layer is the largest contribution for all five books. That reflects ten years of exposure to the U.S. equity market and says nothing about skill. The three layers below it are where the books differ.

Berkshire's market exposure delivered +260%; Kraft Heinz alone took back 19 points of stock-specific return

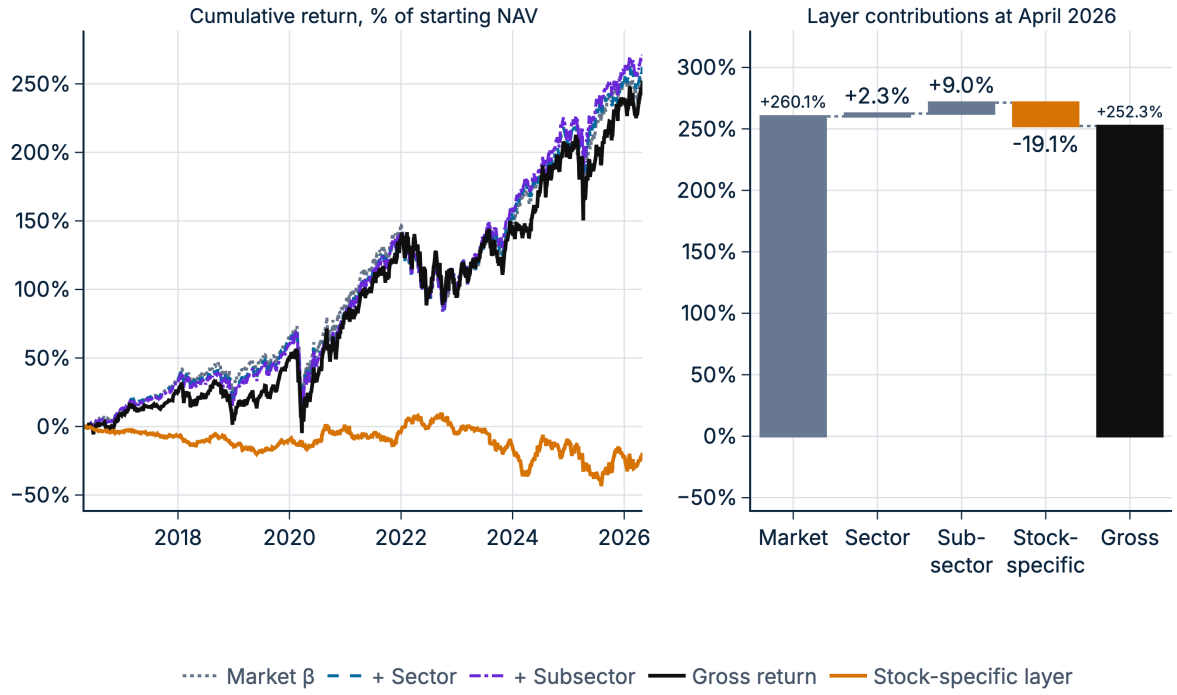


Figure C3. Berkshire Hathaway, cumulative return and composition, May 2016 to April 2026. The same panel for Pershing Square, Lone Pine, Tiger Global and Baupost is at riskmodels.org.

Table C3. Berkshire Hathaway: contribution to cumulative return by position, May 2016 to April 2026, in % of starting NAV. The bar in the last column is scaled to Berkshire's total return of +252.3% over the same window as Figure C3 (a full-width bar would be the whole book). Red bars are negative.

Position	Market	Sector	Subsector	Stock-specific	Total	Share of the +252.3% return
Apple	+98.2	+13.0	+14.1	+10.6	**+135.8**	
American Express	+35.2	–5.7	+2.0	+3.5	**+35.1**	
Bank of America	+31.7	–5.1	–1.8	+0.6	**+25.4**	
Coca-Cola	+10.5	+0.8	–0.2	+6.3	**+17.4**	
Chevron	+6.3	+7.1	+0.3	–1.5	**+12.2**	
Alphabet	+2.7	–0.4	0.0	+5.2	**+7.5**	
Moody's	+9.1	–0.9	+0.4	–1.7	**+6.9**	
Chubb	+1.0	–0.6	0.0	+2.9	**+3.2**	
Delta Air Lines	+1.5	–0.5	–0.4	–1.3	**–0.7**	
Wells Fargo	+11.3	–3.1	+0.1	–9.9	**–1.7**	
Kraft Heinz	+9.3	–1.0	–1.1	–18.7	**–11.5**	
All other positions (159)	+43.5	–1.2	–4.4	–15.0	**+22.8**	
Total, all positions	**+260.1**	**+2.3**	**+9.0**	**–19.1**	**+252.3**	

Apple alone supplied 136 of the book's 252 points, 54% of the ten-year return, and 98 of those 136 points are market-layer return. American Express and Bank of America added another 60 between them. Kraft Heinz is the only position that lost more than two points.

How to read Figure C3. The left panel plots the book's cumulative return in percent of starting NAV. The grey dotted line is what the book earned from its market exposure alone. The blue dashed line adds the sector layer, the purple dash-dot line adds the subsector layer, and the black line is the gross return. The amber line is the distance between the black line and the purple line: the stock-specific layer. The right panel stacks the same four layers at the end of the window: market +260.1%, sector +2.3%, subsector +9.0% and stock-specific –19.1% add to the gross return of +252.3%, and the dotted connector links the end of the gross line to that total. Berkshire's gross line sits below its market line for most of the decade, and the amber line shows the shortfall accumulating in the stock-specific layer. The thematic lines stay close to the market line because large offsetting sector positions net out, as described below.

Where the layers came from. The layer totals hide offsetting positions, and the offsets are what an investor report would discuss. For holding *i* and layer *k* the contribution is the same sum taken one position at a time:

$$C_{i,k} = \sum_{t=1}^{T} \frac{\mathrm{NAV}_{t-1}}{\mathrm{NAV}_0}\, w_{i,t}\, r_{i,t}^{k} \tag{5}$$

where $w_{i,t}$ is the drifted, renormalized weight at the start of the day and $r^{k}_{i,t}$ is the part of the stock's return the model attributes to layer *k*. For the sector layer that part is the return of the stock's sector factor beyond the market; for the subsector layer it is the return of the subsector factor beyond the sector. A position can therefore subtract through a thematic layer while the stock itself rose, and a position held when NAV was several times its starting value contributes several times weight × return. The same weighting is why Tiger's market layer (+560.2%) exceeds its gross return (+462.2%): the later market gains accrue to a larger NAV, and the negative subsector and stock-specific layers bring the total down. Contributions are in percent of starting NAV. Weights quoted with Figure D4 are average weights over the window in the whole mapped book, including names classified through the point-in-time taxonomy.

Appendix D. Sector and stock cascades for the five books

Figure D4. Sector cascades: where each book made and lost money, one manager at a time

How to read Figure D4. Each chart splits one book by sector. Widths add to 100% of the filed book, and bars run from the best-returning sleeve on the left to the worst on the right, so the chart reads as a cascade. Bar height is the annualized return of the sector's holdings while they were held, split into the four layers: market beta, sector, subsector and stock-specific. Each layer is the sleeve's cumulative layer contribution scaled by the ratio of its annualized to its cumulative return, so the layers add to the bar; the scaling is an approximation. The black tick marks each sleeve's net return and the dashed orange line is the whole book. Bars touch, so widths read off the horizontal axis. Labels for narrow bars sit above the chart on leader lines, and sleeves under 3% of the book carry a number keyed to a note below the chart. Layer slices under one point are too thin to see. Sleeves held for less than the full window are annualized over their holding span, so the dashed line differs from the weight-average of the ticks. Because height is an annualized return, the area of a bar for a sleeve held through the whole window approximates that sector's contribution to the book's annual return; partial-window sleeves require a separate contribution calculation. Sectors under 1.5% of the book are combined as Other. Pershing Square, which holds about eight names, is drawn by stock and ordered by average weight. Widths are average weights in the whole mapped book and include names assigned through the model's point-in-time sector taxonomy; the weights quoted in the text use the same basis. Holdings no longer in the current classification, mostly delisted names, are 0.6% of Berkshire's average book, 6% of Lone Pine's, 8% of Tiger's, 11% of Pershing's and 26% of Baupost's. May 2016 to April 2026.

Technology, a third of the book, returned 25% a year against 13.5% for the whole book; staples returned 2%

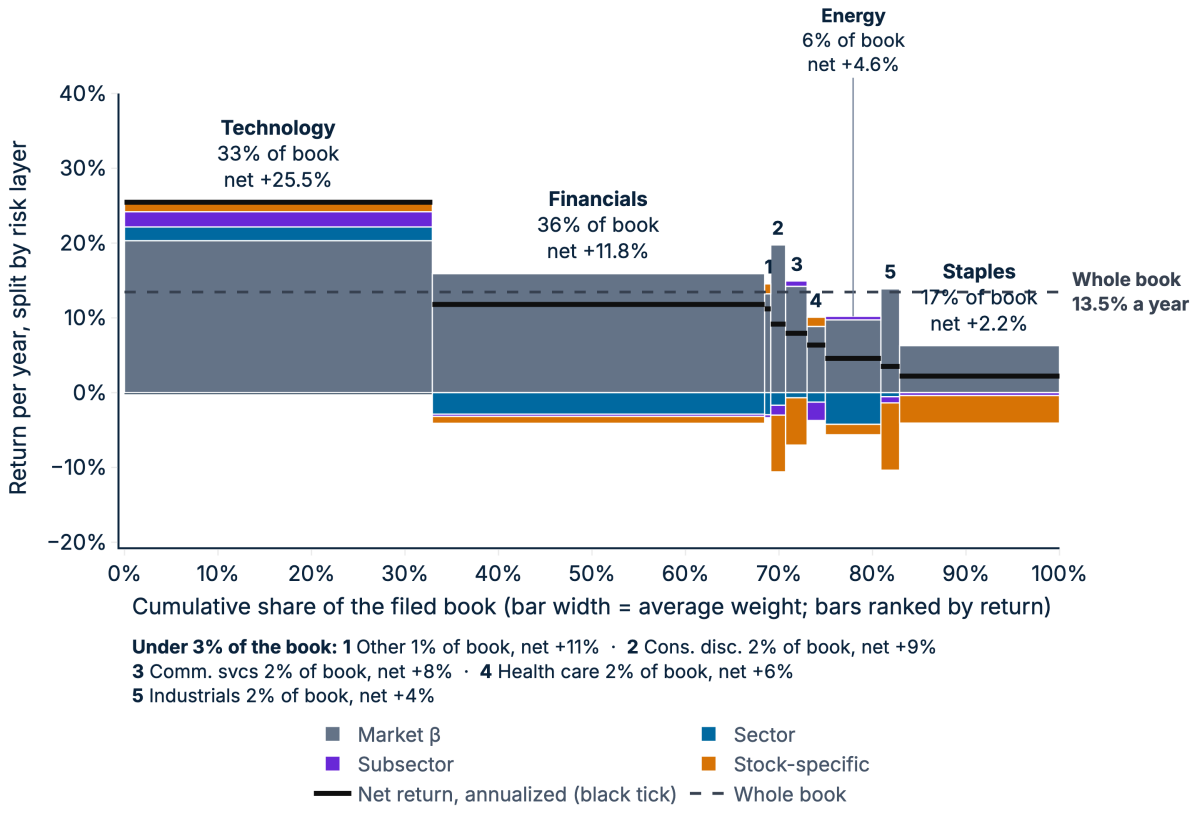


Figure D4a. Berkshire Hathaway, sector cascade, May 2016 to April 2026.

Berkshire held financials at an average of 36% of the filed book, and the sector layer from those holdings cost 18.2% (largest names: American Express –5.7%, Bank of America –5.1%, Wells Fargo –3.1%) because the financial sector trailed the market. Technology (Apple, +13.0%) and energy (largest names: Chevron +7.1%, Occidental +3.1%) added 14.2% and 9.5%, more than offsetting the financials, so the sector layer nets to +2.3%. In the subsector layer, Apple's place in the large-cap technology cohort added 14.1%, while banks (–2.2%), food and beverage (–1.3%, Kraft Heinz –1.1%) and airlines in the transportation cohort (–1.1%) subtracted. The stock-specific layer of –19.1% is Kraft Heinz (–18.7%), Wells Fargo (–9.9%) and Occidental (–5.2%), against Apple (+10.6%), Coca-Cola (+6.3%), Alphabet (+5.2%) and American Express (+3.5%).

Among the large sleeves only technology beat the book: 24% a year against 16.9%

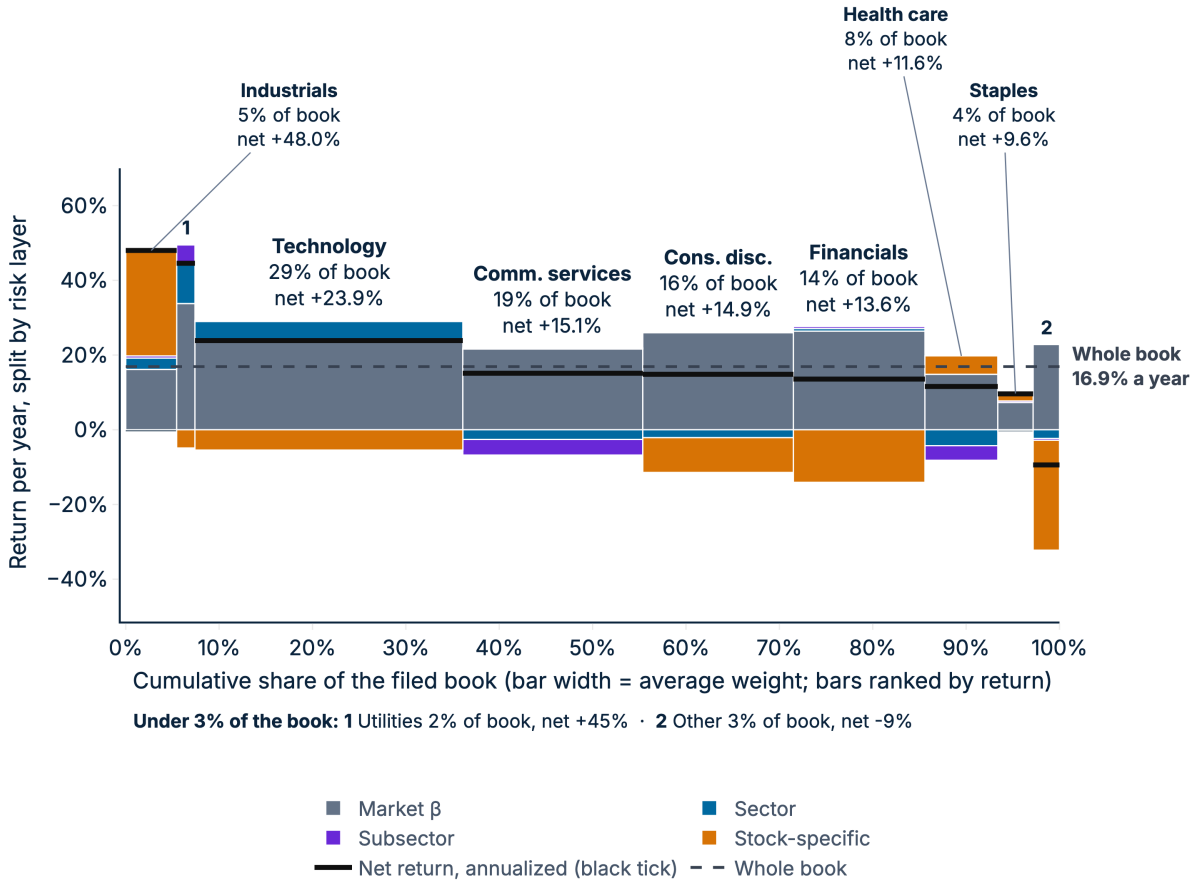


Figure D4b. Lone Pine Capital, sector cascade, May 2016 to April 2026.

Lone Pine held technology at an average of 29% of the book, and the technology sector layer added 27.2% (largest names: Microsoft +5.6%, Adobe +3.7%, Shopify +2.7%). Utilities added 5.5% through the power generators Vistra (+3.5%) and Talen (+1.7%). Communication services (–9.3%, led by Snap at –6.6% and Meta at –3.1%) and consumer discretionary (–6.5%, Amazon –3.1%) subtracted. In the subsector layer, software cost 25.9% (Shopify –7.1%, Workday –5.1%, ServiceNow –4.6%) and the internet platform cohort, which sits under communication services, cost 10.4% (Snap –8.2%), while semiconductor holdings added 19.5% across the two semiconductor cohorts (Taiwan Semiconductor +8.7%, ASML +4.2%, Entegris +3.1%) and brokers and alternative managers added 4.9% (KKR +2.4%, LPL +1.6%, Ares +1.1%). The stock-specific layer of –64.1% was led by Farfetch (–15.7%), Block (–15.4%), StoneCo (–9.9%) and Twitter (–8.0%), against AppLovin (+19.8%), Philip Morris (+11.6%), UnitedHealth (+9.7%) and Spotify (+8.5%).

Technology returned 25% a year; communication services and consumer discretionary, 54% of the book, trailed the 18.9% book return

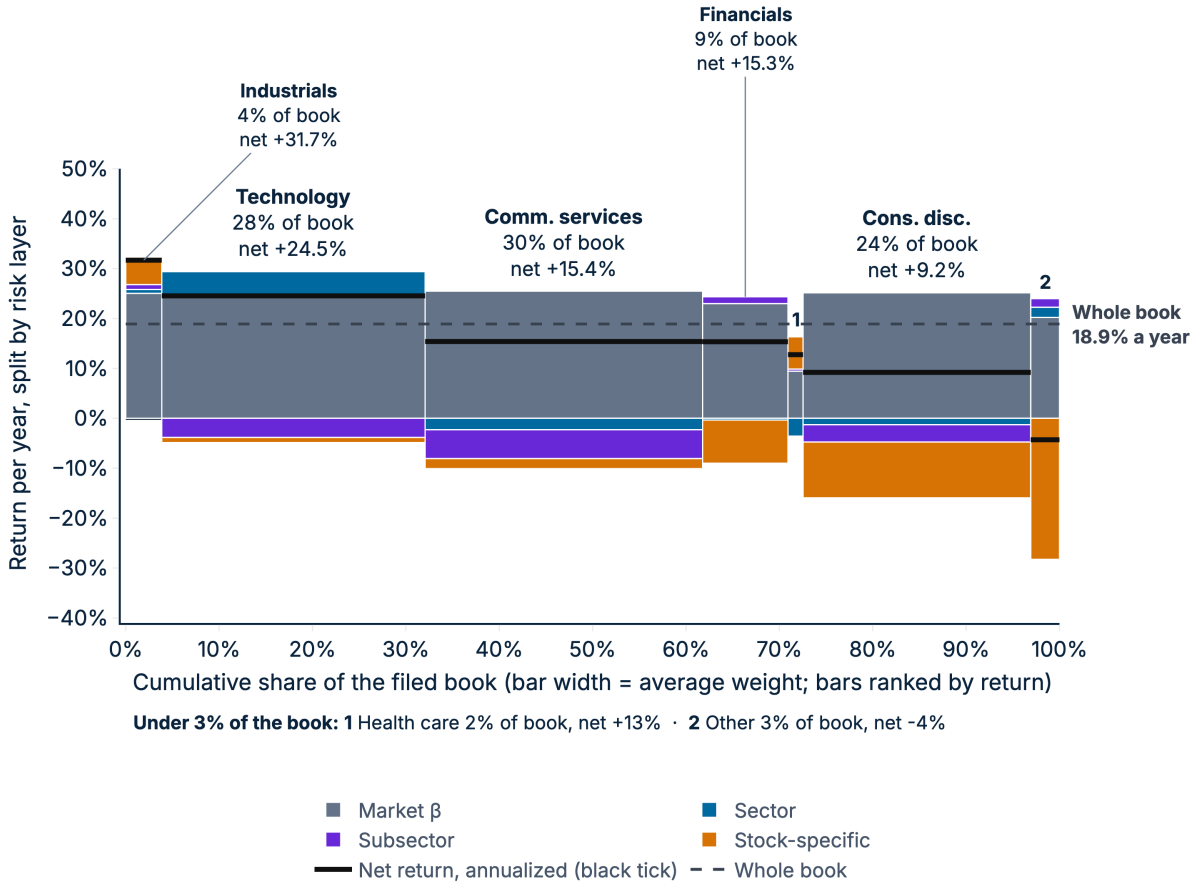


Figure D4c. Tiger Global, sector cascade, May 2016 to April 2026.

Tiger Global averaged 23% of the filed book in software and 23% in internet platforms. The software cohort cost 65.3% (largest names: CrowdStrike –11.0%, ServiceNow –6.9%, DocuSign –5.7%) and the internet cohort cost 30.1% (Sea Limited –18.7%, Reddit –3.8%, Meta –3.5%), with retail costing 9.9% (JD.com –13.1%, Carvana +5.5%). Semiconductors added 22.1% across the two cohorts (Taiwan Semiconductor +6.8%, NVIDIA +5.5%, Lam Research +5.0%), and Sunrun alone added 14.7% through the clean energy cohort. In the sector layer, technology added 28.2% (Microsoft +12.1%, Snowflake +4.5%), while communication services (–4.8%, with Meta +8.2% and Sea –7.6%) and health care (–4.8%, Eli Lilly –2.5%, UnitedHealth –2.0%) subtracted. The stock-specific layer of –38.6% is Carvana (–23.3%), Zoom (–15.4%), Flutter (–12.9%) and Peloton (–9.5%), against Alphabet (+29.0%), Amazon (+20.9%), Meta (+14.7%) and CrowdStrike (+13.4%).

Consumer discretionary was 61% of the book and returned 17.5% a year, ahead of the 15.1% book

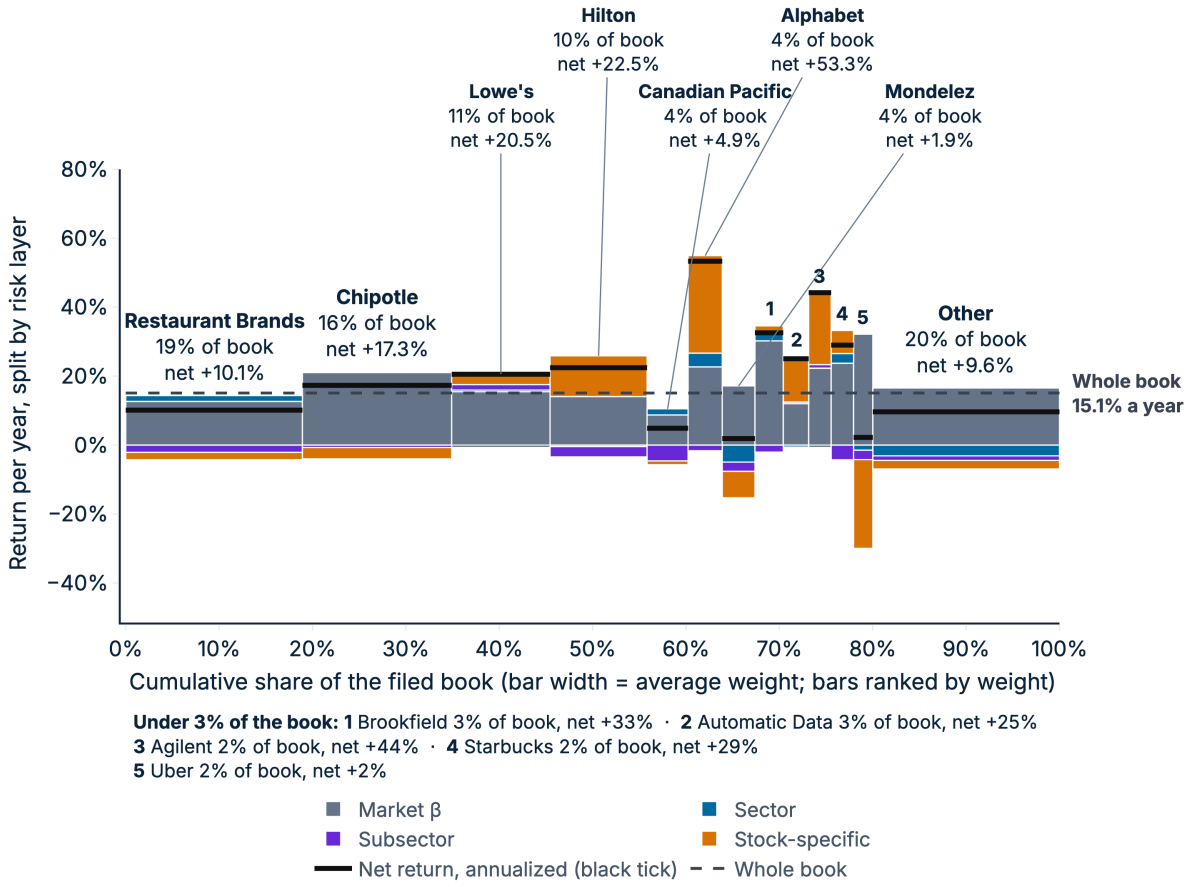


Figure D4d. Pershing Square, cascade by stock, ordered by average weight, May 2016 to April 2026. A book of about eight names is drawn by stock instead of by sector; issuers under 1.5% of the book are pooled as Other holdings.

Pershing Square averaged 61% of the filed book in consumer discretionary and 49% in the leisure and entertainment subsector, which holds Hilton, Restaurant Brands and Chipotle. That subsector layer cost 19.5% (Hilton −10.1%, Restaurant Brands −4.8%, Chipotle −4.2%). Internet platforms (Meta, Alphabet, Uber) cost 7.2% and the transportation cohort, through Canadian Pacific, cost 5.3%. The sector layer nets to −0.4%, with consumer discretionary adding 4.9% (Lowe's +3.5%, Restaurant Brands +3.1%, Amazon −4.4%). The stock-specific layer of +22.5% runs the other way from the subsector layer: Hilton (+28.0%), Chipotle (+20.5%) and Alphabet (+28.2% across its two listings) beat their cohorts by wide margins, against Uber (−15.7%), Howard Hughes (−22.0% across its two entities) and Netflix (−9.3%).

Baupost's red sleeves, 31% of the book, lost money; the whole book returned 5.1% a year

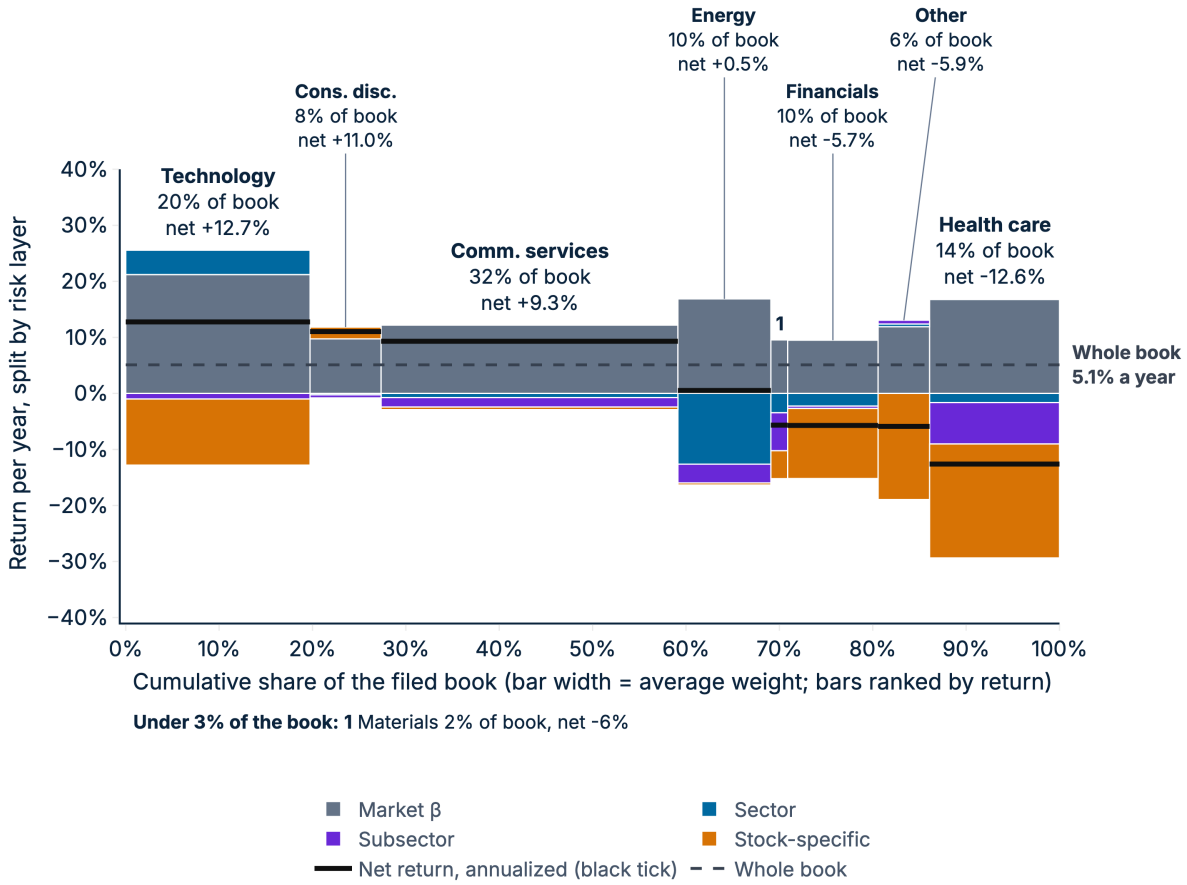


Figure D4e. Baupost Group, sector cascade, May 2016 to April 2026.

Baupost held communication services at an average of 32% of the filed book and energy at 10%. Energy was the largest sector detractor (–10.7%: Cheniere –8.8%, Antero –2.5%), and the communication services cohort cost 4.1% (Viasat –3.9%), with software (–2.5%) and internet (–2.0%) smaller. Cheniere shows how the layers separate: its energy-sector layer cost 8.8% while its stock-specific layer added 6.5%. The stock-specific layer of –77.8% was concentrated in Viasat (–41.2%), Qorvo (–16.3%), Clarivate (–11.2%) and PG&E (–10.4%), against Alphabet (+15.1%), eBay (+9.7%), Cheniere (+6.5%) and New Oriental (+5.3%).

Two readings follow. The subsector layer subtracts for four of the five books (Berkshire is the exception at +9.0%), and the largest subsector detractors are cohorts that lagged their sector: software, internet platforms, leisure, retail, transportation, banks and communication services. The stock-specific layer is negative for four of the five books (Pershing Square is the exception at +22.5%), and a few positions dominate it. Kraft Heinz alone (–18.7%) is about equal to Berkshire's entire stock-specific layer (–19.1%), and Viasat (–41.2%) is more than half of Baupost's (–77.8%). These residual layers belong to a return decomposition; manager skill remains a separate question.

Without the market layer, 67% to 95% of four of the five books had negative active return; only 29% of Pershing Square's book did

Figure D4f. The same cascades with the market layer removed. Each bar is the sector, subsector and stock-specific layers of one sleeve, annualized as in Figure D4 and ranked from best to worst on that active return, with the market layer left out. The black tick is the net active return. Sleeves are the same, and so are the widths. The share of the book with negative active return is 92%

for Baupost, 89% for Lone Pine, 95% for Tiger Global and 67% for Berkshire, against 29% for Pershing Square. The tallest bars are small sleeves (Lone Pine's industrials +32%, Other –32%), so the width of a bar matters as much as its height. The five charts are also in `sector_cascade_active.gif`.

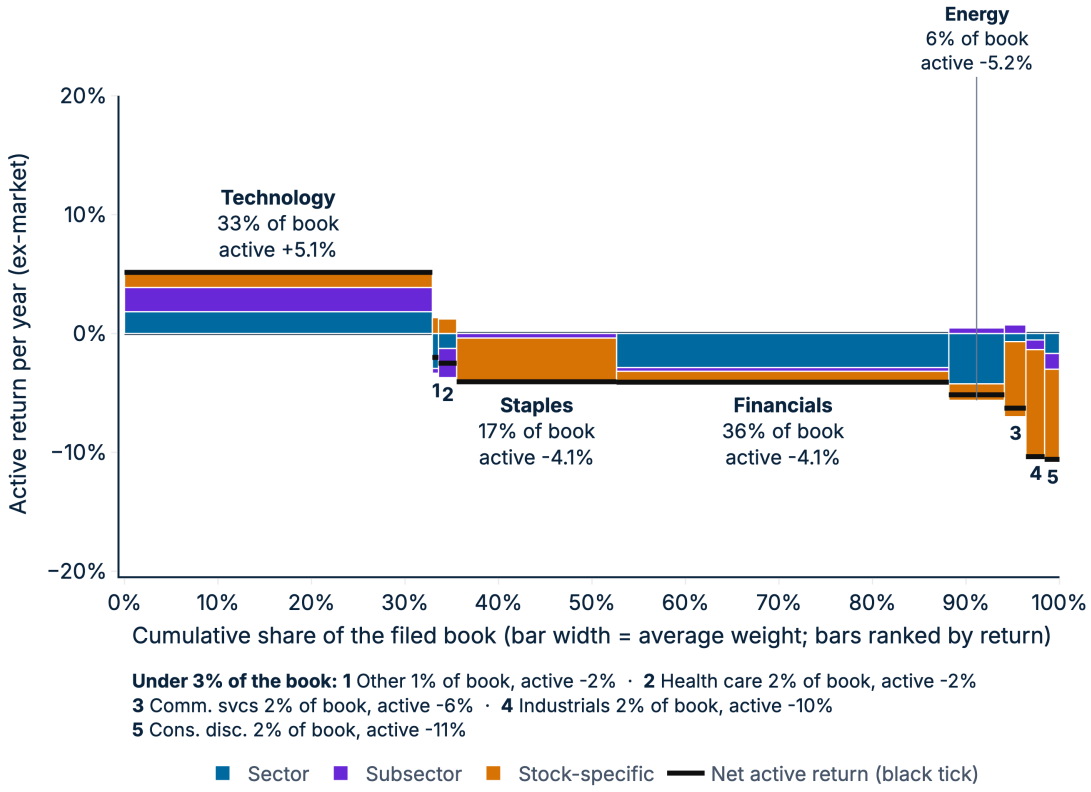


Figure D4f, Berkshire Hathaway: sector, subsector and stock-specific return by sleeve, market layer removed.

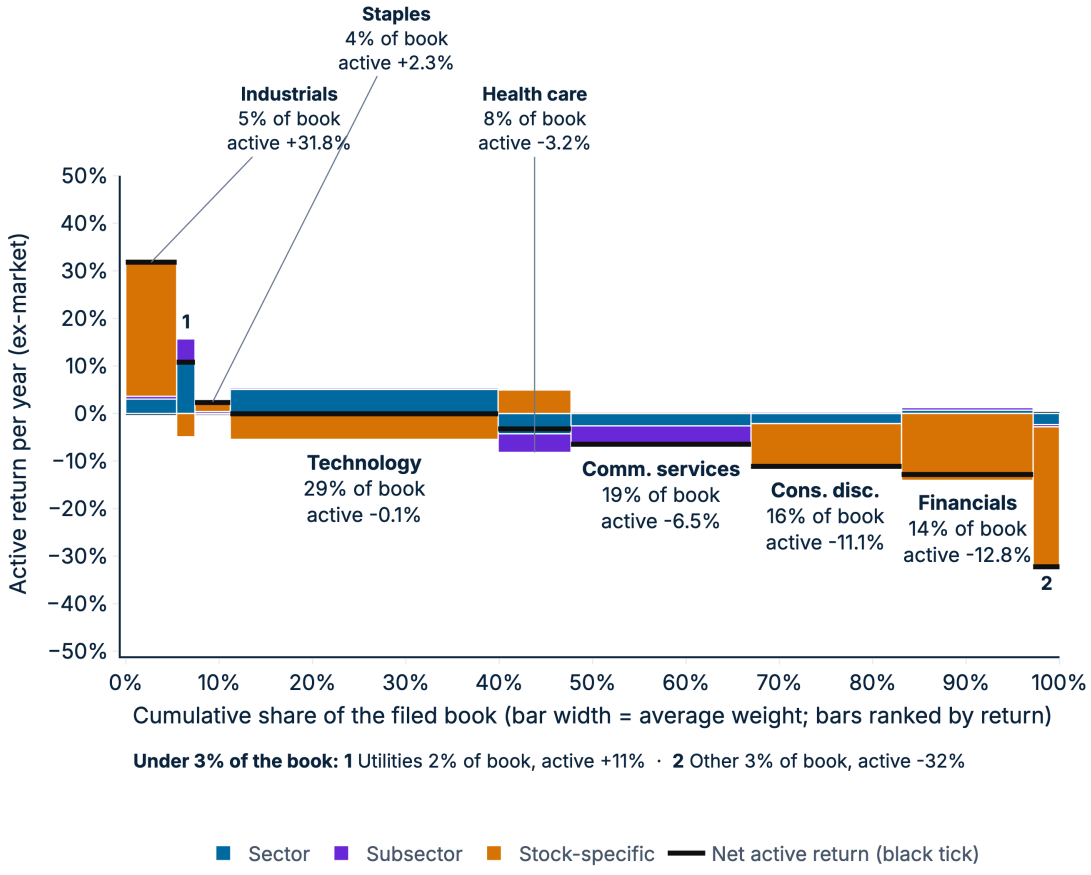


Figure D4f, Lone Pine Capital: sector, subsector and stock-specific return by sleeve, market layer removed.

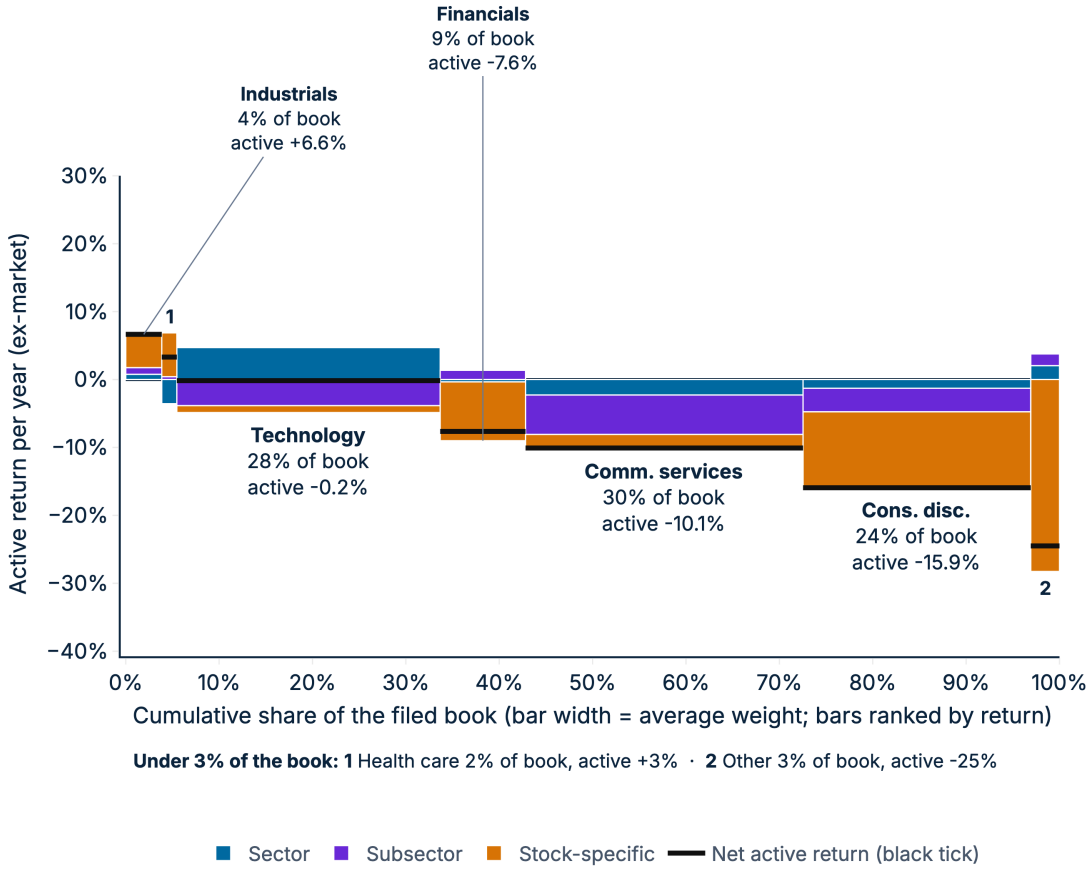


Figure D4f, Tiger Global: sector, subsector and stock-specific return by sleeve, market layer removed.

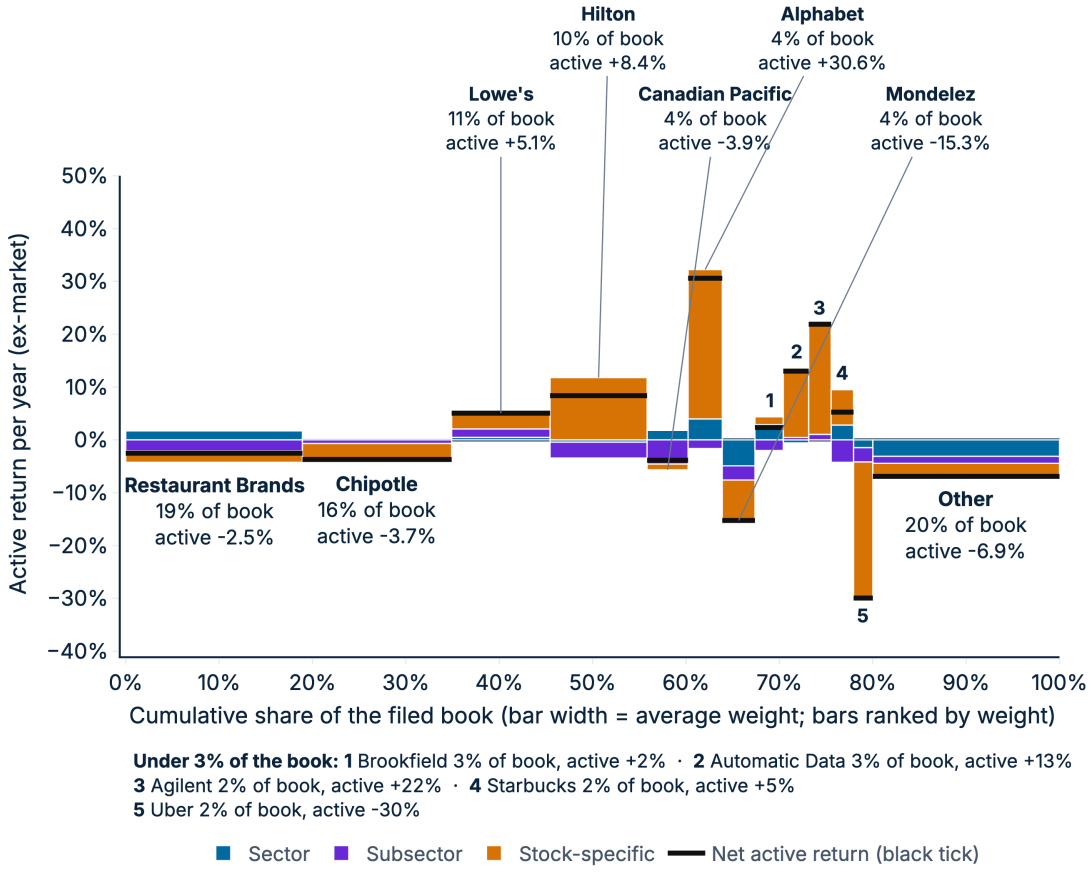


Figure D4f, Pershing Square: sector, subsector and stock-specific return by stock, market layer removed.

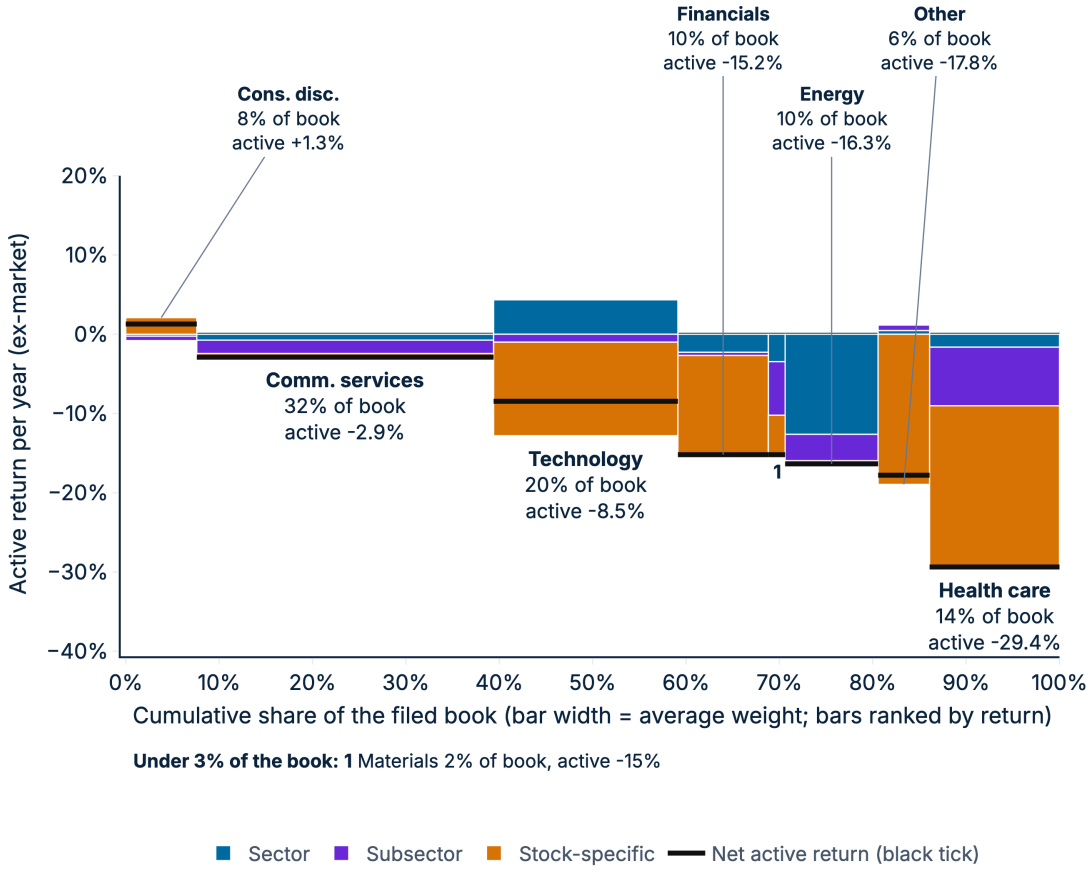


Figure D4f, Baupost Group: sector, subsector and stock-specific return by sleeve, market layer removed.

> **Framework scope: the mapped Baupost sleeve**
>
> The decomposition covers Baupost's **mapped U.S. equity sleeve**. Realized gross over the window is **+64.2%**; the attributed market layer alone sums to **+175.4%**, with offsetting sector (–8.6%), subsector (–24.8%) and stock-specific (–77.8%) contributions. Only 84% of Baupost's book value is priced on the ERM3 stock axis, the lowest of the five, so the stock-specific figure describes the priced part of the book. The pattern applies layered attribution arithmetic to that mapped sleeve; a firm-level verdict lies outside its scope. Baupost's broader opportunistic strategy can include structures, geographies and instruments outside this universe.
>
> For Klarman, the Appendix B **risk attribution** describes the structure of the mapped sleeve and should be read with the return attribution.

Risk share and return contribution. Lone Pine is the clearest contrast: **79%** of portfolio risk from market beta (Appendix B), yet **–64.1%** of stock-specific cumulative contribution here. The market layer carries the return and the residual layers net negative. Variance share and return contribution answer different questions and can differ in sign.

The risk partition (Appendix B) and the return attribution (Appendix C) describe each manager consistently and answer different questions.

Appendix E. Stock-specific structure, turnover, the filing lag and the NAV check

This section examines the stock-specific layer of each book: its size, how it moves through time, how many names carry it, how fast the book turns over, and how much of it survives the 45-day filing lag. Appendix B and Appendix C report thematic structure. This appendix reports the size and drift of the stock-specific share, name-level turnover and the effect of the filing lag.

THROUGH TIME

A single-window risk partition is a snapshot. This subsection reports how each manager's **stock-specific share of portfolio risk** changes across windows and regimes.

In a rolling year, the stock-specific share of risk has run from 2% to 98% across these books

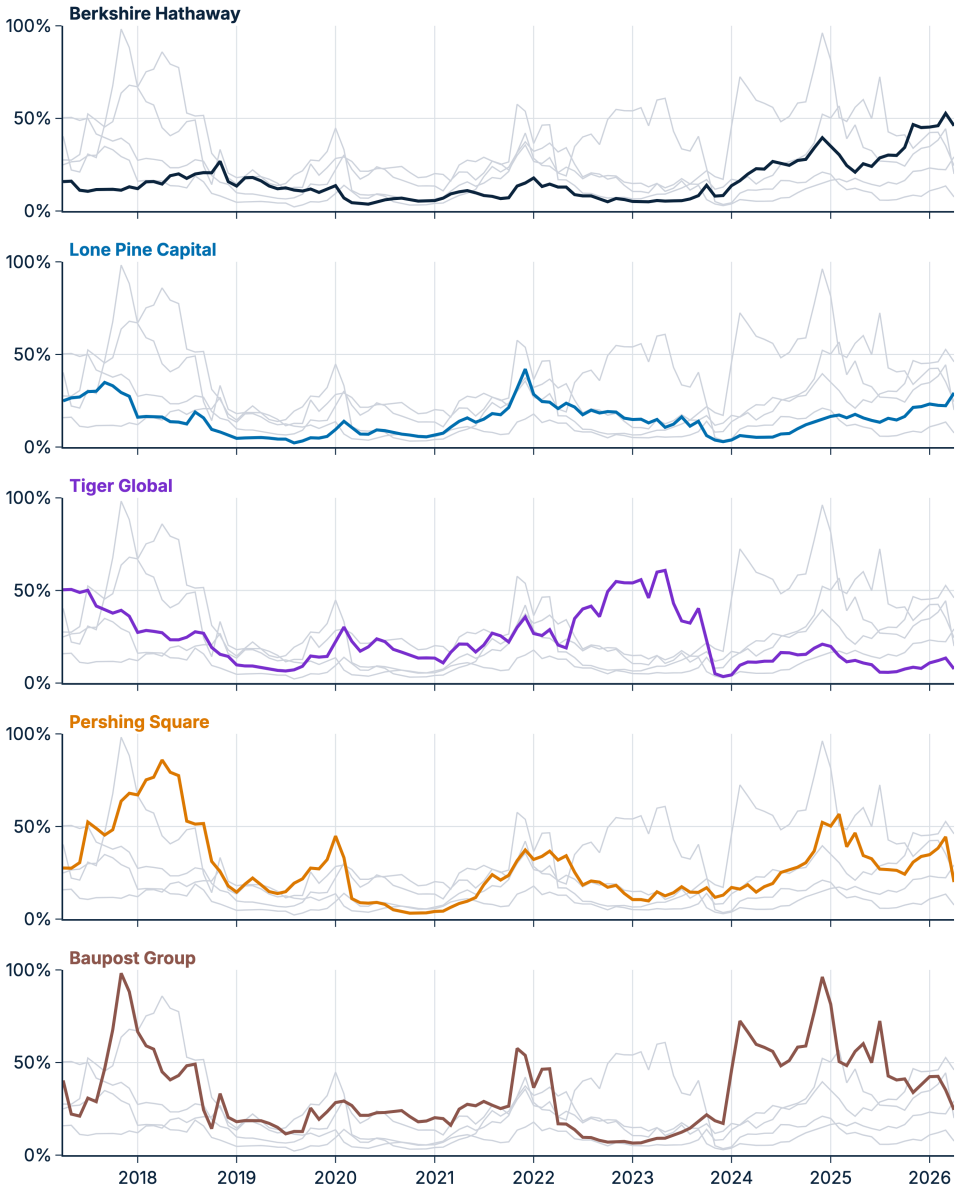


Figure E5. Rolling 12-month stock-specific share of portfolio risk, one panel per manager. The other four managers appear in grey behind each line.

The figure plots a rolling 12-month ratio for each manager: the sum of squared stock-specific monthly returns divided by the sum of squared gross monthly returns, measured around zero. The ratio can exceed 100% in a window where stock-specific and market returns offset each other, and it is a different measure from the full-window variance share in Appendix B. Averages below are over the 109 rolling windows from April 2017 through April 2026.

The stock-specific share of risk is higher in the last third of the decade than the first for Berkshire and Baupost, and lower for Pershing Square and Tiger Global

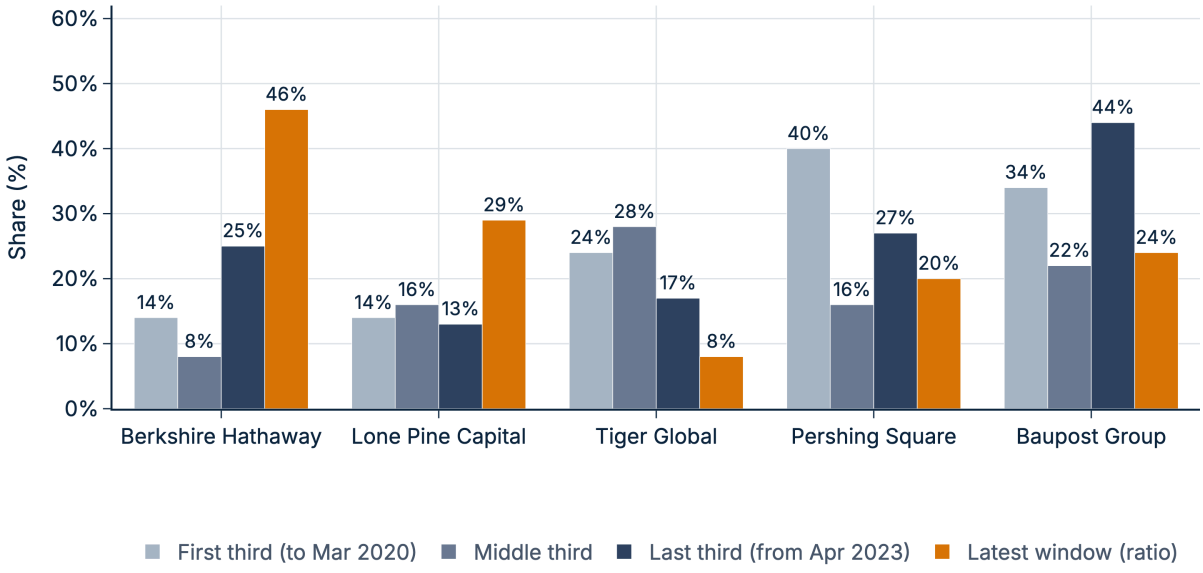


Figure E6. Average rolling 12-month stock-specific share of risk by period. The latest-window bar is the ratio defined above, distinct from the trailing-12-month variance share in Figure E7.

Berkshire's ratio averages 8% in the middle third, with a low of 3.7% in May 2020, and reaches 53% in March 2026 (46% in the latest window). Pershing Square and Baupost record the highest early readings (a maximum of 86% for Pershing in April 2018 and 98% for Baupost in November 2017) and fall to lows of 3.2% (Pershing, October 2020) and 6.5% (Baupost, January 2023). Tiger Global's ratio peaks at 61% in the May 2023 window, falls to 3.5% in December 2023 and stands at 8% in the latest window. Lone Pine's ratio ranges from 2.2% (August 2019) to 42% (December 2021). A twelve-month window holds twelve observations, so single months move these readings materially. The figure shows direction; the level in any one window carries wide error.

STOCK-SPECIFIC INTENSITY

The prior subsection describes the *level* and *drift* of the **stock-specific share of portfolio risk**. Together they rank the five books by the size of the stock-specific share and show whether that share rose or fell over the ten years.

In the latest year Berkshire's stock-specific share rose to 34%, more than double its ten-year level; Tiger Global's fell to 8%

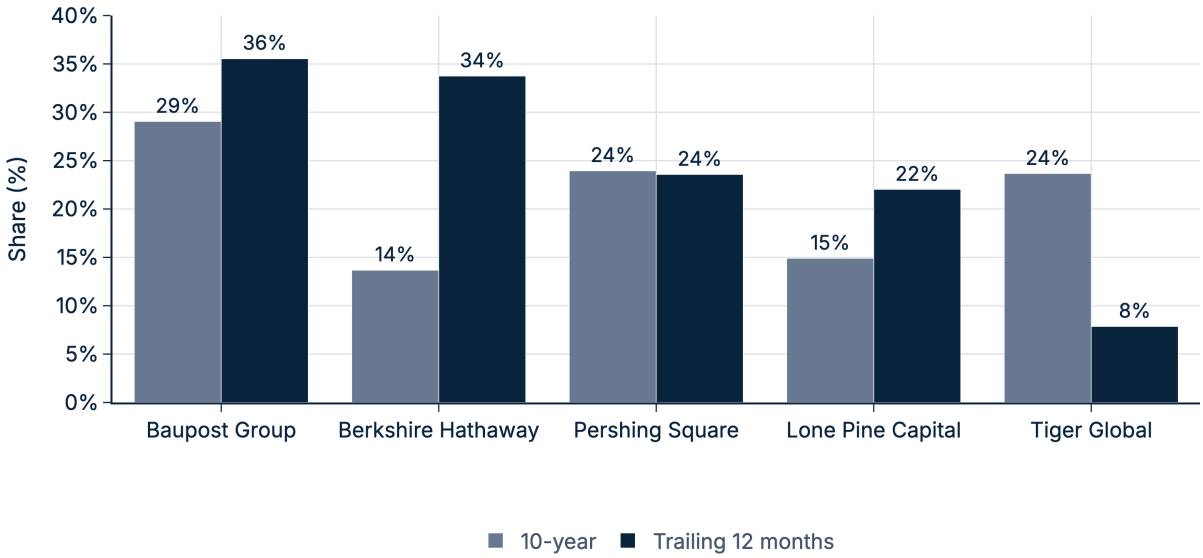


Figure E7. Stock-specific share of portfolio risk, ten-year and trailing 12 months (May 2025 to April 2026), ranked by the trailing share. Both are variance shares of the kind used in Appendix B.

Baupost (29%), Tiger (24%) and Pershing (24%) carry the largest stock-specific shares over the ten years. In the latest twelve months Berkshire's share rises from 14% to 34% and Tiger's falls from 24% to 8%. Baupost rises from 29% to 36%, Lone Pine from 15% to 22%, and Pershing stays within one point of its ten-year level. Twelve observations determine the trailing column, so differences of a few points between managers are within sampling error.

Stock-specific share of portfolio risk and **stock-specific cumulative return contribution** (Appendix C) answer different questions. The first: *how much of the portfolio's volatility was name-driven?* The second: *did the name-driven sleeve add to or subtract from cumulative return?* The two measures are reported separately. Neither is a proxy for total active risk; thematic risk has its own segments in Appendix B and its own columns in Appendix C.

CONCENTRATION: HOW MANY NAMES CARRY THE BOOK

The risk partition (Appendix B) and the return attribution (Appendix C) describe *how return is built*. Concentration describes *how many names carry stock-specific risk*, the bridge between "concentrated manager" and the decomposition above.

Effective N summarizes concentration; Appendix H gives its definition.

Over the quarterly holdings panel from June 2016 through March 2026 (40 quarter-ends; median across quarter-ends):

Pershing Square runs about seven effective positions; Lone Pine about twenty-two

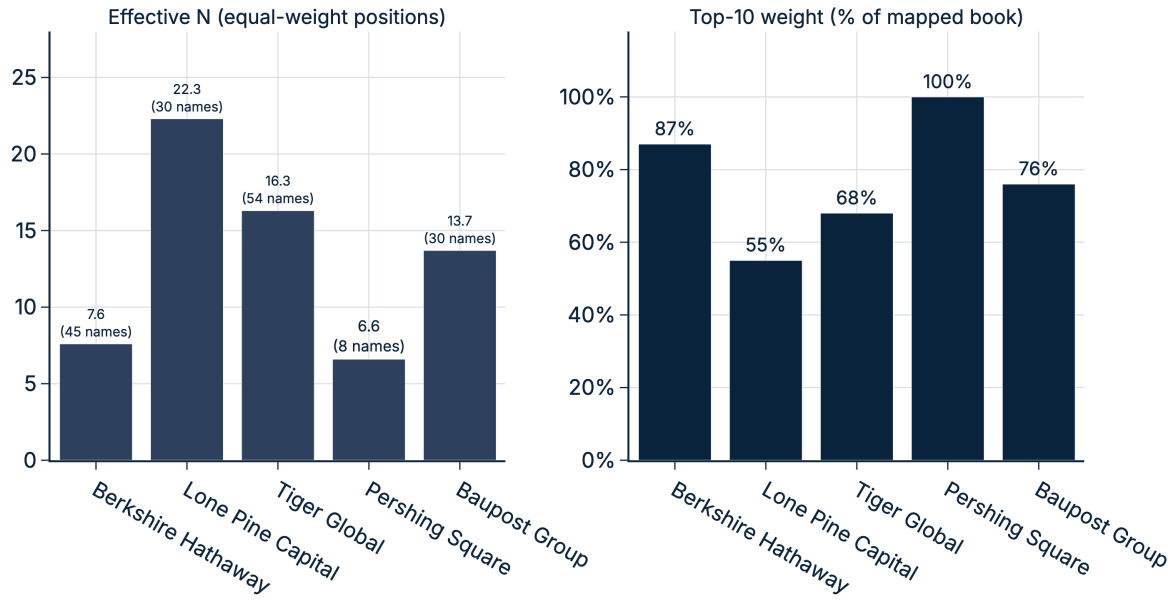


Figure E8. Median effective N and top-10 weight across 40 quarter-ends, June 2016 to March 2026.

Pershing and **Berkshire** have the lowest effective N in the cohort (6.6 and 7.6). Pershing holds a median of 8 mapped names with 100% of weight in the top ten. Berkshire holds a median of 45 names, and its top ten carry 87% of weight.

Baupost (effective N 13.7, 30 names) and **Tiger** (16.3, 54 names) sit in the middle, and **Lone Pine** is the least concentrated of the five (22.3, 30 names, 55% in the top ten). Baupost and Lone Pine hold the same median count of names and differ in effective N because Lone Pine spreads its weights more evenly.

Berkshire's stock-specific share over the latest 12 months (34%) is about two and a half times its ten-year level (14%). Whether a single position accounts for the increase remains outside this holdings panel.

The concentration statistics are available through the RiskModels API (see Section 8 and Appendix F).

TURNOVER

The age of a disclosure depends on turnover: how fast does each manager's book change between filings?

Berkshire turns over 6% of its book in a typical quarter; Lone Pine turns over 22%

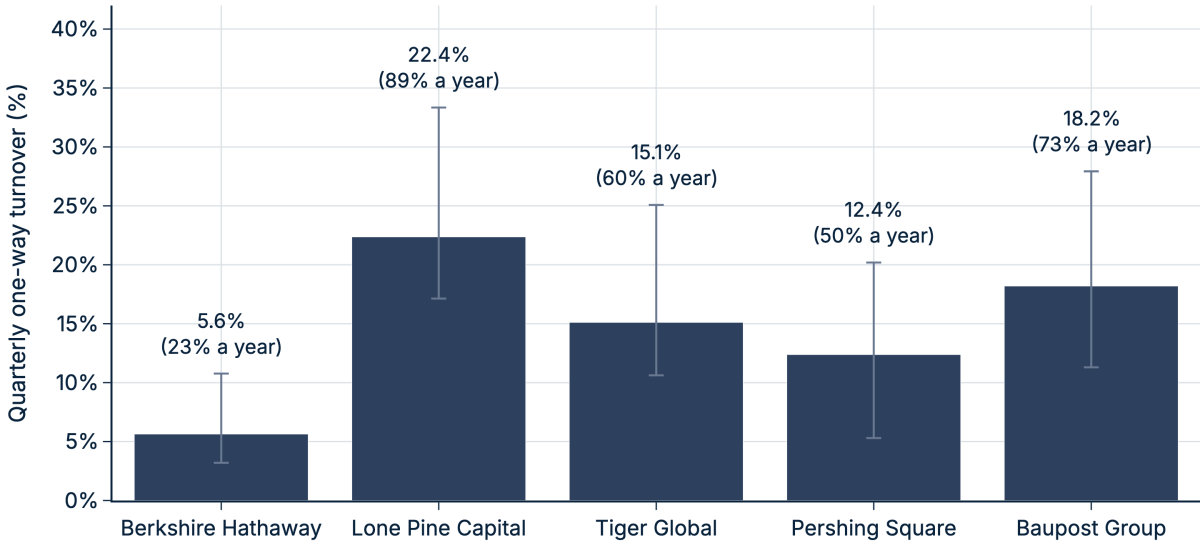


Figure E9. Median quarterly one-way portfolio turnover, June 2016 to March 2026. Median quarterly turnover and its annualized equivalent: Berkshire 5.6% (about 23% a year), Pershing Square 12.4% (50%), Tiger Global 15.1% (60%), Baupost 18.2% (73%), Lone Pine 22.4% (89%).

The window is June 2016 through March 2026 (40 quarter-ends). Berkshire is the slowest-moving book in the group, with about 6% of the portfolio rotating in a typical quarter. That is the operational fingerprint of multi-year holds: the disclosed Q1 book is, by quarter-end, *substantively* the same book. For Tiger Global, Pershing Square, Lone Pine and Baupost the median quarter changes 12%–22% of the book one-way, which is material rebalancing.

This matters for the filing-lag measurement below. The faster the book changes, the less of a given filing still describes the book when it becomes public. The fastest book here rotates about 89% of its holdings a year, so the quarterly disclosure still describes most of the book when it is filed. The next subsection measures what the delay does to a follower's return.

WHAT THE 45-DAY FILING LAG COSTS A FOLLOWER

A 13F records the manager's quarter-end holdings; a follower can begin copying only after the filing becomes public. The filing is due 45 days after quarter-end, so by then the manager has had another six to seven weeks to trade against it. We measure the cost by running each book on two clocks over the same 120 months, May 2016 to April 2026, with the same daily buy-and-hold construction. On the *reported* clock each filing is applied from its report date. On the *available* clock each filing is applied from the first session it could be traded after its public filing date, about two months later. The difference between the two cumulative returns is what waiting for the filing cost, or earned.

Crediting each filing about two months after its report date lowers Tiger Global's ten-year cumulative gross return by 79 points and Lone Pine's by 61, and raises Pershing Square's by 31

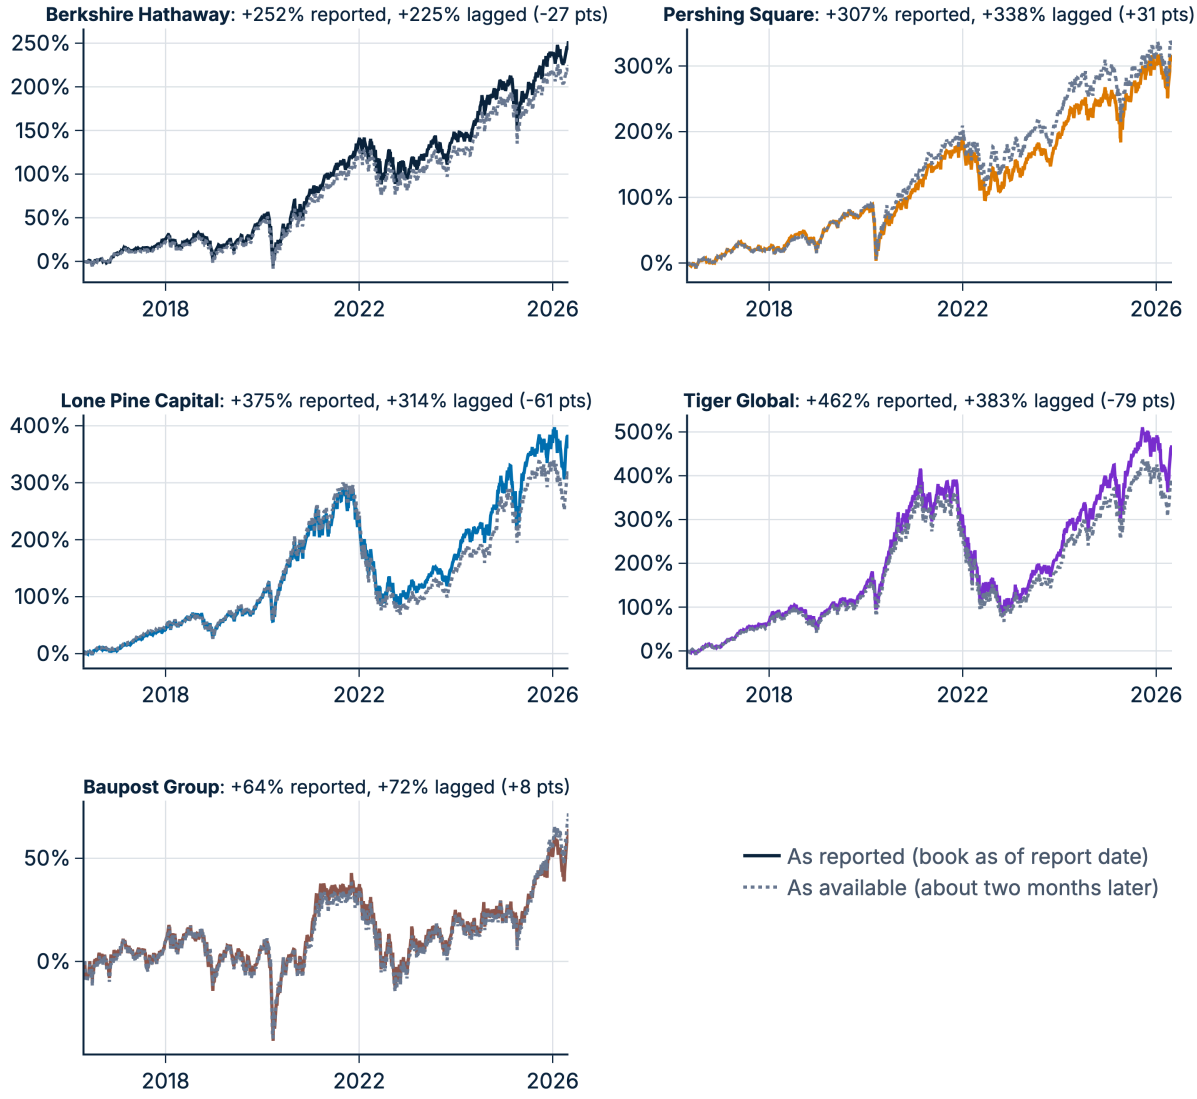


Figure E10. Cumulative gross return, May 2016 to April 2026, of each manager's mapped book as reported and as available about two months after the report date.

Table E4. Effect of the filing lag on cumulative gross return and on each layer's contribution, May 2016 to April 2026, in percentage points of starting NAV. "Change" is available minus reported. The four layer changes add to the total change, up to rounding.

Manager	Reported	Available	Change	Market	Sector	Subsector	Stock-specific
Berkshire Hathaway	+252.3%	+224.9%	**−27.4**	−15.4	−8.3	−2.1	−1.6
Lone Pine Capital	+375.5%	+314.2%	**−61.2**	−27.5	−6.9	−25.2	−1.7
Tiger Global	+462.2%	+382.7%	**−79.5**	−44.6	−3.2	+1.3	−32.9
Pershing Square	+306.8%	+337.8%	**+31.0**	+24.1	−6.4	+15.0	−1.7
Baupost Group	+64.2%	+71.8%	**+7.6**	−5.2	−5.7	+3.2	+15.4

The lag varies by filer. Lagging the book cuts ten-year cumulative gross return by 27.4 points for Berkshire (11% of its return), 61.2 for Lone Pine (16%) and 79.5 for Tiger Global (17%). It adds 31.0 points for Pershing Square (10%) and 7.6 for Baupost (12%). Crediting the books later gives a lower cumulative return for three of the five books and a higher one for two. The layer split in Table E4 shows that most of Tiger Global's 79.5 points is the market layer (-44.6), with the stock-specific layer at -32.9, and that Pershing Square's +31.0 is the market layer (+24.1) and the subsector layer (+15.0). Berkshire, with the lowest turnover of the five, loses 27.4 points, so the lag effect reflects more than turnover.

The source of the change differs by manager. Tiger Global's 79.5 points split between the market layer (–44.6) and the stock-specific layer (–32.9). Lone Pine's 61.2 points sit in the market layer (–27.5) and the subsector layer (–25.2). Pershing Square's gain comes from the market layer (+24.1) and the subsector layer (+15.0); Baupost's comes from the stock-specific layer (+15.4), against market (–5.2) and sector (–5.7) losses. Berkshire's loss is spread across the market (–15.4) and sector (–8.3) layers. The sector layer is negative for all five books, between –3.2 and –8.3 points.

The lag effect differs from the turnover ranking. Berkshire, the slowest book at about 6% of holdings a quarter, gives up 27 points. Pershing Square, at 12%, gains 31. Lone Pine, the fastest at 22%, gives up 61. Distinguishing campaign timing from ordinary dispersion across months for the Pershing Square and Baupost gains would require holdings-level attribution around specific position changes. All columns come from the daily buy-and-hold series on the reported and available clocks.

Risk shares are similar on the two clocks. The stock-specific share of risk on the available clock is within two points of the reported share for all five managers over the ten-year window (Berkshire 14% and 14%, Lone Pine 15% and 16%, Tiger Global 24% and 25%, Pershing Square 24% and 24%, Baupost 29% and 30%), and annualized volatility moves by at most 0.6 points. A single year is less stable: over the most recent 12 months the stock-specific share differs by as much as 3.5 points (Lone Pine, 22.0% reported against 18.5% available). Returns differ more between the clocks than risk shares do: the change in cumulative gross return runs from –79.5 to +31.0 points across the five filers.

DELAYED HOLDINGS AGAINST LIVE FUND NAV

Mutual funds publish a daily NAV, which gives an independent check on the holdings-based method used for the 13F books. For the 200 largest live U.S. equity mutual funds we compared the monthly return implied by each fund's periodic SEC holdings, priced through the same equity risk system with weights drifting between reports, against the fund's reported NAV return over the months where both exist (median 140 months per fund). The holdings are the book as of each report date and are applied without any trading lag, so the comparison tests how well a delayed periodic book explains realized returns.

Across the 200 funds the median monthly correlation is **0.995** (10th percentile 0.981), the median mean absolute monthly gap is **32 basis points** (90th percentile 67 basis points), and the median annualized standard deviation of the gap is 1.7%. The sample includes 57 index funds and 141 actively managed funds; the fund master carries no index flag for the other two. The funds are those of the Funds_DAG qa_top200 run of 9 November 2026. For the 141 actively managed funds the medians are a correlation of 0.994, a gap of 39 basis points and a gap standard deviation of 2.0%. The holdings-implied return is gross and the NAV return is net of fees, so part of every gap is the expense ratio. This is a variant of the return-gap comparison of Kacperczyk, Sialm and Zheng (2008), which subtracts expenses and reads the gap as a signal of unobserved actions; here it tests the return construction. The funds with the largest gaps include Baron Partners (annualized gap standard deviation 12.3%) and Dodge & Cox Stock (8.0%). Their individual diagnoses remain unresolved; unmapped or non-equity positions and trading between reports are candidate explanations.

Delayed periodic holdings track reported fund returns with a median monthly correlation of 0.995

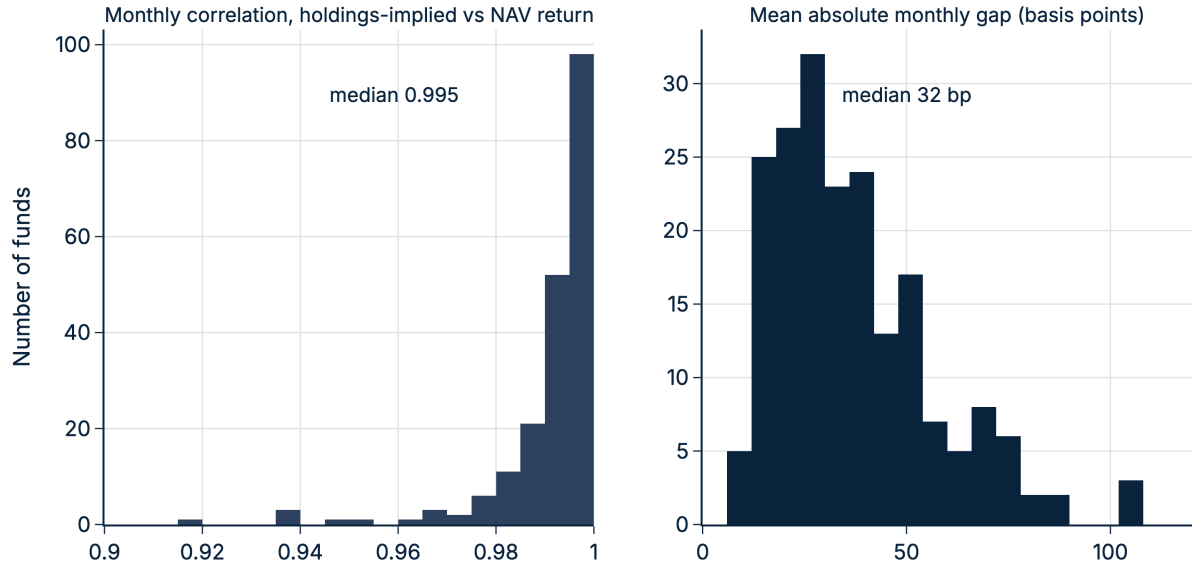


Figure E11. Distribution across the 200 largest live U.S. equity mutual funds of the monthly correlation and the mean absolute monthly gap between holdings-implied and reported NAV returns. NAV returns are distribution-adjusted month-end prices; the run is dated 9 November 2026.

The check covers funds whose books are complete and mostly reported monthly, and it validates the return construction. Its scope excludes whether a quarterly long-only 13F sleeve, which for Baupost maps only 84% of filed value, describes the manager's whole portfolio.

More broadly:

> In this cohort, stock-specific risk shares measured from filings are within two points of the available-at shares over the ten-year window, and cumulative gross returns differ between the two clocks by up to 79.5 points.

Appendix F. Methods, data and API access

- **Universe**: the median attribution uses the 2,801 filers with ten-year histories (May 2016 to April 2026) that meet the mapping-coverage and position-count filters in Section 1. The case studies use five 13F filers: Berkshire Hathaway, Lone Pine Capital, Tiger Global Management, Pershing Square Capital Management and Baupost Group. All file quarterly Form 13F-HR / 13F-HR/A with the SEC; we use the most recent restatement per report date.
- **Window**: cross-manager tables and charts use one **intersection calendar**: **May 2016 through April 2026** (daily returns, aggregated to months for the risk shares). That is the longest span where every filer in the cohort has continuous mapped history with the same start month and end month. Excluding earlier filer-specific months keeps the comparison free of forward-filled returns.

- **Decomposition**: BW ERM3 hierarchical orthogonal L3 risk model (SPY market factor, S&P-mappable U.S. equity universe). The daily components defined in Appendix H aggregate to portfolio returns as

$$R_{k,t} = \sum_{s} w_{s,t}\, r_{k,s,t}$$

  over the manager's mapped sub-portfolio. **Share of portfolio risk** in the narrative is the share of total portfolio return variance attributed to each layer (month-to-month volatility partition). Cumulative return contributions in Appendix C link each layer's daily return, weighted by start-of-day NAV relative to starting NAV, so the four layers add to the compounded gross return.

- **Share of portfolio risk, computation**: for each manager and basis, the monthly layer contributions (market, sector, subsector, and stock-specific, where stock-specific is the style plus idiosyncratic residual after the subsector layer), each built from the daily buy-and-hold layer returns weighted by start-of-day NAV relative to month-start NAV so that the four layers sum to the month's gross return, over the 120 months are each reduced to a sample variance (denominator n–1). Each share is that layer's variance divided by the sum of the four layer variances, so the four shares sum to 100% before rounding. The trailing-12-month columns apply the same calculation to the last 12 monthly observations. Style not absorbed by the traded layers remains in the residual. We report proportional variance shares rather than Euler marginal contributions, Cov(layer, portfolio) / Var(portfolio), so that every share is non-negative and a cross-layer offset cannot hide a layer's standalone variance; the four shares therefore use the sum of layer variances as their denominator rather than gross variance.

- **Cumulative contribution**: the Appendix C table reports each layer's contribution to cumulative return, in percent of starting NAV, on buy-and-hold drift between 13F books on the report-date clock, with priced holdings renormalized (window 2016-05-02 to 2026-04-30). The risk shares, rolling share and lag comparison in Appendix B and Appendix E use monthly layer contributions built from the same daily series; turnover and concentration come from the filed holdings.

- **Filing clocks**: Appendix H defines the report-date and available-at clocks. Their difference, typically about 45 days, is the filing lag measured in Table E4.

- **Turnover**: per-quarter, one-way, market-value-weighted:

$$\text{Turnover}_t = \frac{1}{2} \sum_{s} |w_{s,t} - w_{s,t-1}|$$

  0% means an identical book quarter-over-quarter; 100% means an entirely new book. Mark-to-market value drift between reporting dates contributes a small upward bias which is comparable across managers. Computed on the same June 2016 – March 2026 quarterly holdings panel.

- **Model universe (`uni_mc_3000`)**: the ERM3 return decomposition runs on the 3,000 largest U.S.-listed stocks by market capitalization, defined point in time by a mask. Each day, market capitalization (adjusted price times shares, rebuilt from SEC point-in-time share counts where available) is smoothed with a 23-day rolling median and ranked. Membership uses a hysteresis band, entering at rank 2,700 or better and leaving at rank 3,300 or worse, and is rebalanced monthly. A stock must also pass three gates: at least 113 valid days in the trailing 126, a smoothed price of at least $3, and a dollar-volume rank within the top 6,000. Before the mask is applied, return and price artifacts are scrubbed (returns above 500% are triaged, and level-shift prices such as a $1,000,000 print on a $2.74 stock are removed), and a stock-day with an invalid price, an inconsistent OHLC bar or a failed ticker-shape or data-quality screen is excluded. The build aborts if a set of ten mega-caps drops out or the universe count collapses. Membership is about 2,850 names at its plateau because of the gates, and survivorship coverage is reliable from 2009. A 13F holding is priced on a day when its stock has the full set of ERM3 layer returns, which the model builds on this universe. Holdings outside it (small caps, ADRs without a U.S. listing, non-equity instruments) are dropped and the remaining weights renormalized, which is the coverage figure reported above.

- **Sector sleeves (Figure D4)**: each holding is assigned the ERM3 sector ETF of its current classification. Holdings missing from that table, mostly delisted names, are assigned from the model's point-in-time sector taxonomy, whose eleven sectors map one-to-one to the SPDR sector ETFs. A sleeve is the set of holdings in one sector. Its daily return is the weight-averaged return of its holdings with weights renormalized within the sleeve, the sleeve return is the compounded daily return over the days the sleeve was held, annualized by days held divided by 252, and width is the sleeve's average weight in the renormalized book.

- **Coverage caveat**: the decomposition runs on each manager's *risk-model-mapped* sub-portfolio (names with ERM3 returns in the mapped sleeve). For these five filers, mean mapped **weight coverage** in the daily return series over the analysis window is **84%–98%** of filed book value (Berkshire 98%, Lone Pine 97%, Pershing Square 96%, Tiger Global 91%, Baupost 84%). Foreign ordinary shares, ADRs and other names outside the model universe are excluded from the mapped sleeve. Cross-manager comparison is apples-to-apples on each book's mappable subset; unmapped disclosed dollars and full-firm P&L remain outside the attribution.

- **Reproducing via API**: each manager maps to a canonical filer id (e.g. Berkshire `BW-FILER-CIK0001067983`). `GET /13f/filers/{bw_filer_id}/snapshot` returns the composed filer snapshot (holdings, portfolio diagnostics, `erm3_decomposition` with risk-structure partitions and waterfall when materialized). `GET /13f/filers/{bw_filer_id}/portfolio` with `start_date` / `end_date` supports history and turnover; `GET /13f/filers/{bw_filer_id}/concentration` returns the Appendix E concentration panel. `GET /13f/filers/search` enumerates the universe. For a custom long book, `POST /decompose` or `POST /snapshot` (`type: "portfolio"`) applies the same layers. The Appendix E lag experiment is explicit policy on the client: shift holdings forward ~2 months and re-run decomposition (or `POST` the lagged weights) — the API supplies the model and panels; the bitemporal convention is under your control. See **Programmatic access** below and riskmodels.app/docs/api.

- **Fund cross-check**: for the 200 largest live U.S. equity mutual funds (ranked by latest net assets, equity share of the latest book at least 80%), the monthly return from the migrated daily holdings-based series, compounded to calendar months, is compared with the reported NAV return (distribution-adjusted month-end price) over overlapping months, dropping the first and last partial months. Statistics are the correlation, the mean absolute gap and the annualized standard deviation of the gap per fund, summarized as medians and percentiles across funds. Run of 9 November 2026.

- **Equity-portfolio filter and return construction**: each 13F is filtered to holdings with ERM3 returns and weights are renormalized over the priced holdings. Between filings weights drift with returns; each state is applied from the report date on the report-date basis and from the public filing date on the available-at basis, for at most the carry limit after its report date.

DATA PROVENANCE AND REPRODUCIBILITY

Independent recomputation on 13 November 2026 reproduced the median attribution, ranking statistics, Table 2, ETF-position figures and appendix tables for the five named filers and the 2,801-filer sample. A separate check reproduced the wealth-manager results after correcting market-depth volatility from 4.5% to 4.4%. Both checks start from stored returns, hedge ratios and publicly available filer data. Rebuilding those upstream series, the Form ADV cohort rule and the name-based exclusions remains outside verification. Appendix I's simulation was independently reproduced from daily series on 11 November 2026; its construction and annualization convention are documented there. The recomputation script remains outside the repository.

REFERENCES

Brinson, G., Hood, R. and Beebower, G. (1986), "Determinants of Portfolio Performance," *Financial Analysts Journal* 42(4). Brinson, G. and Fachler, N. (1985), "Measuring Non-U.S. Equity Portfolio Performance," *Journal of Portfolio Management* 11(3). Cariño, D. (1999), "Combining Attribution Effects Over Time," *Journal of Performance Measurement* 3(4). Menchero, J. (2000), "An Optimized Approach to Linking Attribution Effects Over Time," *Journal of Performance Measurement* 5(1). Frongello, A. (2002), "Linking Single Period Attribution Results," *Journal of Performance Measurement* 6(3). Grinblatt, M. and Titman, S. (1989), "Mutual Fund Performance: An Analysis of Quarterly Portfolio Holdings," *Journal of Business* 62(3). Daniel, K., Grinblatt, M., Titman, S. and Wermers, R. (1997), "Measuring Mutual Fund Performance with Characteristic-Based Benchmarks," *Journal of Finance* 52(3). Kacperczyk, M., Sialm, C. and Zheng, L. (2008), "Unobserved Actions of Mutual Funds," *Review of Financial Studies* 21(6). CFA Institute, *Global Investment Performance Standards (GIPS)*.

- **Reproducing this paper**: the holdings, decomposition and concentration panels are available through the API. The daily filer-return series, the universe mask and the scripts behind the figures are available on request.

PROGRAMMATIC ACCESS

The framework in this paper is exposed through RiskModels.app through endpoints for filer search, holdings, concentration, decomposition and benchmark data.

- **Discover filers** — `GET /13f/filers/search` (name, CIK, modelability gate).

- **Manager dossier** — `GET /13f/filers/{bw_filer_id}/snapshot` returns latest holdings, quarterly portfolio history, ERM3 portfolio risk structure, and a latest-period return waterfall (the Appendix B and Appendix C views in JSON). Example: Berkshire `BW-FILER-CIK0001067983`, Pershing Square `BW-FILER-CIK0001336528`.
- **Holdings + history** — `GET /13f/filers/{bw_filer_id}/holdings` and `GET /13f/filers/{bw_filer_id}/portfolio?start_date=&end_date=` for turnover (Appendix E) and custom windows.
- **Concentration summary** — `GET /13f/filers/{bw_filer_id}/concentration` for median effective N, top-5 / top-10 weight share, and weight HHI (Appendix E).
- **Your own book** — `POST /decompose` or `POST /snapshot` with a weighted ticker list; same four-layer decomposition without re-building the factor model.
- **Peer / benchmark context** — `GET /rankings/{ticker}` and `GET /data/benchmark-fit` for stock-level and benchmark framing alongside the manager reads in Appendix E.

Snapshot fields depend on materialization and coverage. The SDK does not promise ETF hedge ratios in the filer snapshot, and no single call returns this paper's complete ten-year measurement. Retain each response's dates and coverage before aggregating positions or comparing periods.

Install: `pip install riskmodels-py` *· API keys at riskmodels.app/get-key · OpenAPI and MCP discovery in the API docs.*

Not investment advice. 13F filings are long U.S. equities only — no shorts, derivatives, cash, or non-U.S. holdings. Cross-manager comparisons are constrained to the universe each manager's 13F discloses, not to the full strategy each firm operates.

Appendix G. The wealth-manager cohort: rule, look-through, concentration and style

The 145 advisers in Section 4 are the SEC-registered advisers that serve mainly individuals, value-weighted, with each quarterly filing credited from the first trading session after its filing date. Return statistics cover May 2016 to April 2026. Structure statistics (lines, effective N, active share, covered share) use each adviser's latest filing, for 30 June to 30 September 2026. The adviser sets allocations and hires managers, and the 13F aggregates client accounts, including accounts run by third-party managers; the measured object therefore spans more than the adviser's own stock selection. The filing contains no short positions, options or non-US holdings, and carries no net asset value. A filed book's return is therefore built from holdings and cannot be compared with what clients earned. Named advisers appear in structural tables with their coverage stated. Lines and coverage use 145 advisers, direct-sleeve concentration 144, and return statistics the 144 advisers with at least 500 return days.

Effective positions sit far below the number of lines in the filing

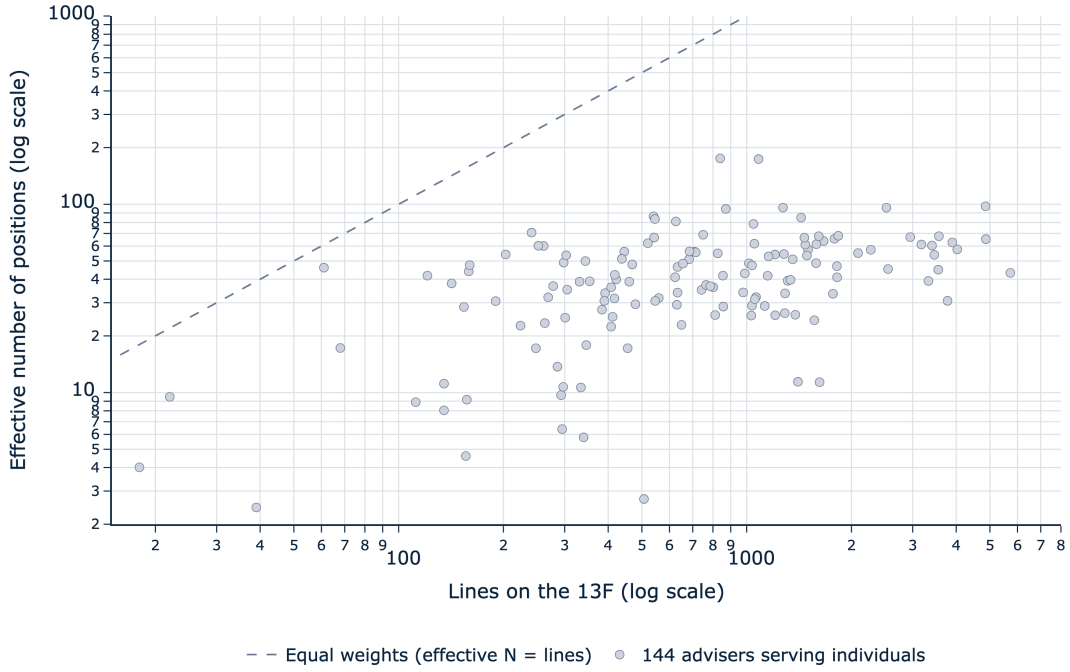


Figure G1. Lines in the latest filing against effective number of positions N = 1 / Σw², direct stock sleeve (positions that map to model stocks, renormalized to their own total), for each adviser in the cohort. Latest filings, 30 June to 30 September 2026.

G1. THE COHORT IS DEFINED BY A RULE ON FORM ADV

The question is which advisers a wealth-management buyer would recognize as the large household firms, selected without a hand-picked list. I used the latest monthly snapshot of SEC-registered advisers in the Form ADV store and kept firms that meet all five conditions in Table G1.

Rule	Condition
R1	Regulatory assets under management (Item 5F(2)(c)) of at least $5 billion
R2	Individuals hold at least 50% of those assets (Item 5D rows a and b: individuals other than high net worth individuals, and high net worth individuals)
R3	Separately managed account clients reported (Item 5K(1) = Y)
R4	CIK in Item 1D appears in the 13F filer master
R5	Registered broker-dealer flag Item 6A(1) = N; excludes platforms, turnkey asset-management providers and robo-advisers on the named exclusion list

Table G1. Cohort rule on the latest monthly Form ADV snapshot of registered advisers. Output: 162 firms, of which 147 have a current 13F filing and 145 have a complete profile.

At a $5 billion minimum the rule returns 162 firms, of which 147 have a current 13F filing in the filer master and 15 have stopped filing. Of the 147, 145 have a complete profile in the model store; the other two are new filers with no holdings store. Among the 145, median regulatory assets are $9.5 billion (quartiles $6.7 and $18.5 billion) and the median individual share is 92%. The exclusion list of platforms and turnkey providers was applied by name. It misses firms with a blank CIK in Form ADV, including Fisher Investments; the current procedure lacks a join on the CRD number. Roll-up filers such as Corient, Mariner and Hightower file on behalf of subsidiaries and may overlap with the filings of their own sub-advisers.

G2. LOOK-THROUGH LIFTS THE COVERED SHARE OF FILED VALUE FROM 39% TO 62% FOR THE MEDIAN ADVISER

A filing line maps to the model if it is a stock in the model universe, the 3,000 US stocks of the uni_mc_3000 universe. ETF and fund lines map through the holdings of the fund. I measured, for the latest filing of each adviser, the share of reported value that sits in direct model stocks, in ETFs in the store, and in lines that cannot be mapped.

The median adviser holds 39% of the filing's value directly in model stocks (quartiles 23% and 69%). ETFs in the store carry a median 15%. After looking through those ETFs to their holdings, 62% of value is attributed to model stocks (quartiles 47% and 82%) and a median 38% remains uncovered. That uncovered share is bonds, private funds, non-US securities and unlisted lines. Every result in Section 4 that uses the stock sleeve is conditional on the covered share.

Look-through raises the covered share from 39% to 62% of filed value for the median adviser

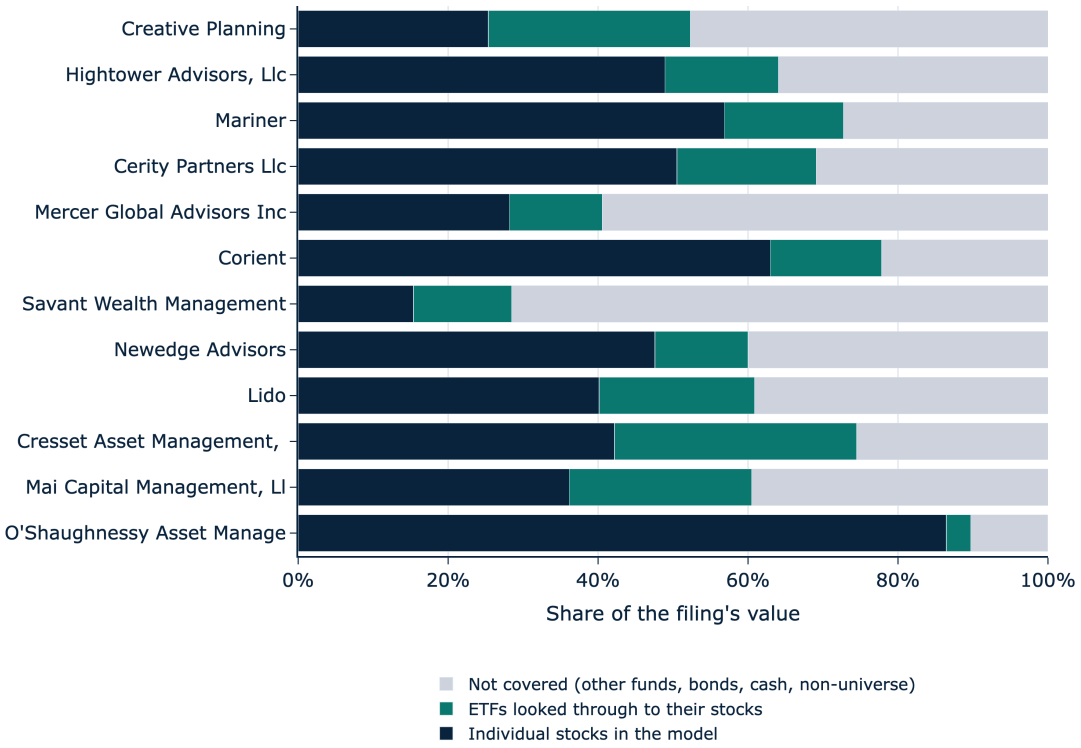


Figure G2. Composition of the latest filing's value for each adviser: direct holdings in model stocks, ETF lines in the store, other ETF lines, and unmapped lines. Value-weighted, filings of 30 June to 30 September 2026.

G3. CONCENTRATION IS MODERATE AND FALLS FURTHER AFTER LOOK-THROUGH

Appendix H defines Effective N and active share. A book with one position at 50% and fifty at 1% has 51 lines and an effective N of 3.9. The active-share benchmark here is the market-cap weights of the model universe (uni_mc_3000); no other benchmark was tested.

Both measures renormalize the sleeve to its own total and therefore describe the covered part of the book rather than the whole filing. The median adviser's direct stock sleeve has an effective N of 40 (quartiles 27 and 56) and an active share of 0.49 (quartiles 0.38 and 0.64). With ETFs looked through, effective N is 55 (40 and 69) and active share 0.32 (0.27 and 0.49). The fall is consistent with the ETF lines in these books holding large, widely owned stocks that the benchmark also holds; that explanation remains untested.

Active share against cap weights drops from 0.49 to 0.32 for the median adviser after look-through

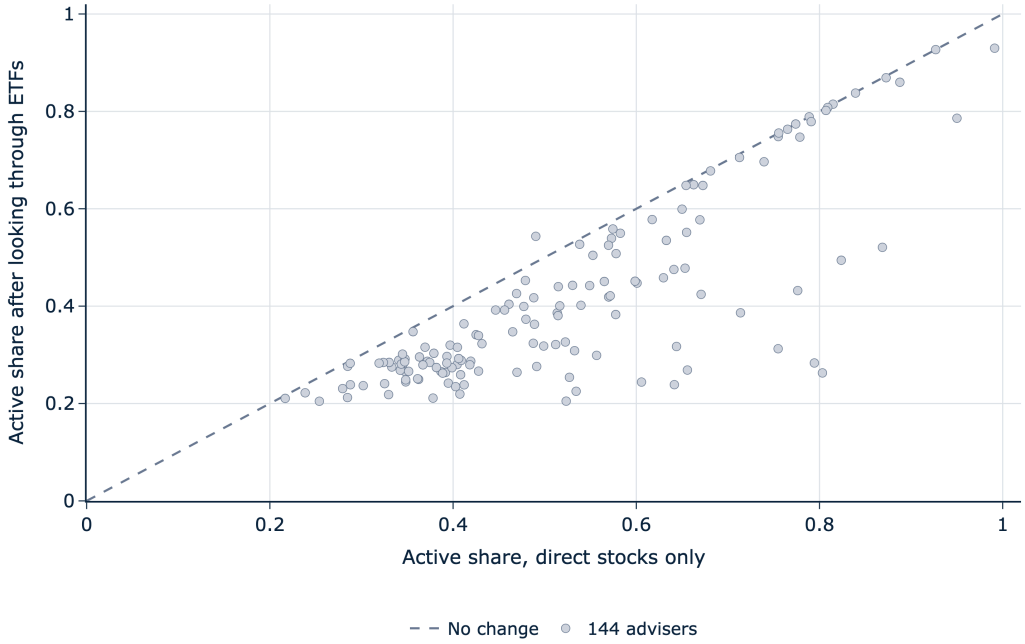


Figure G3. Active share of the stock sleeve (direct) and of the look-through book against market-cap weights of the model universe, by adviser. 30 June to 30 September 2026.

Adviser (12 largest by regulatory assets)	RAUM, $bn	Lines	Effective N, direct / look-through	Active share, direct / look-through	Covered share, direct / look-through	9-box
Creative Planning	296	4,029	58 / 78	0.24 / 0.22	25% / 52%	Large Blend
Hightower Advisors	199	3,568	68 / 75	0.38 / 0.30	49% / 64%	Large Value
Corient	165	2,951	67 / 69	0.34 / 0.29	63% / 78%	Large Blend
Cerity Partners	162	3,178	61 / 66	0.29 / 0.24	51% / 69%	Large Value
Mariner	125	3,328	39 / 47	0.33 / 0.27	57% / 73%	Large Blend
Mercer Global Advisors Inc.	84	3,460	54 / 63	0.28 / 0.23	28% / 41%	Large Value
Cresset Asset Management	79	1,620	11 / 29	0.52 / 0.33	42% / 74%	Large Growth
Savant Wealth Management	56	1,787	66 / 71	0.33 / 0.22	15% / 28%	Large Value
GW&K Investment Management	53	1,082	174 / 175	0.87 / 0.87	98% / 99%	Mid-Cap Blend
Mai Capital Management	52	1,817	47 / 62	0.35 / 0.24	36% / 60%	Large Blend
Moneta Group Investment Advisors	50	857	29 / 83	0.47 / 0.26	14% / 49%	Large Value
AE Wealth Management	50	2,549	45 / 49	0.46 / 0.40	47% / 56%	Large Value

Table G2. Structure of the twelve largest advisers in the cohort. Latest filing (30 June to 30 September 2026); active share against market-cap weights of the model universe. Structural facts only; no return or ranking statistic.

G4. STYLE IS DESCRIBED SEPARATELY FROM THE TRADED HIERARCHY

The benchmark uses no style-ETF hedge. Sector and subsector exposures absorb some style; the remaining style exposure stays in the residual. Style correlations describe that remainder. For each adviser I computed the correlation of the SPY-hedged daily book return with the style ETFs VUG, VTV, IWF, IWD, IWM, MTUM, QUAL, USMV, VLUE, SCHD and MOAT. The median correlation with growth minus value is –0.27 (quartiles –0.54 and +0.03), and 105 of 144 correlations are negative, so most books lean toward value after the market is removed. These correlations use the 753 days from May 2023 to April 2026 on the report-date clock

for 139 of the 145 profiled advisers. The other six have between 314 and 720 days, and one adviser has no correlation. The median correlation with IWM, the small-cap proxy, is 0.25. A holdings-based nine-box classification places 57 advisers in Large Blend, 51 in Large Value, 30 in Large Growth and 7 elsewhere. Across the three large-cap labels, the correlation with growth minus value follows the label in direction: all 51 Large Value books have a negative correlation (median –0.56), 36 of 57 Large Blend books are negative (median –0.11), and 18 of the 29 Large Growth books with a correlation are positive (median +0.13).

Appendix H. Definitions

Effective N. N = 1 / Σ w_i^2, with w_i the value weight of position i in the sleeve.

Active share. $\frac{1}{2}\sum_i |w_i - b_i|$ over the union of the adviser's positions and the benchmark's, with b the market-cap weights of the model universe on the filing's report date.

Available-at clock. A filing is credited from the first trading session after its public filing date. A filing accepted after 4 p.m. Eastern waits one further session.

Report-date clock. A filing is credited from quarter-end, as if its holdings were already known. This is a research convention.

ERM3 layers. Market is exposure to SPY; sector is exposure incremental to market; subsector is exposure incremental to both broader layers. Stock-specific is the return remaining after those three modeled ETF layers, including style and themes the hierarchy does not isolate.

Volatility remaining. The annualized standard deviation of the daily return of the book after the hedge legs, ddof 1, times $\sqrt{252}$.

Turnover. The sum over legs and days of absolute changes in leg weight, divided by days and multiplied by 252, in multiples of book. Cost is turnover times 1 basis point.

Custom benchmark. The long mix of SPY, sector ETFs and subsector ETFs whose weights are the book-level hedge ratios known at the end of the previous day, at a stated depth. This paper measures the book relative to that custom benchmark; standalone benchmark returns (the long ETF mix plus a cash remainder) and comparisons with standard external benchmarks remain outside its scope.

Return relative to the benchmark. The book's return minus the benchmark's return. At market depth it is the SPY-hedged return.

ETF count. The number of ETFs whose time-average absolute weight exceeds 0.5% of the book.

Variance carried. One minus the squared ratio of the volatility remaining after the benchmark to the gross volatility.

Appendix I. Gross and relative return rankings of the 2,801 filers, and a simulation with no return relative to the benchmark

Lone Pine ranks 368th of 2,801 by gross return and 1,095th by return relative to its market-depth benchmark

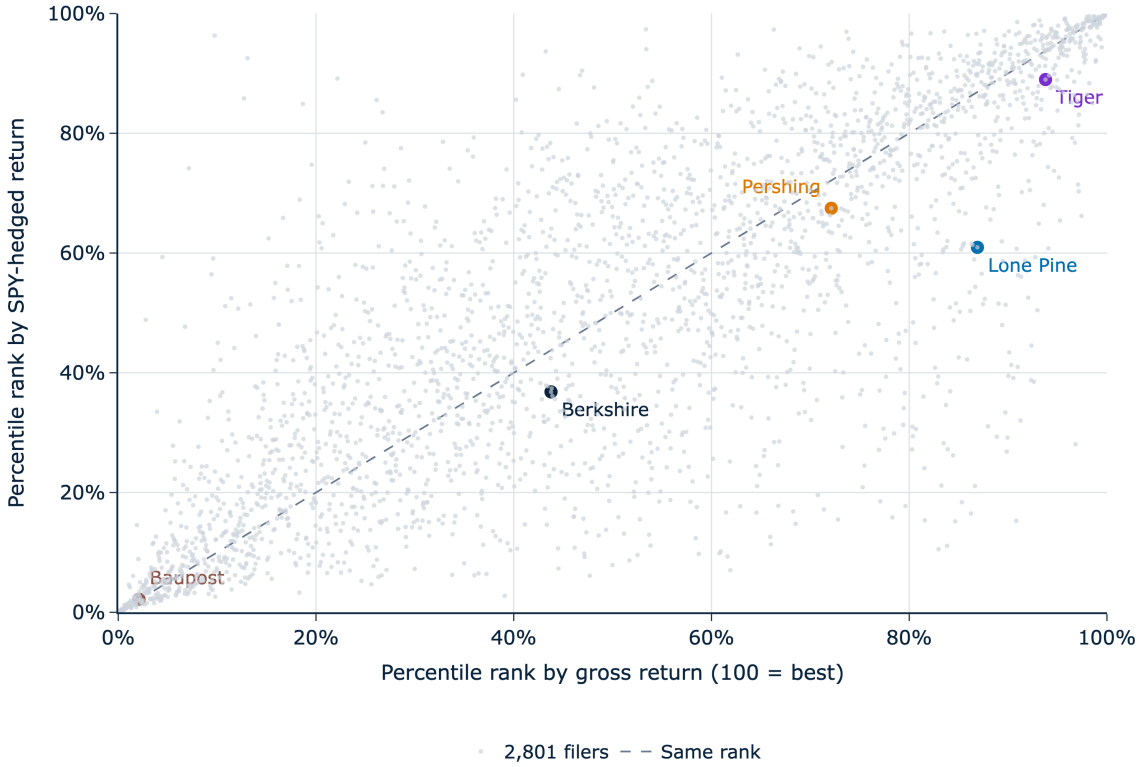


Figure I1. *Percentile rank of each of the 2,801 filers by gross return and by return relative to the market-depth benchmark, May 2016 to April 2026, report-date clock. Points on the dashed line keep their rank.*

The relative return in this appendix is the SPY-hedged return, the return relative to the market-depth benchmark. The Spearman correlation between the two rankings is 0.81. Ranking by relative return changes ranks even in a population with zero return relative to the benchmark, because filers differ in market beta and in idiosyncratic volatility. I simulated such a population: for each filer, daily gross return equal to its beta times the SPY return plus independent daily noise with the volatility of its relative-return series, 200 draws with seed 0 on each filer's usable days, with returns annualized geometrically as in the observed series. The result depends on how the simulated relative-return series is built. When it is the noise alone, which is a constant-beta benchmark, the rank correlation is 0.856 (standard deviation 0.0065, range 0.838 to 0.876) and 10.6% of filers move 25 percentile points or more. When it is the simulated gross return minus the model's time-varying market benchmark, as in the observed series, the rank correlation is 0.841 (0.0071, range 0.820 to 0.863) and 11.6% move. The observed values are 0.807 and 15.0%, outside the simulated range in both versions. The overlap of the top 5% by gross return with the top 5% by relative return is 67.1%, inside the constant-beta simulated range of 65.7% to 81.4%, making that overlap alone insufficient evidence of dispersion in relative returns. With arithmetic annualization of both the observed and the simulated series, the observed correlation is 0.736 and the simulated mean is 0.79 to 0.80, so the observed value lies below the simulated mean under both conventions.

A market-only regression produces rankings similar to the model: its intercept has a 0.91 rank correlation with model-based relative return. Comparing gross and intercept ranks gives 0.817, with 13.4% moving at least 25 percentile points. Sector and subsector results appear in Table 2 and Section 3; risk-normalizing returns moves 22% at least 25 points.