# Does AI in the Investment Process Make Funds Alike? Early Evidence from Mutual Fund and ETF Portfolios*

Kyle Schirmann  Amitoj Singh  Tarun Khanna

November 2, 2026

Preliminary and incomplete. Please do not cite without permission.

**Abstract**

Regulators worry that artificial intelligence will make investors' portfolios homogeneous: if funds apply similar models to similar data, they may buy and sell the same securities at the same time. We study U.S. active mutual funds and ETFs that use AI in their investment process, to screen, rank, or select securities; these funds report their complete holdings every quarter. A large language model reads every AI-related passage in about 330,000 prospectus filings and identifies 132 funds that describe AI performing a specific step of that process. These funds did not converge on the same portfolios, and their trades were less synchronized with one another than those of matched groups of non-adopting funds, although the same tests do detect convergence among funds at different advisers that share a subadviser. Adoption was associated with an increase of 8.5 to 14.5 log points in the number of distinct issuers held, but the added issuers came from the industries, countries, and benchmarks the funds already covered. Portfolios widened without converging.

*JEL codes:* D83, G11, G23, O33. *Keywords:* artificial intelligence; algorithmic monoculture; herding; mutual funds.

*Schirmann: University at Albany, Department of Management, kschirmann@albany.edu (corresponding author). Singh: Harvard Business School, Strategy Unit, amisingh@hbs.edu. Khanna: Harvard Business School, Strategy Unit, tkhanna@hbs.edu. Draft prepared for the NBER Digital Economics and Artificial Intelligence meeting, Spring 2027. Claude (Anthropic) was used to compose and edit the text of the manuscript, with human direction. All errors are our own.

# 1 Introduction

A recurring worry among financial regulators is that the use of artificial intelligence in the investment process will make investors' portfolios homogeneous, especially if funds share models, even without sharing data. If funds feed similar data to similar models, their decisions could produce correlated buying and selling of the same securities and amplify market volatility. The Financial Stability Board raised this concern in an early report on AI in finance (Financial Stability Board 2017), and the International Monetary Fund echoed it in 2024 (International Monetary Fund 2024). Computer scientists call this algorithmic monoculture: when decision-makers rely on shared algorithms and data, errors that would have been idiosyncratic become common (Kleinberg and Raghavan 2021; Bommasani et al. 2022). Similar concerns arise in other domains, such as recommendation systems in digital markets, which may steer consumers toward popular items rather than the long tail (Fleder and Hosanagar 2009; Zielnicki et al. 2025). The concern is with AI as a tool in the decision process, not with AI as an investment theme (for example, a mandate to invest in companies such as Nvidia). Whether adopting AI in the investment process homogenizes investment choices is an empirical question, but one that is rarely answerable, because firms' choices are seldom observed.

Asset management is a setting where these choices can be observed. In the United States, mutual funds and exchange-traded funds (ETFs) must describe their investment process in a prospectus filed under legal liability, and they must amend that description when it materially changes. Since 2018, funds have increasingly disclosed the use of machine learning, natural language processing, and other AI methods in these filings. In April 2022, for example, the Shelton Emerging Markets Fund disclosed that "the team uses data science and machine learning to conduct a deeper dive into each candidate firm to determine the investment merit, suitability for the portfolio, and pertinent risk factors" (Appendix Table A5 lists further examples). Four quarters after that filing, the fund's portfolio had broadened from 32 to 56 corporate issuers. The same funds must also report their complete holdings every quarter.

In this paper, we ask whether adopting AI in the investment process leads adopting funds' portfolios to converge. We are interested in AI as a method, so we distinguish three kinds of AI language in filings. First, a fund can adopt AI as an investment theme, investing in firms in the AI industry; this is not our subject. Second, a fund can include boilerplate AI language in its discussion of general risk factors or describe AI use only in very general terms; such disclosures increased sharply from 2023 to 2026, and they are of secondary interest. Third, and of most interest to us, a fund can describe a concrete role that AI plays in its investment process. With the aid of a commercial large language model, we

read every AI-related sentence in about 330,000 prospectus filings from 2018 to 2026. Our primary treatment indicator, a *described process*, requires that a filing describe AI performing a concrete investment task as at least one input to the security selection decision (that is, buying, selling, or not acting on a security). Of the 718 active funds whose filings first disclosed AI language after their initial filings, only 132 meet this standard; we classify the remaining 586 as *disclosure-only* adopters, most of which adopted in 2023–26 (Figure 1). We link adoption to about 85 million fund-quarter-issuer positions from 2013 to 2026 and estimate staggered-adoption difference-in-differences event studies, with never-adopting active funds as controls.

We reach three main conclusions. First, we find no evidence that adopters converge: introducing AI into security selection did not lead funds to construct portfolios similar to those of other adopters. For each issuer a fund adds, we ask whether AI adopters at other advisers added the same issuer in the same quarter, beyond what the issuer's popularity across all funds predicts. In all six Sun–Abraham designs, this excess co-entry is zero or negative (between $-2.0$ and $-0.4$ percentage points, all $p \geq 0.21$). Similarly, the herding measure of Lakonishok, Shleifer and Vishny (1992) among adopters after adoption is 0.012, against 0.034 in 200 matched groups of never-adopting funds. The same co-entry test does detect convergence where a decision process is shared: among pairs of funds at different advisers that share a subadviser, excess co-entry is 1.7 percentage points, against 0.04 for matched pairs ($p = 0.001$). Beyond naming their subadvisers, however, funds rarely disclose the specific AI tools or vendors they use.

The absence of convergence does not imply the absence of change. Second, adopters searched more widely, consistent with the idea that AI lowers search costs. After adoption, adopting funds held 8.5 to 14.5 log points more distinct issuers than never-adopters; for the median adopter, this is an increase of 10 to 18 names from a pre-adoption median of 114. The estimate is stable across eighteen event-study designs that vary the window, the reference quarter, and the comparison group, and it is robust to two alternative estimators and to randomization inference. Adopters add more names per quarter, mostly as small positions, and they tend to exit more positions within two quarters, which suggests short-lived trial positions.

Third, a finding links the first two: the additional issuers come from the same terrain, in industries and countries, as the fund's prior holdings. This is to be expected, since adopting AI does not change a fund's investment mandate. Nor are the new names smaller or more thinly covered by analysts. Instead, they are mostly constituents of the funds' stated benchmarks, a pattern more commonly associated with closet indexing than with additional search. Curiously, this happens without convergence across funds, even though many funds

share a small number of benchmarks.

Adopting funds hold 1.0 to 1.5 percentage points more of their benchmark's constituents, although their active share against the true index weights does not change. The result is strongest among funds benchmarked to non-U.S. indices and much weaker among funds benchmarked to U.S. indices. Three further results, each resting on small samples, show where the effect arises: it is concentrated among adopting funds that do not delegate management to a subadviser; it appears only among funds whose portfolio managers lack academic training in a technical (STEM) field; and it does not spill over to other funds managed by the same adviser. Finally, adopting AI in the investment process does not appear to affect fund performance, flows, or fees.

The main threat to these estimates is that funds choose when to adopt, and that what we date is disclosure rather than use. A fund that plans to change its portfolio might also write AI into its prospectus, and a fund may have used machine learning for years before saying so. The funds that describe AI also differ from others: 42 percent were quantitative or quantamental before adoption, against 14 percent of never-adopters, and the 102 adopters in the event-study sample sit at 27 advisers, two of which account for 18 each. We address selection with comparison groups that share the plausible confounders. Disclosure-only adopters wrote AI into their prospectuses in the same years and under the same pressures; relative to them, the breadth effect is 6.9 to 12.4 log points. Index funds run by the same adviser share its research infrastructure and filing calendar but cannot act on an AI signal; relative to them, the effect is 7.0 to 12.2 log points. Adviser-by-quarter fixed effects absorb anything common to a firm in a given quarter. Because the amendment that introduces AI often coincides with other changes, we estimate every design with three reference periods and treat any pattern that appears with only one of them as an artifact. Several early results, including apparent gains in flows and alpha, failed this test. What the design cannot rule out is a fund-specific repositioning that begins when the amendment is drafted. Issuer counts dip two to four quarters before adoption, and although the estimates are as large in designs that never use the two quarters before adoption, we describe the breadth effect as following disclosed adoption rather than as caused by the first use of AI.

One reading ties these findings together. As these funds describe it, AI predicts the value of candidates from data the fund already has. That makes it cheaper to evaluate one more name inside the universe the fund already covers, but not to build coverage of new industries, markets, or data. A fund that can evaluate names more cheaply holds more of them, draws them from the universe it already covers, and enters them as small positions, and the marginal names add little to returns. As long as each fund's model is trained on its own data, the extra evaluations are independent across funds, and portfolios do not converge.

This is what Agrawal et al. (2024) call a point solution: cheaper prediction dropped into a decision process whose structure is unchanged. It is also what models of rational inattention predict for added processing capacity, which goes first to the assets an investor already follows (Sims 2003; Van Nieuwerburgh and Veldkamp 2010). The reading also explains why the added names are benchmark constituents, which a manager already knows. Knowing a name is not the same as having evaluated it: a team with limited analyst time leaves much of its benchmark unexamined, and a model that scores every constituent each period makes the next one cheap to consider. A complementary and more mechanical reading is that a model ranking the whole benchmark, paired with an optimizer that controls tracking error, spreads weight across many small positions in benchmark names. Both readings imply wider books inside the same universe and, with fund-specific models, no convergence; Section 6 discusses what separates them.

The findings bound the regulatory concern rather than settle it. That concern centers on two scenarios: many funds relying on a few common models, vendors, or data sets, and signals that react together to news or to market stress. The adopters we study are early, built their own models, managed about one percent of active-fund assets by the first quarter of 2026, and are observed in quarterly holdings. What we show is that adopting in-house AI did not make these funds' portfolios homogeneous. We cannot show that widely shared AI tools would not, or how AI-driven funds trade within a quarter under stress. The tests do carry over to that case: once it is known which funds share a model or a data feed, the same pair statistics can be computed from holdings.

*Related Literature.* The paper contributes to three strands of work. The first asks whether algorithms homogenize choices. Shared recommendation systems can concentrate consumption or diversify it, depending on how they are targeted (Fleder and Hosanagar 2009; Zielnicki et al. 2025). Shared prediction models can make hiring or lending decisions correlated across firms (Kleinberg and Raghavan 2021; Bommasani et al. 2022), although competition among the firms that use such models can favor diverse ones (Raghavan 2026). Regulators have carried the concern into financial markets (Financial Stability Board 2017; International Monetary Fund 2024). In finance, venture capitalists who adopt data technologies tilt toward startups that resemble their historical data (Bonelli 2026), and funds that trade in line with machine-learning signals have been found to outperform (Hu, Rohrer and Zhang 2025). To our knowledge, we provide the first direct test of convergence among firms that adopt AI in a market where every choice is observed, together with a positive control that shows what convergence from a shared decision process looks like.

The second strand concerns how firms adopt AI and how adoption can be measured. Most measures come from job postings, patents, task-exposure indices, or keyword counts

(Goldfarb et al. 2023; Acemoglu et al. 2022; Eloundou et al. 2024; Eisfeldt et al. 2026). These capture capacity, intent, or exposure rather than use in a particular decision. Mandatory process disclosures, read by a language model, let us observe use inside a single decision and grade what the AI is said to do. They also separate substantive from symbolic adoption (Westphal and Zajac 1994). Generic AI disclosure is cheap talk in the sense of Farrell and Rabin (1996), and we show that it has no portfolio consequence; studies that count it as adoption will therefore understate the effects of substantive use.

The third strand is the economics of search and digital technology. Lower search and evaluation costs can expand the variety of options considered or concentrate attention on a familiar core (Stigler 1961; Evans 2008; Van Alstyne and Brynjolfsson 2005; Reimers and Waldfogel 2021). When books became cheaper to produce, the number of titles rose, and the question became where in the quality distribution the extra titles fell (Reimers and Waldfogel 2026). Organizational theory frames the same issue as whether search is local or distant (March 1991; Levinthal and March 1993; Katila and Ahuja 2002). For the same decision-makers, we measure the number of options considered, their distance from what was already held, and their overlap across firms. Research on AI in asset management has concentrated on return forecasting and performance (Gu, Kelly and Xiu 2020; Cao, Jiang, Wang and Yang 2024; Abis 2020; Chen, Sialm and Xu 2026; Bonelli and Foucault forthcoming); it has not measured what AI does to the scope of a portfolio.

The paper proceeds as follows. Section 2 describes what funds do with AI and how they disclose it. Section 3 describes the data and how we measure adoption, and Section 4 sets out the empirical strategy. Section 5 reports the results, beginning with convergence. Section 6 discusses them, and Section 7 concludes. The appendix covers the data linkage, the classification codebook, the robustness checks, and responses to further critiques.

# 2 Institutional Background

## 2.1 AI in the Investment Process

An actively managed fund employs an investment adviser, which may in turn delegate the portfolio to a subadviser, to choose securities within a mandate set out in the prospectus. The mandate specifies a universe (for example, U.S. large-capitalization stocks, emerging-market equities, or high-yield bonds), a benchmark index against which performance is reported, and a method. The method is where AI enters. In the filings we read, funds describe AI in roles that fall along a spectrum. At one end, AI is a research tool: "The Adviser uses artificial intelligence ('AI') tools as part of its investment research process." Further along,

it is one input among several to a ranking: "The Sub-Adviser's quantitative process uses a combination of relative, intrinsic, and artificial intelligence models to rank companies within each economic sector." At the far end, the model selects the holdings: "The AI model is generally updated monthly and typically selects between 60 and 190 equity securities that exhibit strong value characteristics."[1]

Three features of these descriptions matter for what follows. First, the AI is almost always a prediction tool applied to a universe the fund has already defined. It ranks, screens, or forecasts candidates drawn from the mandate, using data the adviser already collects. We find no descriptions of AI used to discover new asset classes or markets. Second, the models are mostly built in-house. When we extract the supplier of the AI from each passage, only a handful of funds name an outside vendor, and once funds under common management are set aside, no two funds at different advisers name the same one in describing their own process; three cross-adviser pairs share a vendor named in an index methodology (Section 4.2).[2] Third, a human remains in the loop. Many filings say so explicitly, and even funds whose models select the holdings describe a manager who evaluates the output. Agrawal et al. (2018) describe this configuration as cheaper prediction paired with unchanged judgment.

The universe to which the prediction is applied is an institution in its own right. An active fund's mandate names the market it invests in and the index it is measured against, and most of what the fund does is organized around that index. Analysts cover the index's sectors, data are licensed for its markets, risk is managed as deviation from its weights, and performance is reported against it. Moving outside the index is possible, but it is costly along each of these dimensions, and clients see it as tracking error.[3] This is what we mean below by a fund's covered universe, and it is why the benchmark is the natural place to look for the names that a cheaper screen adds.

Whether AI could make funds' portfolios homogeneous depends on how much of this machinery funds share. The industry has one important channel for sharing: delegation. A fund's sponsor may hire a subadviser to run the portfolio, and the same subadviser often runs similar mandates for several sponsors. In the N-CEN census, 5,619 active funds list a subadviser; there are 1,025 distinct subadvisers, and 394 of them serve funds at two or more

[1] All quotations are verbatim from N-1A-family filings and appear, with their classification, in Appendix Table A4. The last two roles, together with an intermediate role in which AI is the primary input to selection, are the three that qualify as a described process in Section 3.2.

[2] The best-known exception is an ETF whose subadviser "leverages IBM's Watson AI to conduct an objective, fundamental analysis of U.S.-listed common stocks"; it is an exception precisely because it names a shared platform (Appendix Table A5). Section 5.6 describes the vendor extraction.

[3] Active share, the share of a portfolio that deviates from its benchmark, is a standard summary of how far a manager departs from the index (Cremers and Petajisto 2009; Petajisto 2013). The median equity fund in our sample holds 16 percent of its corporate book in names outside its stated benchmark (Section 3.3).

advisers. Two funds at different sponsors that share a subadviser therefore share a team, a model, and a research process. This is the kind of common input that the homogenization argument envisions for AI, and we use these pairs as a positive control in Section 5.1. AI vendors, by contrast, are barely visible in the filings: each of the few suppliers named in a fund's own process serves a single fund complex.

The funds that describe such a process are a small part of the industry. By the first quarter of 2026, the described-process adopters in our estimation sample managed about one percent of the assets of active funds, and funds carrying generic AI language close to five percent. Our question is therefore not whether AI has already reshaped the market. It is what happened to the portfolios of the firms that adopted AI first, and whether those portfolios moved together.

## 2.2 Disclosure Requirements

Our ability to date adoption rests on the disclosure regime for registered funds. A fund registered under the Investment Company Act of 1940 files its prospectus on Form N-1A. Item 4 of the form requires a summary of the fund's principal investment strategies, and Item 9 a fuller statement of how the adviser selects investments. The prospectus is updated at least annually through a post-effective amendment and between updates through supplements.[4] A fund that begins to use a new method in its process is expected to say so in its next amendment, and each filing carries a date. Because every fund files the same form, a classifier that reads one fund's language applies equally to another's.

The disclosures have also become more numerous and less informative over time, which is why we cannot take them at face value. Through 2022, a fund that mentioned AI in its process almost always said what it did with it. From 2023, the language changed. Large fund complexes added risk factors about AI, statements that the adviser "may utilize AI technologies, including machine learning models and generative AI, to improve operational efficiency and in connection with research," and reassurances that "AI tools are not used to make autonomous investment and/or operational decisions." Such language typically appears in every fund of a complex at once. In the last quarter of 2025, for example, 94 active First Trust funds first carried such language about their process in the same filing cycle. Two regulatory events frame the period. In July 2023, the SEC proposed rules on

[4]Annual updates and non-material changes are filed under Rule 485(b) and become effective on filing; material changes are filed under Rule 485(a) and reviewed by the SEC staff; interim changes are filed as supplements under Rule 497. We read the initial registration (N-1A), the post-effective amendments (485APOS, 485BPOS), and the supplements and summary prospectuses (497, 497K). Misstatements in a registration statement expose the fund and its adviser to liability under Sections 11 and 12 of the Securities Act, which is what makes the strategy text more than marketing copy.

conflicts of interest in advisers' use of predictive data analytics (U.S. Securities and Exchange Commission 2023), and in March 2024 it settled its first enforcement actions against advisers for overstating their use of AI, a practice it called "AI washing" (U.S. Securities and Exchange Commission 2024).[5] For our purposes, the generic wave is both a nuisance and an opportunity. It is a nuisance because it must be separated from adoption, and an opportunity because the funds that joined it form a natural placebo group: they wrote AI into a prospectus in the same years as the adopters without changing what they did.

# 3 Data and Measurement

The treatment comes from the text of fund prospectuses, and the main outcomes come from the holdings that funds report to the SEC. Commercial data classify the securities held and supply returns, fees, benchmarks, and analyst information. Appendix A lists every source and describes the linkage.

## 3.1 Funds and Holdings

The universe is every U.S. open-end fund and ETF series that registers on Form N-1A and filed Form N-PORT between 2021Q3 and 2026Q2, excluding money market funds; we follow these funds' holdings from 2013 (CRSP) and from 2019Q4 (N-PORT). Mutual funds and ETFs file the same forms, and index funds stay in the sample as a placebo group. Form N-PORT discloses each fund's complete portfolio every quarter from 2019Q4. Before that, we use the CRSP Survivor-Bias-Free U.S. Mutual Fund holdings, which reach back to 2013 for most of our funds. Form N-CEN classifies each series as an index fund or ETF and names its adviser and any subadvisers, which tells us whether a fund is run in-house. Funds are linked to CRSP's returns and fee data through dated share-class tickers. The link covers 82.0 percent of series and 93.4 percent of net assets.[6]

A fund's *long book* consists of its positions with a positive share of net assets whose payoff profile is not short. Its *corporate book* contains equity, corporate bonds, loans, and notes. We aggregate positions to the issuer so that a fund's stock and bonds of the same company count as one issuer. To identify the issuer we use, in order, the Compustat GVKEY, the

[5]The 2024 settlements concerned two advisers, Delphia and Global Predictions, that had advertised AI-driven processes the Commission found they did not have. The episode is one reason we define the treatment by what the text says the AI does rather than by the presence of AI vocabulary. Appendix E shows that funds that described AI hedged their language somewhat more after 2023 but did not retract it, and that the break predates the enforcement.

[6]Most unlinked series are insurance and variable-annuity funds without tickers. For these funds, returns and flows come from N-PORT and expense ratios from the prospectus fee tables.

GLEIF ultimate-parent LEI, the issuer's own LEI, the six-digit CUSIP, and, failing these, the security itself.[7] The resulting position file has 85 million fund-quarter-issuer rows, and Table 1 summarizes the funds it covers. Our main breadth outcome is the logarithm of one plus the number of distinct issuers in the corporate book. We report effects in log points, which for the magnitudes here are close to percentage changes (10 log points is a 10.5 percent increase). A fund-quarter enters the estimation sample if corporate positions make up at least 25 percent of the long book and no more than 20 percent of corporate weight lacks an industry group. On the 30,500 fund-quarters observed in both holdings sources (2019Q4 to 2020Q4), issuer counts correlate at 0.92, with a median ratio of 1.00. Quarter fixed effects absorb the 2.2 percent step where funds switch from CRSP to N-PORT.

## 3.2 Measuring AI Adoption from Prospectus Text

We keep every sentence that mentions artificial intelligence, machine learning, deep learning, neural networks, natural language processing, or large language models. We attribute each sentence to a fund using the filing's series headers and the nearest preceding fund name.[8] Adoption is dated on amendments of funds already in operation, so a fund's own launch filing never counts. Some funds have no AI-free filing that restates the investment process (at least 10,000 characters, so that short supplements do not count) before their first AI filing. We treat these funds as left-censored and exclude them from the event studies rather than assign them a date.[9]

A large language model (Anthropic's Claude) labels each distinct sentence in three passes. The first pass separates four kinds of mention: *process* use, in which the adviser or its models use AI in the investment process; *thematic* exposure, in which AI is what the fund invests in; risk disclosure; and other mentions. The second pass separates *specific* use, a concrete task performed with AI, from *generic* statements about AI tools in support of research or about limits on AI's role. The third pass reads each process sentence in its passage. It records the type of statement; the weakest role the passage clearly supports for AI in selection (research support, one input, primary input, or AI that selects holdings); the tasks performed; and

[7] Compustat North America and Global provide each issuer's GVKEY, point-in-time GICS classification, and characteristics. Holdings are matched by CUSIP, ISIN, LEI, and GLEIF ultimate parent, and Bloomberg fills in industries for bonds and loans that Compustat cannot classify. After these steps, 95.7 percent of corporate holdings by weight carry a GVKEY and 96.6 percent an industry.

[8] A filing often covers many funds of one trust. The SGML header lists the series the filing covers and marks each as new or existing, which lets us separate a fund's launch from a later amendment. Text that names no fund is attributed to all existing series in the filing.

[9] The earliest adoption among the 132 adopters in Figure 1 is dated 2021Q1. The SEC's structured risk/return summaries, available from 2010, contain almost no description of AI in the investment process before 2018 and identify no adopter that the full text misses, so starting the full-text collection in 2018 costs little.

the techniques and data used.[10] The *described-process* treatment requires a specific sentence whose passage describes a fund or index process with AI as at least one input to selection.[11] Funds whose AI language never meets this standard are disclosure-only adopters.

Table 2 shows how the corpus falls into these cells. Most AI language has nothing to do with the investment process. Of 41,070 distinct sentences that mention AI, 2,018, about 5 percent, describe the adviser or its models using AI; the rest describe AI as a theme, a risk, or in passing. The second and third passes agree with each other as they should. Where the third pass finds research support, the second pass calls 4 percent of passages specific; where it finds AI as one input, the primary input, or the selector of holdings, the share rises to 77, 92, and 98 percent. The exception, and the reason we require both passes, is a set of 1,228 generic-risk passages stating that a fund's models, "including" AI, may perform poorly. The second pass calls 38 percent of these specific, and without the third pass they would count as adoption.

The labels come from a proprietary model. The prompts and the cached labels are part of the replication files, so the classification can be audited even if the model cannot be rerun identically.

## 3.3 Benchmarks and Other Measures

Each fund's stated benchmark is Bloomberg's primary benchmark field where available, and the prospectus benchmark otherwise. For the 20 U.S. equity indices that funds name, we observe the constituent list at each quarter-end from 2019Q4 (19 lists come from Bloomberg; the Russell 3000 is the union of the Russell 1000 and 2000), and we match 99.8 percent of constituents to our issuer keys.[12] The constituent history of MSCI indices was not available to us. For MSCI EAFE, Emerging Markets, and ACWI ex USA, membership instead comes from an index fund that tracks the index and is run by a family other than the benchmarked fund's adviser, so that no fund is measured against its own family's holdings. MSCI ACWI combines the ACWI ex USA tracker with S&P 500 constituents, and MSCI World uses the S&P Global 1200 list. Index weights are not available for most indices, so in place of active share (Cremers and Petajisto 2009) we measure its floor: the share of the fund's corporate book held in names that are not constituents of its benchmark. The median equity fund

[10] Identical sentences receive the same label wherever they appear. Labels are cached by prompt version, so a change to the model, prompt, or output schema forces relabelling rather than mixing versions. Appendix Table A3 lists the labels, and Appendix Table A4 quotes one example from each cell.

[11] The index branch matters for index funds, which enter only the placebo designs. Among the 132 active adopters, 12 have an index-methodology passage in their adoption filing, and only one qualifies through such a passage alone; the rest also describe AI in their own process.

[12] Thirteen of the 20 are Russell indices; the others are S&P and Nasdaq indices and the Dow Jones Industrial Average. The MSCI indices are EAFE, Emerging Markets, ACWI, ACWI ex USA, and World.

holds 16 percent of its book outside its benchmark. For the 13 Russell indices, we also have FTSE Russell's constituent weights at every quarter-end, which give true active share for Russell-benchmarked funds (Appendix D.7 describes the benchmark construction and its validation).

Three further measures complete the data. First, analyst coverage is measured on each fund's equity holdings from IBES, as the number of analysts with an outstanding recommendation on the issuer at each quarter-end, with Bloomberg's year-end analyst counts as a check; consensus ratings come from Bloomberg and IBES at each quarter-end. Second, manager education comes from the latest prospectus biography of each portfolio manager, read by the same model. Of 11,437 named managers, 5,041 state a degree, and 17.7 percent of these hold a STEM degree. Third, each fund's strategy text before adoption is classified as quantitative, quantamental, fundamental, rules-based, or allocation.

## 3.4 Adoption Patterns

Figure 1 shows the cumulative number of active funds that have adopted, by type. Described-process adoption is a thin stream that grows steadily from 2021 and jumps in 2025, when 71 funds first described an AI process. Disclosure-only adoption is almost entirely a 2023–26 phenomenon: 585 of the 586 such funds first mentioned AI in 2023 or later, most of them in family-wide filing cycles. In other words, before generative AI reached the public in late 2022, a fund that mentioned AI in its process almost always said what it did with it; afterwards, most did not.

Table 1 describes the adopters in their last screened quarter before adoption and compares them with never-adopting active funds. Described-process adopters are broader to begin with. The median adopter held 114 issuers, against 79 for the median never-adopter fund-quarter in the same calendar quarters, and spanned 22 industry groups and 12 countries, against 20 and 9. Adopters are somewhat smaller (median net assets of $242 million against $317 million), cheaper (expense ratio of 0.79 against 0.88 percent), more often ETFs (30 against 17 percent), and about as often subadvised (43 against 50 percent). Disclosure-only adopters resemble the never-adopters on most dimensions except size: they are larger, which reflects their origin in large fund complexes. The adopters are also concentrated. The 132 described-process adopters sit at 35 advisers, the largest of which accounts for 20 percent of them, so we cluster inference by adviser throughout. A further 61 active funds already described AI in the earliest full prospectus we read. Most of these are AI-native launches rather than adoptions; the best known is the AI Powered Equity ETF, launched in 2017 with a stock-selection process built on IBM Watson. We exclude these funds because they have no pre-period. Panel B of Table 1

gives the estimation samples. Adopters enter only if they have at least one pre-adoption quarter of holdings that passes the corporate-book screen. Of the 132 described-process adopters, 8 have no holdings before adoption and 22 have holdings but no quarter that passes the screen, which leaves 102 in the Sun–Abraham sample and 100 in the TWFE sample; Panel A describes 100 of the 102, because two have screened quarters only more than four quarters before adoption. The samples include 281 to 309 disclosure-only adopters, against about 6,800 never-adopters, in 217,000 to 226,000 fund-quarters.

# 4 Empirical Strategy

## 4.1 Estimator and Research Designs

Adoption is staggered across 2021–2026, and treatment effects may vary by cohort and with time since adoption. In that setting, two-way fixed effects can weight cohort effects in misleading ways (de Chaisemartin and D'Haultfœuille 2020; Goodman-Bacon 2021; Roth et al. 2023). Our main estimator is therefore the interaction-weighted event study of Sun and Abraham (2021):

$$y_{it} = \alpha_i + \lambda_t + \sum_g \sum_{k \neq r} \beta_{gk} \, \mathbf{1}[G_i = g] \, \mathbf{1}[t - g = k] + \varepsilon_{it}, \tag{1}$$

where $y_{it}$ is the outcome of fund $i$ in quarter $t$, $G_i$ is the fund's adoption quarter for each adopter type, $r$ is the reference quarter, and event times beyond the window are binned. The cohort-specific effects $\beta_{gk}$ are averaged across cohorts, with weights equal to each cohort's share of treated observations at event time $k$. Never-adopting active funds are the control group. Described-process and disclosure-only adopters enter as separate treated groups in the same regression, so the difference between them is estimated directly. We report three statistics: the average post-adoption effect over event quarters 0 to 2 or 0 to 6; the average of the pre-adoption coefficients from quarter $-8$ onward, excluding the reference quarter and any dropped quarters; and the difference between the two adopter types.

Inference is clustered by adviser, because adoption is often an adviser-level decision and the treated advisers are few. For the primary outcome, we resample advisers with replacement and re-estimate the event study in each of 200 draws, so that the p-values reflect how much the estimate depends on which advisers adopted. For the many secondary outcomes, we report wild cluster restricted bootstrap p-values (Webb weights, 999 replications), which hold the treated advisers fixed. For the primary outcome these p-values are smaller than the resampling p-values (Appendix D.4), so secondary results that are significant only near

the 5 percent level should be read with that in mind.[13] Performance outcomes are estimated on all active fund-quarters from 2019Q4, without the corporate-book screen, and with event time from −4, because factor-model alphas begin in 2019; their samples therefore differ from those of the breadth outcomes. Two-way fixed-effects (TWFE) comparisons use a single post-adoption indicator per adopter type, with adopters' observations restricted to event times −8 to +8. The text reports the Sun–Abraham estimates. TWFE estimates appear in the tables, and the text discusses them only where they disagree with the Sun–Abraham estimates or where TWFE is the only estimator we use: the leave-one-adviser-out checks, randomization inference, a pre-period lead, and a split of post-adoption event time.

Crossing two windows, three reference choices, and three comparison designs gives eighteen designs. The three comparison designs are as follows. The *disclosure* design uses never-adopters alone as controls; it is so called because disclosure-only adopters are its placebo contrast. The *family* design adds the index funds of adopters' advisers, dated at the adviser's first described-process adoption.[14] The *within* design adds adviser-by-quarter fixed effects. The three reference choices are quarter −1, quarter −2, and quarter −3 with quarters −2 and −1 dropped. We call an effect robust when it keeps its sign and significance across the six Sun–Abraham variants of a design, and we say so when it does not. This criterion is descriptive, not a correction for multiple testing. The six variants share the same data and are far from independent, and the paper examines many outcomes. We therefore organize the evidence around one primary outcome, the number of distinct issuers held, for which we also report randomization inference and sensitivity bounds. We call secondary results that are significant in only some variants suggestive, and we do not carry them into the abstract or the conclusions.

## 4.2 Measuring Convergence

The convergence tests need outcomes that distinguish shared inputs from shared tastes. Funds buy the same names for many reasons that have nothing to do with AI, among them index reconstitutions, initial public offerings, and simple popularity, so a raw overlap in additions is

[13] With 26 to 35 treated clusters among about 780, conventional cluster-robust standard errors over-reject (Davidson and MacKinnon 2010). Several post-adoption event-time cells rest on one or a few advisers, and a procedure that keeps the treated advisers fixed cannot learn from such cells how much the estimate depends on them. The figures use the analytical clustered covariance for their confidence bands; the tables and text use bootstrap p-values.

[14] Index funds cannot act on an AI signal because their holdings follow an index. An index sibling that "adopts" on the adviser's date is therefore a placebo for everything the adviser did in that quarter that touched all of its funds: data vendors, research systems, reorganizations, and the filing calendar itself. Index funds with their own AI language are dropped from the sibling group, and quarter effects are estimated separately for active and index funds.

not evidence of a common model. Our first measure corrects for this. For each name a fund adds in a quarter, we record whether a described-process adopter at a different adviser adds the same name in the same quarter, and we average over the fund's new names to obtain its co-entry share. We then compute the share that would be expected if the other adopters drew their additions at random from all active funds' additions that quarter, scaled by the number of names the reference set adds per fund. The difference is the fund's *excess co-entry share*. A positive value means the fund's additions coincide with other adopters' additions more than popularity predicts. We also compute a weight-based version, a rate version, and versions that exclude names involved in an index event, that split names inside and outside the fund's benchmark, and that use the equity book only. Each is estimated in the same event-study designs as the breadth outcomes. The analogous share computed against not-yet adopters and against disclosure-only adopters serves as a placebo.

The second measure moves to the level of the stock. The herding statistic of Lakonishok, Shleifer and Vishny (1992) compares, for each stock and quarter, the share of the group's trades that are purchases with the share expected under independent trading, and averages the absolute imbalance across stock-quarters. We compute it for adopters' trades after adoption, for the same funds' trades before adoption, and for disclosure-only adopters. We compare each with the statistic for 200 groups of never-adopting funds matched on Lipper class, quintile of trading activity, and net-assets tercile in the same quarters. The third measure takes whole portfolios: the cosine similarity between a fund's issuer weights and the aggregate issuer weights of described-process adopters at other advisers, again estimated in the event-study designs.

The co-entry measure requires entry in the same quarter, which is a strict notion of convergence: two funds whose models agree could add the same names a quarter apart. The similarity measure compares levels of holdings rather than additions, so it would register slower convergence of that kind. For the herding statistic, decisions are netted within adviser, so a family that adds a name in three funds counts as one buyer. A stock-quarter enters when at least three advisers in the group trade the stock, and a group-quarter enters when at least five stocks qualify. In the adopters' post-adoption quarters, the number of qualifying stocks ranges from 6 to 760, with a median of 57, and the pooled statistic weights each quarter equally. Each placebo group replaces each adopter with a never-adopting fund at a never-adopting adviser in the same Lipper class, quintile of trading activity (entries plus exits), and net-assets tercile. Because adopters span U.S., international, and bond mandates, we also compute the similarity measure against AI users in the same Lipper asset class only.

Finally, we construct a positive control. For pairs of active funds at different advisers that share a subadviser, we compute the pair's excess co-entry, averaged over both directions, and

compare it with matched control pairs, drawn from the same quarters, that share nothing. We sample 1,500 of the 124,603 such pairs, at 178 advisers, with 13 matched control pairs each. We report both the level difference and a difference-in-differences around the quarter in which the later fund of the pair first used the shared subadviser. A shared subadviser, which means a common management team, is the strongest form of shared input. This control therefore establishes that the measure detects a shared decision process. It does not establish that the measure would detect the weaker correlation that two independent teams licensing the same model or data feed might produce. The closest we can come to that case are pairs of funds that share an index provider (354 pairs at 7 advisers) and the three cross-adviser pairs that name the same AI vendor in an index methodology, which we report alongside.

## 4.3 Threats to Identification

*Selection into Adoption.* Funds that describe AI differ from other funds (Table 1): they are broader, more often quantitative, and concentrated at 35 advisers. We choose comparison groups that share the plausible confounders. Disclosure-only adopters made the same decision to write AI into a prospectus in the same years. Index funds of the same adviser share its infrastructure and calendar but follow an index. Adviser-by-quarter and quarter-by-style fixed effects absorb shocks common to a firm or a style. Removing the 52 never-adopters whose strategy text describes machine-learning-like methods without AI vocabulary changes nothing.[15]

*Disclosure versus Use.* What we date is the first filing that describes an AI process, not the first use of AI. A fund that used machine learning before writing it into its prospectus is dated late. This spreads its true post-adoption quarters into the pre-period and attenuates the estimates toward zero. A fund that described AI in its earliest full prospectus is excluded as left-censored rather than mis-dated (61 active funds, Section 3.4). The disclosure-only group is a placebo for the decision to write AI language into a prospectus, which large complexes made in family-wide filing cycles. It is not a placebo for a fund-level plan to change the portfolio. That plan remains the leading alternative explanation, which is why the timing evidence matters.

*Timing.* The prospectus amendment that introduces AI can coincide with other changes, and issuer counts dip two to four quarters before adoption. We therefore estimate every design with three reference periods and treat a pattern that appears with only one of them as an artifact. Several early results, including apparent gains in flows and alpha, were of this

[15] These are 0.4 percent of the never-adopting funds in the universe, identified by the same language model from strategy and risk text of 2019 onward: funds that describe algorithms, sentiment, or alternative data without using an AI term (Appendix D.3).

kind (Section 5.7). Launches never count as adoptions. The breadth effect is also the same for adopters with and without a manager arrival in the adoption quarter or the six quarters before it (32 and 64 funds; difference −2.1 log points, p = 0.71).

*Measurement.* The treatment could measure marketing rather than practice. The described-process standard requires a passage that assigns AI a role in selection, and the generic wave serves as the placebo for language without practice. Misclassification moves true adopters into the disclosure-only group and attenuates the contrast. Appendix E takes up further critiques, including survivorship, counting artifacts, and the concentration of adopters at few advisers.

# 5 Results

We present convergence first (Section 5.1). We then turn to how much funds searched and where (Sections 5.2 to 5.5), which explains the answer, and finally to who broadened and whether it paid (Sections 5.6 and 5.7). The text reports the headline estimates; the sub-tests behind each statement are in Appendix D.1.

## 5.1 Convergence

If AI tools make funds' portfolios homogeneous, the effect should show first in what funds add. Figure 2 shows the three measures of Section 4.2 around adoption. The first is the excess co-entry of a fund's new names with other adopters' additions, which shared inputs raise and independent models leave at zero. The second is the Lakonishok, Shleifer and Vishny (1992) herding measure among adopters' trades in each stock. The third is the cosine similarity of a fund's issuer weights to the aggregate book of other AI users, which would register convergence of whole portfolios. Appendix Table A15 reports the co-entry and herding tests design by design.

Before reading the adopters' result, we check what the test can and cannot find. Pairs of funds at different advisers that share a subadviser share a decision process, because the same team informs both portfolios. Among 1,500 such pairs sampled from the 124,603 in the data, excess co-entry is 1.7 percentage points, against 0.04 for matched control pairs that share nothing (p = 0.001). Weaker shared inputs are another matter. Pairs that share only an index provider show no excess co-entry beyond their controls (p = 0.73), and neither do the three cross-adviser pairs that name the same AI vendor in an index methodology (p = 0.31). The test is therefore well powered against the case in which AI amounts to a common manager, and of unknown power against the case in which it amounts to a common data feed.

Described-process adopters show no such convergence. Panel (a) of Figure 2 traces their excess co-entry share around adoption: flat before adoption and at or below zero through the fifth quarter after it. The average post-adoption effect is between −2.0 and −0.4 percentage points across the six Sun–Abraham designs (all p ≥ 0.21), and it does not turn positive when index reconstitutions are excluded, when the test is run inside or outside the fund's benchmark, when it is restricted to the equity book, or when it is restricted to fund-quarters with several new names. If anything, the raw overlap with other adopters' additions *falls*, by 4 to 7 percentage points, because adopters add names that fewer other funds are buying.

Panel (b) moves to the stock level. Among adopters after adoption, the herding measure is 0.012, against 0.034 in 200 groups of never-adopting funds matched on class, trading activity, and size, a difference of more than three placebo standard deviations (two-sided p = 0.005). Adopters trade less in step than comparable funds, not more.[16] Panel (c) turns to whole portfolios, which would also register convergence that is slower than same-quarter co-entry. The similarity of a fund's issuer weights to the aggregate portfolio of AI users at other advisers changes by 0.3 to 1.0 percentage points after adoption (all p ≥ 0.26), and by 0.3 to 1.3 points against AI users in the fund's own Lipper asset class (all p ≥ 0.29), where portfolios are comparable.

The scope of the finding follows from what the adopters shared, which was little. Once pairs under common management are removed, no two funds at different advisers name the same AI vendor in describing their own process, and funds whose process reads text or alternative data show no more co-entry than funds that use only market and fundamental data. The positive control establishes that a shared decision process would have been visible. Whether a shared data feed would be visible is a question for a later period, and one that the same pair tests can answer when such sharing appears.

## 5.2 Portfolio Breadth

What did change is how much funds considered. Figure 3 shows the event study for the number of distinct issuers held. Described-process adopters' issuer counts are flat before adoption, except for a dip of up to five log points two to four quarters earlier. They then rise from the adoption quarter onward, to about 10 log points by the second quarter and 15 to 27 log points after four or more quarters. Disclosure-only adopters, the funds that added AI language without describing a process, move little.

[16] The statistic is $|p_{jt} - \bar{p}_t| - \mathbb{E}[\,|p_{jt} - \bar{p}_t|\,]$, averaged over stock-quarters, where $p_{jt}$ is the share of the group's trades in stock $j$ and quarter $t$ that are purchases, $\bar{p}_t$ is the average of $p_{jt}$ across stocks, and the expectation is taken under independent trading, in which each trader in the stock buys with probability $\bar{p}_t$. Subtracting it removes the imbalance expected by chance when few traders are active in a stock (Section 4.2).

Table 3 reports the average post-adoption effect in all eighteen designs. In the disclosure design with reference quarter $-1$, adopters hold 8.7 log points more issuers over the first three quarters and 12.8 over the first seven. Across the eighteen designs, the range is 8.5 to 14.5 log points, with every p at most 0.035. For the median adopter, which held 114 issuers before adoption, this is 10 to 18 additional names. The pre-period averages are small and insignificant ($-3.6$ to $+1.4$ log points, all $p \geq 0.29$). The placebo contrasts are as large as the main effect: 6.9 to 12.4 log points relative to disclosure-only adopters ($p \leq 0.03$ in all six designs) and 7.0 to 12.2 relative to the adviser's own index funds ($p \leq 0.06$, and below 0.05 in five of six). The disclosure-only group pools funds whose filings describe AI only as research support with funds whose AI language is boilerplate. Neither half broadens on its own (research support $-0.2$ to $+2.9$ log points, all $p \geq 0.45$; boilerplate $+1.1$ to $+3.4$, all $p \geq 0.08$), so the pooled comparison does not hide an effect of AI used in research (Appendix D.1). A broader treatment that counts any specific AI sentence, including the generic-risk boilerplate that the third pass removes, gives 5.5 to 8.8 log points, the attenuation that misclassified adopters would produce. The effect is in equity holdings, which is largely a statement about who adopts: 89 of the 102 adopters held essentially no corporate debt before adoption, so whether AI broadens bond books is a question this sample cannot answer.

*Entry and Exit of Positions.* Following each position at the fund-issuer-quarter level (Table 4, upper panel) shows what breadth is made of. The number of names a fund adds in a quarter rises by 24 to 58 log points, a quarter to three-quarters more additions for a fund that adds about eight a quarter, and short-lived positions become somewhat more common, a suggestive result. Per-name survival, the size at which a new name enters, and the scaling-up of names that survive do not change, and the added names are mostly small positions.[17] More trials, more early reversals, and no change in what happens to the trials that work: the added names are ones about which the fund was previously close to indifferent. Within two years, adopters show no turn from trying new names to building positions in the ones they have.

## 5.3 Robustness and Pre-Adoption Trends

Alternative estimators agree (Appendices D.3 and D.4). Callaway and Sant'Anna (2021) estimates on the same samples have the same sign but are smaller and less precise, for reasons Appendix D.4 explains: 6.4 log points over event quarters 0 to 2 ($p = 0.004$) and 10.8 over 0 to 6 ($p = 0.06$). Stacked difference-in-differences, randomization inference, dropping any

[17]The share of names below 0.5 percent of corporate weight rises by 4.2 to 7.1 percentage points from a mean of 37 percent in all six designs (Table 5). For the median adopter, with $242 million in assets, the weight that moves into such names amounts to roughly $7 million.

one treated adviser, and quarter-by-investment-style fixed effects all give effects of similar size (Appendix D.1). The effect is broad rather than carried by a few funds: 64 percent of adopters broaden. Nor is it a feature of quantitative funds that happened to add AI language: fundamental adopters broaden by 10.6 to 15.9 log points, at least as much as quantitative ones.

The main qualification is the pre-period dip visible in Figure 3. A joint test of the pre-period coefficients rejects ($p = 0.002$), and a reader could see the post-adoption rise as a recovery from the dip or as the continuation of a trend. Three facts weigh against these readings. First, the designs that never use quarters $-2$ and $-1$ give 10.3 to 13.7 log points, as large as the designs that do. Second, the rise continues well past the level from which the dip began. Third, the Rambachan and Roth (2023) bounds show the effect surviving post-adoption violations of parallel trends up to 0.4 times the largest pre-period change, as well as a linear continuation of the pre-trend. What we cannot rule out is a repositioning that begins when the prospectus amendment is drafted and continues after it is filed. For that reason, we date the broadening to the disclosed adoption of AI rather than to its first day of use, and the causal reading rests on the comparison groups rather than on the pre-period alone.

## 5.4 Distance of Added Securities

If the extra names had come from new territory, convergence would have been possible even with fund-specific models, because funds extending into the same unfamiliar regions would meet there. They did not. Figure 4 puts the breadth result next to three measures of terrain. Issuers held rise, but the number of GICS industry groups barely moves ($+1.2$ to $+2.3$ log points, significant only with reference $-1$), and sub-industries rise by less than issuers (4.8 to 7.7 log points), which is what happens when a fund adds names inside sectors it already holds. The number of countries rises late and noisily, but only because funds add more names: the share of new names in a country the fund did not already hold does not rise.

The added names themselves (Table 4, lower panel) are no more distant than the names the fund was already adding by any measure we have. They are no more often new to the adviser, new to the fund, or in an industry group or country the fund did not hold. They are no smaller, younger, or more research-intensive than the names kept, no farther in product markets (Hoberg and Phillips 2016), and no more thinly covered by analysts, and their consensus ratings show no robust pattern (Appendix D.1). The broadening also stays inside the adopting fund: summing holdings across all funds of an adviser, or of a corporate group,

shows no broadening once launches of new AI funds are excluded.[18]

## 5.5 Position Relative to the Benchmark

Where, then, did the extra names come from? Table 5 locates them relative to each fund's stated benchmark for the 87 described-process adopters with a benchmark we can observe. The count of benchmark constituents held rises by 8.8 to 14.4 log points in all six designs ($p \leq 0.002$), as much as the count of all issuers, and by 10.7 log points relative to disclosure-only adopters ($p = 0.001$). Funds come to hold 1.0 to 1.5 percentage points more of their benchmark's constituents, from a mean of 18 percent. The count of names outside the benchmark rises too, but significantly in only two of six designs. Among never-adopters, funds that hold more names grow faster outside their benchmark than inside it, so breadth alone would have added more names outside. The extra names are therefore mostly benchmark constituents, although we cannot say that adopters avoid names outside the benchmark. The result is strongest among the 34 adopters benchmarked to MSCI indices, international and global equity funds, and suggestive rather than established among the 53 benchmarked to U.S. indices (Appendix D.7). Breadth itself is smaller among the U.S.-benchmarked funds. One reading is that a model lowers evaluation costs most where a human team's coverage is thinnest, in foreign markets; but MSCI membership is proxied by tracker funds whose sampled holdings can differ from the index, so we do not treat the difference between the two groups as a finding, although a constituent list from another index provider gives the same result.

Holding more benchmark names moves a portfolio toward its index, and Figure 5 shows by how much. After adoption, the weight held outside the benchmark falls by 0.8 to 1.9 percentage points from a median of 16 percent, significantly in four of six designs; the path is flat before adoption and steps down after it, and among never-adopters the number of names held is unrelated to that weight. True index weights show how far the move goes. For the 32 adopters benchmarked to a Russell index, active share against FTSE Russell's constituent weights does not change robustly ($-0.5$ to $-1.0$ percentage points). Funds add benchmark names without moving their weights toward the index's weights.

*Closet Indexing.* A move toward the benchmark raises the question whether adopters drift toward closet indexing, holding a portfolio close to the index while charging active fees (Cremers and Petajisto 2009; Petajisto 2013). The drift we measure is small: active share falls

[18] Including newly launched AI funds produces a firm-level gain that reflects the new funds' holdings rather than a change at existing funds. The firm-level books sum every fund of the adviser, so they are sensitive to quarters in which some of the adviser's funds have not yet filed. 2026Q2 is such a quarter, and we drop it from the firm and group tests.

by at most one percentage point from a median of 80 percent among Russell-benchmarked equity funds, and expense ratios do not change (Table 6). Two limits apply. True active share is available only for Russell-benchmarked funds, not for the international funds where the benchmark result is strongest. And a fund that decided to hug its benchmark when it rewrote its prospectus would produce the same pattern, which is the fund-level repositioning that the comparison groups cannot rule out (Section 4.3).

## 5.6 Heterogeneity and the Organization of Portfolio Management

Figure 6 splits described-process adopters along several dimensions, each estimated jointly against never-adopters and disclosure-only adopters (Appendix Table A11 gives the estimates, differences between groups, and TWFE estimates for the splits in Figure 6). The splits divide about a hundred funds at a few dozen advisers, and none was specified before the data were examined, so we present them as exploratory. The sharpest split is between funds that run their own process and funds that delegate it. In-house adopters broaden by 11.5 to 19.8 log points, significantly in all six designs (57 funds at 15 advisers), and subadvised adopters by 2.9 to 9.5, significantly in one (45 funds at 18 advisers), although the difference is itself significant in only one design. The early adopters of 2021–22, nine funds at five advisers, show the largest effects (21.4 to 28.8 log points), which partly reflects the longer post-adoption windows of that cohort. No difference by vehicle, by what the AI is described as doing, or by the data it reads is robust, and funds whose adoption filing contains more AI sentences broaden somewhat more, although that difference is significant in only one design (Appendix D.1).

Teams whose managers state no STEM degree before adoption broaden by 8.3 to 12.5 log points, significantly in five of six designs (47 funds at 11 advisers). Teams with at least one STEM-trained manager show no effect (0.4 to 3.5 log points, none significant; 26 funds at 10 advisers). The difference is significant in three of seven designs, counting TWFE, and team education does not predict who adopts. One reading is that AI supplies screening capacity that technically trained teams already had. The groups rest on about ten advisers each, so we offer this as a hypothesis rather than a finding.

Visa filings show where technical capacity sits at adopting firms, and it is not in the portfolio-manager role. Employers that sponsor H-1B workers file labor condition applications with the Department of Labor, which record the job title and occupation; we match employers to advisers by name and count applications for data-science and AI positions (Appendix Table A13). At the advisers of described-process adopters, these are mostly operations research analysts, data scientists, business-intelligence analysts and statisticians, with titles

in data science and analytics. None carries a portfolio-manager title, and none of the 61 applications for portfolio-manager positions at these advisers is a data-science role. The hiring also came first. Before 2019, adopters' advisers filed a data-science application in 5.5 percent of quarters, against 1.6 percent for advisers without an adopting fund ($p = 0.034$); 14 of the 16 that ever filed one did so before adoption; and after adoption they hire no more than the advisers of disclosure-only adopters (difference 2.4 percentage points, $p = 0.52$). Taken together with the training result, this suggests that AI reached these funds through data-science staff working alongside the portfolio managers, and broadened the portfolios of teams that lacked such training themselves. The evidence is limited: applications cover only visa-sponsored positions, an application is not a hire, and the comparison above matches employers by brand name, which attaches a parent's hiring to all of its advisers; on exact legal names, only 7 of the 35 adopting advisers file any such application.

The tool also changes how the work is organized, in the direction the labor economics of prediction would suggest (Agrawal et al. 2019; Freund and Mann 2025). Following each fund's named management team from 2019Q4 (95 described-process adopters), issuers held per named manager rise by 14 to 24 log points across the six designs (all $p \leq 0.030$), and the share of adopters run by a single named manager rises by 8 to 13 percentage points.[19] The saving is within the fund: advisers whose Form ADV brochure first describes AI in investment analysis add investment staff rather than shed it.

## 5.7 Performance

We find no robust performance consequence, and Table 7 shows why the reference-period discipline matters here. Over event quarters 0 to 2 with quarter $-1$ as reference, adopters' quarterly net flows rise by 2.4 percent of assets after adoption ($p = 0.005$), and their net alpha by 1.25 percent per quarter ($p = 0.066$). But their pre-period flows are also high relative to that quarter. Moving the reference to quarter $-2$, or dropping quarters $-2$ and $-1$, turns both estimates negative or null: flows of $-1.8$ and $-1.2$ percent ($p \geq 0.29$), and alpha of $-0.7$ and $+0.4$ percent ($p \geq 0.34$). What looked like a reward for adoption was a single unusual quarter before it. Table 6 completes the picture: peer-adjusted returns, value added (Berk and van Binsbergen 2015), assets under management, expense ratios, the return gap of Kacperczyk, Sialm and Zheng (2008), market timing, and information ratios do not move under any reference.

Nor do adopters pick names that go on to do better, as one would expect if the added

[19] The Voya Small Company Portfolio, whose 2022 prospectus states that "the Sub-Adviser uses quantitative methods, including artificial intelligence ('AI') models, to select securities," went from two named managers to one within four quarters, while its issuer count rose from 124 to 169.

names are marginal ones. Risk-adjusted returns of the added names over the next quarter and the next four quarters, and their returns relative to the names the fund drops, show no robust change, and the factor profile of the book does not move (Appendix D.1). No subgroup of adopters earns alpha that survives the reference checks and a correction for the number of splits (Appendix D.8).

Power is limited, and we say so. Post-adoption windows are short, and the treated funds are few. A sharper test asks whether AI relaxes decreasing returns to scale (Pástor et al. 2015): if AI lowers the cost of covering more names, adopters' alphas should fall less as they grow. The change in the size–alpha slope after adoption is small and insignificant, but the smallest change the test can detect is about 60 basis points of quarterly alpha per log unit of size, so the performance results rule out only large effects.[20] Evidence from outside asset management suggests that one payoff to AI may lie in stability rather than in levels: among U.S. firms, AI investment is associated with less volatile sales, earnings, and cash flows (Babina et al. 2025), a dimension that our tests of returns, flows, and costs do not address.

# 6 Discussion

*Breadth without Homogenization.* The homogenization argument has two steps: AI adoption spreads, and adopters' decisions become correlated. In this market and period, the first step happened on a small scale and the second did not happen at all. The funds that put AI to work hold more names. The names they add overlap with other AI users' additions no more than popularity predicts, their trades are less synchronized than those of comparable funds, and their portfolios are no more similar to other AI users' portfolios than before. AI funds' returns also comove less than those of non-AI funds (Chen, Sialm and Xu 2026). The evidence suggests why breadth and homogenization came apart. The filings describe models built in-house, and the added names came from each fund's own covered universe, so the extra evaluations appear to have been fund-specific. Homogenization requires shared models or data, which were rare in this period, and the positive control shows that the test detects a shared input where funds have one. Competition may also work against convergence when tools are similar. Active managers are rewarded for outputs that differ from their rivals', since a fund that holds what its rivals hold cannot outperform them, and in such markets diverse

[20] Two related checks are in Appendix E. First, adopting funds are not less likely to close: the post-adoption exit hazard rises by 0.79 percentage points per quarter for described-process adopters ($p = 0.15$) and by 0.74 for disclosure-only adopters ($p = 0.013$), so differential survival cannot explain the gap in breadth between them. Second, late generic adopters are not underperformers looking for a label. In a hazard model of disclosure-only adoption, trailing peer-adjusted returns, alpha, and flows predict nothing ($p \geq 0.15$), while adoption by the fund's own family ($p \leq 0.006$) and adviser size ($p < 0.001$) do.

models are favored over a common one (Raghavan 2026). As vendors and data products spread across fund complexes, the same pair statistics can be computed from holdings alone for any shared input, once the sharing is known. Current fund disclosures name custodians and auditors but not model or data suppliers, so the limitation lies in the disclosure, not in the test.

*Alternative Explanations.* Several accounts that do not involve AI could produce more names inside a fund's existing terrain, and the data distinguish some of them. If the breadth effect were a quantitative shop widening its screen for reasons unrelated to AI, it should be confined to quantitative adopters; instead, fundamental adopters broaden at least as much (10.6 to 15.9 log points, against 4.9 to 10.0 for quantitative adopters). A related account locates the change in portfolio construction rather than in search: a model that ranks the whole benchmark, combined with an optimizer that manages tracking error, would spread weight across many small benchmark positions. Two facts fit it imperfectly, although neither rules it out. Such construction is most natural in quantitative processes, yet fundamental adopters broaden at least as much; and risk budgeting against the benchmark should move a fund's weights toward the index's, yet active share against true Russell weights falls by at most a percentage point in the event-study designs, a check available only for Russell-benchmarked funds. If it were benchmark-hugging induced by inflows, flows and assets should rise with it; flows do not rise once the reference quarter is moved, and assets do not change. If it were cheaper trading or diversification for its own sake, names outside the benchmark should rise with names inside; they rise less, and less robustly. What the data cannot distinguish is a fund-specific decision to reposition the portfolio, taken when the prospectus was rewritten, from the use of AI that the rewrite describes. The comparison groups rule out such decisions at the adviser level, not at the fund level.

*Evaluation Costs and Discovery Costs.* Search has two costs: the cost of evaluating a candidate once it is in view, and the cost of bringing new candidates into view. As these funds use it, AI lowers the first and leaves the second unchanged. Benchmark constituents show the difference: they are the names a manager knows best, but knowing a name is not the same as having evaluated it. That is the point-solution pattern of Agrawal et al. (2024): cheap prediction added to a decision process whose structure, here the fund's covered universe, stays fixed. It is also what rational-inattention models predict for added processing capacity, which goes first to assets the investor already follows (Sims 2003; Van Nieuwerburgh and Veldkamp 2010; Kacperczyk et al. 2016). Distant search should rise where AI connects bodies of information that were not previously joined, not where it ranks names within a fixed screen. Funds whose AI reads text or alternative data do not search farther than funds using numeric models, so on current evidence even text-reading AI ranks names within a

fixed screen.

*Breadth within the Covered Universe.* The pattern resembles what cheaper production did to books, where the question became whether the extra titles were valuable (Reimers and Waldfogel 2026). Here the extra holdings are more draws from the distribution the fund already samples: the same industries, regions, and coverage strata, and mostly its benchmark. In other settings, cheaper access to a larger collection concentrated use on a popular core (Evans 2008; Fleder and Hosanagar 2009); for active funds, the benchmark plays that part. We cannot yet tell whether wider search toward the benchmark is cheap diversification or a drift toward closet indexing (Cremers and Petajisto 2009; Petajisto 2013). Adopters hold more benchmark names, but their active share against true Russell weights falls by at most a percentage point in the event-study designs (Section 5.5).

*Substantive and Symbolic Adoption.* The 2023–26 wave of generic AI disclosure has no measurable consequence for portfolios, and it follows family-wide rollouts at large advisers rather than performance. Verifying AI claims is costly for investors and regulators (Goldfarb and Tucker 2019), and studies that code adoption from announcements, hiring, or keyword counts will mix the two kinds of disclosure. Reading what a filing claims about a specific decision separates them. The broader definition cuts the breadth estimate by a third and would have blurred the convergence test with funds that changed nothing.

# 7 Conclusion

Does AI in the investment process make fund portfolios homogeneous? Using prospectus text read by a language model and complete holdings records, we find that it did not, in the first years of its use in U.S. active funds. Funds that put AI to work in their process did not add the same names as other AI users, did not herd more than comparable funds, and did not grow more similar to one another. These tests do detect convergence from a shared decision process where one exists. What the adopting funds did do was search more widely. They held more securities, added and dropped names faster, and, most clearly among international funds, came to hold more of their own benchmark. The extra search went into the terrain each fund already covered, and, with models that the filings describe as each fund's own, broader portfolios did not become more homogeneous.

The result is bounded in five ways. First, it concerns the adoption of AI, not the sharing of it. The adopters built their own models; outside pairs under common management, no AI vendor named in a fund's own process served funds at two advisers; and the one weak shared input we observe, a common index provider, produces no co-entry. The power of the test against shared data feeds is therefore not established. Second, we date disclosure, not

use, and the disclosure-only placebo controls for the decision to write about AI, not for a fund-level plan that coincides with it. Third, the treated funds are few and concentrated, and most adopted within the last three years, so the long-horizon estimates rest on a handful of advisers and the heterogeneity results are exploratory. Fourth, a pre-adoption dip complicates the cleanest reading of timing, and the performance nulls rule out only large effects. Fifth, we observe what funds say about the data their AI reads, not the data themselves. Prospectuses describe types of data, such as fundamentals, prices, text, or alternative data, and rarely name a dataset or a supplier. Funds at different advisers could therefore feed the same commercial or proprietary data into their models without our knowing it. The tests register such sharing only through its consequences for holdings. The absence of excess co-entry says that whatever data adopters shared did not move them toward the same securities in this period, not that they shared none. The reading that each fund's model ran on its own data is an inference from the filings, not an observation.

Three questions follow. First, does wider search pay over horizons longer than the two years we observe? Second, do the teams that gained breadth eventually learn to search farther, or does the cheaper screen lock in the familiar terrain? Third, as AI models and data products become shared across fund complexes, will the correlation that regulators fear appear, or will funds continue to point shared tools at different universes? The answer will depend less on how much the tools can process than on whether the firms that use them come to look in the same places.

# References


**Abis, S.**, "Man vs. Machine: Quantitative and Discretionary Equity Management," 2020. Working paper, Columbia Business School.

**Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo**, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

**Agrawal, A., J. Gans, and A. Goldfarb**, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press, 2018.

_____, **J. S. Gans, and A. Goldfarb**, "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction," *Journal of Economic Perspectives*, 2019, *33* (2), 31–50.

_____, _____, **and** _____, "Artificial Intelligence Adoption and System-Wide Change," *Journal of Economics & Management Strategy*, 2024, *33* (2), 327–337.

**Arellano, M. and O. Bover**, "Another Look at the Instrumental Variable Estimation of Error-Components Models," *Journal of Econometrics*, 1995, *68* (1), 29–51.

**Babina, T., A. Fedyk, A. He, and J. Hodson**, "Artificial Intelligence Makes Firm Operating Performance Less Volatile," *AEA Papers and Proceedings*, 2025, *115*, 35–39.

**Berk, J. B. and J. H. van Binsbergen**, "Measuring Skill in the Mutual Fund Industry," *Journal of Financial Economics*, 2015, *118* (1), 1–20.

**Bommasani, R., K. A. Creel, A. Kumar, D. Jurafsky, and P. Liang**, "Picking on the Same Person: Does Algorithmic Monoculture Lead to Outcome Homogenization?," in "Advances in Neural Information Processing Systems 35 (NeurIPS 2022)" 2022, pp. 3663–3678.

**Bonelli, M.**, "Data-Driven Investors," *Review of Financial Studies*, 2026, *39* (7), 1909–1969.

_____ **and T. Foucault**, "Does Big Data Devalue Traditional Expertise? Evidence from Active Funds," *Journal of Finance*, forthcoming. SSRN Working Paper 4527672.

**Callaway, B. and P. H. C. Sant'Anna**, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

**Cao, S., W. Jiang, J. Wang, and B. Yang**, "From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses," *Journal of Financial Economics*, 2024, *160*, 103910.

**Chen, S., C. Sialm, and D. Xu**, "The Growth and Performance of Artificial Intelligence in Asset Management," NBER Working Paper 35273, National Bureau of Economic Research 2026.

**Cremers, K. J. M. and A. Petajisto**, "How Active Is Your Fund Manager? A New Measure That Predicts Performance," *Review of Financial Studies*, 2009, *22* (9), 3329–3365.

**Davidson, R. and J. G. MacKinnon**, "Wild Bootstrap Tests for IV Regression," *Journal of Business & Economic Statistics*, 2010, *28* (1), 128–144.

**de Chaisemartin, C. and X. D'Haultfœuille**, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

**Eisfeldt, A. L., G. Schubert, M. B. Zhang, and B. Taska**, "Generative AI and Firm Values," *Journal of Finance*, 2026. Forthcoming; published online 8 October 2026. Earlier version: NBER Working Paper 31222 (2023).

**Eloundou, T., S. Manning, P. Mishkin, and D. Rock**, "GPTs Are GPTs: Labor Market Impact Potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308.

**Evans, J. A.**, "Electronic Publication and the Narrowing of Science and Scholarship," *Science*, 2008, *321* (5887), 395–399.

**Farrell, Joseph and Matthew Rabin**, "Cheap Talk," *Journal of Economic Perspectives*, 1996, *10* (3), 103–118.

**Financial Stability Board**, "Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications," Technical Report, Financial Stability Board, Basel November 2017.

**Fleder, D. and K. Hosanagar**, "Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity," *Management Science*, 2009, *55* (5), 697–712.

**Freund, L. B. and L. F. Mann**, "Job Transformation, Specialization, and the Labor Market Effects of AI," Discussion Paper 117/25, RFBerlin November 2025.

**Goldfarb, A. and C. Tucker**, "Digital Economics," *Journal of Economic Literature*, 2019, *57* (1), 3–43.

_____, **B. Taska, and F. Teodoridis**, "Could Machine Learning Be a General Purpose Technology? A Comparison of Emerging Technologies Using Data from Online Job Postings," *Research Policy*, 2023, *52* (1), 104653.

**Goodman-Bacon, A.**, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

**Gu, S., B. Kelly, and D. Xiu**, "Empirical Asset Pricing via Machine Learning," *Review of Financial Studies*, 2020, *33* (5), 2223–2273.

**Henriksson, R. D. and R. C. Merton**, "On Market Timing and Investment Performance II," *Journal of Business*, 1981, *54* (4), 513–533.

**Hoberg, G. and G. Phillips**, "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy*, 2016, *124* (5), 1423–1465.

**Hu, X., M. Rohrer, and H. Zhang**, "Active Machine-Learning-Based Trading and Mutual Fund Performance," 2025. SSRN Working Paper 5571002.

**International Monetary Fund**, "Advances in Artificial Intelligence: Implications for Capital Market Activities," in "Global Financial Stability Report: Steadying the Course: Uncertainty, Artificial Intelligence, and Financial Stability," Washington, DC: International Monetary Fund, October 2024, chapter 3.

**Kacperczyk, M., C. Sialm, and L. Zheng**, "Unobserved Actions of Mutual Funds," *Review of Financial Studies*, 2008, *21* (6), 2379–2416.

**_____, S. Van Nieuwerburgh, and L. Veldkamp**, "A Rational Theory of Mutual Funds' Attention Allocation," *Econometrica*, 2016, *84* (2), 571–626.

**Katila, R. and G. Ahuja**, "Something Old, Something New: A Longitudinal Study of Search Behavior and New Product Introduction," *Academy of Management Journal*, 2002, *45* (6), 1183–1194.

**Kleinberg, J. and M. Raghavan**, "Algorithmic Monoculture and Social Welfare," *Proceedings of the National Academy of Sciences*, 2021, *118* (22), e2018340118.

**Lakonishok, J., A. Shleifer, and R. W. Vishny**, "The Impact of Institutional Trading on Stock Prices," *Journal of Financial Economics*, 1992, *32* (1), 23–43.

**Levinthal, D. A. and J. G. March**, "The Myopia of Learning," *Strategic Management Journal*, 1993, *14* (S2), 95–112.

**March, J. G.**, "Exploration and Exploitation in Organizational Learning," *Organization Science*, 1991, *2* (1), 71–87.

**Pástor, L., R. F. Stambaugh, and L. A. Taylor**, "Scale and Skill in Active Management," *Journal of Financial Economics*, 2015, *116* (1), 23–45.

**Petajisto, A.**, "Active Share and Mutual Fund Performance," *Financial Analysts Journal*, 2013, *69* (4), 73–93.

**Raghavan, M.**, "Competition and Diversity in Generative AI," *Management Science*, 2026. Articles in Advance, published online October 16, 2026.

**Rambachan, A. and J. Roth**, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

**Reimers, I. and J. Waldfogel**, "Digitization and Pre-Purchase Information: The Causal and Welfare Impacts of Reviews and Crowd Ratings," *American Economic Review*, 2021, *111* (6), 1944–1971.

**_____ and _____**, "AI and the Quantity and Quality of Creative Products: Have LLMs Boosted Creation of Valuable Books?," Working Paper 34777, National Bureau of Economic Research 2026.

**Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe**, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

**Sims, C. A.**, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

**Stigler, George J.**, "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

**Sun, L. and S. Abraham**, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

**Treynor, J. and K. Mazuy**, "Can Mutual Funds Outguess the Market?," *Harvard Business Review*, 1966, *44* (4), 131–136.

**U.S. Securities and Exchange Commission**, "SEC Proposes New Requirements to Address Risks to Investors From Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers," Press Release 2023-140 July 2023. July 26, 2023. https://www.sec.gov/newsroom/press-releases/2023-140.

_____, "SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence," Press Release 2024-36 March 2024. March 18, 2024 (Delphia (USA) Inc. and Global Predictions Inc.). https://www.sec.gov/newsroom/press-releases/2024-36.

**Van Alstyne, M. and E. Brynjolfsson**, "Global Village or Cyber-Balkans? Modeling and Measuring the Integration of Electronic Communities," *Management Science*, 2005, *51* (6), 851–868.

**Van Nieuwerburgh, S. and L. Veldkamp**, "Information Acquisition and Under-Diversification," *Review of Economic Studies*, 2010, *77* (2), 779–805.

**Westphal, J. D. and E. J. Zajac**, "Substance and Symbolism in CEOs' Long-Term Incentive Plans," *Administrative Science Quarterly*, 1994, *39* (3), 367–390.

**Zielnicki, K., G. Aridor, A. Bibaut, A. Tran, W. Chou, and N. Kallus**, "The Value of Personalized Recommendations: Evidence from Netflix," arXiv preprint 2511.07280 2025. Presented at the NBER Digital Economics and AI meeting, February 2026.

# Tables

**Table 1:** Summary statistics

Active funds in the estimation sample. Adopters are measured in the quarter before adoption; never-adopters over the same calendar quarters.

| | Described-process adopters | | Disclosure-only adopters | | Never-adopters | |
|---|---|---|---|---|---|---|
| | (100 funds) | | (268 funds) | | (6,755 funds; 96,225 fund-quarters) | |
| | Mean | Median | Mean | Median | Mean | Median |
| *Panel A. Portfolio and fund characteristics* | | | | | | |
| Issuers held | 168 | 114 | 196 | 94 | 202 | 79 |
| GICS industry groups held | 20 | 22 | 19 | 21 | 18 | 20 |
| GICS sub-industries held | 59 | 58 | 58 | 46 | 53 | 43 |
| Countries held | 13 | 12 | 13 | 10 | 12 | 9 |
| Weight of the ten largest issuers (%) | 35 | 34 | 31 | 27 | 34 | 31 |
| Corporate share of the long book (%) | 92 | 98 | 87 | 97 | 89 | 97 |
| Net assets ($ million) | 1,128 | 242 | 2,478 | 452 | 2,250 | 317 |
| Expense ratio (%) | 0.79 | 0.79 | 0.84 | 0.85 | 0.87 | 0.88 |
| *Fund attributes (share of funds, %)* | | | | | | |
| Exchange-traded fund | 30 | | 18 | | 17 | |
| Any subadviser (N-CEN) | 43 | | 46 | | 50 | |
| *Panel B. Estimation samples (Section 4)* | | | | | | |
| Funds in the TWFE sample | 100 | | 281 | | 6,757 | |
| Funds in the Sun–Abraham sample | 102 | | 309 | | 6,844 | |
| Advisers (TWFE sample) | 26 | | 44 | | 772 | |
| Advisers (Sun–Abraham sample) | 27 | | 44 | | 778 | |
| Fund-quarters, 2013–2026 | 216,939 (TWFE); 226,281 (Sun–Abraham) | | | | | |

*Notes:* Fund-quarters pass the corporate-book screen (corporate positions at least 25 percent of the long book, at most 20 percent of corporate weight without an industry group). Adopters are active funds that are not left-censored (an AI-free filing restating the investment process precedes their first AI filing), measured in their last screened quarter within four quarters before adoption; funds without such a quarter are not shown. Never-adopter fund-quarters are restricted to the calendar quarters in which those pre-adoption quarters fall. Issuer, industry and country counts use the corporate book. Net assets from CRSP where linked, else N-PORT; expense ratios from CRSP, else the prospectus fee table. Panel B counts funds and fund-quarters in the event-study samples of Section 4; adopters enter with at least one pre-adoption quarter, left-censored adopters are excluded, and TWFE trims adopters to event time −8 to +8. Advisers are primary-adviser CRD numbers.

**Table 2:** Taxonomy of AI language in fund prospectuses

| Cell | Distinct texts | Occurrences | % specific | Intensity | Use in the treatment |
|---|---|---|---|---|---|
| *Panel A. Sentences (passes 1 and 2)* | | | | | |
| Thematic exposure: AI is what the fund invests in | 4,813 | 13,615 | | | Not adoption |
| Risk disclosure: AI as a risk to the fund or its holdings | 9,179 | 96,788 | | | Not adoption |
| Other mentions (names, definitions, unrelated uses of the terms) | 25,060 | 89,388 | | | Not adoption |
| Process use: the adviser or its models use AI | 2,018 | 12,958 | | | Pass 2 |
| Specific: a concrete investment task performed with AI | 1,358 | 7,986 | | | Pass 3 |
| Generic: AI tools in support, efficiency, or limits on AI's role | 660 | 4,972 | | | Disclosure only |
| Total | 41,070 | 212,749 | | | |
| *Panel B. Passages around process sentences (pass 3)* | | | | | |
| Fund process, research support (1) | 283 | 1,569 | 4 | 1 | Disclosure only |
| Fund process, one input to selection (2) | 699 | 1,421 | 77 | 2 | Described process if specific |
| Fund process, primary input (3) | 1,032 | 2,114 | 92 | 3 | Described process if specific |
| Fund process, AI selects holdings (4) | 530 | 1,088 | 98 | 4 | Described process if specific |
| Fund process, role unclear | 357 | 756 | 9 | – | Disclosure only |
| Index methodology, research support (1) | 4 | 18 | 61 | 1 | Disclosure only |
| Index methodology, one input to selection (2) | 161 | 381 | 92 | 2 | Described process if specific |
| Index methodology, primary input (3) | 795 | 1,873 | 98 | 3 | Described process if specific |
| Index methodology, AI selects holdings (4) | 115 | 256 | 100 | 4 | Described process if specific |
| Index methodology, role unclear | 43 | 84 | 48 | – | Disclosure only |
| Generic risk | 1,228 | 2,961 | 38 | 0 | Disclosure only |
| Limitation of AI's role | 50 | 169 | 35 | 0 | Disclosure only |
| Not AI use | 167 | 268 | 29 | 0 | Not adoption |
| Total | 5,464 | 12,958 | 62 | | |
| *Panel C. Active funds by type of first adoption (Figure 1)* | | | | | |
| Described process | | 132 | | | Treated |
| Disclosure only, process at a weaker role (research support) | | 277 | | | Placebo |
| Disclosure only, boilerplate (risk or limitation language) | | 309 | | | Placebo |

*Notes:* Labels from the current version of each pass. Panel A counts distinct sentence texts and their occurrences across filings and attributed series in N-1A-family filings from July 2018 to September 2026. Panel B counts distinct passages (a process sentence with up to 700 characters of context before it and 400 after it) and their occurrences; "% specific" is the share of occurrences whose sentence the second pass labels specific. Intensity is the role of AI in selection on a scale from 1 (research support) to 4 (AI selects holdings); statements that are not process descriptions score 0. A filing describes a process when a specific sentence sits in a fund-process or index-methodology passage with intensity of at least 2. Panel C counts active funds that are not left-censored (an AI-free filing restating the investment process precedes their first AI filing), by the type of their first AI filing; the event-study samples in Table 1 additionally require pre-adoption holdings.

**Table 3:** Effect of adoption on issuers held

Average post-adoption effect of described-process adoption on log(1 + issuers), in log points × 100 [adviser pairs bootstrap p].

| Design | Window | Reference | Described-process adopters | Contrast | Pre-period |
|---|---|---|---|---|---|
| *vs never-adopters; contrast: minus disclosure-only adopters* | | | | | |
| Disclosure | 0–2 | −1 | +8.7 [<0.005] | +6.9 [0.020] | −1.6 [0.635] |
| Disclosure | 0–2 | −2 | +10.7 [<0.005] | +9.6 [0.010] | +0.8 [0.500] |
| Disclosure | 0–2 | −3, drop −2, −1 | +10.3 [0.005] | +9.1 [0.025] | +1.4 [0.390] |
| Disclosure | 0–6 | −1 | +12.8 [<0.005] | +9.9 [0.020] | −1.6 [0.635] |
| Disclosure | 0–6 | −2 | +14.5 [<0.005] | +12.4 [0.010] | +0.8 [0.500] |
| Disclosure | 0–6 | −3, drop −2, −1 | +13.7 [<0.005] | +11.3 [0.030] | +1.4 [0.390] |
| *with index funds of the same adviser; contrast: minus index siblings* | | | | | |
| Family | 0–2 | −1 | +8.7 [<0.005] | +7.0 [0.045] | −1.6 [0.575] |
| Family | 0–2 | −2 | +10.7 [<0.005] | +9.4 [0.010] | +0.8 [0.540] |
| Family | 0–2 | −3, drop −2, −1 | +10.3 [0.005] | +8.5 [0.060] | +1.4 [0.460] |
| Family | 0–6 | −1 | +12.8 [0.005] | +10.4 [0.025] | −1.6 [0.580] |
| Family | 0–6 | −2 | +14.5 [0.005] | +12.2 [0.030] | +0.8 [0.540] |
| Family | 0–6 | −3, drop −2, −1 | +13.7 [0.010] | +10.8 [0.040] | +1.4 [0.460] |
| *adviser × quarter fixed effects; contrast: minus disclosure-only adopters* | | | | | |
| Within adviser | 0–2 | −1 | +8.5 [0.010] | +6.3 [0.090] | −3.6 [0.290] |
| Within adviser | 0–2 | −2 | +10.9 [0.005] | +10.0 [0.035] | −0.8 [0.560] |
| Within adviser | 0–2 | −3, drop −2, −1 | +10.6 [0.020] | +10.0 [0.050] | −0.8 [0.730] |
| Within adviser | 0–6 | −1 | +11.8 [0.010] | +8.2 [0.210] | −3.6 [0.310] |
| Within adviser | 0–6 | −2 | +13.8 [0.005] | +12.1 [0.090] | −0.8 [0.560] |
| Within adviser | 0–6 | −3, drop −2, −1 | +13.2 [0.035] | +10.8 [0.110] | −0.8 [0.725] |

*Notes:* Stitched holdings panel 2013–2026; event window from −8. Pre-period is the average of the pre-adoption event quarters from −8 other than the reference (and dropped) quarters. 102 described-process adopters. p-values from an adviser pairs bootstrap: advisers resampled with replacement and the Sun–Abraham regression re-estimated in each of 200 draws (p = share of draws whose deviation from the estimate is at least as large as the estimate). Appendix D.4 compares this inference with the wild cluster restricted bootstrap.

**Table 4:** Additions, exits and characteristics of added names

Described-process adopters vs never-adopters.

| Outcome | TWFE | Sun–Abraham range (6) | Significant | Reading |
|---|---|---|---|---|
| *Volume of search: changes* | | | | |
| log(1 + names added) | +0.235 [0.033] | +0.236 to +0.579 | 6 of 6 | More trials |
| New names / names held | +0.030 [0.024] | +0.032 to +0.052 | 6 of 6 | More turnover of names |
| Exits after $\leq$ 2 quarters / exits | +0.066 [0.033] | +0.015 to +0.101 | 4 of 6 | Quicker pruning |
| log(1 + names dropped) | +0.120 [0.235] | +0.192 to +0.391 | 6 of 6 | Pre-trend; suggestive |
| *Direction of search: unchanged* | | | | |
| New names new to the adviser | −0.003 [0.800] | −0.067 to −0.006 | 1 of 6 | Not more novel |
| New names first ever held by the fund | +0.016 [0.296] | −0.008 to +0.036 | 0 of 6 | Null |
| New names in a new industry group | −0.003 [0.723] | −0.010 to +0.016 | 0 of 6 | Null |
| New names in a new country | −0.020 [0.058] | −0.029 to −0.006 | 1 of 6 | Null |
| Size tilt of new names (log market cap) | −0.057 [0.488] | −0.137 to +0.019 | 0 of 6 | Null |
| Analysts covering new names (log), minus kept | −0.004 [0.903] | −0.061 to +0.147 | sign flips | Reference-sensitive |
| New names an analyst starts to cover within 4 quarters, minus kept | +0.008 [0.670] | −0.095 to +0.063 | 1 of 6 | Null |
| Industry groups held (log) | −0.001 [0.957] | +0.012 to +0.023 | 2 of 6 | Barely moves |
| *Depth: unchanged* | | | | |
| Survival of new names to 4 quarters | +0.021 [0.524] | −0.021 to +0.128 | 2 of 6 | Fragile |
| Scale-up of survivors | +0.073 [0.273] | −0.086 to +0.312 | 0 of 6 | Null |
| Weight in names held $\geq$ 8 quarters | −0.037 [0.319] | −0.046 to −0.009 | 2 of 6 | Null |

*Notes:* Coefficient [wild cluster restricted bootstrap p]; log outcomes are log differences (0.10 = 10 log points). Sun–Abraham ranges cover windows 0–2 and 0–6 and references −1, −2 and −3 with −2, −1 dropped; "significant" counts designs with $p < 0.05$. Analyst outcomes come from IBES and use each fund's equity holdings only: the number of analysts with an outstanding recommendation on the issuer at the end of the quarter before the trade, and an analyst's first recommendation on the issuer in the four quarters after it. Bloomberg's year-end analyst counts give the same verdicts. Industry groups: Callaway–Sant'Anna estimates range from −0.1 to +0.6 log points (Appendix D.4).

**Table 5:** Concentration of the book and position relative to the benchmark

Described-process adopters vs never-adopters [wild cluster restricted bootstrap p].

| Outcome | TWFE | Sun–Abraham range (6) | Significant | Minus disclosure-only | Minus index siblings | Never-adopter slope |
|---|---|---|---|---|---|---|
| *Shape of the book* | | | | | | |
| log(1 + issuers) | +0.080 [0.022] | +0.087 to +0.145 | 6 of 6 | +0.099 [0.000] | +0.104 [0.001] | 1 |
| log effective N (entropy) | +0.048 [0.152] | +0.058 to +0.104 | 6 of 6 | +0.078 [0.000] | +0.060 [0.001] | 0.89 |
| Evenness (entropy / ln N) | −0.006 [0.181] | −0.009 to −0.004 | 5 of 6 | −0.004 [0.049] | −0.009 [0.003] | −0.006 |
| Share of names < 0.5% of book | +0.039 [0.018] | +0.042 to +0.071 | 6 of 6 | +0.051 [0.008] | +0.065 [0.002] | 0.22 |
| Weight in names < 0.5% of book | +0.028 [0.018] | +0.027 to +0.039 | 6 of 6 | +0.035 [0.016] | +0.028 [0.040] | 0.14 |
| *Inside and outside the stated benchmark (primary sample: U.S. indices from the provider, MSCI indices from a tracker of another family)* | | | | | | |
| log(1 + issuers), this subsample | +0.079 [0.058] | +0.089 to +0.144 | 6 of 6 | +0.098 [0.000] | +0.108 [0.000] | 1 |
| log(1 + benchmark constituents held) | +0.089 [0.040] | +0.088 to +0.144 | 6 of 6 | +0.107 [0.001] | +0.093 [0.000] | 0.81 |
| log(1 + names outside the benchmark) | +0.061 [0.304] | +0.030 to +0.106 | 2 of 6 | −0.044 [0.525] | +0.045 [0.511] | 1.05 |
| Constituents held minus names outside (log) | +0.028 [0.670] | +0.023 to +0.082 | 0 of 6 | +0.151 [0.122] | +0.049 [0.451] | −0.23 |
| Share of the benchmark's constituents held | +0.013 [0.041] | +0.010 to +0.015 | 6 of 6 | +0.010 [0.109] | +0.004 [0.201] | 0.10 |
| New names that are benchmark constituents | +0.021 [0.296] | −0.001 to +0.031 | 0 of 6 | +0.029 [0.306] | +0.025 [0.412] | |
| Weight outside the benchmark | −0.010 [0.111] | −0.019 to −0.008 | 4 of 6 | −0.019 [0.024] | −0.020 [0.005] | 0.011 |
| *Funds benchmarked to U.S. indices only (provider membership throughout)* | | | | | | |
| log(1 + issuers), this subsample | +0.001 [0.993] | +0.029 to +0.067 | 6 of 6 | +0.027 [0.121] | +0.055 [0.003] | 1 |
| log(1 + benchmark constituents held) | +0.006 [0.940] | +0.034 to +0.070 | 6 of 6 | +0.032 [0.107] | +0.040 [0.046] | 0.78 |
| log(1 + names outside the benchmark) | +0.045 [0.561] | +0.024 to +0.083 | 2 of 6 | −0.027 [0.668] | +0.109 [0.234] | 1.05 |
| Share of the benchmark's constituents held | +0.006 [0.422] | +0.006 to +0.009 | 6 of 6 | +0.002 [0.499] | +0.003 [0.303] | 0.10 |
| New names that are benchmark constituents | −0.000 [0.980] | −0.046 to +0.014 | 1 of 6 | −0.006 [0.801] | −0.103 [0.126] | |
| Weight outside the benchmark | −0.004 [0.354] | −0.011 to −0.004 | 3 of 6 | −0.013 [0.093] | −0.015 [0.017] | 0.009 |
| Constituents held minus names outside (log) | −0.038 [0.568] | −0.029 to +0.028 | 0 of 6 | +0.059 [0.371] | −0.068 [0.475] | −0.26 |

*Notes:* Sun–Abraham ranges cover windows 0–2 and 0–6 and references −1, −2 and −3 with −2, −1 dropped; "significant" counts designs with $p < 0.05$. Contrast columns: described-process adopters minus disclosure-only adopters (disclosure design) and minus the adviser's index funds (family design), Sun–Abraham, window 0–6, reference −1. Never-adopter slope: coefficient on log(1 + issuers) among never-adopters with fund and quarter fixed effects, the change breadth alone would imply per log point. Benchmark rows use no index weights. Primary sample: 87 adopters with pre- and post-adoption data; U.S. benchmarks from the index provider's quarter-end constituent lists, MSCI benchmarks from the holdings of an index fund that tracks the index and is run by a family other than the benchmarked fund's adviser (Section 3.3); pre-period averages insignificant in every row ($p \geq 0.08$). U.S. benchmarks only: 53 adopters; pre-period averages insignificant with reference −2 or after dropping quarters −2 and −1 ($p \geq 0.06$), significant with reference −1 for constituents held and the share held ($p = 0.02, 0.01$). Weight outside the benchmark: equity funds.

**Table 6:** Performance and costs after adoption

Described-process adopters vs never-adopters, Sun–Abraham window 0–2 unless noted [wild cluster restricted bootstrap p].

| Outcome | Reference −1 | Reference −2 | Ref. −3, drop −2, −1 |
|---|---|---|---|
| Primary-model alpha, net, per quarter | +1.25% [0.066] | −0.71% [0.340] | +0.43% [0.552] |
| Peer-adjusted return, per quarter | +0.70% [0.097] | +0.21% [0.144] | +0.40% [0.225] |
| Net flow / lagged assets, per quarter | +2.40% [0.005] | −1.75% [0.299] | −1.15% [0.325] |
| log assets under management | −0.011 [0.693] | −0.009 [0.865] | +0.003 [0.980] |
| Value added, asinh $M | +0.40 [0.198] | −0.02 [0.951] | +0.37 [0.498] |
| Expense ratio, CRSP (window 0–6) | −0.01 pp [0.080] | −0.00 pp [0.887] | +0.00 pp [0.776] |
| Return gap, per quarter (window 0–6) | −0.11% [0.609] | −0.01% [0.982]† | +0.27% [0.254] |
| Treynor–Mazuy timing coefficient (window 0–6) | +0.29 [0.531] | −0.08 [0.782]† | +0.36 [0.089] |
| Information ratio, 12 months (window 0–6) | +0.05 [0.102] | −0.00 [0.970]† | −0.04 [0.478] |
| Flow–performance slope, adopter interaction (TWFE) | +0.013 [0.632] against a never-adopter slope of 0.044 | | |

*Notes:* † window 0–2. Return gap: gross fund return minus the return of the previous quarter-end holdings, for funds whose stock holdings cover at least 80 percent of assets. Expense ratios are estimated within the CRSP series; combining CRSP with prospectus fee tables produces spurious changes when a fund's data source switches at the adoption filing.

**Table 7:** Performance after adoption under three reference quarters

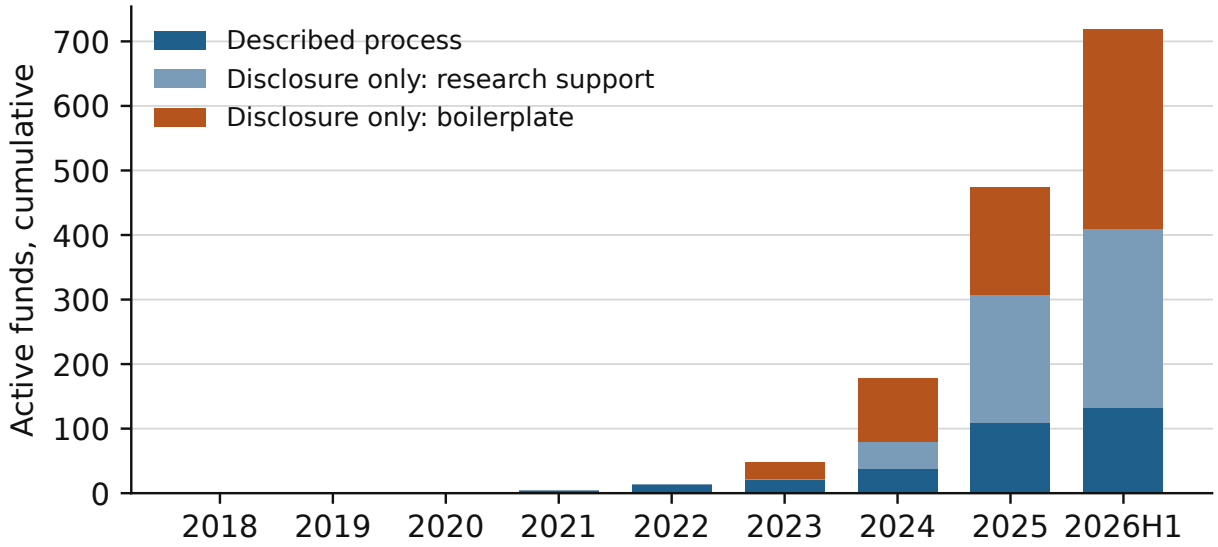
Sun–Abraham, window 0–2, against never-adopting active funds [wild cluster restricted bootstrap p].

| Estimate | Reference −1 | Reference −2 | Ref. −3, drop −2, −1 |
|---|---|---|---|
| *Primary-model alpha, net, per quarter* | | | |
| Described process, after adoption | +1.25% [0.066] | −0.71% [0.340] | +0.43% [0.552] |
| Described process, before adoption | +0.65% [0.208] | −0.87% [0.265] | −0.30% [0.811] |
| Disclosure only, after adoption | −0.56% [0.037] | +0.46% [0.223] | +0.28% [0.283] |
| Described process minus disclosure only | +1.80% [0.011] | −1.16% [0.052] | +0.15% [0.841] |
| *Peer-adjusted return, per quarter* | | | |
| Described process, after adoption | +0.70% [0.097] | +0.21% [0.144] | +0.40% [0.225] |
| Described process, before adoption | +0.60% [0.182] | −0.53% [0.041] | −0.06% [0.930] |
| Disclosure only, after adoption | +0.06% [0.740] | +0.07% [0.642] | +0.03% [0.837] |
| Described process minus disclosure only | +0.64% [0.129] | +0.14% [0.515] | +0.37% [0.236] |
| *Net flow / lagged assets, per quarter* | | | |
| Described process, after adoption | +2.40% [0.005] | −1.75% [0.299] | −1.15% [0.325] |
| Described process, before adoption | +2.94% [0.020] | −2.87% [0.176] | −3.20% [0.030] |
| Disclosure only, after adoption | +0.70% [0.412] | +0.41% [0.530] | +0.21% [0.816] |
| Described process minus disclosure only | +1.70% [0.100] | −2.16% [0.264] | −1.37% [0.396] |
| *log assets under management* | | | |
| Described process, after adoption | −0.011 [0.693] | −0.009 [0.865] | +0.003 [0.980] |
| Described process, before adoption | −0.088 [0.326] | −0.080 [0.126] | −0.180 [0.066] |
| Disclosure only, after adoption | +0.012 [0.783] | +0.017 [0.722] | −0.000 [0.998] |
| Described process minus disclosure only | −0.023 [0.643] | −0.026 [0.665] | +0.003 [0.967] |

*Notes:* After adoption: average effect over event quarters 0 to 2. Before adoption: average of the pre-period coefficients relative to the reference quarter (quarters −4 to −2, excluding the reference; quarter −4 alone in the drop variant). Performance outcomes use all active fund-quarters from 2019Q4, without the corporate-book screen, and event time from −4, so the samples differ from those of the breadth outcomes. Described-process and disclosure-only adopters enter the same regression as separate treated groups. With quarter −1 as reference, the pre-period flows of described-process adopters are high relative to that quarter; the post-adoption gains in flows and alpha do not survive moving the reference.

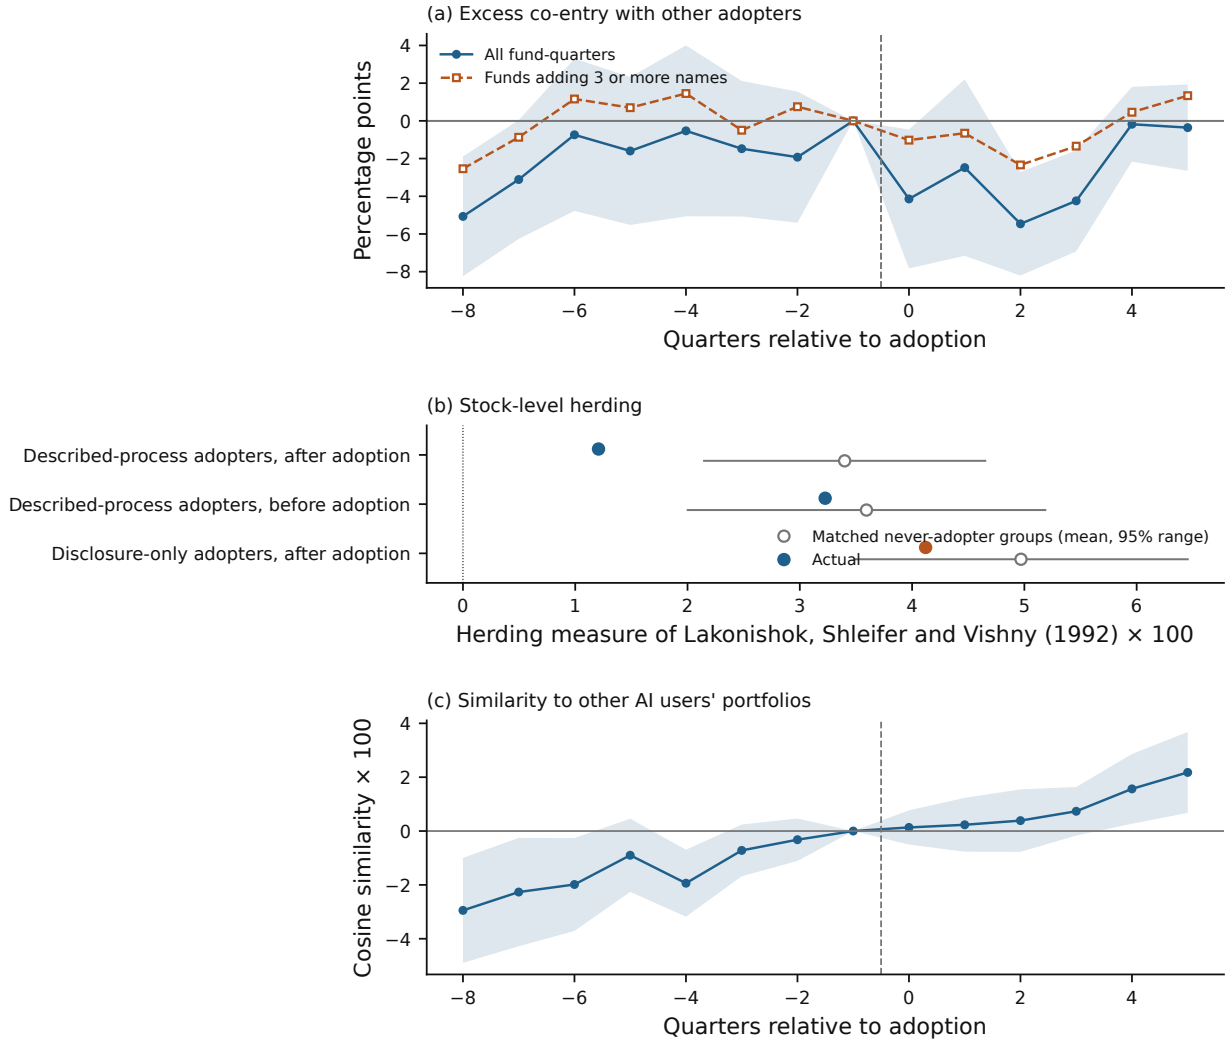
# Figures

**Figure 1:** Cumulative AI adoption among active funds, by type



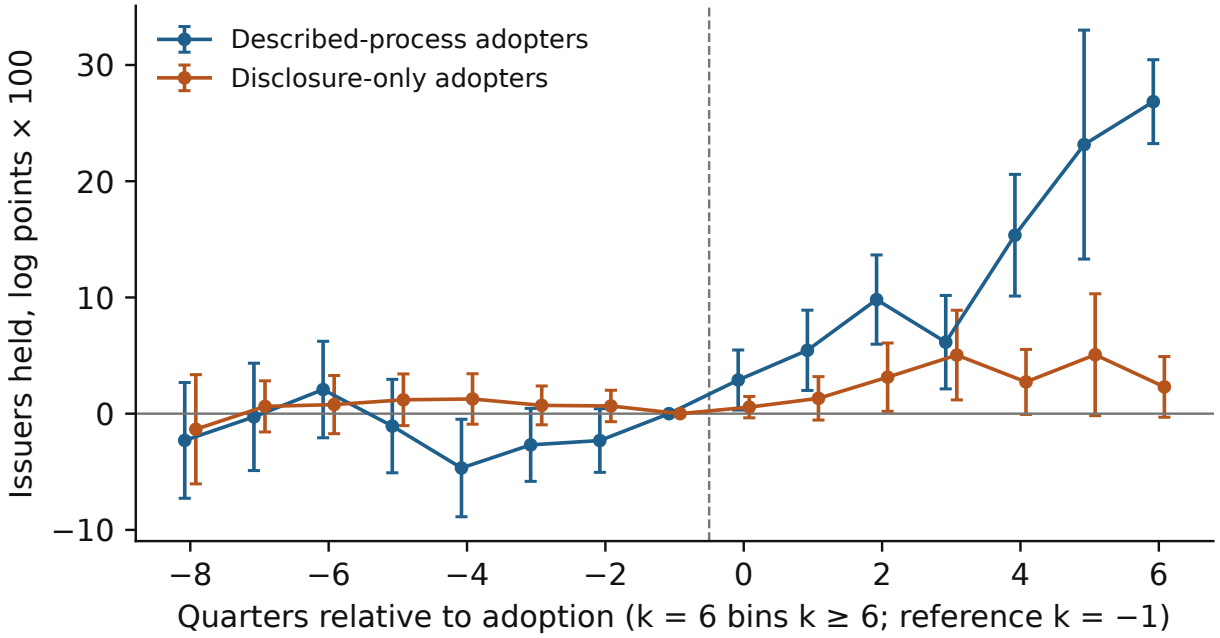

*Notes:* Active funds that are not left-censored (an AI-free filing restating the investment process precedes their first AI filing), counted from the year their filings first mention AI; bars show the cumulative number at each year-end. Described process: the filing describes AI performing a concrete investment task as at least one input to selection. Disclosure only, research support: a described process at a weaker role. Disclosure only, boilerplate: risk or limitation language only. 2026H1: through June 30, 2026. Totals at the end of 2026H1: 132 / 277 / 309. No described-process adoption is dated before 2021; the 2020 bar is one disclosure-only fund.

**Figure 2:** Co-entry, herding and portfolio similarity around adoption



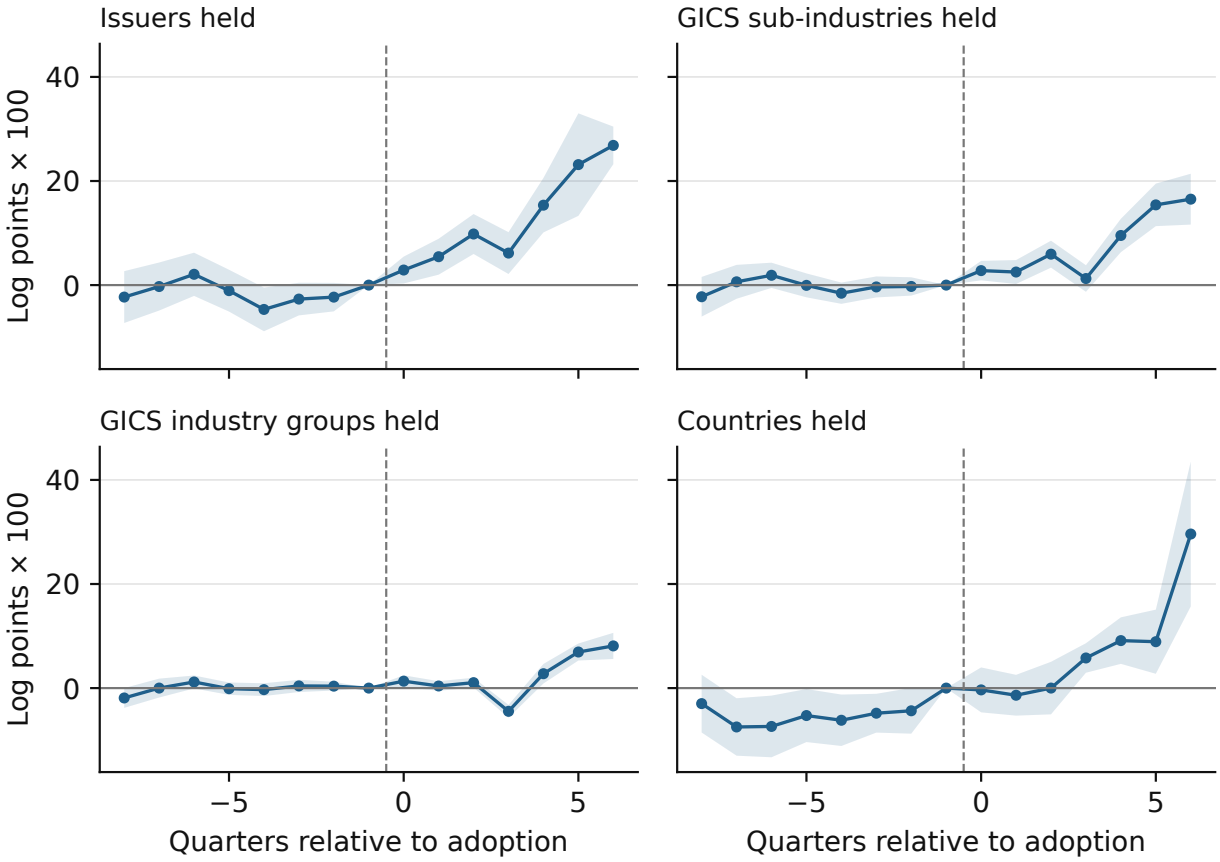

*Notes:* Described-process adopters, active funds. Panel (a): Sun–Abraham event-study coefficients (disclosure design, reference quarter −1) for the excess co-entry share, the share of a fund's new names that a described-process adopter at another adviser adds in the same quarter minus the share expected from the names' popularity among all active funds, in percentage points with 95% intervals from the analytical adviser-clustered covariance (shaded). The dashed series restricts the sample to fund-quarters in which the fund adds at least three names. Quarters 6 and later are pooled into one coefficient that is included in the post-adoption averages but not plotted, because it rests on a few early adopters at six to eight advisers; for the excess co-entry share it is +13.7 points, driven by funds that add one or two names in a quarter (+3.4 among fund-quarters with at least three new names). The positive control, pairs of funds at different advisers that share a subadviser, is reported in Section 5.1. Panel (b): the Lakonishok, Shleifer and Vishny (1992) herding measure among adopters' trades, pooled over quarters, (filled circles; disclosure-only adopters in orange) against the mean of 200 groups of never-adopting funds matched on Lipper class, trading activity and size (open grey circles, with ±1.96 placebo standard deviations); before adoption pools event quarters −8 to −1, after adoption k ≥ 0. Higher values mean that the group's purchases and sales of a stock are more one-sided than independent trading would produce. Panel (c): cosine similarity of the fund's corporate issuer weights to the aggregate portfolio of described-process adopters at other advisers, Sun–Abraham coefficients as in panel (a), quarters through 5 plotted. Table A15 reports every design.

**Figure 3:** Issuers held around adoption



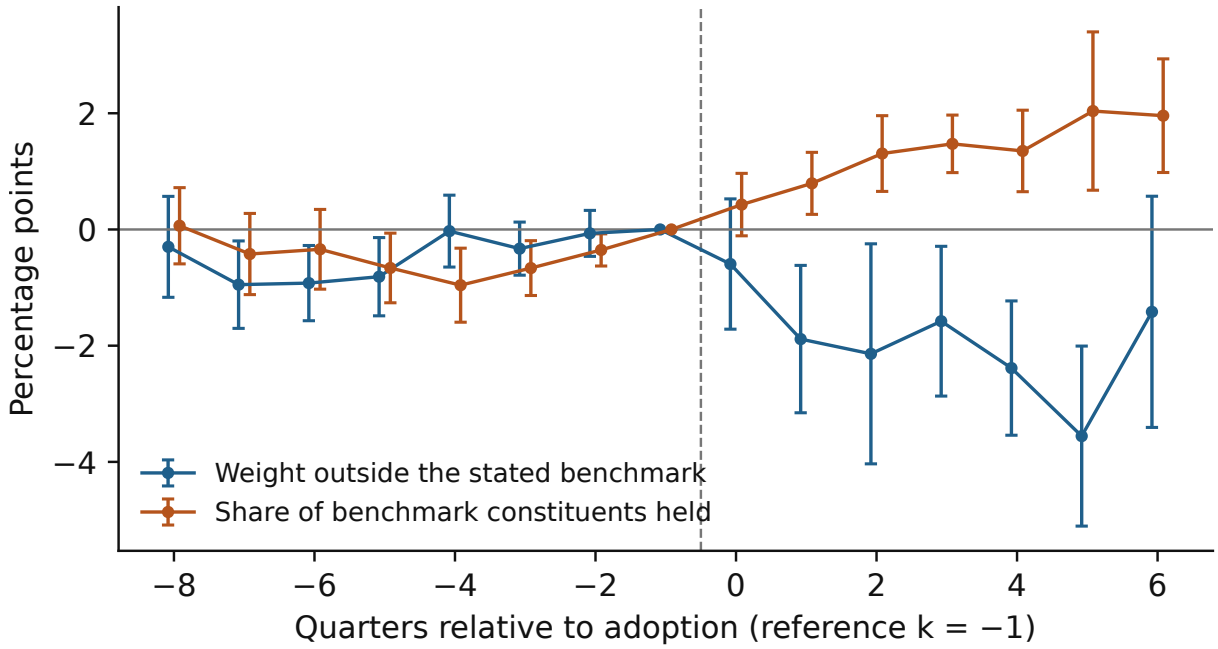

*Notes:* Sun–Abraham event study, disclosure design, event window to k = 6 (binned), reference quarter −1. Points are cohort-weighted effects in log points × 100 with 95% intervals from the analytical clustered covariance, which understates uncertainty where a cohort-quarter cell rests on one adviser; the bootstrap p-values in Table 3 do not have this problem. Late event times rest on few advisers (10 at k = 6). Table A6 lists the coefficients.

**Figure 4:** Issuers, industries and countries held around adoption



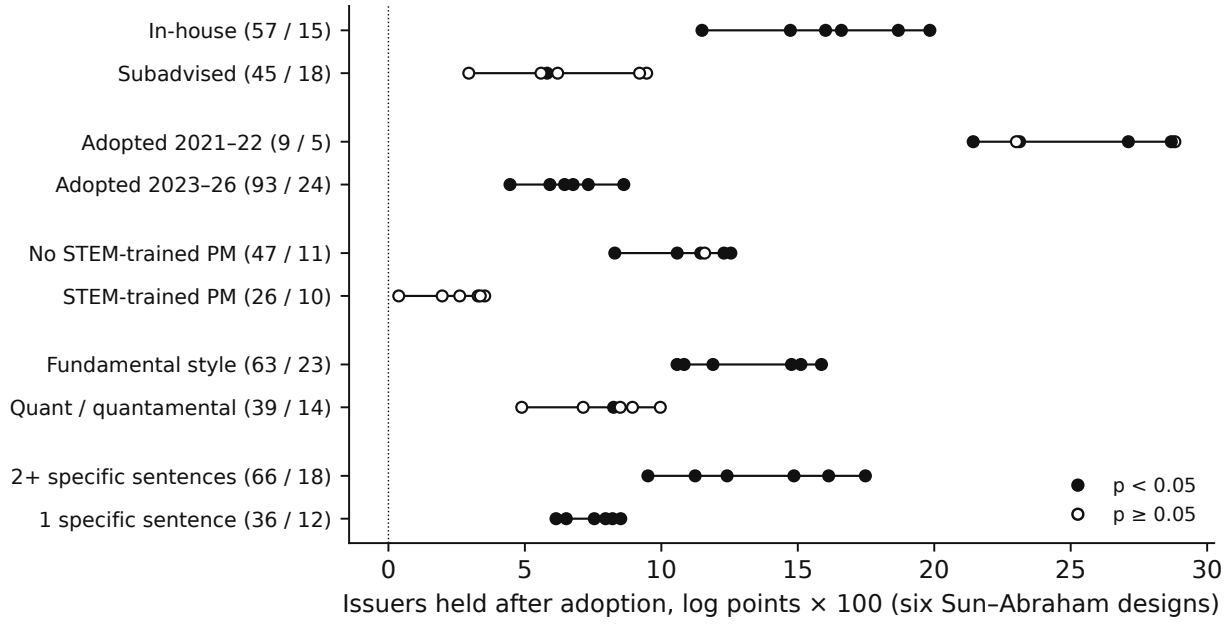

*Notes:* Sun–Abraham event-study coefficients for described-process adopters relative to never-adopting active funds (disclosure design, window to k = 6 with k ≥ 6 binned, reference quarter −1), in log points × 100, with 95% intervals from the analytical adviser-clustered covariance. Each panel is a separate regression on the same sample: the number of distinct corporate issuers held, and the numbers of GICS sub-industries, GICS industry groups and countries those issuers span. Axes are shared.

**Figure 5:** Position relative to the benchmark around adoption




*Notes:* Sun–Abraham, window to k = 6, reference −1, described-process adopters in the primary benchmark sample (87 for constituents held, 83 for weight outside the benchmark), percentage points with 95% intervals from the analytical clustered covariance. Weight outside the benchmark: share of the fund's corporate book in names that are not constituents of its stated benchmark (equity funds). Benchmark membership from the index provider's quarter-end constituent lists for U.S. indices and from a tracker run by another fund family for MSCI indices; no index weights are used.

**Figure 6:** Breadth effect by adopter characteristics




*Notes:* Sun–Abraham post-adoption effects on issuers held (log points × 100) in each of the six designs (windows 0–2 and 0–6; reference quarter −1, −2, or −3 with quarters −2 and −1 dropped). Filled circles: wild cluster restricted bootstrap $p < 0.05$; open circles: $p \geq 0.05$; the line spans the six estimates. Each pair is one regression in which the two groups and disclosure-only adopters are separate treated types against never-adopters. Labels give funds / advisers in the window 0–6 sample; an adviser with funds in both groups of a split is counted in each, so adviser counts can sum to more than the 27 advisers in the sample. Manager education from prospectus biographies, measured the quarter before adoption; funds with no education stated are a separate group (not shown). Table A11 reports the estimates and the differences between groups.

# Appendix

## A Data Sources and Linkage

Table A1 lists each source, what we take from it, the period used and how it is linked to the fund-quarter panel. The regulatory filings carry most of the weight: the treatment comes entirely from the text of fund prospectuses, and the main outcomes come from the holdings funds report to the SEC. Commercial data classify the securities in those holdings, supply returns and fees, and provide the benchmark and analyst information. Coauthor access to additional WRDS libraries supplied FTSE Russell constituent weights, IBES analyst data, the CRSP–Compustat link and MFLINKS (used only to check our fund links).

*Fund Text.* The treatment is read from SEC EDGAR's N-1A-family filings: initial registrations (N-1A), post-effective amendments (485APOS and 485BPOS), and prospectus supplements and summary prospectuses (497 and 497K). We collect every such filing by the registrants in our universe from July 2018 to September 2026, about 330,000 documents, and read the statutory prospectus and the statement of additional information in full. Adoption is dated through 2026Q2, the last quarter with holdings data; the sentence and passage counts in Table 2 cover the full corpus. The same filings supply the portfolio-manager items and biographies used for manager teams and education, and their headers identify which series a filing covers and whether each series is new or existing, which is how we separate a fund's launch from a later amendment. The SEC's risk/return summary data sets, structured XBRL extracts of prospectus Items 2 to 4 available from 2010Q4, provide each fund's principal-strategy text before adoption, the benchmarks named in the prospectus, and fee tables. They also bound what the 2018 start of the full-text collection can miss. The summaries contain almost no description of AI in the investment process before 2018, and every fund in which they identify such a description is also identified in the full text.

*Holdings and Fund Characteristics.* Form N-PORT discloses each fund's complete portfolio as of a date in each quarter, and the SEC's N-PORT data sets, from which we take the filings, begin in 2019Q4. For each holding they report the asset category, issuer type, identifiers (CUSIP, ISIN, LEI), country, payoff profile and fair-value level, along with the fund's net assets and monthly returns and flows. Where a fund amends a filing we keep the latest version for each report date. Before 2019Q4 we use the CRSP Survivor-Bias-Free U.S. Mutual Fund holdings. Form N-CEN, the funds' annual census filing, classifies each series as an index fund or ETF and lists its investment adviser and any subadvisers by CRD number. These lists identify a fund's primary adviser and whether the fund is run in-house or by a subadviser. A small number of N-CEN filings missing from the SEC's structured data sets

are parsed from EDGAR directly. Form ADV supplies each adviser's registration history, main office, control persons, employees in investment advisory functions (Item 5.B(1)) and regulatory assets under management, and its Part 2 brochure describes the adviser's methods of analysis (Item 8), which we read with the same language-model passes as the fund text. The Department of Labor's Office of Foreign Labor Certification publishes every H-1B labor condition application with the employer's name, the job title and the occupation code (fiscal years 2016 to 2026); we match employers to advisers by exact normalized name and, within finance, by brand name, and flag data-science and AI positions by occupation and title.

Funds are linked to CRSP through share-class tickers taken from filing headers and N-CEN and matched to CRSP's dated ticker histories, because tickers are reused after funds close. The link covers 82.0 percent of series and 93.4 percent of net assets in the universe; most unlinked series are insurance and variable-annuity funds without tickers. For these, returns and flows come from N-PORT, and expense ratios from the prospectus fee tables.

*Securities and Issuers.* Compustat North America and Global provide each issuer's GVKEY, point-in-time GICS classification, market value, age, R&D intensity and S&P index membership. Holdings are matched to Compustat by CUSIP, ISIN, the issuer's LEI, its six-digit CUSIP and, through GLEIF's relationship records (a current snapshot, so ultimate parents are not point in time), its ultimate parent's LEI; identifiers are checked for valid check digits first, because N-PORT carries placeholder values. GLEIF also provides headquarters countries, and OpenFIGI resolves instrument identifiers and security types, which separates shares of other funds from operating companies. For corporate holdings that Compustat cannot classify, mostly bonds and loans of private or foreign issuers, we obtained GICS codes and countries of risk from Bloomberg and, for bonds, mapped the issuer to the GVKEY of a listed parent. After these steps 95.7 percent of corporate holdings by weight carry a GVKEY and 96.6 percent an industry. Bloomberg also supplies each fund's primary benchmark, index constituents, the number of analysts with a recommendation on each issuer at each year-end (a check on the IBES counts), and consensus ratings at each quarter-end from 2019Q4. IBES supplies the analyst counts used in the analysis: the number of analysts with an outstanding recommendation on each issuer at each quarter-end, analysts' first recommendations on an issuer, and consensus ratings; IBES tickers are linked to GVKEY through the CRSP–IBES link, the CRSP–Compustat link, CUSIPs and SEDOLs. Product-market similarity between issuers comes from the text-based network industries of Hoberg and Phillips (2016), and distances between issuers and advisers use ZIP-code centroids from the U.S. Census gazetteer.

*Benchmarks and Factors.* Membership of U.S. benchmarks is each index provider's list of constituents at every quarter-end, obtained from Bloomberg and checked against Datastream; MSCI membership comes from tracker funds of another family, and FTSE Russell's constituent

weights give true active share for Russell-benchmarked funds; we use no other index weights (Section 3.3 and Appendix D.7). Equity factors (market, size, value, profitability, investment and momentum) come from Kenneth French's data library through WRDS. Bond factors are long-short returns of iShares bond ETFs from the CRSP mutual fund file: aggregate bonds, term (long minus short Treasuries), credit, high yield, emerging-market debt, municipal bonds and inflation protection. Each fund is evaluated against the model for its Lipper asset class, and peer-adjusted returns compare it with other funds in its Lipper class.

# B Outcome Definitions

Figure A1 shows how the sources in Table A1 are joined: every source is linked through one of four entity identifiers (the fund's SEC series ID, the adviser's CRD number, the security's identifiers, and the issuer key built from them) or by month (factor returns), and the resulting tables are merged on series and quarter. Table A2 defines the main outcomes used in the paper. Section 3.1 describes the holdings data, the construction of issuer books and the sample screen.

# C Classification Codebook

Table A3 lists the labels of the three classification passes described in Section 3.2. The described-process treatment requires a specific sentence whose passage is a fund-process or index-methodology statement with AI at least one input to selection. Disclosure-only adopters are split into those whose filings describe a process at a weaker role, such as research support, and those with only risk or limitation language (Figure 1).

Table A4 gives one example from each cell, quoted verbatim: three sentences that stop at the first or second pass, three process sentences that the third pass profiles as something other than a described process (a fund that weights companies by their AI revenues, which is not AI use; a model-risk factor that the second pass had called specific, which is generic risk; and a statement that AI is not used for autonomous decisions, which is a limitation), and four fund processes of increasing intensity, from AI tools in research (which does not qualify) to a model that selects between 60 and 190 securities each month, and one index methodology that ranks companies with natural language processing.

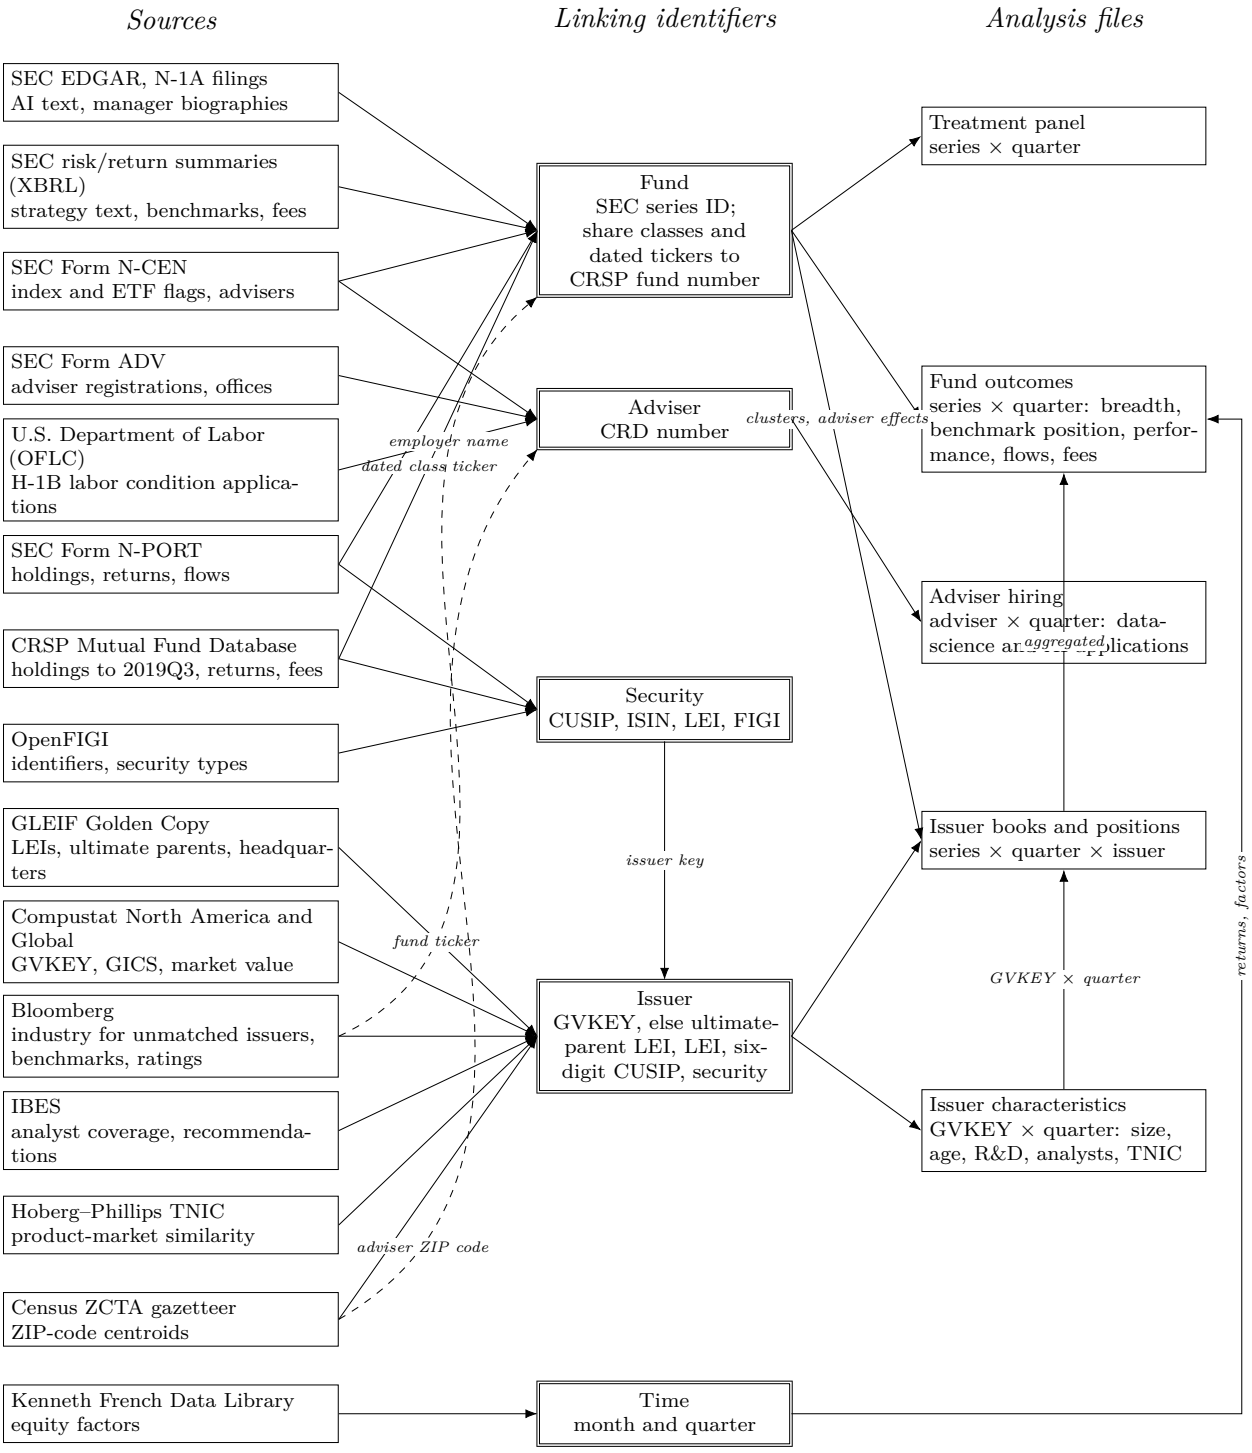
Table A5 gives concrete cases: the earliest dated and a left-censored adopter, funds that broadened after adopting in-house, a fund whose team shrank to one manager, AI supplied by an outside firm or an index provider, and a generic disclosure from a family-wide rollout.

**Table A1:** Data sources

| Source | Content used | Period | Link to the panel |
|---|---|---|---|
| *Treatment and fund text* | | | |
| SEC EDGAR, N-1A-family filings (485BPOS, 485APOS, 497, 497K, N-1A) | Full text of about 330,000 prospectuses and statements of additional information: AI sentences, portfolio-manager items and biographies, series headers identifying new and existing series | July 2018–September 2026 | Series ID from the filing header |
| SEC risk/return summary data sets (XBRL) | Principal-strategy text before adoption, prospectus benchmarks, fee tables | 2010Q4–2026Q2 | Series and class ID |
| *Funds and holdings* | | | |
| SEC Form N-PORT data sets | Quarterly holdings with asset category, issuer type, identifiers, country, payoff profile and fair-value level; net assets; monthly returns and flows | 2019Q4–2026Q2 | Series ID |
| SEC Form N-CEN data sets and EDGAR N-CEN filings | Index and ETF flags; adviser and subadviser CRD numbers | 2018–2026 | Series ID |
| CRSP Survivor-Bias-Free U.S. Mutual Fund Database | Holdings before N-PORT; monthly returns, total net assets, expense ratios, turnover, Lipper classes; returns of iShares bond ETFs for the bond factors | Holdings 2013Q1–2019Q3; returns 2017–2026 | Class ticker to CRSP fund number, dated; portfolio number to series |
| SEC Form ADV | Adviser registrations, main office, control persons, advisory employees (Item 5.B(1)), regulatory assets under management; Part 2A brochures (Item 8 methods of analysis) | 2013–2026 | Adviser CRD |
| U.S. Department of Labor, OFLC labor condition application disclosure files | Employer, job title and occupation of H-1B applications; data-science and AI positions | Fiscal years 2016–2026 | Employer name to adviser (exact name; brand name within finance) |
| *Securities and issuers* | | | |
| Compustat North America and Global | GVKEY, point-in-time GICS, market values, fundamentals, S&P index membership | 2012–2026 | CUSIP, ISIN, LEI, issuer CUSIP |
| GLEIF Golden Copy | Legal entity identifiers, ultimate-parent relationships, headquarters country | Snapshot, September 2026 | LEI, ISIN |
| OpenFIGI | Instrument identifiers and security types | – | CUSIP, ISIN |
| Bloomberg | GICS and country of risk for issuers outside Compustat; funds' primary benchmarks; index constituents; year-end analyst counts (check on IBES) and quarter-end consensus ratings | 2019–2026 | FIGI, ISIN, CUSIP; fund ticker |
| IBES | Analysts with an outstanding recommendation at each quarter-end; first recommendations; consensus ratings | 2012–2026 | IBES ticker to GVKEY (CRSP–IBES and CRSP–Compustat links, CUSIP, SEDOL) |
| Hoberg–Phillips TNIC | Product-market similarity between firms | 1988–2023 | GVKEY |
| *Benchmarks and factors* | | | |
| Bloomberg index member lists | Quarter-end constituents of the 19 U.S. equity benchmarks (membership only) | 2019Q4–2026Q2 | ISIN, CUSIP, FIGI |
| Datastream index lists (via WRDS) | Constituents of the S&P 500, Russell 2000 and 3000 (cross-check), STOXX Global 1800 ex Americas and S&P Global 1200 (MSCI substitutes) | 2019Q4–2026Q2 | ISIN, CUSIP |
| Index funds tracking MSCI indices (N-PORT holdings) | Membership of MSCI benchmarks, from a tracker outside the benchmarked fund's family (primary benchmark measure for MSCI-benchmarked funds) | 2019Q4–2026Q2 | Series ID |
| MSCI website | Current constituents of five MSCI indices (validation of the trackers) | June 2026 | Security name |
| FTSE Russell index weights (via WRDS) | Quarter-end constituent weights of 13 Russell indices (true active share for Russell-benchmarked funds) | 2019Q3–2026Q3 | CUSIP, ISIN to issuer key |
| Kenneth French data library (via WRDS) | Market, size, value, profitability, investment and momentum factors | 2017–2026 | Month |
| U.S. Census ZCTA gazetteer | ZIP-code centroids for distances between issuers and advisers | 2023 | ZIP code |

*Notes:* Periods are those used in the paper. Alphas use factor loadings estimated over up to 36 preceding months (at least 24), so they begin in 2019. The GLEIF relationship file is a current snapshot, so ultimate parents are not point in time.

**Figure A1:** Data sources and linkage




*Notes:* Sources are listed in Table A1. Funds are linked to CRSP through share-class tickers matched to CRSP's dated ticker histories, because tickers are reused. Holdings are aggregated to the issuer using the first available identifier in the order listed. A fund's equity and debt of the same issuer share one issuer key. Dashed lines are secondary links: Bloomberg's benchmark fields are keyed by fund ticker, and adviser offices are located by ZIP code. Labor condition applications are matched to advisers by employer name (exact name, and brand name within finance). The fund-level analysis files are merged on series and quarter; adviser hiring is merged on the adviser's CRD number and quarter.

## Table A2: Outcome definitions

| Outcome | Definition |
|---|---|
| *Convergence (fund-quarter unless noted)* | |
| Excess co-entry share | Share of the fund's new names that a described-process adopter at another adviser also adds in the same quarter, minus the share expected if those adopters drew their additions at random from all active funds' additions that quarter, scaled by the number of names they add per fund (Section 4.2). |
| Herding (group-quarter) | Lakonishok, Shleifer and Vishny (1992) statistic: the absolute difference between the share of the group's trades in a stock that are purchases and its cross-stock mean, minus its expectation under independent trading, averaged over stock-quarters; trades netted within adviser; stocks traded by at least three advisers. |
| Similarity to other AI users | Cosine similarity between the fund's corporate issuer weights and the aggregate issuer weights of described-process adopters at other advisers. |
| *Breadth (fund-quarter, corporate book)* | |
| Issuers held | log(1 + number of distinct issuers). |
| Sub-industries, industry groups held | log(1 + number of distinct GICS sub-industries or industry groups held); unclassified issuers are not counted. |
| Names below 0.5% of the book | Share of issuers, and their combined weight, with less than 0.5 percent of the corporate book. |
| Benchmark constituents held | log(1 + number of the stated benchmark's constituents held), from the index provider's quarter-end constituent list for U.S. indices and from the holdings of an index fund of another family for MSCI indices (MSCI World: the S&P Global 1200 list; MSCI ACWI: the ACWI ex USA tracker plus S&P 500 constituents); names outside the benchmark analogously. |
| Share of the benchmark's constituents held | Constituents held divided by the benchmark's number of constituents. |
| *Positions (fund–issuer–quarter; consecutive quarters only)* | |
| New name | An issuer held at the end of quarter $t$ but not at the end of $t-1$. |
| Names added, names dropped | log(1 + number of new names), log(1 + number of issuers held at $t-1$ but not at $t$). |
| New names / names held | Number of new names divided by the number of issuers held. |
| New to the adviser | Share of new names that no other fund of the same primary adviser held at $t-1$. |
| First ever held | Share of new names the fund never held in its earlier observed quarters (requires four prior quarters). |
| New industry group, new country | Share of new names whose GICS industry group, or country, was not in the fund's book at $t-1$. |
| Survival to four quarters | Share of new names still held at $t+4$. |
| Scale-up of survivors | log of the median ratio of weight at $t+4$ to weight at $t$ among new names still held. |
| Quick exits | Share of exits from positions held two quarters or less. |
| Weight in long-held names | Weight in issuers held for at least eight consecutive quarters. |
| Size tilt of new names | Weighted mean log market capitalization of new names minus that of continuing names. |
| Analysts covering new names, minus kept | log(1 + IBES analysts with an outstanding recommendation at the previous quarter-end), mean over new names minus mean over continuing names, equity holdings only. |
| *Position relative to the benchmark (equity funds; constituent membership)* | |
| Weight outside the benchmark | Share of the fund's corporate book in issuers that are not constituents of its stated benchmark (Bloomberg's primary benchmark, else the prospectus benchmark) at the quarter-end; a lower bound on active share. |
| Active share, Russell weights | $1 - \sum_i \min(w_{fi}, w_{bi})$ over renormalized corporate issuer books, with $w_{bi}$ FTSE Russell's quarter-end weights; Russell-benchmarked equity funds. |
| *Management team (fund-quarter; named portfolio managers in the prospectus)* | |
| Issuers per named manager | log of the number of issuers held divided by the number of named portfolio managers. |
| Single named manager | 1 if the prospectus names one portfolio manager. |
| *Adviser hiring (adviser-quarter; H-1B labor condition applications)* | |
| Data-science or AI application | A certified application whose occupation code is data scientist, statistician, operations research analyst or computer and information research scientist, or whose job title names data science, machine learning or AI; dated by quarter of receipt and matched to advisers by employer name. |
| Any data-science or AI application | 1 if the adviser has at least one such application in the quarter; before 2019, the adviser's share of quarters with one. |
| *Performance and costs (per quarter, sums of monthly values)* | |
| Primary-model alpha | Net return minus factor exposures from rolling 36-month regressions estimated before each month; six-factor equity model (market, size, value, profitability, investment, momentum) for equity funds, bond and municipal-bond factor models for taxable and municipal bond funds, a mixed model otherwise. |
| Peer-adjusted return | Net return minus the prior-assets-weighted mean return of the other funds in the same Lipper class and month. |
| Net flow | Dollar flow over the quarter divided by total net assets at the end of the previous quarter (CRSP, else N-PORT). |
| Value added | Gross primary-model alpha times prior-month total net assets (Berk and van Binsbergen 2015); inverse hyperbolic sine of $ millions. |

**Table A3:** Codebook for the language-model classification passes

| Pass and unit | Field | Labels |
|---|---|---|
| 1. Sentence | Type of AI mention | *Process*: the adviser or its models use AI in the investment process. *Thematic*: AI is what the fund invests in. *Risk*: AI as a risk factor. *Other*. |
| 2. Process sentence | Specificity | *Specific*: a concrete investment task performed with AI. *Generic*: AI tools in support, efficiency language, or statements limiting AI's role. |
| 3. Passage around a process sentence | Statement type | Fund process; index methodology; generic risk; limitation (the purpose is to limit AI's role); not AI use. |
| | Role in selection | The weakest role clearly supported: research support; one input to selection; primary input; AI selects holdings. Recorded for fund-process and index-methodology statements only. |
| | Tasks, techniques, data | Multiple labels each. Tasks include screening, selection or ranking, forecasting, portfolio construction, risk monitoring, trading, text analysis and general research. Techniques include machine learning, deep learning, natural language processing, reinforcement learning and generative models. Data: alternative, text, market and fundamental. |
| | Provenance, coverage, commitment | In-house or third-party model; whole portfolio or a sleeve; does use or may use. |

*Described-process treatment:* a specific sentence (pass 2) whose passage is a fund-process or index-methodology statement with AI at least one input to selection (pass 3). *Disclosure-only adopters:* funds whose AI language never meets this standard.

**Table A4:** Examples of AI language and its classification

| | Sentence (verbatim) | Pass 1 | Pass 2 | Pass 3 statement | Intensity (role) | Tasks | Classification |
|---|---|---|---|---|---|---|---|
| 1 | The Index focuses on companies that supply the specialized hardware, cooling systems, and energy management solutions necessary to support ultra-dense AI compute workloads. | Thematic | – | – | – | – | Not adoption |
| 2 | Regulatory frameworks governing AI's use, particularly concerning data privacy and protection, are evolving rapidly. | Risk | – | – | – | – | Not adoption |
| 3 | The Manager uses a number of artificial intelligence ("AI") tools to facilitate and enhance its operations, including its investment research processes, and will continue to. explore and expects to deploy other tools in future. | Process | Generic | – | – | – | Disclosure only |
| 4 | In determining weights, the Adviser considers each company's market share and revenue share from the sale or production of LLM products and related technologies. | Process | Generic | Not AI use | 0 | – | Not adoption |
| 5 | Portfolios that are actively managed, in whole or in part, according to a quantitative investment model (including models that utilize forms of artificial intelligence, such as machine learning) can perform differently from the market, based on the investment model and [. . . ] | Process | Specific | Generic risk | 0 | – | Disclosure only |
| 6 | The Manager operates the "human in the loop" principle; AI tools are not used to make autonomous investment and/or operational decisions for any Fund. | Process | Generic | Limitation | 0 | – | Disclosure only |
| 7 | The Adviser uses artificial intelligence ("AI") tools as part of its investment research process. | Process | Generic | Fund process | 1 (Research support) | general research | Disclosure only |
| 8 | The Sub-Adviser's quantitative process uses a combination of relative, intrinsic, and artificial intelligence models to rank companies within each economic sector. | Process | Specific | Fund process | 2 (One input) | selection ranking | Described process |
| 9 | The artificial intelligence model uses 250+ features engineered from a broad range of data such as fundamentals, analyst sentiment, prices and market activity, short interest and calendar effects to make relative forecasts of the stocks of the Index. | Process | Specific | Fund process | 3 (Primary input) | forecasting | Described process |
| 10 | The AI model is generally updated monthly and typically selects between 60 and 190 equity securities that exhibit strong value characteristics, such as those noted above, and have the greatest potential to achieve income and capital appreciation for inclusion in [. . . ] | Process | Specific | Fund process | 4 (AI selects) | selection ranking | Described process |
| 11 | In constructing the Underlying Index, the Index Provider first applies a proprietary natural language processing algorithm to the eligible universe, which seeks to identify and rank companies involved in the field of solar technology based on filings, disclosures and other [. . . ] | Process | Specific | Index method. | 3 (Primary input) | screening, selection ranking, text analysis | Described process |

*Notes:* Sentences are quoted verbatim from N-1A-family filings; long sentences are cut after 40 words, and page numbers picked up by text extraction are removed. Rows 1–3 stop at the first or second pass. Rows 4–11 are profiled in their passage (the surrounding text is not shown); the first pass labelled row 4 a process sentence, and the passage shows it describes the weighting of AI-related companies, not the use of AI. Intensity is the role of AI in selection (0 for statements that are not process descriptions). "Described process" requires a specific sentence in a fund-process or index-methodology passage with intensity of at least 2.

**Table A5:** Illustrative adopting funds

| Case | Fund (filing) | What the filing says | What we observe |
|---|---|---|---|
| Earliest dated adoption | James Alpha Yorkville MLP Portfolio, Saratoga Advantage Trust (497, Dec. 31, 2018) | "This evaluation includes advanced linear and non-linear numerical techniques, including artificial intelligence." | Described-process adoption dated 2018Q4, the earliest in the filing corpus; the fund filed no Form N-PORT, so it is outside the universe and not among the 132 adopters. |
| Left-censored | Catalyst Multi-Strategy Fund, Mutual Fund Series Trust (485APOS, Aug. 31, 2018) | "The Trading Advisor employs multiple unique trading models ... namely, relative value, trend following, pattern recognition, machine learning ... and calendar spread strategies." | AI language in the first filing we read; dropped from the event studies. |
| In-house broadening | Shelton Emerging Markets Fund (485BPOS, Apr. 29, 2022) | "... the team uses data science and machine learning to conduct a deeper dive into each candidate firm to determine the investment merit, suitability for the portfolio, and pertinent risk factors." | Adopter 2022Q2, run in-house; issuers held 32 in the quarter before adoption, 56 four quarters after. |
| In-house broadening (ETF) | ERShares NextGen Entrepreneurs ETF (485APOS, Jan. 22, 2021) | "The EF model incorporates a bottom-up investment orientation, powered by artificial intelligence ('AI'), that includes investment criteria such as management attributes, sector, growth, value, leverage, market capitalization (size), momentum, and geographic orientation." | Adopter 2021Q1, run in-house; issuers held 50 before, 92 four quarters after. |
| Fewer managers | Voya Small Company Portfolio (485BPOS, Apr. 25, 2022) | "The Sub-Adviser uses quantitative methods, including artificial intelligence ('AI') models, to select securities and to support portfolio trading." | Adopter 2022Q2 (affiliated subadviser); named managers 2 before, 1 four quarters after; issuers 124 to 169. |
| Third-party AI supplier | AI Powered Equity ETF (EquBot, subadviser) | "As an IBM Global Entrepreneur company, EquBot leverages IBM's Watson AI to conduct an objective, fundamental analysis of U.S.-listed common stocks and real estate investment trusts ..." | Supplier extracted as EquBot (model: EquBot Model); one of few funds whose AI comes from a named vendor or subadviser. |
| Index methodology | Global X E-commerce ETF (2018 filings) | "In constructing the Solactive E-commerce Index, the Index Provider first applies a proprietary natural language processing algorithm to the eligible universe ..." | Index-methodology passage; left-censored. |
| Generic disclosure (rollout) | First Trust Income Opportunities ETF (485BPOS, Dec. 29, 2025) | "The Fund, the Advisor and the Fund's service providers may utilize AI technologies, including machine learning models and generative AI, to improve operational efficiency and in connection with research, among other purposes." | Disclosure-only adopter; one of 94 active First Trust funds whose first AI process language is dated 2025Q4. |

*Notes:* Cases chosen to illustrate the classification and the main results, not as a representative sample. Quotations are from the fund's own N-1A-family filings (ellipses ours); dates are filing dates. "Issuers held" counts corporate issuers in the fund's N-PORT book in the quarter before adoption and four quarters after. Named managers are from the prospectus portfolio-manager item, carried forward between filings. Supplier names come from the language-model extraction in Section 5.6.

# D Additional Estimates

## D.1 Details of the Results in Section 5

This subsection collects the sub-tests summarized in Section 5, in the order of that section.

*Convergence.* The difference-in-differences of the positive control around the later fund's first use of the shared subadviser is positive ($p = 0.027$; $p = 0.003$ with index-event names excluded). Pairs that share only an index provider show excess co-entry of 0.19 points, against 0.12 for their controls ($p = 0.73$), and the three cross-adviser pairs that name the same AI vendor in an index methodology show +0.65 points ($p = 0.31$). Excluding names involved in an index event, so that two funds adding a stock because it entered the Russell 2000 do not count as converging, leaves the excess co-entry share between −0.9 and +0.1 percentage points (all $p \geq 0.46$). The excess co-entry share for names inside and outside the fund's benchmark, and for the equity book alone, is never significantly positive; where it is significant, it is negative. Because the share is a ratio over the names a fund adds, a fund that adds a single name scores either far above or far below its expectation. Restricting the test to fund-quarters with at least three new names gives −1.2 to +1.5 percentage points (all $p \geq 0.31$), and restricting it to fund-quarters with at least five gives −2.1 to +0.5 (one design significant, and negative). Adopters' co-entry with funds that will adopt later but have not yet, and with disclosure-only adopters, shows no robust change, and randomization inference that reassigns adoption to placebo funds gives $p = 0.15$ for the TWFE estimate of the excess share. The raw overlap falls because adopters add less popular names, and the expected share falls with it, which is why the excess is zero. In the herding test, the difference between adopters and their matched groups is 2.2 points. Before adoption, the same funds' measure was 0.032, indistinguishable from that of their matched groups (0.036), and disclosure-only adopters after their adoption stand at 0.041, against 0.050. We do not lean on the before-and-after comparison for adopters, which rests on eleven pre-adoption quarters and which a timing permutation leaves insignificant (two-sided $p = 0.23$). For portfolio similarity, a single post-adoption indicator in TWFE picks up 1.8 points ($p = 0.046$), but the pre-period coefficients in the event studies are of the same size and negative (−1.1 to −1.6 points, $p = 0.05$ to 0.09): the series drifts upward through the whole window, before adoption as much as after, so whatever similarity AI users share is not acquired at adoption.

*Breadth and Positions.* In the within-adviser design, which compares adopters with the adviser's other funds, the contrast with disclosure-only adopters is 6.3 to 12.1 log points but less precise (p between 0.035 and 0.21). Splitting the disclosure-only adopters into the two groups of Figure 1 (121 research-support funds at 16 advisers and 188 boilerplate funds at 31 in the estimation sample), the effect relative to research-support funds alone is 6.8 to 13.5 log

points (all p $\leq$ 0.003). Research-support funds add no more names ($-$2.1 to +9.8 log points, none significant), their new-name share moves by $-$0.9 to +1.3 percentage points (significant in one design), and their sub-industry count does not change ($-$0.8 to +0.8). Without First Trust, the largest adviser in that group, research-support funds broaden by +0.2 to +3.4 log points (all p $\geq$ 0.27), and the contrast is 6.0 to 13.2 (all p $\leq$ 0.008). These p-values are from the wild cluster restricted bootstrap. Counting only issuers whose stock a fund holds, adopters broaden by 6.5 to 13.2 log points across the six Sun–Abraham designs (all p $\leq$ 0.020); among the pure-equity funds, issuers rise by 8.2 to 14.5 log points (all p $\leq$ 0.033). New names make up 3.2 to 5.2 percentage points more of the names held, and the share of exits that close positions held two quarters or less rises by 1.5 to 10.1 percentage points from a base of 30 percent, a suggestive result (significant in four of six designs). When we split post-adoption event time at three quarters in TWFE, names added rise further from early to late (+22 log points, p = 0.025), and so do names dropped (+17, p = 0.017), while weight added to continuing names and weight in names held two years or more do not rise.

*Robustness.* Stacked difference-in-differences gives 8.1 log points (p = 0.037). Randomization inference that reassigns adoption to placebo funds or placebo advisers rejects at p $\leq$ 0.006. Dropping any one treated adviser leaves the TWFE effect between 6.6 and 10.1 log points, and quarter-by-investment-style fixed effects give 8.4 to 14.1 (all p $\leq$ 0.001). Quantitative adopters broaden by 4.9 to 10.0 log points (significant in one of six designs), against 10.6 to 15.9 for fundamental adopters (all six); the difference is significant in two of six designs, and TWFE puts fundamental adopters at 7.9 log points and quantitative adopters at 8.3 (difference p = 0.94). A TWFE lead over quarters $-$4 to $-$2 is $-$3.3 log points (p = 0.042).

*Distance of Added Securities.* Industry groups change by $-$0.1 to +0.6 log points in Callaway–Sant'Anna estimates. The number of countries held rises by +4.4 to +11.7 log points, significant in three of six designs. Among the names a fund adds, the share that no sibling fund of the same adviser held in the previous quarter does not rise ($-$6.7 to $-$0.6 percentage points across the six designs, significant in one, from a base of 44 percent), nor does the share the fund had never held before (71 percent at baseline), the share in an industry group the fund did not hold (7 percent), or the share in a country it did not hold (6 percent). The market-capitalization tilt of new names ranges from $-$14 to +2 log points and is significant in no design. New names are no more often recent listings, and their distance from the adviser's office does not change robustly. Measured on equity holdings at the end of the quarter before the trade, new names carry more IBES analysts than kept names after adoption only with quarter $-$1 as reference (14 to 15 log points), and the difference vanishes when the reference is moved; the added names are also no more often ones that an analyst

begins to cover within the following year. Bloomberg's year-end analyst counts give the same answer, and none of 25 consensus-rating outcomes is robust, whether the ratings come from Bloomberg or from IBES. At the level of the adviser, issuers held do not rise once launches of new AI funds in the adoption quarter are excluded (TWFE, 4.4 log points, $p = 0.76$). In the 19 corporate groups with a described-process adopter, the group's combined book gains 5.0 log points in TWFE ($p = 0.57$), and the Sun–Abraham estimates are negative. The numbers of industry groups, sub-industries, and countries held do not change at either level.

*Benchmarks.* Among the 87 adopters with an observed benchmark, breadth matches the full sample (8.9 to 14.4 log points), and constituents held rise by 9.3 log points relative to the adviser's own index funds ($p < 0.001$). The count of names outside the benchmark rises by 3.0 to 10.6 log points, and the difference between growth inside and outside the benchmark is not itself significant. The weight held outside the benchmark falls by 1.9 points relative to disclosure-only adopters ($p = 0.024$) and by 2.0 relative to the adviser's index funds ($p = 0.005$). Russell active share falls in two of six designs. About a fifth of Russell-benchmarked equity funds' fund-quarters lie below the 60 percent level that Cremers and Petajisto (2009) use to identify closet indexers. Among the 53 adopters benchmarked to U.S. indices, breadth itself is small (TWFE +0.1 log points, insignificant; 2.9 to 6.7 in the Sun–Abraham designs), and constituents held rise by 3.4 to 7.0 log points in the six event-study designs but not in the TWFE, stacked, or randomization-inference estimates. The benchmark results hold across estimators, when funds advised by BlackRock are excluded, and with benchmark-by-quarter fixed effects that absorb index reconstitutions.

*Heterogeneity.* A separate classification of who supplies the AI, which the language model extracts from the passages that name a third party, gives the same answer as the in-house split: funds whose AI is the adviser's own broaden by 8.8 to 16.7 log points in all six designs. The much larger 2023–26 cohort broadens by 4.5 to 8.6 log points (significant in all six). ETFs and mutual funds do not differ, and no difference by what the AI is described as doing, or by the data it reads, is significant in more than two of six designs. Funds whose adoption filing contains two or more specific AI sentences broaden by 9.5 to 17.5 log points, and funds with a single sentence by 6.1 to 8.5, both significantly in all six designs. This dose-response pattern is consistent with the text measuring something real about the process, although the difference between the two groups is significant in only one design.

*Performance.* Adopters' flow-performance relation is no steeper than other funds'. For each name a fund adds or drops, we compute its return over the next quarter and the next four quarters, net of the average return of stocks in the same region, size decile, and momentum quintile. With quarter $-1$ as reference, the added names' next-quarter return rises by 3.2 to 3.5 percentage points after adoption ($p = 0.08$ to 0.14). However, the same comparison

shows a gap of 2.3 points before adoption ($p = 0.027$), and the rise disappears under the other two references. Over four quarters, the estimates range from $-4.1$ to $+4.8$ points, with no p-value below 0.18, and added names do not outperform the names the fund drops ($-4.9$ to $+5.5$ points across the six designs, none significant). Shifts in the momentum of added names and in the profitability rank of the book amount to one to three percentile points, and they either predate adoption or appear equally among disclosure-only adopters. Splitting adopters by fund, team, and AI attributes, two cells pass the reference checks: 24 recently launched ETFs, whose outperformance predates their adoption filing, and 35 funds whose filings name text or alternative data. With sixteen splits and Benjamini–Hochberg q-values of 0.25 and 0.27, we do not read either as a finding (Appendix D.8).

## D.2 Event-Time Coefficients and Alternative Estimators

Table A6 lists the coefficients plotted in Figure 3. Described-process adopters' issuer counts are within three log points of the reference quarter from eight to six quarters before adoption, dip by up to five log points between quarters $-4$ and $-2$, and rise from the adoption quarter onward. Disclosure-only adopters' coefficients stay within about five log points of zero throughout. Tables A7 and A8 report the alternative estimators and the sensitivity analysis summarized in Section 5.3.

## D.3 Robustness of the Breadth Effect

Tables A9 and A7 report the alternative estimators summarized in Section 5.3. Callaway–Sant'Anna estimates with never-treated controls are 6.4 log points over event quarters 0 to 2 ($p = 0.004$) and 8.7 to 10.8 over 0 to 6 ($p = 0.015$ cell-weighted, 0.055 for the simple mean); they rise when the base period is moved one or two quarters earlier to allow for anticipation, to 7.6 and 8.6 log points over 0 to 2 and 9.6 and 10.1 cell-weighted over 0 to 6 (Appendix D.4). Stacked difference-in-differences gives 8.1 log points ($p = 0.037$), and 7.4 when quarters $-2$ and $-1$ are dropped ($p = 0.086$). Randomization inference rejects in three of four procedures: reassigning adoption to placebo funds or placebo advisers gives $p \leq 0.006$, and permuting adoption dates across advisers rather than funds gives $p = 0.079$. Dropping any one treated adviser leaves the TWFE effect between 6.6 and 10.1 log points. Controlling for investment style with quarter-by-style fixed effects gives 7.7 log points in TWFE ($p = 0.030$; 8.4 to 14.1 in the Sun–Abraham designs, Section 5.3), and dropping the 52 control funds whose strategy text describes an ML-type technique without AI vocabulary (0.4 percent of the never-adopting funds in the universe) leaves the estimate unchanged.

The effect is broad rather than carried by a few funds. For each adopter we compute a

**Table A6:** Event-study coefficients for issuers held

Sun–Abraham, disclosure design, window to $k = 6$ (binned), reference quarter $-1$; log points $\times$ 100.

| Event quarter $k$ | Described process | (s.e.) | Disclosure only | (s.e.) |
|---|---|---|---|---|
| $-8$ | $-2.3$ | (2.5) | $-1.3$ | (2.4) |
| $-7$ | $-0.3$ | (2.4) | $+0.6$ | (1.1) |
| $-6$ | $+2.1$ | (2.1) | $+0.8$ | (1.3) |
| $-5$ | $-1.1$ | (2.1) | $+1.2$ | (1.1) |
| $-4$ | $-4.7$ | (2.1) | $+1.3$ | (1.1) |
| $-3$ | $-2.7$ | (1.6) | $+0.7$ | (0.9) |
| $-2$ | $-2.3$ | (1.4) | $+0.7$ | (0.7) |
| $-1$ | (reference) | | (reference) | |
| 0 | $+2.9$ | (1.3) | $+0.6$ | (0.5) |
| 1 | $+5.5$ | (1.8) | $+1.3$ | (1.0) |
| 2 | $+9.8$ | (2.0) | $+3.1$ | (1.5) |
| 3 | $+6.2$ | (2.1) | $+5.0$ | (2.0) |
| 4 | $+15.4$ | (2.7) | $+2.7$ | (1.4) |
| 5 | $+23.2$ | (5.0) | $+5.1$ | (2.7) |
| 6+ | $+26.8$ | (1.8) | $+2.3$ | (1.3) |

*Notes:* These are the coefficients plotted in Figure 3. Standard errors are the analytical adviser-clustered standard errors, which understate uncertainty where a cohort-quarter cell rests on one adviser; inference in the text uses the adviser pairs bootstrap for the primary outcome (Table 3) and, for HonestDiD, a leave-one-adviser-out covariance.

calendar-matched difference-in-differences against never-adopters over the same quarters. The median adopter's estimate is 8.7 log points (90 percent interval 1.8 to 17.4, adviser-clustered bootstrap), close to the mean of 9.0, and 64 percent of adopters broaden. The distribution has a long right tail: a few funds, mostly small, roughly double their holdings or more. Trimming the one to ten most extreme adopters from each end leaves the TWFE estimate between 6.8 and 8.8 log points ($p \leq 0.018$), and Sun–Abraham estimates without the seven largest responders are 6.0 to 9.7 ($p \leq 0.014$). Removing only the largest responders lowers the TWFE estimate further (4.0 log points without the top seven, $p = 0.21$), as removing the upper tail of any positive effect must.

On the pre-period, a lead term over quarters $-4$ to $-2$ is $-3.3$ log points ($p = 0.042$). The Rambachan and Roth (2023) smoothness tolerance of 0.25 log points per quarter squared ($M = 0.0025$ in log units) is about five times that of the broader treatment definition on the same data. Across designs (Table A8), the relative-magnitudes breakdown is 0.55 for the window of two quarters, 0.4 for six quarters, and 0.05 when quarter $-3$ is the reference and quarters $-2$ and $-1$ are treated as anticipation; the smoothness breakdowns range from 0.0025 to 0.005.

Table A9 collects the main alternative estimators and specifications.

**Table A7:** Alternative estimators for issuers held

Described-process adopters vs never-adopters; log points × 100.

| Estimator and target | Estimate | 95% interval | $p$ |
|---|---|---|---|
| *Callaway–Sant'Anna, never-treated controls* | | | |
| base $g-1$, mean ATT($e$), $e = 0$–2 | +6.4 | [+3.2, +9.6] | 0.004 |
| base $g-1$, mean ATT($e$), $e = 0$–6 | +10.8 | [−0.3, +21.8] | 0.055 |
| base $g-1$, cell-weighted, $e = 0$–6 | +8.7 | [+2.5, +14.9] | 0.015 |
| base $g-2$, mean ATT($e$), $e = 0$–2 | +7.6 | [+3.0, +12.2] | 0.004 |
| base $g-2$, mean ATT($e$), $e = 0$–6 | +11.4 | [−0.2, +22.9] | 0.052 |
| base $g-2$, cell-weighted, $e = 0$–6 | +9.6 | [+2.8, +16.4] | 0.015 |
| base $g-3$, mean ATT($e$), $e = 0$–2 | +8.6 | [+2.2, +15.0] | 0.012 |
| base $g-3$, mean ATT($e$), $e = 0$–6 | +11.5 | [−1.4, +24.5] | 0.080 |
| base $g-3$, cell-weighted, $e = 0$–6 | +10.1 | [+2.0, +18.2] | 0.024 |
| *Stacked difference-in-differences, cohort stacks with clean never-treated controls* | | | |
| Base | +8.1 | [+0.5, +15.7] | 0.037 |
| Dropping event quarters −2, −1 | +7.4 | [−0.9, +15.7] | 0.086 |
| With a lead term (−4 to −2) | +6.8 | [−0.7, +14.4] | 0.078 |
| *Randomization inference for the TWFE estimate (+8.0 log points)* | | | |
| Placebo funds | | | 0.001 (studentized 0.037) |
| Placebo advisers | | | 0.006 (studentized 0.063) |
| Permuted adoption dates, funds | | | 0.009 (studentized 0.179) |
| Permuted adoption dates, advisers | | | 0.079 (studentized 0.100) |

*Notes:* Callaway–Sant'Anna group-time effects are estimated on the Sun–Abraham sample with 999 adviser pairs-bootstrap replications (Appendix D.4); the base period is the quarter before adoption ($g-1$) or one or two quarters earlier, allowing for anticipation. Stacked estimates use a window of $g-8$ to $g+8$. Randomization inference re-estimates the TWFE model for each draw; $p$ is the share of placebo coefficients at least as large in absolute value, and the studentized version uses the adviser-clustered $t$.

**Table A8:** Sensitivity of the breadth effect to violations of parallel trends

Rambachan and Roth (2023) bounds for the Sun–Abraham average post-adoption effect on issuers held.

| Design | Estimate | s.e. | Pre-trend Wald $p$ | Breakdown $\bar{M}$ | Breakdown $M$ |
|---|---|---|---|---|---|
| Window 0–6, reference −1 | +12.8 | 3.4 | 0.002 | 0.4 | 0.0025 |
| Window 0–6, reference −3, −2 and −1 as anticipation | +13.6 | 6.0 | 0.003 | 0.05 | 0.00325 |
| Window 0–2, reference −1 | +8.7 | 2.4 | 0.002 | 0.55 | 0.005 |

*Notes:* Estimates in log points × 100; covariance from a leave-one-adviser-out jackknife of the event-time vector. $\bar{M}$: relative-magnitudes restriction, under which the post-adoption violation is at most $\bar{M}$ times the largest change between consecutive pre-adoption quarters. $M$: smoothness restriction on the second difference of the violation, in log units (log points ÷ 100) per quarter squared. The breakdown value is the largest value on the grid at which the robust 95% interval excludes zero.

**Table A9:** Robustness of the breadth effect

Issuers held, described-process adopters.

| Specification | Estimate | p / interval |
|---|---|---|
| Two-way fixed effects (event time −8 to +8) | +8.0 lp | 0.022 |
| TWFE with adviser × quarter FE | +8.1 lp | 0.017 |
| Callaway–Sant'Anna, ATT over e = 0–2 | +6.4 lp | [3.2, 9.6], p = 0.004 |
| Callaway–Sant'Anna, cell-weighted over e = 0–6 | +8.7 lp | [2.5, 14.9], p = 0.015 |
| Stacked difference-in-differences | +8.1 lp | [0.5, 15.7], p = 0.037 |
| Randomization inference: placebo funds / advisers | +8.0 lp | p = 0.001 / 0.006 |
| Jackknife, dropping one treated adviser | +6.6 to +10.1 lp | min to max |
| Quarter × investment-style FE (TWFE) | +7.7 lp | 0.030 |
| Quant / quantamental adopters (TWFE) | +8.3 lp | 0.037 |
| Fundamental adopters (TWFE) | +7.9 lp | 0.083 |
| Drop controls with ML language but no AI terms (TWFE) | +8.0 lp | 0.025 |
| Broader treatment (any specific AI sentence), 18 designs | +5.5 to +8.8 lp | 0.001–0.014 |
| HonestDiD, relative magnitudes breakdown $\bar{M}$ | 0.40 | window 0–6, jackknife covariance |
| HonestDiD, smoothness breakdown M | 0.0025 | broader treatment: 0.0005 |
| Joint pre-trend Wald test | | p = 0.002 |

*Notes:* lp = log points × 100. Callaway–Sant'Anna rows use the adviser pairs bootstrap of Appendix D.4. Quant and fundamental adopters do not differ (difference +0.4 lp, p = 0.94). Style is classified from each fund's strategy text in force before adoption; 42 percent of described-process adopters are quantitative or quantamental, against 14 percent of never-adopters.

## D.4 Callaway–Sant'Anna Estimates and Inference

The main text reports Sun–Abraham estimates, with p-values for the primary outcome from an adviser pairs bootstrap. We checked both choices against alternatives.

*Inference.* The wild cluster restricted bootstrap holds the set of treated advisers fixed and perturbs their residuals. For the primary outcome it gives p-values of at most 0.009 for the effect of adoption and at most 0.005 for the placebo contrasts in the disclosure and family designs (0.029 to 0.104 for the contrast with disclosure-only adopters in the within design). Resampling advisers with replacement and re-estimating the event study gives at most 0.035 for the effect across all eighteen designs, 0.010 to 0.030 for the contrast with disclosure-only adopters in the disclosure design, and 0.010 to 0.060 for the contrast with the adviser's index funds. Several event-time cells after adoption rest on one or a few advisers, and a procedure that keeps the advisers fixed cannot learn from those cells how much the estimate depends on which advisers adopted. We therefore report the resampling p-values for the primary outcome (Table 3). For secondary outcomes the tables report wild-bootstrap p-values, which on the evidence of the primary outcome are optimistic. The multiplier bootstrap that is the default for Callaway–Sant'Anna estimators keeps the treated advisers fixed in the same way; on the primary outcome its standard error for the average over quarters 0 to 6 is about a third of the resampling standard error, and we do not report it.

*Estimator.* Table A10 reports Callaway and Sant'Anna (2021) estimates for the main

outcomes, computed on exactly the sample of each Sun–Abraham estimate and with the same adviser resampling. Signs agree for every outcome with a significant Sun–Abraham estimate; three reversals (excess co-entry with at least three new names, the size tilt of new names and peer-adjusted returns) are within about a log point of zero under both estimators, and names outside the benchmark, positive but significant in only two of six Sun–Abraham designs, is −2.0 over quarters 0 to 2 and +1.4 over 0 to 6 under Callaway–Sant'Anna. For the outcomes with significant effects (issuers, equity issuers, sub-industries, benchmark constituents held and the share of the benchmark held) and for the placebo contrasts, the Callaway–Sant'Anna estimates are smaller, by up to a quarter, and less precise over the longer window: issuers held rise by 6.4 log points over event quarters 0 to 2 ($p = 0.004$) and by 10.8 over 0 to 6 ($p = 0.055$), and the contrast with disclosure-only adopters is 5.2 ($p = 0.019$) and 8.6 ($p = 0.117$). Over the first three quarters, issuers, equity issuers, sub-industries, benchmark constituents held, the share of the benchmark held and both placebo contrasts are significant under every base period, while the weight outside the benchmark is not; over seven quarters most of these estimates are imprecise. The convergence results are null or negative under both estimators, and the performance results show the same reference-quarter pattern as in Table 7. The difference in precision is concentrated in the later event quarters: over quarters 0 to 2 the standard error of the effect on issuers is 1.6 log points for Callaway–Sant'Anna and 2.4 for Sun–Abraham, while over quarters 0 to 6 it is 5.6 against 3.5. Two features of the estimators produce this. Sun–Abraham compares each event quarter with all of a cohort's pre-adoption quarters and estimates the common quarter effects from never-adopters' full histories, which is efficient when parallel trends hold over the whole window. Callaway–Sant'Anna compares each later quarter with the single quarter before adoption, which requires parallel trends only between those two quarters but leaves the other pre-adoption quarters unused. Noise in that one base quarter therefore passes directly into every later cell, and the late cells rest on the 2021–22 adopters at a handful of advisers: their Callaway–Sant'Anna cohort average is 28 log points against 5 for the 2023–26 cohorts, in line with Figure 6. Not-yet-adopting controls and outcome-regression adjustment for size, breadth, ETF status and asset class before adoption change the estimates for issuers by less than one log point.

## D.5 Heterogeneity

Table A11 gives the Sun–Abraham ranges behind Figure 6, the differences between groups, and the TWFE estimates. Table A12 repeats the split by managers' training in the six Sun–Abraham designs. The difference between teams with and without a STEM-trained manager has the same sign in all seven designs and is significant at the 5 percent level in

## Table A10: Callaway–Sant'Anna estimates for the main outcomes

Described-process adopters vs never-adopting active funds; × 100 (log points, percentage points or percent per quarter) [adviser pairs bootstrap $p$].

| Outcome | SA, 0–6 | CS, 0–2 | CS, 0–6 | CS, pre-period | CS with covariates, 0–6 | CS significant |
|---|---|---|---|---|---|---|
| *Breadth* | | | | | | |
| Issuers held (log) | +12.8 | +6.4 [0.004] | +10.8 [0.055] | −0.5 [0.857] | +10.9 | 3 of 6 |
| minus disclosure-only adopters | | +5.2 [0.019] | +8.6 [0.117] | | | |
| minus the adviser's index funds | | +6.1 [0.021] | +8.8 [0.137] | | | |
| Sub-industries held (log) | +7.7 | +3.8 [0.012] | +6.0 [0.118] | +0.6 [0.726] | +6.1 | 3 of 6 |
| Industry groups held (log) | +2.3 | +0.6 [0.363] | +0.6 [0.842] | +0.2 [0.843] | +0.6 | 0 of 6 |
| Equity issuers held (log) | +10.5 | +5.8 [0.012] | +10.3 [0.066] | −1.7 [0.470] | +10.3 | 3 of 6 |
| *Convergence* | | | | | | |
| Excess co-entry share | −0.4 | −4.4 [0.016] | −1.0 [0.626] | −2.0 [0.366] | −1.3 | 1 of 6 |
| Excess co-entry share, ≥ 3 new names | −0.0 | −0.8 [0.692] | +0.5 [0.745] | +0.5 [0.664] | +0.4 | 0 of 6 |
| Cosine similarity to other AI users | +0.7 | +0.4 [0.410] | +0.7 [0.393] | −1.8 [0.045] | +0.7 | 0 of 6 |
| Overlap with other AI users | +0.3 | +0.1 [0.539] | +0.3 [0.260] | −0.6 [0.191] | +0.3 | 0 of 6 |
| *Positions and added names* | | | | | | |
| Names added (log) | +40.2 | +15.5 [0.175] | +36.0 [0.034] | −3.1 [0.712] | +36.0 | 3 of 6 |
| New names / names held | +4.6 | +3.4 [0.111] | +4.3 [0.037] | −0.7 [0.319] | +4.3 | 2 of 6 |
| Exits after ≤ 2 quarters / exits | +10.1 | +6.4 [0.020] | +9.1 [0.085] | −2.2 [0.411] | +10.3 | 2 of 6 |
| New names new to the adviser | −6.7 | −2.6 [0.068] | −4.1 [0.175] | −1.6 [0.298] | −4.4 | 0 of 6 |
| Size tilt of new names | +0.5 | −0.6 [0.941] | −0.8 [0.942] | +0.5 [0.957] | −1.1 | 0 of 6 |
| Analysts covering new names, minus kept | +14.7 | +7.2 [0.208] | +8.8 [0.239] | +1.1 [0.696] | +8.9 | 0 of 6 |
| *Benchmark* | | | | | | |
| Share of the benchmark's constituents held | +1.3 | +0.9 [0.006] | +1.3 [0.008] | −0.5 [0.434] | +1.3 | 6 of 6 |
| Benchmark constituents held (log) | +12.1 | +8.2 [0.012] | +11.5 [0.058] | −1.5 [0.547] | +11.6 | 3 of 6 |
| Names outside the benchmark (log) | +5.3 | −2.0 [0.739] | +1.4 [0.857] | −8.0 [0.035] | +1.4 | 0 of 6 |
| Weight outside the benchmark | −1.9 | −1.5 [0.084] | −2.4 [0.047] | −0.7 [0.030] | −2.3 | 1 of 6 |
| Active share, Russell weights | −0.9 | −1.7 [0.131] | −1.2 [0.156] | +2.2 [0.057] | −1.2 | 0 of 6 |
| *Management team and performance* | | | | | | |
| Issuers per named manager (log) | +23.8 | +15.0 [0.014] | +19.0 [0.117] | +1.4 [0.776] | +19.6 | 3 of 6 |
| Single named manager | +10.8 | +7.0 [0.037] | +11.0 [0.088] | −1.0 [0.449] | +10.9 | 3 of 6 |
| Net primary-model alpha, per quarter | +1.3 | +1.2 [0.098] | +1.0 [0.172] | +0.9 [0.145] | +1.0 | 0 of 6 |
| Peer-adjusted return, per quarter | +0.1 | +0.9 [0.150] | −0.1 [0.676] | +0.3 [0.433] | −0.2 | 0 of 6 |
| Net flow / lagged assets, per quarter | +2.6 | +2.0 [0.014] | +2.7 [0.006] | +2.3 [0.014] | +2.8 | 2 of 6 |

*Notes:* SA: Sun–Abraham point estimate on the same sample (window 0–6, reference quarter −1); its inference is in the main tables. CS: Callaway and Sant'Anna (2021) group-time effects with never-adopting funds as controls and the quarter before adoption as base, averaged over event quarters 0–2 or 0–6 (equal weight per event quarter); pre-period: average over quarters −8 to −2 relative to the same base. Covariates: outcome-regression adjustment for log net assets, log issuers, ETF status and Lipper asset class in the base quarter (point estimate only). CS significant: number of six specifications (base quarter $g-1$, $g-2$ or $g-3$, window 0–2 or 0–6) with $p < 0.05$ and the sign of the 0–6 estimate. Each CS estimate uses exactly the Sun–Abraham sample of its outcome. $p$-values: advisers resampled with replacement and every group-time cell re-estimated, 999 draws.

three.

**Table A11:** Heterogeneity of the breadth effect

Issuers held, log points × 100.

| Group | Sun–Abraham range (6) | Significant | TWFE [$p$] | Funds | Advisers |
|---|---|---|---|---|---|
| In-house | +11.5 to +19.8 | 6 of 6 | +12.4 [0.001] | 57 | 15 |
| Subadvised | +2.9 to +9.5 | 1 of 6 | +2.4 [0.705] | 45 | 18 |
| *In-house minus subadvised* | +5.7 to +13.7 | 1 of 6 | +10.0 [0.184] | | |
| Adopted 2021–22 | +21.4 to +28.8 | 4 of 6 | +29.8 [0.040] | 9 | 5 |
| Adopted 2023–26 | +4.5 to +8.6 | 6 of 6 | +4.4 [0.258] | 93 | 24 |
| *2021–22 minus 2023–26 adopters* | +16.7 to +21.4 | 4 of 6 | +25.4 [0.048] | | |
| No STEM-trained PM | +8.3 to +12.5 | 5 of 6 | +11.3 [0.031] | 47 | 11 |
| STEM-trained PM | +0.4 to +3.5 | 0 of 6 | −6.3 [0.903] | 26 | 10 |
| *No STEM-trained PM minus STEM-trained PM* | +4.8 to +11.9 | 2 of 6 | +17.6 [0.037] | | |
| Fundamental style | +10.6 to +15.9 | 6 of 6 | +7.9 [0.083] | 63 | 23 |
| Quant / quantamental | +4.9 to +10.0 | 1 of 6 | +8.3 [0.037] | 39 | 14 |
| *Fundamental minus quant / quantamental* | +1.9 to +7.6 | 2 of 6 | −0.4 [0.938] | | |
| 2+ specific sentences | +9.5 to +17.5 | 6 of 6 | +9.9 [0.073] | 66 | 18 |
| 1 specific sentence | +6.1 to +8.5 | 6 of 6 | +4.6 [0.062] | 36 | 12 |
| *2+ minus 1 specific sentence* | +3.4 to +9.0 | 1 of 6 | +5.3 [0.466] | | |

*Notes:* Each pair of rows is one regression in which the two groups of described-process adopters and the disclosure-only adopters enter as separate treated types against never-adopters. Sun–Abraham ranges cover windows 0–2 and 0–6 and references −1, −2 and −3 with −2, −1 dropped; "significant" counts designs with wild cluster restricted bootstrap $p < 0.05$ (Webb weights, 999 replications, clustering by adviser). TWFE restricts adopters to event times −8 to +8. Funds and advisers are those in the Sun–Abraham sample for window 0–6; an adviser with funds in both groups of a split is counted in each. Figure 6 plots the six Sun–Abraham estimates for each group. Funds whose managers' education is not stated form a third group in the education regression (not shown).

**Table A12:** Breadth effect by managers' technical training

Issuers held, log points × 100 [wild cluster restricted bootstrap $p$].

| Design | No STEM-trained PM | STEM-trained PM | Difference |
|---|---|---|---|
| TWFE (−8 to +8) | +11.3 [0.031] | −6.3 [0.903] | +17.6 [0.037] |
| Sun–Abraham, 0–2, ref. −1 | +8.3 [0.025] | +3.5 [0.220] | +4.8 [0.284] |
| Sun–Abraham, 0–2, ref. −2 | +10.6 [0.007] | +3.3 [0.217] | +7.3 [0.123] |
| Sun–Abraham, 0–2, ref. −3, drop −2, −1 | +11.4 [0.003] | +3.4 [0.149] | +8.1 [0.021] |
| Sun–Abraham, 0–6, ref. −1 | +11.6 [0.069] | +2.6 [0.232] | +9.0 [0.091] |
| Sun–Abraham, 0–6, ref. −2 | +12.5 [0.006] | +2.0 [0.295] | +10.6 [0.062] |
| Sun–Abraham, 0–6, ref. −3, drop −2, −1 | +12.3 [0.003] | +0.4 [0.724] | +11.9 [0.018] |

*Notes:* Teams are classified at the quarter before adoption from prospectus biographies. 47 funds at 11 advisers have no STEM-trained manager and 25 funds at 10 advisers have at least one (TWFE estimation sample). Difference: no STEM minus STEM, as in Table A11.

Table A13 reports the data-science hiring discussed in Section 5.6, with employers matched to advisers by brand name and, more strictly, by exact legal name. Before 2019 the advisers of described-process adopters filed data-science applications more often than advisers without an adopting fund under the brand match, but not under the exact match, where they filed none; under both, most adopting advisers that ever filed one did so before adoption, and hiring after adoption does not differ from that at disclosure-only adopters' advisers.

## D.6 Positions, Benchmarks, and Performance

Table A14 reports each position-level outcome in Table 4 design by design, except the analyst rows. Table A15 reports the synchronized-entry tests of Section 5.1: co-entry of new names with other adopters beyond the names' popularity, and stock-level herding against matched groups of never-adopters. Table A16 compares estimators for the benchmark outcomes in the primary sample and among U.S.-benchmarked funds. Table A17 extends Table 6 to the window of six quarters: flows rise, and alpha rises marginally ($p = 0.054$), only with quarter −1 as the reference, and no outcome changes significantly once the reference quarter is moved.

**Table A13:** Data-science hiring at adopting advisers

H-1B labor condition applications (U.S. Department of Labor) matched to advisers.

| | Brand match | Exact name |
|---|---|---|
| *A. Share of quarters with a data-science or AI application, through 2018Q4 (percent)* | | |
| Advisers of described-process adopters (35) | 5.5 | 0.0 |
| Advisers of disclosure-only adopters (41) | 4.0 | 2.0 |
| Advisers without an adopting fund (1,687) | 1.6 | 0.3 |
| Described-process minus no adopting fund | +3.9 [0.034] | −0.3 [0.000] |
| Described-process minus disclosure-only | +1.5 [0.577] | −2.0 [0.004] |
| *B. Any data-science or AI application, event study (percentage points)* | | |
| Described-process adopters, quarters −7 to −2 | +1.9 [0.402] | +1.8 [0.073] |
| Described-process adopters, quarters 0 to +4 | +5.4 [0.066] | +1.6 [0.274] |
| Disclosure-only adopters, quarters 0 to +4 | +3.0 [0.189] | +0.4 [0.716] |
| Difference, quarters 0 to +4 | +2.4 [0.517] | +1.1 [0.509] |
| *C. Timing and occupations at the advisers of described-process adopters* | | |
| Advisers with any data-science or AI application | 16 | 7 |
| of which the first precedes adoption | 14 | 5 |
| Data-science and AI applications | 1,294 | 42 |
| with a data-science or machine-learning title (percent) | 21 | 10 |
| with a quantitative-research title (percent) | 2 | 26 |
| with a portfolio-manager title | 0 | 0 |
| Applications with a portfolio-manager title, all occupations | 61 | 41 |
| of which data-science or AI | 0 | 0 |

*Notes:* An application is data-science or AI if its occupation code is data scientist, statistician, operations research analyst or computer and information research scientist, or its job title names data science, machine learning or AI. Exact name: the employer's normalized name equals an adviser's legal, business or N-CEN name. Brand match adds employers in finance (NAICS 52) whose first word is an adviser's brand word, attaching parent-level hiring to every adviser of the brand. Adopters are the primary advisers of the 132 described-process adopters (and, separately, of disclosure-only adopters), dated at their earliest adoption; advisers with a left-censored or index-fund adopter are dropped from the comparison group. Panel A: adviser means over quarters of receipt through 2018Q4; bootstrap $p$ in brackets. Panel B: Sun–Abraham, adviser and quarter fixed effects, reference eight or more quarters before adoption, so the pre-period shows any build-up; wild cluster restricted bootstrap $p$ (Webb weights, 999 replications). Panel C: occupations from job-title keywords. The applications cover visa-sponsored positions only, and an application is not a hire.

**Table A14:** Position-level outcomes by design

Described-process adopters vs never-adopters; coefficient [wild cluster restricted bootstrap $p$].

| Outcome | TWFE | SA ref. −1 | SA ref. −2 | SA drop −2, −1 | SA 0–6 range | Significant |
|---|---|---|---|---|---|---|
| New names new to the adviser | −0.0034 [0.800] | −0.0610 [0.009] | −0.0190 [0.474] | −0.0128 [0.072] | −0.0671 to −0.0056 | 1 of 6 |
| New names first ever held by the fund | +0.0161 [0.296] | −0.0081 [0.732] | +0.0081 [0.720] | +0.0125 [0.686] | −0.0003 to +0.0365 | 0 of 6 |
| New names in a new industry group | −0.0028 [0.723] | +0.0114 [0.235] | +0.0013 [0.881] | −0.0068 [0.638] | −0.0098 to +0.0158 | 0 of 6 |
| New names in a new country | −0.0203 [0.058] | −0.0188 [0.158] | −0.0277 [0.231] | −0.0067 [0.706] | −0.0289 to −0.0061 | 1 of 6 |
| New names / names held | +0.0304 [0.024] | +0.0403 [0.003] | +0.0317 [0.003] | +0.0473 [0.001] | +0.0344 to +0.0522 | 6 of 6 |
| log(1 + names added) | +0.2350 [0.033] | +0.2644 [0.014] | +0.2364 [0.002] | +0.4262 [0.004] | +0.3646 to +0.5788 | 6 of 6 |
| log(1 + names dropped) | +0.1200 [0.235] | +0.1972 [0.001] | +0.1923 [0.000] | +0.3287 [0.005] | +0.2559 to +0.3911 | 6 of 6 |
| Survival of new names to 4 quarters | +0.0208 [0.524] | +0.0967 [0.006] | −0.0211 [0.583] | +0.1177 [0.253] | −0.0172 to +0.1275 | 2 of 6 |
| Scale-up of survivors | +0.0730 [0.273] | +0.2091 [0.337] | +0.1008 [0.357] | −0.0861 [0.508] | −0.0859 to +0.3117 | 0 of 6 |
| Exits after $\leq 2$ quarters / exits | +0.0657 [0.033] | +0.0757 [0.009] | +0.0152 [0.393] | +0.0660 [0.019] | +0.0211 to +0.1012 | 4 of 6 |
| Weight in names held $\geq 8$ quarters | −0.0373 [0.319] | −0.0249 [0.052] | −0.0085 [0.319] | −0.0156 [0.177] | −0.0465 to −0.0281 | 2 of 6 |
| Size tilt of new names (log market cap) | −0.0566 [0.488] | +0.0190 [0.672] | +0.0009 [0.997] | −0.0463 [0.635] | −0.1368 to +0.0051 | 0 of 6 |

*Notes:* SA columns: Sun–Abraham, window 0–2, with reference quarter −1, −2, or −3 after dropping quarters −2 and −1. The range covers the same three references with window 0–6. "Significant" counts Sun–Abraham designs with $p < 0.05$. Shares are fractions; log outcomes are in log points. Positions are fund–issuer–quarter rows of the corporate book; a name is new in the first quarter it is held after a quarter in which it was not.

## Table A15: Co-entry and herding tests

Panel A: fund level, percentage points [wild cluster restricted bootstrap $p$]; described-process adopters vs never-adopters.

| Outcome | TWFE | Sun–Abraham range (6) | Significant | Minus disclosure-only | Minus index siblings |
|---|---|---|---|---|---|
| *Main treatment (described process)* | | | | | |
| Share of new names also added by an adopter at another adviser | −5.5 [0.019] | −6.6 to −4.4 | 2 of 6 | −5.5 [0.055] | −6.8 [0.034] |
| Expected share given the names' popularity | −3.4 [0.079] | −6.1 to −2.4 | 4 of 6 | −5.8 [0.013] | −2.2 [0.043] |
| Excess co-entry share (count) | −2.1 [0.262] | −2.0 to −0.4 | 0 of 6 | +0.2 [0.849] | −4.5 [0.201] |
| Excess co-entry share (weight) | −1.4 [0.528] | −1.4 to +0.5 | 0 of 6 | −0.1 [0.962] | −5.2 [0.075] |
| excluding index-event names | −1.2 [0.565] | −0.9 to +0.1 | 0 of 6 | +0.5 [0.630] | +0.3 [0.765] |
| new names inside the benchmark | −3.8 [0.231] | −6.8 to +0.9 | 1 of 6 | +0.1 [0.946] | −7.6 [0.488] |
| new names outside the benchmark | −4.6 [0.099] | −14.3 to −7.1 | 2 of 6 | −17.0 [0.026] | −21.3 [0.022] |
| *Broader treatment (any specific AI sentence)* | | | | | |
| Share of new names also added by an adopter at another adviser | −5.8 [0.005] | −6.3 to −3.2 | 6 of 6 | −3.2 [0.094] | −3.6 [0.013] |
| Expected share given the names' popularity | −3.8 [0.023] | −4.7 to −2.1 | 5 of 6 | −3.0 [0.162] | −0.9 [0.700] |
| Excess co-entry share (count) | −2.0 [0.032] | −1.7 to −1.1 | 0 of 6 | −0.1 [0.927] | −2.7 [0.265] |
| Excess co-entry share (weight) | −1.6 [0.094] | −1.1 to −0.1 | 0 of 6 | +0.7 [0.680] | −2.3 [0.210] |
| excluding index-event names | −2.1 [0.012] | −2.6 to −0.6 | 0 of 6 | −0.8 [0.713] | −2.3 [0.194] |
| new names inside the benchmark | −2.0 [0.037] | −4.0 to +2.0 | 2 of 6 | +6.8 [0.010] | −1.8 [0.619] |
| new names outside the benchmark | −4.2 [0.009] | −11.7 to −1.8 | 4 of 6 | −6.3 [0.431] | −16.6 [0.004] |

Panel B: stock-level herding (Lakonishok, Shleifer and Vishny 1992) against 200 matched never-adopter groups.

| Group | Quarters | $H$ | $H$, matched groups | Difference | $p$ (upper) |
|---|---|---|---|---|---|
| *Main treatment* | | | | | |
| Adopters after adoption | 21 | 0.012 | 0.034 | −0.022 | 1.000 |
| Adopters before adoption ($k = -8$ to $-1$) | 11 | 0.032 | 0.036 | −0.004 | 0.672 |
| Disclosure-only adopters after adoption | 11 | 0.041 | 0.050 | −0.009 | 0.891 |
| *Broader treatment* | | | | | |
| Adopters after adoption | 24 | 0.020 | 0.039 | −0.019 | 1.000 |
| Adopters before adoption ($k = -8$ to $-1$) | 17 | 0.027 | 0.044 | −0.017 | 1.000 |
| Disclosure-only adopters after adoption | 6 | 0.037 | 0.045 | −0.008 | 0.955 |

*Notes:* Panel A: for each name a fund adds (GVKEY-keyed issuers), whether a described-process adopter at another adviser also adds it that quarter; the expected share uses the name's popularity among all active funds outside the fund's adviser, scaled by the number of names the reference set adds per fund. Index-event names are those added by at least 1 percent of eligible index funds, or newly in a tracked benchmark. The Sun–Abraham range covers windows 0–2 and 0–6 with references −1, −2 and −3 after dropping −2 and −1; "Significant" counts designs with $p < 0.05$. Contrast columns: described-process adopters minus disclosure-only adopters (disclosure design) and minus the adviser's index funds (family design), Sun–Abraham, window 0–6, reference −1. Panel B: the Lakonishok–Shleifer–Vishny measure of net buying imbalance across advisers per stock-quarter (stocks traded by at least three advisers), averaged over quarters, against 200 groups of never-adopters at advisers without adopters, matched on Lipper class, trading activity and size; $p$ is the share of matched groups with herding at least as high (a one-sided test of more herding than matched funds). Benchmark membership from the index provider.

**Table A16:** Benchmark outcomes across estimators

Described-process adopters [p]; window 0–6 unless noted.

| Outcome | Sun–Abraham (6 designs) | TWFE | Stacked | Callaway–Sant'Anna | RI $p$ (4 schemes) | HonestDiD $\bar{M}$ |
|---|---|---|---|---|---|---|
| *Panel A: primary sample (87 adopters)* | | | | | | |
| log(1 + constituents held) | +0.088 to +0.144 (6 of 6) | +0.089 [0.040] | +0.089 [0.041] | +0.115 [0.058] | 0.002–0.046 | 0.50 |
| log(1 + names outside) | +0.030 to +0.106 (2 of 6) | +0.061 [0.304] | +0.062 [0.295] | +0.014 [0.857] | | |
| Constituents held minus names outside | +0.023 to +0.082 (0 of 6) | +0.028 [0.670] | +0.027 [0.676] | +0.101 [0.269] | | |
| Share of constituents held (pp) | +0.96 to +1.54 (6 of 6) | +1.27 [0.041] | +1.28 [0.038] | +1.31 [0.008] | 0.008–0.084 | 0.50 |
| Weight outside the benchmark (pp) | −1.94 to −0.84 (4 of 6) | −1.04 [0.111] | −1.03 [0.141] | −2.36 [0.047] | 0.022–0.202 | 0.40 |
| *Panel B: U.S. benchmarks only (53 adopters)* | | | | | | |
| log(1 + constituents held) | +0.034 to +0.070 (6 of 6) | +0.006 [0.940] | +0.006 [0.943] | +0.029 [0.683] | 0.784–0.869 | 0.10 |
| log(1 + names outside) | +0.024 to +0.083 (2 of 6) | +0.045 [0.561] | +0.046 [0.544] | +0.052 [0.601] | | |
| Constituents held minus names outside | −0.029 to +0.028 (0 of 6) | −0.038 [0.568] | −0.040 [0.556] | −0.023 [0.767] | | |
| Share of constituents held (pp) | +0.57 to +0.93 (6 of 6) | +0.58 [0.422] | +0.60 [0.383] | +0.65 [0.091] | 0.176–0.441 | 0.15 |
| Weight outside the benchmark (pp) | −1.11 to −0.37 (3 of 6) | −0.42 [0.354] | −0.42 [0.345] | −1.10 [0.073] | 0.234–0.686 | 0.35 |

*Notes:* Sun–Abraham: range over windows 0–2 and 0–6 and references −1, −2 and −3 with −2, −1 dropped (number with $p < 0.05$). Stacked: one stack per adoption cohort with never-adopters, stack-by-fund and stack-by-quarter effects. Callaway–Sant'Anna: never-adopters as controls, average of event times 0–6, adviser-cluster bootstrap (Callaway and Sant'Anna 2021). RI: randomization inference for the TWFE coefficient (placebo funds, placebo advisers, permuted fund and adviser adoption dates). HonestDiD: the largest relative-magnitude bound $\bar{M}$ at which the 95% interval for the average effect over event times 0–6 excludes zero (Rambachan and Roth 2023), with the analytical clustered covariance (the leave-one-adviser-out covariance is degenerate here because some event-time cells are identified by a single adviser), which makes the bounds optimistic.

**Table A17:** Performance and costs after adoption, window 0–6

Described-process adopters vs never-adopters, Sun–Abraham [wild cluster restricted bootstrap $p$].

| Outcome | Reference −1 | Reference −2 | Ref. −3, drop −2, −1 |
|---|---|---|---|
| Primary-model alpha, net, per quarter (%) | +1.31 [0.054] | −0.95 [0.336] | +0.30 [0.628] |
| Peer-adjusted return, per quarter (%) | +0.11 [0.656] | −0.13 [0.472] | +0.16 [0.567] |
| Net flow / lagged assets, per quarter (%) | +2.64 [0.000] | −1.28 [0.420] | −0.71 [0.524] |
| log assets under management | −0.039 [0.343] | −0.040 [0.439] | −0.044 [0.499] |
| Value added, asinh $M | +0.38 [0.304] | −0.10 [0.789] | +0.22 [0.608] |
| Expense ratio, CRSP (pp) | −0.01 [0.087] | −0.00 [0.894] | +0.00 [0.762] |

*Notes:* Companion to Table 6, which reports window 0–2. Alpha uses rolling 36-month betas estimated before each month. Flows are net of returns. Value added follows Berk and van Binsbergen (2015).

## D.7 Benchmark Measures and Additional Benchmark Results

Except for the Russell indices, we use membership only. Index weights were not available to us for most indices, and the obvious substitute, the holdings of an exchange-traded fund that tracks the index, would tie the benchmark to one market participant whose fund family includes several of our adopters. The substitute would have mattered little for membership: the iShares ETFs hold 99.2 to 100 percent of their index's constituents in the median quarter for 17 of the 19 indices, but only 89 percent for the Russell 3000 and 20 percent for the Russell 2500, whose ETF holds part of the index through other funds. No constituent history was available to us for MSCI indices: neither our LSEG and WRDS subscriptions nor MSCI's website, which publishes only the current lists, carries it. For funds benchmarked to MSCI indices we therefore take each index's membership from an index fund that tracks it and is run by a different family than the benchmarked fund's adviser (MetLife and Transamerica funds for EAFE; Xtrackers, or iShares for Xtrackers' own funds, for Emerging Markets and ACWI ex USA; ACWI as the ACWI ex USA tracker plus the S&P 500's constituents; World as the S&P Global 1200), so that no fund is measured against its own family's holdings. This combination, provider lists for U.S. indices and cross-family trackers for MSCI indices, is our primary benchmark measure; Section 5.5 also reports results for U.S.-benchmarked funds alone, whose membership comes entirely from the index provider. The trackers agree with the iShares funds on 92 to 99 percent of members, and match MSCI's current lists as closely as the replicating iShares funds do. A further check uses other providers' developed-market lists from Datastream: the STOXX Global 1800 ex Americas for EAFE and the S&P Global 1200 for World. Datastream's lists also confirm the Bloomberg membership, agreeing on 99.8 to 100 percent of S&P 500, Russell 2000 and Russell 3000 constituents.

These benchmark results hold across estimators (Appendix Table A16). For constituents

held, stacked difference-in-differences gives 8.9 log points ($p = 0.041$), Callaway–Sant'Anna 8.2 over event quarters 0 to 2 ($p = 0.012$) and 11.5 over 0 to 6 ($p = 0.058$), and randomization inference rejects in all four procedures ($p \leq 0.012$; 0.046 when adoption dates are permuted across advisers). The share of constituents held is significant in every estimator (Callaway–Sant'Anna $p = 0.008$) except that one randomization procedure ($p = 0.084$). Pre-period averages are insignificant in every design ($p \geq 0.08$), and Rambachan and Roth (2023) bounds show the effect on constituents held surviving post-adoption trend violations up to 0.5 times the largest pre-period change (0.4 to 0.75 across variants), similar to the headline estimate. Dropping one adviser at a time leaves constituents held between 11.0 and 14.3 log points in the window to six quarters ($t \geq 6.4$) and the share held between 1.1 and 1.5 points ($t \geq 3.9$). Without every BlackRock-advised fund, treated or not (69 adopters remain), constituents held rise by 9.1 to 13.7 log points and the share held by 0.8 to 1.2 points, both in all six designs, and the headline breadth estimate on the full sample holds in all eighteen Sun–Abraham designs: 8.1 to 13.4 log points against never-adopters (81 adopters) and 8.4 to 12.9 with adviser-by-quarter effects (60 adopters), all $p \leq 0.018$; relative to the adviser's index funds it is 4.6 to 10.5 ($p \leq 0.02$ in five of six designs). Giving each benchmark its own quarter effects, so that index reconstitutions and benchmark-specific trends are absorbed, leaves constituents held at 8.7 to 14.2 log points and the share held at 1.0 to 1.6 points (all six designs). Counting only equity holdings, so that a fund's bonds of an index company do not count as a constituent held, changes nothing (U.S. benchmarks: 3.1 to 6.9 log points, all six designs).

The results are weaker among the 53 adopters benchmarked to U.S. indices alone, whose membership comes entirely from the index provider. Breadth itself is smaller there (2.9 to 6.7 log points in the Sun–Abraham designs, nothing in TWFE). Constituents held rise by 3.4 to 7.0 log points and the share held by 0.6 to 0.9 points in all six Sun–Abraham designs, names outside in two, but TWFE, stacked and randomization-inference estimates are insignificant, and the bounds tolerate trend violations only up to 0.1 times the largest pre-period change. On these funds the source of membership does not matter: reading it from the tracking iShares ETF instead gives nearly the same estimate (5.4 log points for constituents held in the window to six quarters, against 5.6). The larger estimates in the full benchmark sample come from the 34 adopters benchmarked to MSCI indices, international and global equity funds, which is where the breadth effect is also largest. With another provider's list for EAFE instead of a tracker (World uses the S&P Global 1200 list in both measures; Emerging Markets and ACWI funds drop out) (69 adopters), constituents held rise by 5.2 to 9.8 log points (all six designs) and names outside in none.

For the 13 Russell indices, FTSE Russell's quarter-end constituent weights match our

issuer keys for 99.98 percent of index weight, and FTSE's member lists agree with Bloomberg's at a Jaccard index of 0.996 to 1.000 in the median quarter (0.980 to 0.997 at the minimum). Against these weights, the median Russell-benchmarked equity fund (2,180 funds) has an active share of 0.80. The iShares ETF holdings are within 0.05 to 0.7 percent of the FTSE weights for every Russell index except the Russell 2500, whose ETF holds 58 percent of the index's weight.

## D.8 Performance of Adopter Subgroups

The pooled alpha estimate of Table 6 could hide a group of adopters for whom AI pays. Table A18 reruns that specification with the described-process adopters split into cells by one attribute at a time: the fund's vehicle, asset class and size, its adviser's size and data-science hiring, whether it is subadvised, its managers' training, its strategy before adoption, and what the adoption filing says the AI does, who supplies it and what data it reads. With sixteen splits and two to three cells each, several estimates with $p < 0.05$ at reference $-1$ are expected by chance, so we ask that a cell keep its sign with $p < 0.10$ when the reference quarter is moved to $-2$ and when quarters $-2$ and $-1$ are dropped.

One group passes for alpha: the 24 adopting ETFs, at 9 advisers, whose net alpha is 4.5 percent per quarter with reference $-1$, 1.9 with reference $-2$ and 2.9 in the drop variant (3.4, 1.1 and 1.8 in the window of six quarters). The funds whose filings name text or alternative data (35 funds at 11 advisers) pass as well. Of the 34 described-process adopters that are ETFs, 25 adopted in 2025 or 2026; the largest groups are Fidelity (7 funds), State Street (5), WisdomTree (5), BlackRock (4) and Rayliant (3), and several are thematic funds. The evidence is weak. The Benjamini–Hochberg $q$ is 0.25, and the pre-adoption coefficients are already 1.4 percent per quarter ($p = 0.08$), so these funds were outperforming before their adoption filing. No other cell, including the managers' training, the role and supplier of the AI and the adoption cohort, is significant under both alternative reference quarters; a few are significant under one of them, and several change sign. For peer-adjusted returns (Table A19), more cells pass the reference checks, but every one except the ETFs and the text-or-alternative-data cell has a significant pre-adoption coefficient, which is outperformance that predates adoption rather than an effect of it. We do not read the table as evidence that AI pays for any identifiable group.

## Table A18: Alpha by adopter characteristics

Primary-model alpha, net, percent per quarter [wild cluster restricted bootstrap $p$]. Columns give the reference quarter.

| Cell | Funds / adv. | −1 | −2 | drop | −1 | −2 | drop | Pre (−1) | $q$ |
|---|---|---|---|---|---|---|---|---|---|
| | | Window 0–2 | | | Window 0–6 | | | | |
| *All described-process adopters* | | | | | | | | | |
| Pooled (Table 6) | 96 / 26 | +1.25 [0.07] | −0.71 [0.34] | +0.43 [0.55] | +1.31 [0.05] | −0.95 [0.34] | +0.30 [0.63] | +0.65 [0.21] | |
| *Vehicle* | | | | | | | | | |
| ETF† | 24 / 9 | +4.52 [0.01] | +1.87 [0.10] | +2.89 [0.02] | +3.37 [0.01] | +1.05 [0.06] | +1.83 [0.02] | +1.39 [0.08] | 0.25 |
| Mutual fund | 72 / 18 | +0.61 [0.37] | −1.36 [0.25] | −0.40 [0.70] | +0.89 [0.22] | −1.46 [0.35] | −0.29 [0.72] | +0.65 [0.29] | 0.47 |
| *Asset class* | | | | | | | | | |
| Equity | 80 / 23 | +1.46 [0.09] | −0.72 [0.39] | +0.37 [0.60] | +1.56 [0.08] | −1.01 [0.37] | +0.12 [0.86] | +0.71 [0.26] | 0.27 |
| Other | 16 / 7 | +0.08 [0.50] | −0.01 [0.92] | +0.06 [0.66] | −0.09 [0.42] | −0.17 [0.33] | −0.03 [0.81] | +0.15 [0.40] | 0.58 |
| *Fund size before adoption* | | | | | | | | | |
| Above active median | 44 / 7 | +0.82 [0.45] | −1.25 [0.02] | −0.21 [0.94] | +1.17 [0.34] | −1.44 [0.03] | +0.07 [0.93] | +0.77 [0.28] | 0.54 |
| Below | 48 / 25 | +1.88 [0.11] | +0.27 [0.57] | +1.25 [0.20] | +1.75 [0.17] | +0.14 [0.81] | +0.73 [0.35] | +0.48 [0.58] | 0.27 |
| *Adviser size (Form ADV)* | | | | | | | | | |
| Above active median | 64 / 9 | +0.58 [0.17] | −0.79 [0.03] | +0.19 [0.76] | +0.79 [0.14] | −0.90 [0.13] | +0.47 [0.57] | +0.44 [0.36] | 0.29 |
| Below | 32 / 17 | +2.55 [0.12] | −0.06 [0.94] | +1.07 [0.21] | +2.18 [0.15] | −0.64 [0.46] | +0.14 [0.85] | +1.03 [0.44] | 0.27 |
| *Management* | | | | | | | | | |
| In-house | 56 / 16 | +0.62 [0.10] | −0.21 [0.33] | +0.16 [0.80] | +1.06 [0.10] | +0.14 [0.46] | +0.50 [0.51] | +0.11 [0.84] | 0.27 |
| Subadvised | 40 / 14 | +2.14 [0.26] | −0.97 [0.07] | +0.72 [0.34] | +1.64 [0.28] | −1.87 [0.04] | −0.07 [0.92] | +1.47 [0.27] | 0.35 |
| *Adviser data-science hiring 2016–21* | | | | | | | | | |
| Yes | 49 / 9 | +0.90 [0.17] | −1.65 [0.01] | +0.85 [0.33] | +1.16 [0.02] | −1.77 [0.03] | +1.02 [0.11] | +0.70 [0.21] | 0.29 |
| No | 47 / 17 | +1.86 [0.14] | +0.65 [0.27] | −0.01 [0.99] | +1.71 [0.15] | +0.44 [0.40] | −0.29 [0.63] | +0.58 [0.44] | 0.27 |
| *Team before adoption: STEM degree* | | | | | | | | | |
| Yes | 20 / 11 | +2.08 [0.09] | +2.45 [0.05] | +1.92 [0.25] | +1.14 [0.31] | +1.81 [0.16] | +1.13 [0.28] | +0.04 [0.94] | 0.27 |
| No | 51 / 12 | −0.40 [0.58] | −0.22 [0.49] | +0.30 [0.79] | −0.36 [0.63] | +0.12 [0.70] | +0.60 [0.56] | −0.48 [0.36] | 0.64 |
| Education not stated | 25 / 11 | +3.54 [0.14] | −2.97 [0.37] | −0.13 [0.62] | +4.18 [0.12] | −3.22 [0.34] | −0.07 [0.82] | +3.94 [0.22] | 0.27 |
| *Team before adoption: STEM PhD* | | | | | | | | | |
| Yes | 4 / 2 | +9.66 [0.08] | +9.86 [0.06] | +5.10 [0.06] | +5.16 [0.09] | +6.58 [0.06] | +1.51 [0.08] | +2.06 [0.11] | 0.27 |
| No | 67 / 18 | −0.45 [0.55] | −0.15 [0.64] | +0.39 [0.67] | −0.39 [0.60] | +0.22 [0.49] | +0.59 [0.50] | −0.49 [0.39] | 0.62 |
| *Team before adoption: MBA* | | | | | | | | | |
| Yes | 41 / 11 | −0.15 [0.83] | +0.49 [0.47] | +0.23 [0.83] | −0.46 [0.64] | +0.34 [0.56] | +0.23 [0.81] | −0.46 [0.66] | 0.88 |
| No | 30 / 10 | +0.70 [0.08] | +0.99 [0.10] | +1.39 [0.56] | +0.17 [0.20] | +1.10 [0.06] | +0.76 [0.50] | −0.12 [0.81] | 0.27 |
| *Strategy before adoption* | | | | | | | | | |
| Quantitative / quantamental | 28 / 11 | +1.39 [0.04] | −0.54 [0.22] | +1.72 [0.08] | +1.89 [0.11] | −0.77 [0.13] | +1.95 [0.14] | +0.05 [0.76] | 0.27 |
| Other | 68 / 23 | +1.20 [0.15] | −0.75 [0.35] | −0.13 [0.84] | +1.15 [0.16] | −1.03 [0.28] | −0.33 [0.65] | +0.92 [0.21] | 0.27 |
| *Adoption date* | | | | | | | | | |
| 2021–22 | 12 / 8 | +1.12 [0.19] | −5.21 [0.18] | −1.58 [0.05] | +2.51 [0.02] | −4.04 [0.40] | −0.30 [0.06] | +4.27 [0.04] | 0.30 |
| 2023 or later | 84 / 20 | +1.01 [0.18] | +0.24 [0.47] | +0.68 [0.47] | +0.72 [0.34] | −0.13 [0.81] | +0.40 [0.65] | +0.39 [0.43] | 0.29 |
| *Specific process sentences at adoption* | | | | | | | | | |
| One | 38 / 11 | +0.17 [0.89] | −0.21 [0.68] | −1.69 [0.51] | +0.22 [0.88] | −0.59 [0.60] | −1.70 [0.45] | +0.57 [0.58] | 0.92 |
| Two or more | 58 / 18 | +1.88 [0.01] | −0.94 [0.42] | +1.55 [0.04] | +1.89 [0.01] | −1.24 [0.46] | +1.19 [0.05] | +0.68 [0.15] | 0.25 |
| *Process sentence in the principal strategy* | | | | | | | | | |
| Yes | 61 / 24 | +2.04 [0.04] | +0.41 [0.29] | +1.75 [0.05] | +2.21 [0.03] | +0.39 [0.46] | +1.51 [0.14] | +0.60 [0.42] | 0.27 |
| No | 35 / 4 | −0.04 [0.93] | −1.96 [0.21] | −1.56 [0.47] | +0.30 [0.77] | −2.50 [0.21] | −1.19 [0.43] | +0.71 [0.75] | 0.93 |
| *Described AI role* | | | | | | | | | |
| Primary input or selects | 23 / 13 | +2.22 [0.23] | −3.30 [0.07] | +0.35 [0.50] | +2.83 [0.18] | −2.83 [0.11] | +0.57 [0.41] | +2.96 [0.22] | 0.33 |
| One input | 73 / 14 | +0.79 [0.32] | +0.72 [0.06] | +0.51 [0.59] | +0.57 [0.43] | +0.59 [0.17] | +0.40 [0.59] | +0.08 [0.83] | 0.42 |
| *AI supplier* | | | | | | | | | |
| In-house | 76 / 21 | +1.10 [0.11] | +0.16 [0.60] | +0.35 [0.67] | +1.14 [0.12] | +0.03 [0.92] | +0.15 [0.86] | +0.24 [0.65] | 0.27 |
| Third party / not stated | 20 / 9 | +2.16 [0.07] | −3.43 [0.05] | +0.86 [0.02] | +2.44 [0.32] | −3.51 [0.11] | +1.10 [0.11] | +2.62 [0.04] | 0.27 |
| *Data named* | | | | | | | | | |
| Text or alternative data† | 35 / 11 | +1.50 [0.03] | +2.34 [0.01] | +2.19 [0.06] | +1.05 [0.08] | +2.03 [0.02] | +1.94 [0.15] | −0.45 [0.52] | 0.27 |
| None named | 61 / 17 | +1.09 [0.25] | −2.39 [0.18] | −0.65 [0.70] | +1.49 [0.17] | −2.44 [0.26] | −0.56 [0.66] | +1.20 [0.17] | 0.35 |

*Notes:* The specification of Table 6 (Sun–Abraham, all active fund-quarters, disclosure-only adopters as a separate treated group, never-adopters as controls; N-PORT panel from 2019Q4, event time from −4, no corporate-book screen, so cell sizes differ from the breadth splits of Table A11 and need not be subsets of them) with described-process adopters split into cells by one attribute at a time; attributes defined for every fund (vehicle, asset class, sizes, management, data-science hiring) also get quarter-by-attribute fixed effects. Team attributes refer to the quarter before adoption; the AI attributes to the adoption filing. "drop": reference −3 with quarters −2 and −1 dropped. Pre: mean of the coefficients for quarters −4 to −2, window 0–2, reference −1. $q$: Benjamini–Hochberg over the cells of the outcome in that design. †: $p < 0.05$ with reference −1 and the same sign with $p < 0.10$ with reference −2 and in the drop variant, in either window. Gross alpha gives the same cells.

**Table A19:** Peer-adjusted returns by adopter characteristics

Peer-adjusted return (Lipper class), percent per quarter [wild cluster restricted bootstrap $p$].

| Cell | Funds / adv. | −1 | −2 | drop | −1 | −2 | drop | Pre (−1) | $q$ |
|---|---|---|---|---|---|---|---|---|---|
| | | Window 0–2 | | | Window 0–6 | | | | |
| *Vehicle* | | | | | | | | | |
| ETF† | 33 / 12 | +3.70 [0.04] | +2.41 [0.00] | +1.54 [0.00] | +2.43 [0.05] | +1.84 [0.02] | +1.43 [0.02] | +2.29 [0.10] | 0.15 |
| *Adviser size (Form ADV)* | | | | | | | | | |
| Below† | 35 / 19 | +1.31 [0.00] | +0.91 [0.01] | +1.12 [0.01] | +0.55 [0.21] | +0.25 [0.29] | +0.47 [0.16] | +0.65 [0.03] | 0.04 |
| *Management* | | | | | | | | | |
| Subadvised† | 45 / 17 | +1.83 [0.01] | +0.68 [0.02] | +0.85 [0.08] | +0.52 [0.09] | −0.10 [0.62] | +0.51 [0.18] | +1.30 [0.03] | 0.10 |
| *Team before adoption: STEM degree* | | | | | | | | | |
| Yes† | 24 / 12 | +3.68 [0.04] | +2.54 [0.02] | +1.18 [0.04] | +1.55 [0.05] | +1.56 [0.08] | +0.81 [0.18] | +2.56 [0.01] | 0.15 |
| *Team before adoption: MBA* | | | | | | | | | |
| Yes† | 46 / 12 | +1.30 [0.02] | +0.86 [0.01] | +0.42 [0.08] | −0.03 [0.88] | −0.00 [0.97] | −0.10 [0.58] | +0.67 [0.05] | 0.10 |
| *Data named* | | | | | | | | | |
| Text or alternative data† | 41 / 13 | +2.00 [0.02] | +1.60 [0.00] | +1.16 [0.01] | +0.80 [0.03] | +0.95 [0.01] | +0.73 [0.02] | +0.66 [0.29] | 0.10 |

*Notes:* As Table A18, for the return in excess of the prior-assets-weighted return of the fund's Lipper-class peers; only cells marked † are shown. Cell sizes differ from Table A18 because peer-adjusted returns need no factor-model history.

# E Responses to Potential Concerns

The design rests on comparisons a skeptical reader can question. This section takes the main critiques in turn, states the evidence that answers each, and says where the answer is incomplete.

*Selection into Adoption.* Funds that describe AI in their process differ from other funds: they are more often quantitative (42 percent against 14 percent of never-adopters) and concentrated at a few dozen advisers. If the same plans that led a fund to adopt also led it to broaden, the estimates would confound the two. We answer with comparison groups chosen to share the plausible confounders: disclosure-only adopters, which wrote AI language into their prospectuses in the same years and under the same pressures; index funds of the same adviser, which share its research infrastructure, data vendors and filing calendar; adviser-by-quarter and quarter-by-style fixed effects; and the removal of never-adopters whose strategy text describes machine-learning-like methods without AI vocabulary, which changes nothing. These comparisons rule out confounders common to an adviser, a style or a period. They cannot rule out a fund-specific plan that coincides with adoption, which is why the pre-period evidence matters.

*Pre-Adoption Trends.* Issuer counts dip two to four quarters before adoption, and a joint test of the pre-period coefficients rejects ($p = 0.002$). A reader could see the post-adoption rise as recovery from the dip, or as the continuation of a trend. Three facts weigh against those readings. The estimates are as large in the designs that never use the two quarters before adoption (10.3 to 13.7 log points) as in those that do, so they are not driven by the

trough. The rise continues well past the level from which the dip began: adopters' issuer counts are 15 to 27 log points above the reference quarter after four or more quarters, against a dip of at most five. And the Rambachan and Roth (2023) bounds show that the effect survives a linear continuation of the pre-trend. What we cannot exclude is a repositioning that begins when the prospectus amendment is drafted and continues after it is filed.

*Cheap Talk in Prospectus Language.* If funds describe AI to market themselves, the treatment would measure marketing rather than practice. The measurement is built to separate the two: the described-process standard requires a passage that assigns AI a role in selection, and the generic AI language that spread across large fund families in 2023–26 has no portfolio consequence. A related worry is that language responds to enforcement rather than practice: on 18 March 2024 the SEC settled its first cases against advisers for overstating their use of AI (U.S. Securities and Exchange Commission 2024). Funds that already described AI did use more hedged language afterwards (the share of their process passages that are risk or limitation statements rose by 8.9 percentage points relative to their previous annual filing, $p = 0.019$), but in levels a break at 2023Q3, the quarter of the SEC's predictive-analytics proposal, is as large as one at 2024Q1 (5.8 against 3.1 percentage points, neither significant), so the drift predates the enforcement, and funds did not retract their claims: the rate at which a fund drops all process language from its next annual filing shows no break.

*Coincident Changes.* A new portfolio manager, a rewritten strategy or a fund launch could coincide with the prospectus amendment that introduces AI. Adoption is dated only on amendments of funds already in operation, so launches never count, and funds whose first filing already describes AI are excluded rather than assigned an adoption date. Manager arrivals do spike just before adoption, but the breadth effect is the same for adopters with and without an arrival in the adoption quarter or the six quarters before it (32 adopters with such an arrival against 64 without; difference $-2.1$ log points, $p = 0.71$, although each half alone is imprecise).

*Counting Artifacts.* Issuer counts depend on how holdings are matched to issuers and on which holdings source is used. The CRSP and N-PORT counts agree closely where both exist, the step where funds switch sources is absorbed by quarter fixed effects, and identifiers are validated before any match. The effect is also broad rather than the product of a few funds whose counts jumped: the median adopter broadens by 8.7 log points, 64 percent of adopters broaden, and trimming the most extreme adopters on either side leaves the estimate between 6.8 and 8.8 log points (Section 5.3 and Appendix D.3). Nor do bond holdings inflate the count: 89 of the 102 adopters hold essentially no corporate debt, and the effect among them is as large as in the full sample (Section 5.2). Analyst coverage and ratings are measured on

equity holdings only, so a bond position does not inherit the issuer's equity coverage; the restriction changes no conclusion.

*Concentration among Advisers.* The 132 described-process adopters sit at 35 advisers. Inference is clustered by adviser throughout; for the primary outcome the $p$-values come from an adviser pairs bootstrap that resamples the treated advisers (Appendix D.4), and wild-bootstrap $p$-values for secondary outcomes are likely optimistic; dropping any one treated adviser leaves the estimate between 6.6 and 10.1 log points; and randomization inference that reassigns adoption to placebo advisers rejects at $p = 0.006$. The heterogeneity results rest on about ten advisers per group, and we present them as hypotheses.

*Survivorship.* If adopters that failed were closed, the adopters that remain could look broader. Adopting funds are not less likely to close: in a quarterly hazard model with quarter and Lipper-class fixed effects and controls for size, age, flows and past returns, the post-adoption exit probability rises by 0.79 percentage points for described-process adopters ($p = 0.15$) and by 0.74 for disclosure-only adopters ($p = 0.013$), so the two do not differ. The exits are mostly liquidations of thematic ETFs, and equal exit rates mean that differential attrition cannot explain the gap in breadth between the two adopter types.

*Statistical Power.* The performance nulls could reflect low power. We agree, and do not claim that AI does not pay. The post-adoption windows are short and the treated funds few. One sharper test asks whether AI relaxes decreasing returns to scale (Pástor et al. 2015): if AI lowers the cost of covering more names, adopters' alphas should fall less as they grow. Estimated with recursive demeaning (as in Arellano and Bover 1995), the change in the size–alpha slope after adoption is small and insignificant, but the smallest detectable change is about 60 basis points of quarterly alpha per log unit of size, so the performance results rule out large effects only. One recent paper finds that equity funds whose trades align most with machine-learning signals outperform (Hu, Rohrer and Zhang 2025); its measure classifies funds by what they trade, ours by the process they describe, so the two need not capture the same funds.

*External Validity.* Most treated funds adopted in 2023–26 and have short post-adoption windows, and the early adopters of 2021–22, whose effects are largest, described machine-learning processes before generative AI was widely available. Whether language-model tools change the distance of search as well as its breadth remains open.