Internet Appendix:

The Rise of Alternatives

Juliane Begenau* Pauline Liang† Emil Siriwardane‡

Abstract

In this Internet Appendix, we provide: (i) detailed descriptions of all datasets used, along with additional facts about the U.S. pension sector; (ii) robustness results using imputed missing target shares based on the method of Calvet et al. (2009); (iii) an extension of the portfolio choice model incorporating Bayesian learning and consultant effects; (iv) further evidence related to belief-based and risk-seeking explanations; (v) the complete set of results from our simulation of a mean-variance model for the aggregate pension portfolio; (vi) a discussion of alternative explanations; and (vii) full derivations of the mean-variance model presented in the main text.

*Begenau: Stanford GSB and NBER and CEPR. E-mail: begenau@stanford.edu
†Liang: Stanford GSB. E-mail: pliang20@stanford.edu
‡Siriwardane: Harvard Business School and NBER. E-mail: esiriwardane@hbs.edu

A Data

A.1 PPD

A.1.1 Sample Description and Variable Definitions

In this subsection, we discuss how we define variables in the PPD and filter the data. Asset class definitions are not standardized across CAFRs. For example, some pensions classify publicly traded real estate investment trusts (REITs) as public equities and others consider them real estate. We adjust PPD allocations manually based on information in CAFRs in order to harmonize asset class definitions across time and pension systems. As a general rule, we consolidate private credit and private equity into a single asset class and consider natural resources and infrastructure as investments into commodities. Real estate includes core real estate and private-equity real estate, but does not include REITs as they are included in public equities. Real estate and commodities are then further aggregated into an asset class that we label as real assets. Alternatives include private equity and credit, real assets, hedge funds, and other alternatives. Risky investments are then defined as everything outside of cash and fixed income. The alternative-to-risky share is defined as the portfolio share of alternatives scaled by the portfolio share of risky investments. The bulk of the adjustments we make to PPD result in reallocations within alternatives (e.g., from miscellaneous alternatives to hedge funds) and thus have a negligible impact on the overall alternative and alternative-to-risky shares.

We further process the raw PPD data as follows. First, we exclude plans that are missing information on either the market value of assets or the value of liabilities under GASB 25 standards. Second, we screen observations for plan p and fiscal year t based on the sum of actual portfolio weights (A_{pt}) and the sum of target portfolio weights (T_{pt}). Specifically, we drop observations if both $|A_{pt} - 1|$ and $|T_{pt} - 1|$ are greater than 0.05. If $|A_{pt} - 1| > 0.05$ but $|T_{pt} - 1| \leq 0.05$, we replace actual weights with target weights, and vice versa. Third, we retain data between 2001 and 2021, though in our cross-sectional analyses we often start in 2002 to increase the size of the cross-section. After aggregating to the pension system level, the resulting panel has 3,128 system-year observations. Annual dates are based on pension plan fiscal years, not calendar years, and the fiscal year for most plans begins in July and ends in June.

The code used to clean the PPD data is available in a public GitHub repository.[1] Relative to the repository's output, we make two additional adjustments. During our sample, Idaho PERSI reports broad targets to equities that include both public and private equity. The PPD records the total value as the public equity target and sets private equity to zero. To maintain asset-level granularity, we replace target shares with actuals for this plan. We apply the same adjustment for the NY State and Local Retirement System in FY 2001–2003, since PPD reports target shares that are not disclosed in the annual reports listed on the PPD website. Our conclusions are not sensitive to these choices.

In general, missing target shares occur infrequently in the PPD data—they are available for 89% of pension systems by count and 94% by assets. The decision to substitute missing target shares with actual shares follows from the fact that most institutional investors have explicit rebalancing policies that are triggered when actual shares drift too far from target. For public pensions, this rebalancing policy is typically de-

[1] The repository can be found at: https://github.com/esiriwardane/ppd-cleaning-public. We use the PPD 2023-07-28 release for this paper.

scribed in what is known as an investment policy statement (NASRA, 2020). For example, CALPERs—the largest U.S. pension by assets—currently allows for a range of 7 pp around its target allocation of 37 pp to public equities.[2] Rebalancing policies of this nature are found across pensions of all sizes. Consider a median-sized pension in our sample, such as the Oklahoma Public Employees Retirement System (PERS). In their latest CAFR, Oklahoma PERS writes:

> Asset allocation is the key determinant of return. Therefore, the commitments to asset allocation targets are mantained through a disciplined rebalancing program.

Accordingly, in their most recent annual report, Oklahoma PERS further reports a target allocation to U.S. equities of 40 pp, with actual allocations during the year never falling below 41.2 pp or exceeding 43.0 pp.[3] Similarly, the Charlotte Firefighters Retirement System—one of the smallest pensions in our sample—states in their investment policy statement that domestic equity exposure can fluctuate 5 pp around its target of 38 pp.[4] In addition to explicit rebalancing rules, public pensions have strong incentives to keep their actual portfolios close to target, as the target portfolio determines the weights used in constructing performance benchmarks.

A.1.2 Basic Summary Statistics

Here, we further discuss the basic properties of the PPD data, as summarized in Table 1. Owing to the fact that the PPD adds pension plans as their CAFRs become available, the number of pension systems in our sample steadily increases through time, starting with 165 in the early part of the sample and ending with 208 in the late part. In more recent years, these pension systems cover 27 million members, 37% of whom are retired. The share of retired members has also increased steadily over time, which reflects broader aging trends in the population. From 2016 to 2020, the pensions in our sample managed just under $4 trillion of total assets on average. The value of these assets was 73% of total pension liabilities, where liabilities are measured according to GASB 25 accounting standards. Many observers have argued that this level of funding is overstated (Novy-Marx and Rauh, 2009; Brown and Wilcox, 2009). The reason is that under GASB 25, future pension promises are discounted using each plan's assumed long-run return on assets. The national average assumed asset return – equivalently, the liability discount rate – over the last five years of the sample was 7.2%. This liability discount rate is likely too high because pension liabilities are essentially a form of government debt that have strong creditor rights, making them safer than the assets in which pensions invest. Interestingly, the average annual realized return on U.S. pension assets has exceeded the assumed return for only one of the five-year subsamples reported in Table 1.

A.1.3 Decomposing Fixed Income

The PPD provides a dataset called "PensionCreditRating", which contains dollar holdings by credit rating for a number of ppd_id-year pairs. The "PensionCreditRating" only begins in 2004 and ends in 2018 and

[2] See https://www.calpers.ca.gov/docs/total-fund-investment-policy.pdf.
[3] See https://www.opers.ok.gov/wp-content/uploads/2024/11/2024_ACFR_URSJJ-Investment.pdf.
[4] https://charlottefireretire.com/wp-content/uploads/2024/04/Investment-Policy-April-2024-final-approved-by-Board.pdf

is not well-populated for a number of pensions in our sample. It is for this reason that we do not use it to further decompose the fixed income share for individual pension plans.

Nevertheless, the dataset can be used to get a sense of the aggregate share of fixed income in high-yield debt. To this end, we clean the data as follows. The column names in this dataset contain three pieces of information: (i) credit rating; (ii) whether the amount listed is in par or market values; and (iii) the rating agency. For instance, the column "Aa_MktVal_M" contains the market value of holdings for debt securities rated Aa by Moody's. PPD transcribes these data from pension annual reports and so the column names are exactly as they appear in the annual report. This means there are many columns in the dataset and we assign each to an investment-grade or high-yield bucket. For example, one pension may list all of its holdings under "Aa_MktVal_M" whereas another may further separate government debt holdings under "Aa_GovBond_MktVal_M". In both cases, we designate the holdings as being investment-grade. We compute the share of investment-grade holdings (IG-share) according to Moody's if the total amount of bonds rated by Moody's exceeds the amount rated by S&P and vice versa. We then aggregate investment-grade and high-yield holdings to the aggregate level by year, including only years where there are at least twenty plans (2005-2018).

Figure A1 plots the share of high-yield debt relative to fixed income (including cash) and relative to the whole portfolio. To compute the latter, we scale the share of high-yield debt relative to all holdings in the "PensionCreditRating" data by the fixed-income share from the main PPD dataset. This is because total holdings in "PensionCreditRating" do not line up with total fixed income holdings in the main PPD data, which is part of why we don't use this dataset for our cross-sectional analysis. The plot shows that the share of high-yield debt was relatively stable from 2005 to 2018. The average share of high-yield relative to fixed income is around 30% and its portfolio share is around 8%.

A.2 National Pension Data

As mentioned in the main text, we supplement plan-level information from the PPD with two datasets that are available from the U.S. Census Bureau. The first is the Annual Survey of Public Pensions (ASPP). The ASPP is designed to provide information on revenues, expenditures, membership information, and financial assets for all public employee pension plans administered by state and local governments across the United States. Survey years in the ASPP begin in July and end in June, which accords with the fiscal years for most U.S. plans. Survey respondents are asked to provide responses via a standardized form, though some instead provide their CAFR. In these cases, Census Bureau staff fill out the survey based on the information in the CAFRs. The Census Bureau aggregates the data to the state and national levels, and we obtained the aggregated data directly from the Bureau's website. These data are available from 1993 onward.

The second and related dataset from the Census Bureau is the Quarterly Survey of Public Pensions (QSPP). It provides much of the same information as the ASPP, but is administered quarterly to only the hundred largest pensions in the country. These pensions comprise roughly 90% of total U.S. public pension assets. Nationally aggregated data based on the QSPP is available from the Census Bureau from 1968 onward. The definition of asset classes in the QSPP materially changed in 2019, so to maintain consistency in the time-series we only use asset allocation data through 2018.

A comparison to the ASPP reveals that the PPD is representative of the overall U.S. public pension system. As one way to measure representativeness, we compare total pension assets in each year of the PPD sample to the total amount of cash and investments listed in ASPP. For each of the five-year spans listed in Table 1, the PPD covers no less than 85% of pension assets listed in the ASPP. From 2006 onwards, it covers over 90% of total U.S. public pension assets. Moreover, according to the ASPP, U.S. public pensions had 33.7 million members in 2020, whereas plans in the PPD had 27.4 million members (81% of ASPP). We show below that asset allocation data in the PPD also lines up well with the ASPP and QSPP at the national level.

In addition to validating the PPD, the ASPP and QSPP are useful because they begin much earlier than the PPD and decompose asset allocations into U.S. sovereign debt, federally-sponsored agency debt, and corporate bonds. This breakdown lets us study the composition of fixed income exposures in more detail at the national level. Though the PPD does not provide such a breakdown, it does allow for a coarse decomposition of fixed income allocations into those with low or high credit ratings. The PPD also breaks out allocations to alternative asset classes like private equity, whereas the ASPP and QSPP did not separately report the major alternative asset classes prior to the 2019 survey wave.

PPD coverage is also fairly large in comparison to the broader U.S. pension system, defined as the total amount of assets held by all private- and public-sector pension funds. This definition includes both defined benefit and defined contribution pensions, but excludes the value of social security benefits. Information on total pension assets is taken from Table L.117 of the Financial Accounts of the United States that is published by the Federal Reserve. As Table 1 shows, the pensions in the PPD sample comprise roughly one-quarter of all U.S. pension assets, though this number has declined in recent years with the growth of defined contribution pensions.

A.3 Capital Market Assumptions

Consultants' reported beliefs come from Capital Market Assumptions (CMAs) that all major consultancies produce each year. We obtain reports from 14 consultants under a non-disclosure agreement. For most of these consultants, we compile a history of their beliefs dating back to the early 2000s. While we cannot disclose specific consultant names, together they oversee about 78% of pension assets in the U.S..

As mentioned in Section 3.1.2, CMA reports contain consultants' forecasts on expected returns, volatilities, and correlation matrices of different asset classes. We collect information on eight asset classes: cash, fixed income, public equity, private equity, real estate, hedge funds, private debt, and infrastructure. The latter five categories fall under the umbrella of alternative assets. Consultants often break down these categories into subcategories. We use U.S. core bonds as a benchmark for fixed income and U.S. stock as a benchmark for public equity. If U.S. stock data are not available, we use large-cap domestic equity instead. For real estate, we look at private-label real estate, excluding publicly traded securities such as REITs. In the reports, hedge funds are occasionally referred to as absolute return. When multiple subcategories exist for hedge funds, we aim to identify the most inclusive category.

Until recently, most consultants did not provide forecasts for private debt and infrastructure. Recall that in the PPD holdings data, these asset classes are included in our definitions of private equity and real assets,

respectively (see Section 2.1). Because of incomplete historical CMA data availability, we exclude private debt and infrastructure from our analysis in Section 3.1.2. Practically speaking, this means that we proxy for beliefs about all real assets using only real estate. Similarly, we proxy for beliefs about private equity (including private credit) using only private equity.

Finally, we perform two adjustments in the data to obtain a consistent panel. First, in earlier years, expected returns are presented as arithmetic returns. We convert them to geometric returns using the formula $\mu_{Geometric} = \mu_{Arithmetric} - 0.5\sigma^2$. Second, while expected returns and volatilities are well populated, CMA reports may lack correlation matrices in some years. In such instances, we rely on the correlations provided in the consultant's report from the preceding years. Our underlying assumption is that correlations are relatively stable over time and, therefore, do not undergo abrupt changes between years.

Using the information from CMAs, we derive the implied alpha and beta for each alternative asset class. We first compute the implied beta as $\beta = corr \cdot \frac{\sigma_A}{\sigma_E}$, where σ_A is the volatility of the alternative asset, σ_E is the volatility of public equity, and *corr* is the correlation between the two assets. We compute implied alpha as $\alpha = \mu_A^{excess} - \beta\mu_E^{excess}$, where μ_A^{excess} and μ_E^{excess} are the expected excess returns of the alternative asset and public equity, respectively.

A.4 PENDAT Data

In our analysis of experience effects in Section 3.2, we use data on 1990s portfolio composition from the PENDAT surveys that were administered by the Public Pension Coordinating Council (PPCC). According to the National Association of State Retirement Administrators, the PPCC is:

> ... a coalition of three national associations that represent public retirement systems and administrators: NASRA, the National Council on Teacher Retirement (NCTR) and the National Conference on Public Employee Retirement Systems (NCPERS). Together, these associations represent more than 500 of the largest pension plans in the United States, serving most of the nation's 18+ million state and local government employees.

The PENDAT surveys were conducted in several waves during the 1990s. These surveys asked for a variety of information about each pension system, including cash flow and balance sheet data, portfolio construction, investment restrictions, governance, and more (Zorn, 1997; Mitchell et al., 2001). Each survey wave consists of several files, three of which are used repeatedly in our analysis.

The first is the "system" file, which contains identifying information for each pension (e.g., its name, location, etc.) and the fiscal year to which the survey response corresponds. The second is the "investment" file, which contains among other things the target and actual allocations to various asset classes, as well as the value of each pension's assets (see Figure A2 for an example of a survey question). The third is the "historical" file, which contains actual asset allocation (but not target), total assets, membership information, and rates of return from the three previous surveys. Generally, the system and investment files for a given survey contain data from that fiscal year, whereas the historical file contains data from previous fiscal years. For instance, the system and investment files from the 1994 survey correspond to FY1994, while the

historical file includes data for fiscal years 1990, 1991, and 1992.[5]

Figure A3 gives a sense of coverage in the PENDAT data. Blue bars in the plot show the aggregate amount of non-missing assets in PENDAT for a given survey year and red bars plot the aggregate amount of assets in the ASPP. The figure shows that PENDAT covers virtually all U.S. pension assets during the 1990s.

In Section 3.2, the PENDAT data are used to construct changes in the target share of public equities between 1990-2002 or 1996-2002. In both cases, the 2002 target share is taken from the PPD. The 1996 target share comes from one of three sources: (i) the investment file of the 1996 PENDAT survey; (ii) the average of the target shares from the investment files of the 1994 and 1998 surveys; or (iii) direct FOIA requests with pensions. Through these efforts, target shares for 1996 are available for 72% of pensions by count and 89% by assets. Any remaining missing target shares are replaced with actual shares from the 1996 PENDAT survey. The fact that target shares are more likely to be missing for smaller pensions is one reason why we run Table 5 at the state-level, as this reduces the influence of smaller pensions with missing data.

Similarly, the 1990 target share of public equities is obtained from one of two sources: (i) the investment file of the 1990 PENDAT survey; and (ii) direct FOIA requests with pensions. These collection efforts yielded modest success, providing target shares in 1990 for 32% of pensions by count and 57% by assets. Any remaining missing target shares are replaced with actual shares in 1990, which are sourced from the historical file of the 1994 survey and the investment file of the 1990 survey.

A.4.1 Additional Variables

The PENDAT survey also asks two questions that we use in our analysis. The first is:

> Is your asset allocation: (a) Long-term? (i.e., not changed often with varying economic conditions) or (b) tactically set? (i.e., changed often with varying economic conditions).

This question is clearly related to whether pensions behave as Merton (1973) ICAPM investors or not. The second question we use is:

> Are investment restrictions specified in your state's constitution?

This question is helpful for identifying the states in which changes in the alternative-to-risky share or risky share may be influenced by legal restrictions. When aggregating to the state level, we compute the fraction of each state's assets that are affected by constitutional investment restrictions. For both the first and second questions, we use the 1998 PENDAT survey data when available, as this is the last wave containing answers to either. If data from the 1998 survey is unavailable, we use the 1996 PENDAT survey as an alternative.

A.5 S&P Money Market Directory

In Section 3.3, we compute the average actual (not target) alternative-to-risky and risky share of private-sector institutions in each pension's state. The asset allocation of private-sector institutions is further used

[5] According to the PENDAT user guides, values of -9 indicate missing data and -1 indicate not applicable. We set both instances to missing.

below in Section D.4.3. The data for both analyses calculation come from S&P's Money Market Directory (MMD), which contains information such as sponsor description and type, asset allocation, consultant and other service provider relationships, management rosters, and key contacts.[6] Asset allocation data are specifically contained in "MMASSET" tables that are indexed by year. For each plan, these data provide dollar and percent allocations at a granular level (e.g., large-cap equities), which are then manually assigned to one of three aggregate categories: (i) fixed income; (ii) equities; and (iii) alternatives. As one check of data quality, we compare the sum of holdings for each plan-year in the "MMASSET" table with its total assets in the "MMPLANS" table, the latter of which contains basic information on each plan (e.g., defined benefit vs contribution). Total assets in MMASSET must be within 5% of MMPLANS to be included. Only plans classified as "Defined Benefit Plan", "Endowment Fund", and "Closed or Frozen Defined Benefit Plan" are included. Moreover, these plans must have sponsors that are domiciled in the United States and are either listed as either "Corporation", "Endowment", or "Union".

After computing portfolio shares for each plan-sponsor-year tuple, we aggregate the data in one of two ways. The first is to the state-year level (Section 3.3), for which we require more than five private-sector pensions per cell. The second is to the consultant-year level (see Section D.4.3). For this second purpose, we use the "MMPROV" table to match plans to their consultants. This is not a one-for-one match because investors may employ multiple consultants and the data does not allow us to distinguish between general and specialty consultants. Consultants are identified in MMD by a unique numerical identifier (prov_no) and we obtain consultant names associated with each identifier from the MM85COMP table. Consultant names in MMD are then manually mapped to names in the PPD data for merging purposes. For the purposes of the analysis in Section D.4.3, after matching plans to consultants, we collapse portfolio shares to the consultant-year level by taking simple averages. A consultant must have more than five private-sector clients in a given year to be included in the analysis.

A.6 U.K. Corporate DB Pensions

Data for U.K. corporate DB pensions comes directly from the U.K. Pension Protection Fund's (PPF) Purple Book. The PPF provides a backstop to U.K. private-sector pensions in the event of insolvency, similar to the Pension Benefit Guaranty Corporation in the United States. The PPF was founded in 2004 as a part of the Pensions Act of 2004 and is accountable to the U.K. Parliament. Starting in 2006, the PPF published the Pensions Universe Risk Profile (the "Purple Book"), which contains comprehensive data on scheme funding, size, demographics, and asset allocation. Historical asset allocation for UK corporate pensions is taken directly from the 2022 Purple Book. As of March 2022, the Purple Book covers 5,131 pension schemes and over £1.7 trillion of assets.

Figure 7.3 of the 2022 Purple Book contains actual (not target) allocations broken into equities, bonds, and other investments. Other investments includes property, cash and deposits, insurance policies, hedge funds, annuities, and miscellaneous. Figure 7.5 further breaks equities into one of three categories: (i) U.K. quoted; (ii) overseas quoted; and (iii) unquoted/private. We define alternatives as hedge funds, property,

[6] https://www.spglobal.com/marketintelligence/en/solutions/money-market-directories

miscellaneous, and unquoted/private equities.[7] Fixed income is defined as bonds, cash and deposits, insurance policies, and annuities, the bulk of which is in bonds and cash. The risky share is defined as everything outside of fixed income.

A.7 Endowments

Asset allocation for endowments is taken directly from the National Association of College and University Business Officers (NACUBO). Each year, TIAA and NACUBO conduct and publish a survey (the NTSET) of college and university endowments that contains information like size, returns, and asset allocation. We use the public NTSE tables that are available directly on the NACUBO website.[8] These data start in 2002 and provide aggregate asset allocation data for over 600 institutions, who in 2022 accounted for $807 billion of assets. We manually collect the dollar-weighted average asset allocation for each year and then aggregate it to asset classes that align with the PPD data. Fixed income is defined as NACUBO's fixed income category plus cash. Public equities exactly match the "Equity" category from NACUBO up to 2008 and then are the sum of domestic and international equities thereafter. Alternatives are the residual. The risky share is defined as public equities plus alternatives.

As mentioned in the main text, only actual portfolio shares are available for private-sector institutions, who generally do not publicly release their policy portfolios. However, like public pensions, most private-sector institutions follow rebalancing rules that prevent actual shares from drifting too far from target. For example, according to the 2021 NACUBO annual study of endowments (see Chapter 8, NACUBO (2021)), 92% of endowments had a formal rebalancing policy. Additionally, 87% rebalanced using market value-based rules (i.e., targets and range-based) or a combination of market value-based rules and calendar rules (e.g., quarterly). Among the 5% of endowments that relied solely on calendar-based rules, virtually all rebalanced at least annually, with nearly 65% rebalancing quarterly or more frequently.

B Additional Facts

B.1 The Risky Share

B.1.1 Aggregate Risky Share

Figure 1b shows that the QSPP risky share – defined as non-fixed-income investments – increased dramatically from 1970 to 2000. We now decompose these changes at a more granular level using data from both the ASPP and QSPP. Figure A4a uses the ASPP to plot the share of U.S. Treasuries, U.S. Agency debt, cash, and corporate debt. The plot confirms a decline in exposure to safe assets over the last thirty years. Cash holdings fell from 7.3% in 1993 to 3.9% in 2016, whereas the share of Agency debt has remained relatively stable throughout. Furthermore, the figure shows that the rapid decline in fixed income during the 1990s was driven primarily by a shift out of Treasuries. The share of Treasuries continued to decline after 2000, as did exposure to corporate bonds, albeit at a much slower pace.

[7] The miscellaneous share is a small part of the portfolio. It never exceeds 6% and has been around 1% since 2016.

[8] See https://www.nacubo.org/Research/2022/Public-NTSE-Tables.

Figure A4b uses the QSPP to extend the preceding decomposition of fixed income over a longer window. Unlike the ASPP, the QSPP groups U.S. Treasury and Agency debt into a single category called U.S. Government Sponsored Debt (henceforth, U.S. sponsored debt). The figure shows that the large allocation to fixed income in the early 1970s was primarily through corporate bonds and not U.S. sponsored debt – 56.4% of the total portfolio in 1970 was in corporate bonds compared to 8.6% in U.S. sponsored debt. U.S. public pensions cut their corporate bond allocation by 15.4 pp during the 1970s, which was only partially offset by an increase of 9.1 pp in U.S. sponsored debt. This trend continued into the 1980s, as public pensions further reduced their corporate bond exposure by 20.2 pp and increased their U.S. sponsored debt allocation by 10.5 pp. Because we have no indication of the credit quality of corporate bond exposures prior to 2000, it is difficult to know whether the shift out of corporate bonds and into Treasuries during the 1970s and 1980s led to an increase or decrease in the credit risk of the aggregate portfolio. With that said, we strongly suspect that U.S. pensions did not hold large amounts of high-yield corporate bonds during this period, in which case it is likely that pension risk-taking increased as the overall fixed income share declined.[9] As discussed above, the share of corporate bonds and U.S. government-sponsored debt both declined starting in the 1990s, though the reduction in the share of government-sponsored debt was much sharper compared to corporate bonds. The QSPP data in Figure A4b also further confirms that the decline in fixed income since the 2000s was driven by shift out of both corporate bonds and government debt.

B.1.2 The Risky Share in the Cross-Section

Figure A5a visualizes the cross-sectional distribution of the risky share using PPD data from 2002 onward, showing only even years for readability. Consistent with the aggregate trends in Figure 1a, the risky share has increased for the median pension over time. However, the plot also highlights a fair amount of heterogeneity. For instance, in 2021 the risky share for the 10th percentile pension system was 67% and was 85% for the 90th percentile. This degree of heterogeneity has declined slightly over time, as the spread between the 10th and 90th percentile pension was 21 pp in 2002.

It is unclear from Figure A5a whether U.S. pensions have uniformly increased their risky share over the last two decades or whether some have done so more than others. Figure A5b sheds light on this question by plotting the distribution of changes in the risky share between 2002 and 2021. We start our analysis in 2002 because the number of available systems in the PPD increases sharply in this year. The plot reveals some extreme outliers in changes in the risky share, the most prominent of which is the Texas Municipal Retirement System (TMRS). Prior to 2009, the pension invested entirely in fixed income because of how investment gains and losses were credited to member and city accounts. The passage of House Bill 360 during the 81st Session of the Texas Legislature paved the way for the fund to diversify, leading to a large increase in its risky share over our sample. Another extreme case is the South Carolina Retirement System Investment Commission (RSIC).[10] Prior to 1997, pensions in South Carolina were prohibited by law from investing in public equities. As a result, they entered the 2000s with a risky share of less than 25%, but this number grew to 74% by the end of our sample. Similar legal restrictions suppressed the risky share of

[9] One reason why is that the high-yield debt market in the United States did not mature until the 1980s.

[10] The RSIC manages the assets for the five major state pensions (e.g., South Carolina Retirement System).

the Texas County and District Retirement System (TCDRS) at the turn of the century, but the TCDRS risky share rose rapidly once these restrictions were lifted.[11]

Even without the outliers, there is still meaningful variation in the degree to which pensions have taken on more risk since the 2000s. For example, the 25th percentile pension only increased its risky share by 1 pp whereas the 75th percentile pension increased it by 16 pp. The shift has led to turnover in terms of the pensions that are taking the most risk as of 2021. Table A1a summarizes this turnover by grouping pensions into quartiles based on their risky share in both 2002 and 2021, then computing transition rates between the two groups. Rows in the table are normalized to sum to one. The top left entry of the table shows that about 33% of the pensions who were in the bottom quartile of risk-taking in 2002 remained in there in 2021. While this indicates some degree of persistence in relative risk-taking, there is also meaningful turnover as well: 15% of the pensions in the bottom quartile of risk-taking in 2021 transitioned from the top quartile of risk-taking in 2002. Moreover, as the bottom row of the table shows, 29% of the pensions who were in the top quartile of risk-taking in 2021 were in the bottom quartile in 2002.

B.2 Heterogeneity in the Alternative-to-Risky Share

There are several notable outliers in the alternative-to-risky share distribution in the cross-section of pensions, especially in the later years of the sample. For example, the target alternative-to-risky share for the Pennsylvania Public School Employees Retirement System (PSERS) has fluctuated around 75% since the early 2010s.[12] Similarly, the Indiana Public Retirement System has targeted roughly 70% of its risky investments in alternatives since 2012.[13]

For completeness, Figure A6 shows the cross-sectional distribution of alternatives in the *overall* portfolio (as opposed to the risky portfolio). The degree of heterogeneity in the alternative-to-risky share is mirrored when studying the overall share of alternatives. For example, in 2001, the difference in the overall alternative share between the 10th and 90th percentile pension was only 18 pp, roughly half the analogous spread in 2021.

C Robustness: Missing Target Shares

As mentioned in Section A.1, target shares are rarely missing in the PPD—they are available for 89% of pension systems by count and 94% by assets. In the PENDAT data, target shares are also reasonably populated on an asset-weighted basis in 1996 but are less so in 1990. Throughout the main text, we replace missing target shares with actual shares. For robustness, we now reproduce our main results by instead

[11] See this link for the South Carolina legislature, this link for the TMRS legislature, and this link for the TCDRS legislature. The TCDRS classifies private credit, distressed debt investing, strategic credit, and direct lending as "credit", whereas we classify them as alternative investments based on the description provided in TCDRS annual reports. For instance, strategic credit and direct lending investments are typically structured through limited partnerships.

[12] For example, consider the asset allocation for PSERS in 2016, which can be found on p.78 of its CAFR. Fixed income was 31.9% of the portfolio and alternatives were 48.9% of the overall portfolio (72% of the risky portfolio). We consider absolute return, commodities, master limited partnerships, infrastructure, and risk parity in alternatives.

[13] For example, according to page 48 of the 2020 CAFR, target weights in public equity, private markets, commodities, real estate, absolute return, and risk parity were 22%, 14%, 8%, 7%, 10%, and 12%, respectively.

imputing missing target shares using the method in Calvet et al. (2009).

C.1 Methodology

Using tax data for households in Sweden, Calvet et al. (2009) decompose changes in a household's risky share into a passive component driven by realized returns and an active component that reflects deliberate rebalancing. The general idea is to take existing portfolio weights and grow them based on realized returns. These re-normalized weights are considered passive because they represent the portfolio composition that would result if the investor did not trade. The difference between the observed portfolio weight next period and the calculated passive weight thus captures active portfolio adjustments made by the investor.

In our setting, we implement this idea as follows. Consider a pension that allocates $w_{j,t}$ of its portfolio in asset j at time t, such that $\sum_j w_{j,t} = 1$. Let $r_{j,t\to T}$ denote the net return of j between time t and $T > t$. The passive share at time T for asset j is then defined as:

$$w_{j,T}^{passive} = \frac{w_{j,t}(1 + r_{j,t\to T})}{\sum_k w_{k,t}(1 + r_{k,t\to T})} \tag{1}$$

The active change in j between the two periods is given by:

$$A_{t\to T}^{j} = w_{j,T} - w_{j,T}^{passive}.$$

This procedure can also be extended to construct active changes in the risky share or alternative-to-risky share. For instance, suppose assets $j < K < J$ are classified as risky assets, while $j \geq K$ are considered riskless. The passive risky share at time T is then:

$$\text{Passive Risky Share}_T = \frac{\sum_{j<K} w_{j,T}^{passive}}{\sum_k w_{k,t}(1 + r_{k,t\to T})},$$

and the active change in the risky share can be computed by subtracting the passive risky share from the observed risky share at T.

Calvet et al. (2009) construct $A_{t\to T}^{j}$ using data on individual security holdings. In our data, we do not observe this level of granularity and so j necessarily represents different asset classes (e.g., public equities). Ideally, $r_{j,t\to T}$ would reflect the exact composition of asset class j at time t. For example, suppose a pension's public equity exposure at t consists of 30% Apple stock and 70% Microsoft stock. In this case, $r_{j,t\to T}$ for public equities should, in principle, equal 30% times the return on Apple stock plus 70% times the return on Microsoft stock over the period t to T. However, because we do not observe the specific composition of asset classes, we now describe how we approximate $r_{j,t\to T}$ in the PPD and PENDAT data.

C.1.1 Active Portfolio Changes in the PPD Data

In the PPD, we approximate $r_{j,t\to T}$ for each pension based on its own return on asset class j between t and T. If a pension is missing a return in a given year, we assume it earned the aggregate pension return on

j.[14] Constructing active changes in this manner is conservative because $w_{j,T}^{passive}$ will reflect both true passive realized returns and any changes in the composition of assets within j made between t and T.

C.1.2 Active Portfolio Changes in the PENDAT Data

Unliked the PPD, the PENDAT data does not provide pension-level returns by asset class. To address this, we select benchmark indices for each asset class and assume that each pension earns the benchmark. As a result, cross-pension variation in the passive share of asset j at time T is driven entirely by differences in portfolio shares at time t. The 1990 PENDAT survey reports actual weights in the following asset classes: domestic equities, international equities, real estate equities, fixed income (government plus corporate debt), real estate mortgages, and an "other" category. For the purposes of our analysis, we combine fixed income and real estate mortgages into a single category and use the Barclay's Aggregate Fixed Income Index as its return. For real estate equities, the benchmark we use is the FTSE Nareit U.S. Real Estate Index, which tracks the performance of all U.S. real estate investment trusts (REITs). Based on the PENDAT user guidelines, the "other" category includes cash, so we use the one-month Treasury bill rate from ICE BofA US 1-Month Treasury Bill Index. For domestic equities, we use the CRSP-value weighted return from Ken French's website. For international equities, we use the return on the developed-market (ex-US) index, also from Ken French's website.

Together with actual shares from 1990, returns between 1991 and 2002 are then used to construct the passive share in equities in 2002 at the state level. These passive shares are subsequently subtracted from the actual 2002 shares from the PPD to calculate the active change in public equities between 1990 and 2002, $A^{e}_{1990\to2002}$. We focus on public equities because this is the only asset class in which we need target shares (see Section 3.2).

A similar approach is employed to construct active changes in public equities from 1996 to 2002, $A^{e}_{1996\to2002}$. The only difference is that the 1996 PENDAT survey lists alternatives assets as a separate category. For alternatives, we use the Cambridge Associates U.S. Private Equity index as the benchmark return.

C.1.3 Changes vs Levels

Before presenting the results, we want to clarify how the active portfolio changes from Calvet et al. (2009) are used in these robustness checks. Specifically, we substitute active changes in place of *changes* in target shares when target shares are missing. This substitution affects two analyses. First, it impacts our analysis of experience effects in Table 5. This table examines state-level changes in the alternative-to-risky share from 2002 to 2020 ($\Delta\Omega_{s,2002\to2021}$) in relation to changes in the share of public equities from 1996 to 2002 ($\Delta\text{Equity}_{s,1996\to2002}$) and from 1990 to 2002 ($\Delta\text{Equity}_{s,1990\to2002}$). For these variables, we use active changes in our robustness tests whenever more than 50% of the state's assets are missing target shares in at least one of the years used to construct changes. Our results are robust to different thresholds and to fully using active

[14] If a pension is missing all data for a given year in the PPD, we do not construct active changes.

changes.[15]

Second, substituting active changes in place of changes in target shares affects our analysis of risk-shifting motives in Table 7. This table shows regressions of changes in the alternative-to-risky share from 2002 to 2020 on several proxies for risk-shifting motives. When these regressions are run at the system level, missing changes in target alternative-to-risky shares are replaced with active changes. For the state-level regressions, we fill missing targets using the same substitution rule employed in the robustness check for Table 5 described above.

There is also a final question of how to handle missing *levels* of target shares in the PPD. The best way to understand how we handle these cases is through a simple example. Consider the City of Phoenix Employees Retirement System, which reports actual shares in the PPD starting in 2004 but does not report target shares until 2006. To infer target shares in 2004, we construct active changes $A^{j}_{2004\to2006}$ for each asset class j and then subtract them from the target share of asset class j in 2006. The same procedure is then repeated for missing 2005 target shares.[16] When inferring the level of target shares in this manner, we enforce short-sale and leverage constraints, and rescale imputed target shares so that they sum to one.

We emphasize that this imputation method affects a relatively small number of observations within the PPD sample, as target shares are available for 89% of pension systems by count and 94% by assets.

C.2 Results

Tables A2-A6 reproduce our main cross-sectional results when imputing missing target shares based on Calvet et al. (2009). We do not reproduce any aggregate analyses because missing shares are rare in the PPD on an asset-weighted basis. The following list compares each figure and table to its counterpart in the main text:

- Table A2: the regression coefficients are similar in magnitude to those in Table 3, and all coefficients remain statistically significant.
- Table A3: there continues to be large and statistically significant consultant fixed effects, and the incremental R^2 from adding consultant fixed effects is comparable to Table 4.
- Table A4: the four columns in this table should be compared to columns (1), (3), (4), and (5) in Table 5. Table A4 has one less column because, unlike Table 5, it does not consider the version of $\Delta\text{Equity}_{s,1990\to2002}$ that is constructed by substituting missing target shares with actual shares. The coefficients on $\Delta\text{Equity}_{s,1996\to2002}$ are similar in magnitude across the two tables, as is the pattern of statistical significance and goodness-of-fit (R^2).
- Table A5: the strength of the peer effects across all specifications is close in magnitude to those found in Table 6. Moreover, the coefficients that are statistically significant in Table 6 remain so in Table A5.

[15] A weakness of this general imputation approach is that active portfolio changes from Calvet et al. (2009) can be non-zero even absent changes in target shares. For example, suppose the equity market rises and a pension automatically rebalances toward target. The active change in its equity share will be negative in this case, even if its target share did not change.

[16] In this example, we chose 2006 as the base year because it is the first year in which target shares are available. We have also experimented with using different base years (e.g., 2020) and find that our results are very similar.

- Table A6: Consistent with Table 7, this table shows that proxies for risk-seeking motives provide little explanatory power for long-run changes in the alternative-to-risky share. The sample underlying Table A6 contains one less pension than Table 7 because the Philadelphia Board of Pensions and Retirement lacks data in 2010, and we only construct active changes for pensions that have a continuous sample.

Overall, our core motivating facts and results are extremely similar regardless of how we handle missing target shares.

D How Consultants Shape Pension Beliefs

D.1 Bayesian Learning

In this subsection, we extend the portfolio choice model introduced in Section 2.3 by incorporating a Bayesian learning framework, following Anagol et al. (2021). We first show that the paper's main results are consistent with pension funds updating their beliefs upward in response to consultant signals that are more optimistic than their own priors. Bayesian learning thus provides a natural mechanism through which consultant beliefs influence pension allocation decisions. We then characterize how pension skill—defined by the relative precision of consultant signals—modulates the causal effect of those beliefs on portfolio structure. In the model, consultants have a greater influence on the raw alternative share for lower-skill pensions, although the impact on the alternative-to-risky share is more ambiguous due to induced changes in they risky share. Empirically, we find that consultants exert more influence over smaller pensions and those with more politicized boards, two proxies for relative signal precision. Finally, we use the Bayesian framework to clarify the identifying assumptions needed to estimate the causal effect of consultant beliefs.

D.1.1 Model

The model environment is nearly identical to Section 2.3. The investable universe consists of three assets: a riskless asset, public equities, and alternatives. The log return on each is denoted by r_f, r_E, and r_A, respectively. Each investor p chooses portfolio weights $\omega = [\omega_f, \omega_A, \omega_E]'$ across asset classes to maximize expected utility, $\mathrm{E}\left[(1-\gamma)^{-1}(W_0 R_p)^{1-\gamma}\right]$, where R_p is the return of its portfolio. As before, beliefs about the risk-return characteristics of the risky assets determine the optimal allocation. Returns on public equities and alternatives are assumed to be jointly log-normal, with mean μ and variance-covariance matrix Σ of log excess returns given by:

$$\mu = \begin{bmatrix} \mu_A \\ \mu_E \end{bmatrix}, \Sigma = \begin{bmatrix} \sigma_A^2 & \sigma_{AE} \\ \sigma_{AE} & \sigma_E^2 \end{bmatrix}.$$

All pension investors are assumed to know μ_E, σ_E^2, and σ_{AE}, but they learn about the true return-generating process for alternatives. Using CAPM notation, the true return-generating process for alternatives is given by:

$$r_A - r_f = \alpha + \beta(r_E - r_f) + \eta_A,$$

where η_A is an idiosyncratic shock to alternatives that has a volatility of σ_η and β is the CAPM-slope coefficient on the market excess-return. Both σ_η and β are assumed to be known to pension investors. However, pensions do not know the true α of alternatives. Instead, each pension begins with a normally distributed prior belief about α, characterized by mean $\alpha_{0,p}$ and precision $\tau_{0,p}$:

$$p_{0,p}(\alpha) = N(\alpha_{0,p}, \frac{1}{\tau_{0,p}}).$$

We further assume that each pension p is matched with a consultant $c(p)$, who provides a (noisy) signal about alpha. The signal received by the consultant is given by:

$$\alpha_{c(p)} = \alpha + \varepsilon_p, \quad \varepsilon_p \sim N(0, \tau_p^{-1}),$$

where $\tau_p = s_p \tau_{0,p}$ determines the precision of the consultant's signal. When $s_p > 1$, the pension views the consultant's signal as more precise than its own prior. Thus, as in the Bayesian models of Gervais and Odean (2001), Linnainmaa (2011), and Anagol et al. (2021), larger values of s_p can be interpreted as a lower (relative) skill to evaluate or invest in alternatives.

After receiving the consultant's signal, each pension updates its beliefs using standard Bayesian updating. This means the posterior distribution $p_{1,p}(\alpha) \sim N(\alpha_{1,p}, \tau_{1,p}^{-1})$ of alpha is parametrized by:

$$\alpha_{1,p} = \frac{1}{1+s_p}\alpha_{0,p} + \frac{s_p}{1+s_p}\alpha_{c(p)}, \tag{2}$$

$$\tau_{1,p} = (1+s_p)\tau_{0,p}. \tag{3}$$

Let $\Delta_{c(p)} = \alpha_{c(p)} - \alpha_{0,p}$ denote the difference between consultant signal and prior belief. A positive difference $\Delta_{c(p)} > 0$ means that the consultant is more bullish about alternatives than pension p. Under the posterior distribution, the perceived expected return and risk of alternatives are given by:

$$\begin{aligned}
\mu_{A,p} &= \alpha_{1,p} + \beta\mu_E = \alpha_{0,p} + \frac{s_p}{1+s_p}\Delta_{c(p)} + \beta\mu_E \\
\sigma^2_{A,p} &= \frac{1}{1+s_p}\tau^{-1}_{0,p} + \beta^2\sigma^2_E + \sigma^2_\eta \\
\sigma_{AE} &= \beta\sigma^2_E.
\end{aligned} \tag{4}$$

In Section I.2-I.3, we derive the optimal portfolio weights of alternatives and public equities as a function of the parameters that govern returns:

$$\begin{aligned}
\omega_A &= \frac{1}{\gamma}\frac{\sigma^2_E(\mu_A + \frac{1}{2}\sigma^2_A) - \sigma_{AE}(\mu_E + \frac{1}{2}\sigma^2_E)}{\sigma^2_A\sigma^2_E - \sigma^2_{AE}}, \\
\omega_E &= \frac{1}{\gamma}\frac{\sigma^2_A(\mu_E + \frac{1}{2}\sigma^2_E) - \sigma_{AE}(\mu_A + \frac{1}{2}\sigma^2_A)}{\sigma^2_A\sigma^2_E - \sigma^2_{AE}}.
\end{aligned}$$

Substituting in the posterior beliefs from Equation (4) yields the following optimal portfolio weights:

$$\omega_A = \frac{1}{\gamma}\left(\frac{\alpha_0 + \frac{s_p}{1+s_p}\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2}{\frac{1}{1+s_p}\tau_0^{-1} + \sigma_\eta^2} + \frac{1}{2}\right),$$
$$\omega_E = \frac{1}{\gamma}\left(-\beta \cdot \frac{\alpha_0 + \frac{s_p}{1+s_p}\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2}{\frac{1}{1+s_p}\tau_0^{-1} + \sigma_\eta^2} + \frac{\mu_E}{\sigma_E^2} + \frac{1}{2}(1-\beta)\right). \tag{5}$$

D.1.2 Comparative Statics

How do pensions adjust their portfolios in response to a positive consultant signal? How does their ability to invest—captured by the relative signal precision s_p or their initial signal precision affect their portfolios and the influence of consultants? We now address these questions by deriving comparative statistics within the model.

Learning from Consultants The partial derivatives of the optimal portfolio shares with respect to $\Delta_{c(p)} = \alpha_{c(p)} - \alpha_{0,p}$ are

$$\frac{\partial \omega_A}{\partial \Delta_{c(p)}} = \frac{1}{\gamma}\frac{s_p\tau_{0,p}}{1+(1+s_p)\sigma_\eta^2\tau_{0,p}}, \tag{6}$$
$$\frac{\partial \omega_E}{\partial \Delta_{c(p)}} = -\beta\left[\frac{1}{\gamma}\frac{s_p\tau_{0,p}}{1+(1+s_p)\sigma_\eta^2\tau_{0,p}}\right] = -\beta\frac{\partial \omega_A}{\partial \Delta_{c(p)}}, \tag{7}$$

which reflects the fact that signals from consultants cause substitution between alternatives and public equities. Let $\Omega = \omega_A/(\omega_A+\omega_E)$ denote the alternative-to-risky share. Differentiating Ω with respect to $\Delta_{c(p)}$ yields:

$$\begin{aligned}\frac{\partial \Omega}{\partial \Delta_{c(p)}} &= \frac{\omega_E + \beta\omega_A}{(\omega_A+\omega_E)^2}\cdot\frac{\partial \omega_A}{\partial \Delta_{c(p)}} \\ &= \frac{\omega_E + \beta\omega_A}{(\omega_A+\omega_E)^2}\frac{1}{\gamma}\frac{s_p\tau_{0,p}}{1+(1+s_p)\sigma_\eta^2\tau_{0,p}} > 0.\end{aligned} \tag{8}$$

Intuitively, pensions will increase their alternative-to-risky share when they receive a relatively optimistic signal from their consultant. This result provides a natural microfoundation for how and why consultants have shaped the rise of alternatives.

Initial pension signal precision Next, we analyze how a pension's prior precision changes its portfolio allocations via the following partial derivatives:

$$\frac{\partial \omega_A}{\partial \tau_{0,p}} = \frac{1}{\gamma} \frac{\left(\alpha_{0,p} + \frac{s_p}{(1+s_p)}\Delta_{c(p)} - \frac{1}{2}(1-\beta)\sigma_E^2\right)}{\left(\frac{1}{(1+s_p)}\tau_{0,p}^{-1} + \sigma_\eta^2\right)^2} \frac{1}{\tau_{0,p}^2(1+s_p)},$$

$$\frac{\partial \omega_E}{\partial \tau_{0,p}} = -\beta \frac{\partial \omega_A}{\partial \tau_{0,p}}.$$

In turn, the partial derivative of Ω with respect to prior precision equals:

$$\frac{\partial \Omega}{\partial \tau_{0,p}} = \frac{\omega_E + \beta\omega_A}{(\omega_A + \omega_E)^2} \cdot \frac{1}{\gamma} \frac{\left(\overbrace{\alpha_{0,p}}^{(A)} + \overbrace{\frac{s_p}{(1+s_p)}\Delta_{c(p)}}^{(B)} \overbrace{-\frac{1}{2}(1-\beta)\beta\sigma_E^2}^{(C)}\right)}{\left(\frac{1}{(1+s_p)}\tau_{0,p}^{-1} + \sigma_\eta^2\right)^2} \frac{1}{\tau_{0,p}^2(1+s_p)} \lessgtr 0. \tag{9}$$

Interestingly, even when pensions receive a relatively bullish signal from their consultants (i.e., $\Delta_{c(p)} > 0$), the sign of the derivative in Equation (9) is still ambiguous. For sufficiently low prior beliefs $\alpha_{0,p}$ and large enough perceived equity volatility, a modestly positive signal from consultants may not be enough to outweigh the pension's initial skepticism.

A more interpretable condition follows from rearranging the (5) and substituting into Eq. (9):

$$\frac{\partial \Omega}{\partial \tau_{0,p}} = \frac{\omega_E + \beta\omega_A}{(\omega_A + \omega_E)^2} \cdot \frac{\omega_A - \frac{1}{2\gamma}}{\left(\frac{1}{(1+s_p)}\tau_{0,p}^{-1} + \sigma_\eta^2\right)\tau_{0,p}^2(1+s_p)}.$$

This derivative is only positive only if $\frac{1}{2\gamma} < \omega_A$. To gauge whether this condition is empirically relevant, suppose $\omega_A = 0.09$, its observed aggregate level in 2001, and assume $\gamma = 2$. In this case, $\omega_A = 0.09 << 0.25 = 1/2\gamma$, implying a negative partial derivative. Intuitively, by revealed preference, the low observed share of alternatives in 2001 suggests that pensions were not particularly bullish on alpha at that time. As a result, an increase in signal precision simply reinforces this skepticism, leading to a decline in the alternative-to-risky share. In contrast, when pensions begin with optimistic views about alternatives and hence a high alternative-to-risky share, an increase in confidence causes an even greater reallocation toward alternatives and away from public equities.

Pension skill (relative signal precision) Next, we investigate how the alternative-to-risky share responds to changes in s_p, which captures the precision of the consultant's signal relative to each pension's prior. Following Anagol et al. (2021), we interpret s_p as a proxy for each pension's internal ability to invest or evaluate alternative investments. The basic idea is that pensions high values of s_p are less confident in their own prior and therefore put more weight on their consultant's signal. As before, we first compute the partial

derivatives of the alternative and equity shares with respect to s_p:

$$\frac{\partial \omega_A}{\partial s_p} = \frac{\tau_{0,p}}{\gamma} \cdot \left(\frac{\alpha_{0,p} + (1+\tau_{0,p}\sigma_\eta^2)\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2}{(1+(1+s_p)\tau_{0,p}\sigma_\eta^2)^2} \right), \tag{10}$$

$$\frac{\partial \omega_E}{\partial s_p} = -\beta \cdot \frac{\tau_{0,p}}{\gamma} \cdot \left(\frac{\alpha_{0,p} + (1+\tau_{0,p}\sigma_\eta^2)\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2}{(1+(1+s_p)\tau_{0,p}\sigma_\eta^2)^2} \right) = -\beta \frac{\partial \omega_A}{\partial s_p}. \tag{11}$$

The elasticity of the alternative-to-risky share then follows:

$$\frac{\partial \Omega}{\partial s_p} = \frac{\omega_E + \beta\omega_A}{(\omega_A + \omega_E)^2} \cdot \frac{\tau_{0,p}}{\gamma} \cdot \left[\frac{\overbrace{\alpha_{0,p}}^{(A)} + \overbrace{(1+\tau_{0,p}\sigma_\eta^2)\Delta_{c(p)}}^{(B)} \overbrace{- \frac{1}{2}(1-\beta)\beta\sigma_E^2}^{(C)}}{(1+(1+s_p)\tau_{0,p}\sigma_\eta^2)^2} \right] \lesseqgtr 0. \tag{12}$$

In general, the effect of increasing s_p on the raw alternative and alternative-to-risky shares is ambiguous and depends on the relative magnitudes of the three annotated components in Equation (12). To better understand what's likely to be the empirically relevant case, we can re-write these terms as:

$$\alpha_{0,p} + (1+\tau_{0,p}\sigma_\eta^2)\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2 = \alpha_{c(p)} + \tau_{0,p}\sigma_\eta^2\Delta_{c(p)} - \frac{1}{2}(1-\beta)\beta\sigma_E^2.$$

Now assume that consultants are more bullish than pensions, so that $\Delta_{c(p)} > 0$. In this case, if $|\beta| > 1$, $C > 0$ and so increasing relative precision s_p will always increase the alternative-to-risky share. If instead $0 < \beta < 1$, then C is likely to be negligible in practice. To see why, notice that the volatility of public equities is generally around $\sigma_E \approx 17\%$ (see Figure A8). The C term is maximized at $\beta = 0.5$, which implies a maximum value of roughly $C \approx 0.0036$. This is quite small relative to the other terms in the expression. For example, $\alpha_{c(p)} \approx 0.01 - 0.03$ in the data. This suggests that the alternative-to-risky share is likely increasing with s_p when pensions receive a positive signal from their consultants.

D.1.3 The Causal Impact of Consultants

In Section 3.1, we estimate the causal impact of consultant beliefs on the alternative-to-risky share. In the model, this effect equals:

$$\begin{aligned} \frac{\partial \Omega}{\partial \alpha_{c(p)}} &= \frac{\partial \Omega}{\partial \Delta_{c(p)}} \cdot \frac{\partial \Delta_{c(p)}}{\partial \alpha_{c(p)}}. \\ &= \frac{\partial \Omega}{\partial \Delta_{c(p)}} > 0 \end{aligned}$$

where the final line follows from $\partial \Delta_{c(p)} / \partial \alpha_{c(p)} = 1$. Unsurprisingly, the causal impact of consultant beliefs is generally positive.

In our empirical analysis below, we explore how this causal effect varies with pension skill, s_p. In the

model, this is captured by the following cross-partial derivative:

$$\frac{\partial^2 \Omega}{\partial \alpha_{c(p)} \partial s_p} = \frac{\omega_e + \beta \omega_a}{(\omega_a + \omega_e)^2} \left[\frac{\partial^2 \omega_a}{\partial \alpha_c \partial s_p} - \frac{2(1-\beta)}{(\omega_a + \omega_e)} \cdot \frac{\partial \omega_a}{\partial \alpha_c} \cdot \frac{\partial \omega_a}{\partial s_p} \right] \tag{13}$$

The sign of this expression depends on the term in square brackets:

$$\frac{\partial^2 \omega_a}{\partial \alpha_{c(p)} \partial s_p} - \frac{2(1-\beta)}{(\omega_a + \omega_e)} \cdot \frac{\partial \omega_a}{\partial \alpha_{c(p)}} \cdot \frac{\partial \omega_a}{\partial s_p} \lessgtr 0.$$

Importantly, the sign of the cross-partial of the raw alternative share is unambiguously positive:

$$\frac{\partial^2 \omega_a}{\partial \alpha_{c(p)} \partial s_p} = \frac{\tau_{0,p}}{\gamma} \cdot \frac{(1 + \tau_{0,p} \sigma_\eta^2)}{(1 + (1 + s_p) \tau_{0,p} \sigma_\eta^2)^2} > 0.$$

In other words, the causal effect of consultants on the *raw alternative share* increases with s_p (or decreases with pension skill). To build more intuition, consider two pensions, A and B, that share a consultant and are identical in all respects except that $s_A > s_B$. In the language of Anagol et al. (2021), this means pension A is less skilled than B. If their consultant becomes more bullish, pension A will update its posterior mean belief about alpha ($\alpha_{1,A}$) more than B, as it places more weight on the consultant's signal. At the same time, Equation (4) shows that pension A's posterior variance of alternatives will also fall more than pension B's. Both forces will push pension A to increase its alternative share more than B.

Furthermore, when the consultant becomes more bullish, pension A will positively update its belief about the Sharpe ratio of the overall risky portfolio more aggressively. Consequently, it will also increase its *overall risky share*—potentially enough that its alternative-to-risky share could rise less than pension B's. As this example highlights, the causal impact of consultant beliefs on the *alternative-to-risky share* is therefore theoretically ambiguous.[17] When $\beta > 1$, the effect on the raw alternative share dominates, and the cross-partial $\frac{\partial^2 \Omega}{\partial \alpha_{c(p)} \partial s_p}$ is always positive. One intuition for this result is that high values of beta reduce the diversification benefits of alternatives, which mute the impact of consultants on the perceived Sharpe ratio and hence the risky share. When $0 < \beta < 1$, the relative impact of consultant beliefs on the raw alternative and risky shares depends on the underlying belief parameters in the model. Thus, whether the causal influence of consultants on the alternative-to-risky share increases with s_p is ultimately an empirical question, one that we address in the data below.

D.1.4 Identification with Bayesian Learning

In Section 3.1.2, we use an IV strategy to estimate the causal impact of consultants beliefs on the alternative-to-risky share. There, we used a reduced-form statistical framework to clarify the key identification issues and emphasize that the variation driving our Bartik-style instrument—namely, geographic distance between pensions and consultants—must be orthogonal to pension-specific preferences for alternatives, whether

[17] Formally, this is because the $\partial \omega_a / \partial \alpha_{c(p)}$ term in Equation (13) is always positive. Moreover, as discussed in Section D.1.2, the sign of $\partial \omega_a / \partial s_p$ is ambiguous. Consequently, the sign of $\frac{\partial^2 \Omega}{\partial \alpha_{c(p)} \partial s_p}$ is also ambiguous in general.

those preferences are belief-driven or not. In this subsection, we show how connect that statistical model to the Bayesian extension developed above.

To start, suppose pension p has updated its beliefs about alternatives after receiving it's consultant signal $\alpha_{c(p)}$ and chooses its alternative-to-risky share. We allow for each pension to have an idiosyncratic preference ν_p for alternatives that is orthogonal to beliefs (e.g., agency frictions). This non-belief preference can be interpreted as arising from outside the Bayesian learning model or as implicitly affecting the pension's risk aversion parameter. We can then write the alternative-to-risky share as a function of the posterior mean of alpha, a vector of all other belief-related parameters Ψ_p ($\tau_{0,p}$, s_p, etc.), and its non-belief preferences:

$$\Omega_p = f(\alpha_{1,p}, \Psi_p, \nu_p).$$

To connect this framework to a linear regression model, it is instructive to linearize $f(\cdot)$ around small changes in $\alpha_{1,p}$, Ψ_p, and ν_p:

$$\begin{aligned}\Omega_p &\approx \kappa + \theta\alpha_{1,p} + \chi'\Psi_p + \pi\nu_p. \\ &= \kappa + \underbrace{\frac{\theta s_p}{1+s_p}\alpha_{c(p)}}_{\text{Causal effect of consultants}} + \underbrace{\frac{\theta}{1+s_p}\alpha_{0,p} + \chi'\Psi_p}_{\text{Effect from pension-specific beliefs and learning}} + \underbrace{\pi\nu_p}_{\text{Non-belief motives}},\end{aligned}$$

where we have used Equation (2) to decompose $\alpha_{1,p}$ into its components from the consultant's belief and the pension's prior about the mean of alpha. Here, κ, θ, χ,and π are reduced-form coefficients that depend on the underlying structural parameters of the Bayesian portfolio choice model, including the pension's prior precision and relative signal precision. The first bracketed term in this expression reflects the causal impact of consultant beliefs $\alpha_{c(p)}$ on the alternative-to-risky share. The second set of terms reflects the influence of pension-specific belief and learning parameters. The third term represents any non-belief-based preferences for alternatives, such as agency frictions or liquidity needs.

Next, suppose we run an OLS regression of pension's alternative-to-risky share Ω_p on its consultant belief $\alpha_{c(p)}$. In this case, the population regression coefficient is given by:

$$\begin{aligned}\beta^{OLS} &= \frac{Cov(\Omega_p, \alpha_{c(p)})}{Var(\alpha_{c(p)})} \\ &= \underbrace{\frac{Cov(\theta s_p(1+s_p)^{-1}\alpha_{c(p)}, \alpha_{c(p)})}{Var(\alpha_{c(p)})}}_{\text{Causal effect of consultant}} + \underbrace{\frac{Cov\left(\theta(1+s_p)^{-1}\alpha_{0,p} + \chi'\Psi_p, \alpha_{c(p)}\right)}{Var(\alpha_{c(p)})}}_{\text{Selection on prior belief}} + \underbrace{\frac{Cov(\pi\nu_p, \alpha_{c(p)})}{Var(\alpha_{c(p)})}}_{\text{Selection on non beliefs}}. \quad (14)\end{aligned}$$

This equation highlights several important issues associated with estimating the causal impact of consultant beliefs on the alternative-to-risky share. First, this causal effect clearly depends on the pension-specific Bayesian learning parameters and may therefore vary across subsamples of pensions with different relative signal precision (s_p). In Section D.1.3, we derived comparative statistics to characterize the nature of this variation within the Bayesian model. Second, selection on beliefs can arise not only through shared priors about the mean of alpha ($cov(\alpha_{0,p}, \alpha_{c(p)}) \neq 0$), but also through any correlation between pension-specific

belief or learning parameters and consultant beliefs ($cov(\chi'\Psi_p, \alpha_{c(p)}) \neq 0$). For example, if pensions with high ability (low s_p) select consultants based on $\alpha_{c(p)}$, the OLS regression will be biased and fail to recover the true causal impact of consultants.

In Section (3.1.2), we handle these selection issues using a Bartik-style instrument whose weights are based on the distance between pensions and consultants. As the Bayesian framework makes clear, this instrument will satisfy the exclusion restriction if the weights—equivalently geographical location—are exogenous to all pension belief and learning parameters (e.g., $\alpha_{0,p}$ and s_p), as well as non-belief preferences for alternatives (v_p). In this case, the IV estimand recovers the local average treatment effect of $\alpha_{c(p)}$ on the alternative-to-risky share:

$$\beta^{IV} = E\left[\theta \frac{s_p}{1+s_p}\right]. \tag{15}$$

From here, it is easy to see that heterogenous treatment effects from consultants can be explored using interaction terms that depend on s_p.

D.1.5 Heterogenous Treatment Effects

In what follows, we explore heterogeneous treatment effects by allowing the IV estimate to vary with pensions' relative precision of the consultant signal, s_p. To do so, we modify the baseline panel regression from Section (3.1.2), which we repeat here for convenience:

$$\Omega_{p,t} = \lambda_t + \beta \alpha_{c(p),t} + \Gamma X_{p,t} + \varepsilon_{p,t},$$

where $\Omega_{p,t}$ is the target alternative-to-risky share of pension p in year t and $\alpha_{c(p),t}$ is the reported alpha of p's consultant c in year t. λ_t is a time fixed effect and $X_{p,t}$ is a large vector of controls that is designed to account for any non-belief reasons why pensions may choose a consultant (e.g., agency-driven motives). One simple way to explore heterogenous treatment effects in this setup is to incorporate interaction terms between s_p and $\alpha_{c(p),t}$. These interaction terms can then be instrumented using the interaction of $z_{p,t}$ and s_p. For this analysis, we use the instrument from Section (3.1.2) that is defined based on the location of consultants in 2005.

We use two empirical proxies for the relative signal precision parameter, s_p, which, as discussed above, can be interpreted as the (inverse) of a pension's ability to evaluate or invest in alternatives. The first proxy is overall pension size. Each year, we classify pensions as "small" or "large" based on whether their assets under management (AUM) fall below or above the sample median for that year. The rationale is that larger pensions generally have more internal resources and investment staff, enabling them to conduct more independent due diligence. As a result, they may place less weight on consultant signals, implying a lower s_p. The exact variable that we use in the regression is an indicator for whether each pension is small or large.

The second proxy is based on the composition of pension boards. Following Andonov et al. (2018), we classify pensions as more or less politicized depending on whether the share of board members who are either appointed by an elected official or serve ex officio is above or below the full-sample median.

The motivation for this proxy is that politicized boards may be more reliant on consultants due to greater internal disagreement, limited investment expertise, or a desire to offload responsibility—factors that could all increase the appeal of external guidance and thereby imply a higher value of s_p.

Figure A7 presents the results, with the IV estimates for pensions based on size presented in blue and based on board structure in red. While all coefficients are positive and statistically significant, there is meaningful heterogeneity. In particular, the causal impact of consultant beliefs on the alternative-to-risky share is roughly twice as large for smaller pensions and those with more politicized boards. This finding aligns with the intuition that consultant influence is stronger when clients have low skill and are less confident in their own priors. In the context of the Bayesian model, this pattern is consistent with $\partial\Omega/\partial\alpha_{c(p)}\partial s_p > 0$.

D.2 Decomposing Changes in Consultant-Perceived Alpha

Figure 4a in the main text shows that average consultant's perceived alpha of alternatives relative to public equities has steadily risen since the early 2000. What has driven this increase? To formally decompose the change in $\bar{\alpha}_t$, we start by breaking down the average alpha for any asset class j using the standard CAPM formula:

$$\bar{\alpha}_{j,t} = \bar{\mu}_{j,t} - \overline{rp}_{j,t},$$

where $\bar{\mu}_{j,t}$ is the average expected excess return of j, $\overline{rp}_{j,t} = \overline{\beta_{j,t}\mu_{E,t}}$ is the average risk premium of j, and μ_E is the excess return on public equities. Here, we use $\bar{x}_t$ to denote the average of variable x across consultants at time t. The total change in $\bar{\alpha}_{j,t}$ between two arbitrary periods $t = 0$ and $t = 1$ can then be written as:

$$\Delta\bar{\alpha}_j = \Delta\bar{\mu}_j - \Delta\overline{rp}_j. \tag{16}$$

The change in the risk premium term can be further decomposed be recognizing that:

$$\begin{aligned}\overline{rp}_{j,t} &= \overline{\beta_{j,t}\mu_{E,t}} \\ &= \overline{\beta}_{j,t}\overline{\mu}_{E,t} + \mathbb{C}(\beta_{j,t}, \mu_{E,t}),\end{aligned}$$

where $\mathbb{C}(\beta_{j,t}, \mu_{E,t})$ is the covariance between $\beta_{j,t}$ and $\mu_{E,t}$ across consultants at time t. Taking the difference over time, we obtain:

$$\begin{aligned}\Delta\overline{rp}_j &= \overline{\beta}_{j,1}\overline{\mu}_{E,1} - \overline{\beta}_{j,0}\overline{\mu}_{E,0} + \Delta\mathbb{C}(\beta_j, \mu_E) \\ &= \overline{\beta}_{j,1}(\overline{\mu}_{E,1} - \overline{\mu}_{E,0}) + \overline{\mu}_{E,0}(\overline{\beta}_{j,1} - \overline{\beta}_{j,0}) + \Delta\mathbb{C}(\beta_j, \mu_E),\end{aligned} \tag{17}$$

where the second line comes from adding and subtracting $\overline{\beta}_{j,1}\overline{\mu}_{E,0}$ to the first line and grouping terms. Plugging Equation (17) into Equation (16), then dividing by $\Delta\bar{\alpha}_j > 0$ yields the following decomposition:

$$\begin{aligned}&\Delta\bar{\alpha}_j = \Delta\bar{\mu}_j - \overline{\beta}_{j,1}\Delta\overline{\mu}_E - \overline{\mu}_{E,0}\Delta\overline{\beta}_j - \Delta\mathbb{C}(\beta_j, \mu_E) \\ &\quad\Leftrightarrow \\ &1 = R_j + P_j + D_j + C_p,\end{aligned} \tag{18}$$

where $R_j = \Delta\bar{\mu}_j/\Delta\bar{\alpha}_j$, $P_j = -\overline{\beta}_{j,1}\Delta\overline{\mu}_E/\Delta\bar{\alpha}_j$, and so on. Assuming $\bar{\beta}_{j,t}, \mu_{E,t} > 0$ for all t, Equation (18) says that an increase in the average perceived alpha can come from four sources:

- An increase in average expected excess returns or what we call the "expected return" channel ($\Delta\bar{\mu}_j > 0 \Leftrightarrow R_j > 0$).
- A decline in the average perceived risk premium on public equities or what we call the "public equity risk premium" channel ($\Delta\overline{\mu}_E < 0 \Leftrightarrow P_j > 0$).
- A decline in the average perceived beta or, equivalently, an improvement in diversification. We call this the "diversification" channel ($\Delta\overline{\beta}_j < 0 \Leftrightarrow D_j > 0$).
- A decline in the cross-sectional covariance between beta and the public equity risk premium ($\Delta\mathbb{C}(\beta_j, \mu_E) < 0 \Leftrightarrow C_p > 0$).

Table A7 shows the formal decomposition of changes in alpha based on Equation (18). Each column reports the percentage of the increase in average perceived alpha from Table 2 attributable to the expected return, public equity risk premium, and diversification channels. For example, "Diversification" column corresponds to D_j in Equation (18). A negative value in this column indicates that perceived diversification benefits have worsened since 2003, contributing negatively to the increase in perceived alpha. The rows in Table A7 do not sum to 100% because we omit the covariance term C_p for readability.

Overall, Table A7 confirms that higher expected returns on alternatives relative to public equities—rather than improved diversification benefits—are the primary driver of the rising perceived alpha over time. For all alternatives, virtually all of the increase in average alpha is attributable to higher expected returns, while the effects of a declining public equity risk premium and worsened diversification benefits largely offset each other. At the individual asset class level, the majority of the increase in perceived alpha has also come from an increase in expected returns. For PE and HFs, the decline in the public equity risk premium has also contributed meaningfully to the increase in average alpha. For RE, the increase in expected excess returns is even larger than the rise in perceived alpha, by a factor of over two. However, this increase has been dampened by a rise in the perceived beta of RE with respect to public equities.

For completeness, Figure A8 provides insight into the sources of the small increases in $\bar{\beta}_j$ over time. Figure A8a plots the average perceived volatility ($\bar{\sigma}_j$) for public equities, PE, RE, and HFs in both 2003 and 2021, while Figure A8b shows average correlations with public equities. The results indicate that perceived volatility has remained largely stable, with the largest decline occurring for PE. All else equal, this decline in volatility should have lowered average betas, not increased them. In contrast, perceived correlations with public equities have increased across the board, suggesting that the modest rise in average perceived beta is primarily driven by stronger perceived co-movement with public markets. While it is possible to conduct a formal decomposition similar to Equation (18), we opt not to do so, given that the magnitude of the increase in average beta is small.

D.3 The Causal Impact of Consultants

This subsection elaborates on how we identify the causal effect of consultants in Section 3.1.2 of the main text.

D.3.1 Definition of active consultant

The instrument we use in Section 3.1.2 is based on whether a consultant is active in a given pension's state. We define a consultant being active in pension $p's$ state as of year k if it had employed at least one active "Investment Advisory Representatives" (IARs) in that state. That is, our instrument is *not* defined based on whether the consultancy itself is registered with a specific state. Rather, we define it based on the set of IARs employed by the consultancy. Specifically, we consider a consultant as active in a particular state if it employs at least one IAR who is registered to operate or maintains their home office in that state. Information on the registration and employment status of IARs is derived from an updated version of the dataset used by Egan et al. (2019).[18]

To become an IAR, individuals must pass credentialing exams administered by the Financial Industry Regulatory Authority (FINRA). Although exact requirements can differ by state, most IARs complete at least the Uniform Securities Agent State Law Examination (Series 63) and the Uniform Investment Adviser Law Examination (Series 65). IARs are subject to a fiduciary standard under the Investment Advisers Act of 1940 and must follow SEC-compliant recordkeeping standards. Beyond these federal requirements, IARs must also adhere to specific state registration regulations. According to NASAA website, these state-level requirements are not trivial:

> All states require that investment advisers (RIAs) and investment adviser representatives (IARs) conducting business in the state register with the state's securities regulatory authority (or qualify for an exemption from registration). Approximately 17,500 investment advisers are so registered. Registration standards are complex and you should consult an attorney if you have questions about whether you or your firm must register...
>
> An applicant for registration as an investment adviser representative must file a Form U4, which also has its own completion instructions. The Form U4 must show that the applicant has passed the required competency exams or holds an active, acceptable professional designation.

Overall, these institutional features suggest that there are meaningful barriers to entry for IARs registering in individual states. These costs are evident when looking at the number of states in which consultancies are active, again defined based on employing at least one IAR who is registered to operate or maintains their home office in that state. If such costs were negligible, we would expect most consultancies to be active in all 51 jurisdictions (including the District of Columbia). Instead, the average consultancy is active in about 13 states at some point during our sample, with only a handful maintaining a nationwide presence. We revisit this issue below when discussing the relevance of our instrument below.

[18] We are grateful to Mark Egan for sharing this data with us.

D.3.2 Instrument Validity

For our IV to be valid, the instrument must satisfy the standard conditions of exogeneity and relevance.

Instrument Exogeneity: As discussed in the main text, we use two different weighting schemes to ensure our instrument $z_{p,k,t}$ is plausibly exogenous. The first is to set $k = t$, so that consultant locations and beliefs are measured in the same year. The logic of this weighting scheme is that consultant location in year t is exogenous to contemporaneous pension preferences for alternatives, *after* conditioning on $X_{p,t}$. One might object that consultants have incentives to be geographically close to large pensions given their fees are typically based on AUM. This is precisely why we directly control for size, funding, and several other pension characteristics in $X_{p,t}$.

The second weighting scheme sets $k = 2005$, the year before the shift to alternatives accelerated (Figure 2c). The logic of this choice is that consultant location choices in 2005 are plausibly exogenous to pension-specific preferences for alternatives in year $t > 2005$. The exogeneity condition would fail if, for instance, consultants anticipated which pensions would later have idiosyncratic preferences for alternatives and chose their registration locations accordingly. However, we think this possibility is unlikely.

Instrument Relevance: The second condition the instrument must satisfy is relevance, which relies on the fact that most consultancies employ registered financial advisors in only a limited subset of states. To illustrate this more formally, consider the instrument that sets $k = t$, meaning consultant locations and beliefs are measured in the same year. If all consultants employed registered advisors in every state, the instrument $z_{p,k,t}$ would exhibit no cross-sectional variation and collapse to a time-specific value $z_{c,k,t} = z_t$. In this case, z_t will be fully absorbed by the time fixed effect λ_t in the first-stage of our IV, leading to a weak instrument problem.

Reassuringly, the first-stage F-statistics in columns (3) and (4) of Table 3 indicate that our IV does not suffer from a weak instruments problem. Both F-statistics exceed 10, the conventional threshold below which weak instrument concerns typically arise (Staiger and Stock, 1997; Andrews et al., 2019). This provides strong evidence that there is sufficient variation in where consultants employ registered financial advisors for the purposes of our IV. As discussed above, this result is not surprising given the costs associated with becoming an IAR.

D.4 Variance Decompositions

In this subsection, we use the statistical framework from Section 3.1.2 to develop a lower bound on the amount of cross-sectional variation in Ω_p that is attributable to beliefs, regardless of whether these beliefs are driven by consultants or not. To develop the intuition, we start with the simple case in which there is no selection by pensions of consultants based on non-belief preferences for alternatives. We then relax this assumption and show the bounding argument applies so long as there is no selection on non-beliefs, after conditioning on observables.

D.4.1 Simple Case

In Section D.1.4, we used the Bayesian extension of the mean-variance model to derive a reduced-form relationship between the alternative-to-risky share, consultant beliefs, and pension-specific preferences (prior beliefs and non-belief preferences). For the purposes of this section, it will be helpful to work with the following simplified version of that relationship, though our main point also holds in the more general case:

$$\Omega_p = \theta \alpha_{c(p)} + \kappa \xi_p + \pi \nu_p,$$

where $\alpha_{c(p)}$ is its consultant's belief about alpha and ξ_p is a pension-specific prior belief. ν_p reflects is any non-belief preferences for alternatives, such as those arising from agency-based or risk-seeking incentives.

Assuming ν_p and ξ_p are orthogonal, this reduced-form relationship implies that the total cross-sectional variance of the target alternative-to-risky share Ω_p is given by:

$$\mathbb{V}(\Omega_p) = \mathbb{V}(\theta \alpha_{c(p)}) + \mathbb{V}(\kappa \xi_p) + \mathbb{V}(\pi \nu_p) + 2\mathbb{C}(\theta \alpha_{c(p)}, \kappa \xi_p) + 2\mathbb{C}(\theta \alpha_{c(p)}, \pi \nu_p), \tag{19}$$

where the last two covariance terms reflect selection on beliefs and non-beliefs, respectively. The total amount of variation B attributable to beliefs is therefore given by:

$$B = \frac{\mathbb{V}(\theta \alpha_{c(p)}) + \mathbb{V}(\kappa \xi_p) + 2\mathbb{C}(\theta \alpha_{c(p)}, \kappa \xi_p)}{\mathbb{V}(\omega_p^*)}. \tag{20}$$

In general, it is difficult to estimate B because we don't observe pension-specific beliefs about the risk and return of alternatives. However, it is possible to put a lower bound on B using simple OLS regressions. This point is easiest to see when there no selection ($\tau_B = \tau_{NB} = 0$), so that the covariance terms in Equation (19) and (20) are zero. In this case, a lower bound $\underline{B}$ is given by:

$$B \geq \underline{B} = \frac{\mathbb{V}(\theta \alpha_{c(p)})}{\mathbb{V}(\Omega_p)}.$$

$\underline{B}$ is the simply the R^2 from an OLS regression of Ω_p on $\alpha_{c(p)}$. If $\alpha_{c(p)}$ is not available for all consultants, as is the case in our data, this lower bound can also be obtained by regressing Ω_p on consultant fixed effects.

When there is selection, the R^2 from this fixed effects regression still provides a lower bound on B, provided that there is no selection into consultants based on non-beliefs ($\tau_{NB} = 0$). Intuitively, when there is no selection on non-beliefs, variation in consultant fixed effects can only be driven by variation in consultant beliefs, pension beliefs, or matching based on beliefs. To see this point formally, let $\mathbb{E}_c[\cdot]$ denote the expectation operator, conditional on consultant identity. Consultant fixed effects are then defined as $\lambda_c = \mathbb{E}_c[\Omega_p] = \theta \alpha_{c(p)} + \mathbb{E}_c[\kappa \xi_p] + \mathbb{E}_c[\pi \nu_p]$. When there is no selection on non-beliefs, $\mathbb{E}_c[\nu_p] = \mathbb{E}[\nu_p]$, since knowing the identity of a pension's consultant provides no information about ν_p. This means that $\mathbb{E}_c[\nu_p]$

does not vary in the cross-section of consultants, so that the variance of the fixed effects is given by:

$$\begin{aligned}\mathbb{V}(\lambda_c) &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\mathbb{E}_c[\kappa\xi_p]) + 2\mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]) \\ &\leq \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p),\end{aligned}$$

where the inequality comes from repeated application of the law of total covariance. Specifically, note that:

$$\begin{aligned}\mathbb{V}(\lambda_c) &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\mathbb{E}_c[\kappa\xi_p]) + 2\mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]) \\ &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) - \mathbb{E}[\mathbb{V}_c(\kappa\xi_p)] + 2\mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]) \\ &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) - \mathbb{E}[\mathbb{V}_c(\kappa\xi_p)] + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p) \\ &\leq \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p), \qquad (21)\end{aligned}$$

where the second line uses the law of total variance and the third line uses the law of total covariance. We also take advantage of the fact that:

$$\begin{aligned}\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p) &= \mathbb{C}(\mathbb{E}_c[\theta\alpha_{c(p)}], \mathbb{E}_c[\kappa\xi_p]) + \mathbb{E}[\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p|c)] \\ &= \mathbb{C}(\mathbb{E}_c[\theta\alpha_{c(p)}], \mathbb{E}_c[\kappa\xi_p]) \\ &= \mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]),\end{aligned}$$

where $\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p|c)$ is the conditional covariance given consultant identity. Here, $\mathbb{E}_c[\theta\alpha_{c(p)}] = \theta\alpha_{c(p)}$ because $\alpha_{c(p)}$ is constant within consultant. For the same reason, $\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p|c) = 0$.

Dividing both sides of Equation (21) by $\mathbb{V}(\Omega_p)$ delivers a lower bound for B:

$$\underline{B} = \frac{\mathbb{V}(\lambda_c)}{\mathbb{V}(\Omega_p)} \leq \frac{\mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p)}{\mathbb{V}(\Omega_p)} = B.$$

The bound $\underline{B}$ is simply the R^2 from a regression of ω_p^* on consultant fixed effects. The tightness of the bound depends on the expected variance of ξ_p within each consultant, $\mathbb{E}[\mathbb{V}_c(\kappa\xi_p)]$.

D.4.2 Conditioning on Observables

We modify the variance decomposition by conditioning on a set of pension observables X_p that plausibly capture non-belief driven motives for pensions to invest in alternatives. The OLS regression of Ω_p on consultant beliefs $\alpha_{c(p)}$ and pension observables becomes

$$\Omega_p = a_0 + \beta\alpha_{c(p)} + \Gamma X_p + \varepsilon_p.$$

The population regression coefficients are given by

$$\beta = \theta + \tau_B \kappa + \tau_{NB} \pi,$$
$$\Gamma = \Psi_B \kappa + \Psi_{NB} \pi,$$

where τ_B and Ψ_B are coefficients from a regression of pension-specific belief component ξ_p on $\alpha_{c(p)}$ and X_p, and τ_{NB} and Ψ_{NB} are coefficients from a regression of non-belief-driven component ν_p on $\alpha_{c(p)}$ and X_p. Here, $\tau_{NB} = 0$ means that there is no selection on non-beliefs, *conditioning* on pension observables X_p. This is a weaker assumption than in D.4.1 which assumes no selection on non-beliefs unconditionally.

Let $\mathbb{V}_{X_p}(\cdot)$ denote the variance operator, conditional on pension observables X_p. Similarly, $\mathbb{V}_{c,X_p}(\cdot)$ denotes the variance operator conditional on both X_p and consultant identity c. Conditional expectations and covariances are denoted analogously. The total amount of variation B attributable to beliefs, conditioning on pension observables, is given by

$$B = \frac{\mathbb{V}_{X_p}(\theta \alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa \xi_p) + 2\mathbb{C}_{X_p}(\theta \alpha_{c(p)}, \kappa \xi_p)}{\mathbb{V}_{X_p}(\Omega_p)}$$

Consider the fixed effect regression $\Omega_p = \lambda_c + \Phi X_p + \varepsilon_p$. We wish to prove that $\underline{B} = \frac{\mathbb{V}_{X_p}(\lambda_c)}{\mathbb{V}_{X_p}(\Omega_p)} = \frac{\mathbb{V}_{X_p}(\lambda_c + \Phi X_p)}{\mathbb{V}_{X_p}(\Omega_p)}$ forms a lower bound on the size of belief-driven variation B in the cross-section. While we cannot compute the actual bound B, we can compute $\underline{B}$ from the data.

The proof is similar to the simple case in the previous subsection. The value of $\lambda_c + \Phi X_p$ is given by

$$\begin{aligned}\lambda_c + \Phi X_p &= \mathbb{E}_{c,X_p}[\Omega_p] \\ &= \mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}] + \mathbb{E}_{c,X_p}[\kappa \xi_p] + \mathbb{E}_{c,X_p}[\pi \nu_p].\end{aligned}$$

Assuming ν_p and ξ_p are orthogonal, the general formula for the conditional variance is

$$\begin{aligned}\mathbb{V}_{X_p}(\lambda_c) =& \mathbb{V}_{X_p}(\lambda_c + \Phi X_p) \\ =& \mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}]) + \mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\kappa \xi_p]) + \mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\pi \nu_p]) \\ &+ 2\mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}], \mathbb{E}_{c,X_p}[\kappa \xi_p]) + 2\mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}], \mathbb{E}_{c,X_p}[\pi \nu_p]).\end{aligned}$$

We further assume that there is no selection on non-beliefs, conditioning on X_p. That is, $\tau_{NB} = 0$. Then, v_p does not depend on consultant identity c. We have $\mathbb{E}_{c,X_p}[\pi \nu_p] = \mathbb{E}_{X_p}[\pi \nu_p]$. The term $\mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\pi \nu_p]) = \mathbb{V}_{X_p}(\mathbb{E}_{X_p}[\pi \nu_p])$ becomes zero because conditioning on X_p, there is only one value of $\mathbb{E}_{X_p}[\pi \nu_p]$ and no variation left. Similarly, the second covariance term $\mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}], \mathbb{E}_{c,X_p}[\pi \nu_p]) = \mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}], \mathbb{E}_{X_p}[\pi \nu_p])$ becomes zero. The conditional variance is now

$$\mathbb{V}_{X_p}(\lambda_c) = \mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}]) + \mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\kappa \xi_p]) + 2\mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta \alpha_{c(p)}], \mathbb{E}_{c,X_p}[\kappa \xi_p]).$$

Using the law of total variance, the first term becomes

$$\mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\theta\alpha_{c(p)}]) = \mathbb{V}_{X_p}(\theta\alpha_{c(p)}) - \underbrace{\mathbb{E}_{X_p}[\mathbb{V}_{c,X_p}(\theta\alpha_{c(p)})]}_{=0} = \mathbb{V}_{X_p}(\theta\alpha_{c(p)}),$$

where in the conditional expectation, $\theta\alpha_{c(p)}$ has no variation within consultant c, so certainly no variation conditional on both $c$and X_p. This implies $\mathbb{V}_{c,X_p}(\theta\alpha_{c(p)}) = 0$. Again using the law of total variance, the second term becomes

$$\mathbb{V}_{X_p}(\mathbb{E}_{c,X_p}[\kappa\xi_p]) = \mathbb{V}_{X_p}(\kappa\xi_p) - \mathbb{E}_{X_p}[\mathbb{V}_{c,X_p}(\kappa\xi_p)].$$

Using the law of total covariance, the third term becomes

$$\begin{aligned}\mathbb{C}_{X_p}(\mathbb{E}_{c,X_p}[\theta\alpha_{c(p)}], \mathbb{E}_{c,X_p}[\kappa\xi_p]) &= \mathbb{C}_{X_p}(\theta\alpha_{c(p)}, \kappa\xi_p) - \mathbb{E}_{X_p}[\underbrace{\mathbb{C}_{c,X_p}(\theta\alpha_{c(p)}, \kappa\xi_p)}_{=0}] \\ &= \mathbb{C}_{X_p}(\theta\alpha_{c(p)}, \kappa\xi_p),\end{aligned}$$

where $\mathbb{C}_{c,X_p}(\theta\alpha_{c(p)}, \kappa\xi_p) = 0$ because $\theta\alpha_{c(p)}$ is constant conditional on both c and X_p. Combining all three terms gives us

$$\begin{aligned}\mathbb{V}_{X_p}(\lambda_c) &= \mathbb{V}_{X_p}(\theta\alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa\xi_p) - \mathbb{E}_{X_p}[\mathbb{V}_{c,X_p}(\kappa\xi_p)] + 2\mathbb{C}_{X_p}(\theta\alpha_{c(p)}, \kappa\xi_p) \\ &\leq \mathbb{V}_{X_p}(\theta\alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa\xi_p) + 2\mathbb{C}_{X_p}(\theta\alpha_{c(p)}, \kappa\xi_p) \qquad (22)\end{aligned}$$

Dividing both side of Equation (22) by $\mathbb{V}_{X_p}(\Omega_p)$ yields

$$\underline{B} = \frac{\mathbb{V}_{X_p}(\lambda_c)}{\mathbb{V}_{X_p}(\Omega_p)} \leq \frac{\mathbb{V}_{X_p}(\theta\alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa\xi_p) + 2\mathbb{C}_{X_p}(\theta\alpha_{c(p)}, \kappa\xi_p)}{\mathbb{V}_{X_p}(\Omega_p)} = B.$$

This lower bound $\underline{B}$ is equivalent to the incremental R^2 between the regression with consultant fixed effects only $\Omega_p = \lambda_c + \varepsilon_p$ and the regression with consultant fixed effects and pension controls $\Omega_p = \lambda_c + \Gamma X_p + \varepsilon_p$. The tightness of the bound depends on the conditional expected variance on $\kappa\xi_p$ within each consultant c and pension observable X_p pair.

D.4.3 Selection on Non-Beliefs: Evidence from Private-Sector Pensions

The variance decompositions in Section 3.1.3 rely on the assumption that, conditional on observables, there is no selection of consultants by pensions on non-beliefs. The fact that the variance decompositions are similar when including state-by-time fixed effects supports this conclusion, as does our analysis of the causal impact of consultant beliefs on portfolio composition from Section 3.1.2. To reinforce the validity of this assumption further, we now study the relationship between public- and private-sector pensions who share a consultant. The idea behind this analysis is as follows. Suppose that public- and private-sector clients who

share a consultant have similar alternative-to-risky shares. The two-fund separation theorem of Tobin (1958) points to beliefs as the obvious explanation for this pattern, regardless of whether consultants have a causal impact on pension beliefs. Agency frictions could also drive this result, but then the agency problem would need to be shared by pensions in both the private and public sectors. Importantly, this would immediately rule out frictions that are specific to the governance, accounting, and regulation of public pensions. A similar logic applies if consultant selection is driven by risk-seeking motives. Notably, any agency or risk-seeking motive of this sort would also need to be orthogonal to the controls included in Table 4, including size and pension funding.

Figure A9a visualizes the aforementioned relationship using a binned scatter plot of the average alternative-to-risky share of each consultant's public-sector clients, $\bar{u}^a_{c,t}$, against the average of its private-sector clients, $\bar{p}^a_{c,t}$. The binned scatter plot also controls for a year fixed effect.[19] To compute $\bar{p}^a_{c,t}$, we use data on actual (not target) U.S. private-sector allocations matched to consultants from S&P's Money Market Directory. These data are discussed in detail in Section A.5. Private-sector clients include corporate defined-benefit pensions, endowments, and unions. The figure reveals a strong and positive correlation between the alternative-to-risky share of a consultant's public- and private-sector clients. In the cross-section, a 10 pp increase in the alternative-to-risky share of a consultant's private-sector clients is associated with a 6 pp ($t = 4.15$) increase in the alternative-to-risky share of its public-sector clients.

Figure A9b repeats this exercise using the average risky share of each consultant's public- and private-sector clients, denoted by $\bar{u}^r_{c,t}$ and $\bar{p}^r_{c,t}$, respectively. Unlike the alternative-to-risky share, there is no evidence that clients who share a consultant have similar risky shares. After controlling for a time fixed effect, a regression of $\bar{u}^r_{c,t}$ on $\bar{p}^r_{c,t}$ yields a statistically insignificant coefficient that is near zero and a within-R^2 that is also effectively zero. Figure A9 supports the idea that independent factors drive the risk budgets of public- and private-sector pensions, whereas beliefs about the alpha of alternatives drive the composition of their risky portfolios. This interpretation echoes the findings of Foerster et al. (2017), who find similarly large advisor effects within a sample of Canadian households. They further emphasize the role of beliefs by showing that advisors invest their own portfolios in a similar manner to their clients.

E Belief-Based Explanations: Additional Evidence

E.1 Peer Effects

In the main text, we document that pensions are more likely to invest in alternatives when nearby pensions do as well. The broad mechanism we propose in the paper is that pension beliefs about the risk-return profile of different asset classes are shaped by the actions and beliefs of their peers. This idea draws inspiration from recent research in household finance, which finds that social networks influence housing market expectations (Bailey et al., 2018) and product adoption Bailey et al. (2022).

In this appendix, we elaborate on how peer networks might form in the context of public pensions. One plausible channel is through conference travel, industry events, and professional networking opportunities.

[19] Figure A9a is effectively unchanged if we first orthogonalize $\bar{u}^a_{c,t}$ and $\bar{p}^a_{c,t}$ to the same pension controls used in Table 4 before collapsing to the consultant-year level.

For example, the California Association of Public Retirement Systems (CALAPRS) is a consortium of public pensions in California, including large systems like CALSTRS as well as smaller municipal pensions like the San Jose Federated City Employees' Retirement System. According to its website, CALAPRS:

> ...sponsors educational forums for sharing information and exchanging ideas among trustees and staff of California public retirement systems to enhance their ability to administer public pension plan benefits and manage investments consistent with their fiduciary responsibility.
>
> CALAPRS carries out its mission through an annual conference called the General Assembly, round table meetings for discussion of topics of interest, the annual Principles of Pension Governance courses for Trustees, the Overview, Intermediate, and Advanced Staff Training, Disability Staff Training, and the Management/Supervisory Academy.

Beyond state-level organizations, there are also several national associations where public pension employees and executives exchange information. The two largest are the National Association of State Retirement Administrators (NASRA) and the National Conference on Public Employees Retirement Systems (NCPERS). Both host annual conferences explicitly designed to provide training, discuss investment strategies, connect with service providers, and facilitate networking.[20] For instance, the 2024 NCPERS annual meeting included sessions such as:

- "Economic Overview: A Look at What's Ahead"
- "Are We There Yet? Has the Recovery in Real Estate Values Begun? And What Comes Next?"
- "Asset Allocation: Building Resilient Portfolios".

Much like the American Finance Association (AFA) annual meetings, the NASRA and NCPERS annual conferences rotate locations each year. If the likelihood of attendance increases with geographical proximity, it is easy to imagine how peer networks could form regionally.

The general notion that pension employees form peer networks through conference travel and industry events—and are more likely to do so at nearby ones—has some anecdotal support. As part of its Transparency and Accountability Program, CalPERS publishes a Travel Transparency Report for each fiscal year, detailing the trips taken by its employees. In FY 2023, CalPERS employees took 162 trips classified as "Conference / Forum", amounting to 665 total travel days.[21] These trips exhibit a clear geographic skew, with 38% occurring in California or within their own census division (the Pacific).

There are, of course, many other factors outside of conference and industry events that influence peer network formation. Labor flows between pensions are one such candidate, as are shared educational experiences. These channels could be explored by testing whether peer effects are stronger between pensions that exhibit greater labor mobility or between those whose employees attended the same academic institutions. The travel-related mechanism we laid out about could also be systematically studied by assembling a national database of pension travel records using Freedom of Information Act (FOIA) requests. We view these as interesting avenues for future research.

[20] Information on the annual NASRA and NCPERS meetings can be found here and here.

[21] The report can be found at this link.

There is also a possibility that consultants themselves are the source of peer networks. Note that the consultant-by-census-division-by-year fixed effect in our baseline regression do not necessarily rule this possibility out.[22] To see why, first recall that the workhorse regression we use to study peer effects is given by:

$$\Omega_{p,t} = \lambda_{cdt} + \beta_z n_{pt} + \theta^{'} \mathbf{X}_{pt} + \varepsilon_{pt}.$$

where λ_{cdt} is a consultant-by-time-by-census-division fixed effect and $\mathbf{X}_{pt}$ is a vector of pension observables. The regressor of interest is n_{pt}, which is a inverse-distance-weighted average of the target alternative-to-risky share of *all* U.S. pensions, regardless of whether they share a consultant. Formally, n_{pt} is defined as:

$$n_{pt} \equiv \sum_{k \neq p} \delta_{pk} \times \Omega_{k,t},$$

with weight $\delta_{pk} = d_{pkt}^{-1} \times \left(\sum_{j \neq p} d_{p,j,t}^{-1} \right)^{-1}$. $\Omega_{k,t}$ is the target alternative-to-risky share of pension k in year t and d_{pkt} is the distance (in kilometers) between the headquarters of pension p and pension k in year t.

Next, suppose that proximity-based peer effects only occur among pensions that share a consultant. In this case, β_z will still be positive even when including λ_{cdt}, as the fixed effect only controls for the average effect of consultants on clients in the same census division. In other words, variation in $\Omega_{p,t}$ that exists within a consultant-census-division-year could still be driven by distance to only other pensions that share a consultant (as opposed to all pensions).

The correct way to handle this issue is therefore through the construction of n_{pt}. As a benchmark, column (1) of Table A8 reproduces the estimates from the main paper. In addition, column (2) presents results when omitting the consultant fixed effect and instead including only a time-by-census-division effect. The point estimates are stable across both specifications, though as noted above, this does not rule out the possibility that consultants are the source of peer effects.

Column (3) addresses this issue directly by introducing an auxiliary measure $n_{pt}^{no-share}$ of peer behavior that distance-weights across pensions that do *not* share a consultant with p. If consultants are the sole source of the peer network, then replacing n_{pt} with $n_{pt}^{no-share}$ should drive the regression coefficient β_z to zero. However, the results in Column (2) show this is not the case. The estimated β_z is of comparable magnitude to column (1) and is strongly statistically significant. Furthermore, we cannot reject the null hypothesis that it equals the estimate in column (1). The results in column (2) therefore imply that the peer networks we document are distinct from consultants.

E.2 Heterogeneity Across Different Types of Alternatives

What explains the heterogeneity documented in Figure 3c? In the textbook mean-variance model, differences in the composition of alternative investments across pensions can arise for one of two reasons. The first is rooted in beliefs. In practice, pensions and consultants form beliefs about individual asset classes (e.g., PE) rather than alternatives as a whole. Any disagreement about the risk-return properties of individual asset classes will naturally give rise to heterogeneity in the composition of the target alternative portfolio.

[22] We thank an anonymous referee for raising this possibility to us.

The second is rooted in risk-seeking motives: some pensions are constrained from increasing their risky share (e.g., a leverage constraint), forcing them to instead tilt the composition of their risky (and alternative) investments toward specific types of alternatives. As we show in a three-asset mean-variance model in Section 2.3, these same mechanisms also generate variation in the overall alternative-to-risky share. Here, we are simply extending the logic of that simplified model to allow for a larger number of more granular asset classes.

We employ a simple variance decomposition to explore both potential mechanisms in the PPD data. We run the following panel regression:

$$y_{p,t} = \alpha_t + \lambda_c + \theta RS_{p,t} + \varepsilon_{p,t},$$

where $y_{p,t}$ is the target share of alternatives invested in either PE, RA, or HFs for pension p in year t. α_t is a time fixed effect, λ_c is a consultant fixed effect, and $RS_{p,t}$ are a set of proxies for risk-seeking motives. $RS_{p,t}$ is drawn primarily from Section 4.2 and includes GASB 25 funding level, liability discount rates, and the fraction of retired members in a given year. For completeness, $RS_{p,t}$ also includes measures of liquidity and risk constraints, such as the ratio of required contributions to payroll, the ratio of total spending to payroll, the cash share, the difference between each pension's target and actual risky share (see Section 4.1), and realized 1-year portfolio returns. This regression is similar in spirit to the one used in Section 3.1.3 for analyzing consultant fixed effects, but here it is applied specifically to the composition of alternative investments, rather than to the composition of all risky investments.

Figure A10 presents the variance decomposition results across alternative asset subcategories. The red bars in the plot show the adjusted R^2 when adding risk-seeking proxies $RS_{p,t}$ alongside the time fixed effect. Across all subcategories of alternatives, the incremental R^2 from adding risk-seeking proxies is small and never exceeds 1 pp. The green bars in the plot show that, in contrast, adding consultant fixed effects λ_c to the model dramatically increases the adjusted R^2. The incremental improvement in explanatory power from consultant fixed effects is an order of magnitude greater than that from adding risk-seeking proxies.

We conclude that risk-seeking motives alone are unlikely to explain heterogeneity in the composition of alternative investments across pensions. Instead, they suggest that consultants disagree about the risk-return properties of specific alternatives and advise their clients accordingly. Section H.1.1 further supports this conclusion by showing that pensions are more likely to invest in real assets over private equity when their consultant reports real assets to have a higher alpha.

E.3 Are Pension Beliefs Rational?

In this section, we provide a preliminary assessment of whether beliefs about alpha are rational. Our analysis proceeds in two steps. First, we benchmark consultant beliefs about the beta, alpha, and expected return of alternatives against estimates from historical data and the academic literature. Second, we examine whether the alternative investments held by pensions with high alternative-to-risky shares have outperformed those held by pensions with lower shares. Overall, the evidence is mixed at best, underscoring the need for further research on this question.

E.3.1 Benchmarking Consultant Beliefs

To start, we investigate how consultant beliefs stack up against the historical performance of alternatives. When comparing consultant beliefs about alternatives to historical returns, we will repeatedly estimate CAPM alphas and betas using the following regression:

$$R_{a,t} - R_{f,t} = \alpha + \beta_0(R_{m,t} - R_{f,t}) + \beta_1(R_{m,t-1} - R_{f,t-1}) + \varepsilon_{a,t},$$

where $R_{a,t}$ is the annual return of alternative asset a (in this case PE), $R_{m,t}$ is the annual return of the CRSP value-weighted index, and $R_{f,t}$ is the annual return of the 1-month Treasury bills. The latter two are taken from Ken French's website. The CAPM beta is defined as $\beta = \beta_0 + \beta_1$. Heteroskedasticity robust standard errors are used throughout.

Estimating alphas and betas in this manner has well-known limitations due to the illiquidity and non-linear return profile of many alternative asset classes (e.g., Jurek and Stafford, 2015; Stafford, 2022). As above, the inclusion of lagged market returns partly addresses illiquidity (Dimson, 1979; Korteweg, 2019), but is far from a perfect solution. In the comparisons that follow, it is also important to keep in mind that CMAs typically contain beliefs about PE as an aggregate asset class, one that includes buyout and venture capital. We will often analyze both asset classes, though buyouts the most natural benchmark for CMAs because account for a much larger share of AUM than venture capital.[23]

Private Equity We start by analyzing the historical performance of PE, as measured by the annual return on the Cambridge Associates (CA) U.S. Private Equity Index, a widely used benchmark in both industry and academia (e.g., Korteweg, 2023). Figure A11 reports 95% confidence intervals for estimates of the annual CAPM alpha, beta, and average return of the CA PE index from 1983 to 2020. The plot also overlays how the consultant-perceived value of each performance metric as of 2021. Figure A11a shows that the consultant-perceived alpha of PE is lower than what is suggested by historical returns, though is still within the 95% confidence interval. This relative pessimism arises because the consultant-perceived beta of private equity $\bar{\beta}_{PE,t} \approx 1.2$ is higher than its historical counterpart (Figure A11b), while consultant-perceived expected returns to PE are lower than historical averages (Figure A11c). Together, the results in Figure A11 make it hard to claim that consultant beliefs are wildly optimistic compared to historical returns.

The same appears to be true when comparing consultant beliefs to estimates from the academic literature. Figure 1 from a recent survey of Korteweg (2023) illustrates the distribution of estimated β's across several studies. CAPM betas for buyout range from about 0.75 to 1.5, whereas those for venture funds generally lie between 1 to 2.75.[24] A perceived beta of $\bar{\beta}_{PE,t} \approx 1.2$ does not stand out in comparison to either asset class. Table 1 from an earlier review paper by Korteweg (2019) supports this conclusion, as do estimates from a recent paper by Brown et al. (2025) and a 2023 study on private equity authored by Norges Bank Investment Management (NBIM).[25]

[23] According to Preqin, the average AUM of buyout funds has been twice that of venture funds during the vintages that span sample.

[24] The benchmark in these studies varies but is usually tilted toward large-cap U.S. equities (e.g., S&P 500)

[25] Figure 9a of the NBIM study, which can be found here here, finds betas of VC and PE to be around 1.

In terms of PE alpha, the 2023 NBIM study mentioned above estimates *quarterly* CAPM alphas of around 1% for buyouts and essentially zero for venture (their Figure 9b). In the academic literature, a widely cited paper by Harris et al. (2014) finds that buyout funds from vintages 1984 to 2008 outperformed the S&P 500 by over 300 bps per year, whereas venture funds outperformed in the 1990s but not the 2000s. Thus, even a simple averaging of the buyout and venture from both sources yields alphas of 150-200 bps. Table 1 of Korteweg (2019) also shows that the consultant-perceived alpha of $\bar{\alpha}_{PE,t} \approx 180$ bps is within the very wide range of estimates from several other academic studies.

Real estate We examine the annual historical performance of the Cambridge Associates U.S. Real Estate Index from 1986 to 2020. Figure A11 shows 95% confidence bands for the estimated historical CAPM alpha, beta, and average return of this index, alongside consultant beliefs for real estate in 2021. For all three performance metrics, consultant perceptions are within the 95% confidence interval of historical estimates. For both alpha and beta, consultant perceptions line up fairly well with the historical point estimate, whereas consultants are relatively pessimistic about expected returns compared to historical averages.[26]

Hedge Funds For HFs, we start by evaluating consultant beliefs using the historical return of the Credit Suisse Hedge Fund Index from 1994 to 2020.[27] Figure A11 shows that the consultant beliefs in 2021 about the CAPM alpha, CAPM beta, and expected return of HFs were all within the 95% confidence band of historical estimates. More specifically, the alpha of $\bar{\alpha}_{HF} \approx 190$ bps perceived by consultants closely aligns with the historical point estimate, as does their perceive beta of $\bar{\beta}_{HF} \approx 0.35$. The consultant-perceived expected return on hedge funds, 420 bps, is actually relatively low compared to historical averages. These historical benchmarks are also consistent with Jurek and Stafford (2015), who analyze various methods for risk-adjusting hedge fund returns from 1996 to 2012. According to Table 1 of their study, the HFRI Fund Weighted Composite Index (HFRI) earned an annual return of 550 bps over this period, with a simple CAPM beta of 0.35 and an alpha of 360 bps.

Summary Overall, the preceding analysis does not immediately suggest that consultant beliefs about alpha are entirely unjustified based on academic research or historical returns, at least when using simple single-factor models for risk adjustment. However, this evidence should not be taken as evidence that consultant beliefs are rational. One can reasonably argue that the historical risk-adjusted returns of alternatives are overstated due to non-linearities in returns (Jurek and Stafford, 2015), return smoothing (Couts et al., 2020; Stafford, 2022), and omitted risk factors (Gupta and Van Nieuwerburgh, 2021). Rather than making a claim about the rationality of beliefs, we view this discussion as illustrating how consultants may justify their expectations in practice, as reasonable differences in benchmarking choices, methodologies, and sam-

[26] These results are not directly comparable to Brown et al. (2025), who compute betas and alphas of private real estate funds relative to a REIT index (not public equities).

[27] According to the Credit Suisse website, the Credit Suisse Hedge Fund Index is "is an asset-weighted hedge fund index and includes only funds, as opposed to separate accounts. The index uses the Credit Suisse Hedge Fund Database, which tracks approximately 9,000 funds and consists only of funds with a minimum of US$50 million under management, a 12-month track record, and audited financial statements. The index is calculated and rebalanced on a monthly basis, and reflects performance net of all hedge fund component performance fees and expenses."

ple periods can lead to widely divergent conclusions about alpha. This observation also helps explain why the alternative-to-risky share varies so widely across pensions—when risk and return are hard to pin down, it becomes easier for pensions to "agree to disagree".

E.3.2 Ex-Post Pension Performance

Under the view that beliefs about alpha drive the alternative-to-risky share, full rationality would imply that pensions with higher alternative-to-risky shares should earn higher realized alphas on their alternative investments. A decisive test of this proposition is extremely challenging in our setting for a few reasons. First, evaluating the performance of alternatives raises well-known difficulties around risk-adjustment, which remains an active area of debate in the academic finance literature (see e.g., Korteweg, 2019, 2023). Second, statistical power is limited by the fact that the PPD provides annual pension returns over only a 20-year sample. Beyond sample-size constraints, the illiquid nature of many alternative asset classes poses a challenge to traditional regression-based models of alpha (Jurek and Stafford, 2015; Couts et al., 2020; Stafford, 2022). Most modern approaches, including public market equivalents (Kaplan and Schoar, 2005) and generalized public market equivalents (Korteweg and Nagel, 2016), require more granular cash-flow data and are best-suited for large samples of funds, like those studied in Brown et al. (2025).

With those caveats in mind, we now present some suggestive evidence about the performance of pensions that vary in their use of alternatives. For each pension p that has at least twenty years of data from 2001 to 2021, we calculate the average alternative-to-risky share, denoted by $\bar{\Omega}_p$. Each pension p is then assigned to a quartile $q(p)$ based on $\bar{\Omega}_p$, with $q = 4$ corresponding to the top quartile. The average $\bar{\Omega}_p$ in quartile 1 is around 15% compared to 45% in quartile 4. For pensions in a given quartile q, we compare the alpha of their alternative portfolio with respect to the CRSP value-weighted index, total portfolio net-of-fee returns, the CAPM alpha of the total portfolio, and the portfolio Sharpe ratio, as shown in FigureA12.[28]

Figure A12a illustrates how the estimated alpha of alternatives $\alpha_{a,q}$ varies across quartiles of $\bar{\Omega}_p$. The bars in the plot show the constant in a regression of excess alternative returns $R_{a,p,t}$ on contemporaneous and lagged market excess returns, estimated separately by quartile.[29] The 95% confidence bands are based on standard errors that are clustered by year. The plot provides some evidence that high alternative users have earned larger alphas compared to low alternative users, though we cannot reject the null hypothesis that the two alphas are equal ($p = 0.70$). Moreover, $\alpha_{a,q}$ does not exhibit a clear monotonic pattern across the second, third, and fourth quartile.

Next, we examine the excess net-of-fee performance $\tilde{R}_{p,t}$ for the entire portfolio, measured using net investment income and beginning-of-year assets as in Andonov and Rauh (2022).[30] Figure A12b shows that

[28] Constructing the return on alternatives for each pension is complicated by the fact that returns to a specific type of alternative are not always populated in the PPD. In these cases, we impute missing returns using the aggregate return for a given asset class (e.g., PE). However, we do so separately for pensions based on whether their fiscal year ends in the middle or at the end of the calendar year. This adjustment accounts for differences in return measurement that are due to fiscal year timing.

[29] The inclusion of lagged market returns partly addresses illiquidity (Dimson, 1979; Korteweg, 2019), but is far from a perfect solution. More sophisticated approaches have been developed in recent years (e.g., Korteweg, 2019, 2023; Korteweg and Nagel, 2024), but these methods are considerably more complex and often rely on deal- or fund-level data. In contrast, regression-based estimates of alpha are straightforward to implement, have a long-standing tradition in finance, and are therefore more likely to be used by consultants as an input to their CMAs.

[30] Comparisons of performance across quartile q could be biased if, for instance, pensions in one quartile all report gross returns

low alternative users have, if anything, earned slightly higher average annual returns than high alternative users—a result largely driven by the sustained bull market in public equities following the Great Financial Crisis. Still, we cannot reject the null hypothesis that net-of-fee returns are equal across quartiles of alternative use.

A natural objection to the comparison in Figure A12b is that low and high alternative users may differ in their exposure to systematic risk. Figure A12c addresses this issue by estimating CAPM alphas for the entire portfolio within each quartile, again using contemporaneous and lagged market excess returns. The realized CAPM alpha does not meaningfully differ across quartiles of alternative use. And, unsurprisingly, we cannot reject the null that low and high alternative users have earned equal alphas on their total portfolios over the last twenty years.

Finally, Figure A12d plots the average Sharpe ratio of returns within each quartile of alternative use. The plot shows that Sharpe ratios do materially differ across quartiles of alternative use, with low alternative users having slightly higher Sharpe ratios. This simple fact also challenges the idea that alternatives are primarily used to conceal risk-taking (see Section 6.2).

In summary, we find little evidence that pensions with higher alternative-to-risky shares achieved superior performance on their alternative investments compared to those with lower allocations. Total returns, return volatility, and CAPM alphas also all exhibit little systematic variation across quartiles of alternative use. While these findings could be interpreted as consistent with a narrative in which pension beliefs are not rational, we are cautious about drawing strong conclusions given the relatively short time series and the well-known statistical challenges associated with estimating alpha in this setting (Couts et al., 2020; Stafford, 2022; Korteweg, 2023).

F Risk-Seeking Explanations: Additional Evidence

F.1 Robustness: The Alternative-to-Risky Share

Section 4.2 demonstrates a weak relationship between the alternative-to-risky share and proxies for risk-seeking motives (e.g., underfunding) in the cross-section of pensions. We now present several additional pieces of evidence confirming the robustness of this result.

F.1.1 Long-Run Changes: Model Misspecification

One potential issue with Table 7 is that it imposes a linear relationship between changes in the alternative-to-risky share and contemporaneous changes in funding, hurdle rates, and fund maturity. To directly assess this possibility, Figure A14 presents a series of scatter plots that show the raw data underneath the regressions in Table 7. These plots show no obvious non-linearities, further highlighting how challenging it is to explain the adoption of alternatives with changes in underfunding, fund maturity, or hurdle rates.

while those in another report net returns. To test for this possibility, we follow Andonov and Rauh (2022) and construct portfolio returns based on net investment income scaled by beginning-of-year portfolio value, denoted by $\tilde{R}_{p,t}$. The difference between $\tilde{R}_{p,t}$ and reported annual returns R_{pt} (variable = InvestmentReturn_1yr) serves as natural proxy for whether reported returns are in gross or net terms. Reassuringly, we cannot reject the null hypothesis that $g_{pt} = \tilde{R}_{p,t} - R_{pt}$ is, on average, equal across quartiles q ($p = 0.39$).

F.1.2 Post GFC

Reach-for-yield motives tend to be strongest during periods of low interest rates (Lian et al., 2018; Campbell and Sigalov, 2022). To give these motives the best chance to manifest empirically, we repeat the Table 7 analysis for the post-2008 period, shown in Table A9. We regress the change in each pension's target alternative-to-risky share between 2010 and 2021, $\Delta\Omega_p$, on the contemporaneous change in a proxy for reach-for-yield motives, ΔX_p. We cluster standard errors by state for system-level regressions and use robust errors otherwise.

Column (1) tests whether pensions that became more underfunded after the GFC (based on GASB 25 standards) increased their alternative-to-risky share. The positive β suggests the opposite: pensions with larger funding declines were less likely to shift toward alternatives. However, the effect is small and only marginally significant. The low R^2 further shows that portfolio changes are largely unrelated to funding levels.

Columns (3) and (4) test two other proxies for risk-seeking motives: changes in the nominal hurdle rate and in the fraction of retired members (discussed in Section 4.2). In both cases, the estimated β is statistically insignificant, and the R^2 values remain low.

Columns (5) and (6) examine whether funding levels in 2010 predict later shifts toward alternatives. Again, we find no evidence that changes in the alternative-to-risky share are driven by reach-for-yield motives in the post-GFC.

It is also important to note that these results do not contradict the findings of Mittal (2024) , who shows that newly committed capital to alternatives post-GFC mainly came from underfunded pensions. However, this fact does not necessarily imply that underfunded pensions increased their alternative-to-risky share more aggressively. To see why, consider two pensions, A and B, each with \$100 in assets. A is underfunded; B is well-funded. Initially, both allocate 70% to risky assets, and 30% of their risky portfolio to alternatives, meaning both invest \$21 (\$100 × 70% × 30%) in alternatives. Now suppose, A increases its risky share from 70% to 90%, while keeping its alternative-to-risky share constant at 30%. Pension B keeps its risky share at 70%, but raises its alternative-to-risky share from 30% to 32%. A's alternative investment rises to \$27 (\$100 × 90% × 30%), a \$6 increase, while B's alternative investment rises to \$22.4 (\$100 × 70% × 32%), a \$1.4 increase. Thus, A commits far more new dollars to alternatives than B (\$6 vs. \$1.4), even though B increases its alternative-to-risky share by more (2 percentage points vs. 0 for A).

The key takeaway from this example is that changes in the risky share of underfunded pensions could generate the flows documented by Mittal (2024), even if funding is largely disconnected from the alternative-to-risky share (as in Table A9). While clearly stylized, this example accords with the commonly held view that underfunding has led to increased risk-taking by pensions since the turn of the century (Eaton and Nofsinger, 2004; Pennacchi and Rastad, 2011; Mohan and Zhang, 2014; Andonov, Bauer, and Cremers, 2017; Lu, Pritsker, Zlate, Anadu, and Bohn, 2019).[31]

[31] We also find some support for this conclusion in Section F.3.

F.1.3 Long-Run Changes: Additional Hypotheses

Table A10 explores additional hypotheses regarding why some pensions have increased their alternative-to-risky share more than others. In particular, we explore whether the initial levels of the covariates used in Table 7 can predict subsequent changes in the alternative-to-risky share. For example, this tests whether the most underfunded pensions in 2002 were also those who subsequently moved to alternatives. Column (1) shows that, if anything, the opposite is true. The positive coefficient indicates that better-funded pensions in 2002 were slightly more likely to increase their alternative-to-risky share. Columns (2) to (4) show that none of the initial values of the other variables contain predictive power for changes in the alternative-to-risky share. In all cases, the R^2 are near zero.

Another possibility is that pensions who failed to make actuarial required contributions in the early 2000s were more likely to turn to alternatives. To test this idea, we group pensions into buckets based on the number of times they missed an actuarial required contribution between 2001 and 2005. We then compute the average change in the alternative-to-risky share for pensions in each bucket and plot them along with standard error bands in Figure A15a. The plot shows no discernible relationship between failure to make required contributions and the move to alternatives. In fact, pensions who failed to make these contributions shifted less aggressively to alternatives, though the standard errors are such that we cannot reject the null of no difference across the buckets.

Lastly, we test the hypothesis that larger pensions switched more aggressively into alternatives, perhaps because size lowers the cost of managing an alternatives strategy (e.g., Begenau and Siriwardane (2022)). To do so, we sort pensions into quintiles based on their size and then compute the average change in the alternative-to-risky share within each bin. in Figure A15b. The plots shows some evidence that the smallest pensions shifted less of their risky assets into alternatives, but the magnitude of this difference is relatively small. The standard errors also imply that we cannot reject the null that pensions in each size bin equally shifted to alternatives. Figure A15b also shows ethat size is not a particularly strong predictor of the average level of the alternative-to-risky share.

F.1.4 Panel Analysis of Levels

Our analysis thus far has focused on understanding changes in the alternative-to-risky share. However, the model in Section 2.3 further suggests the level of the alternative-to-risky share should be related to risk-seeking behavior (provided portfolio constraints bind). We explore this hypothesis using variants of the following panel regression:

$$\Omega_{p,t} = \alpha_t + \beta X_{pt} + \varepsilon_{pt} \tag{23}$$

where $\Omega_{p,t}$ is the target alternative-to-risky share of pension p in fiscal year t.[32] α_t in the regression is a time fixed effect, so the coefficient of interest β is not identified from common time-series variation across pensions. X_{pt} is the same set of explanatory variables used in Table 7, measured at time t. Standard errors are double-clustered by state and year.

[32]The fiscal year for most pensions ends in June, though there are several that end in December. These timing differences would matter more for studying cross-sectional variation in returns, though they matter less for studying variation in portfolio composition.

Table A11 summarizes the results. Columns (1) to (4) reveal no meaningful relationship between the level of the alternative-to-risky share and the level of GASB 25 funding (column 1), BEA-adjusted funding (column 2), hurdle rates (column 3), or the fraction of retired members (column 4). In addition, the within-R^2 in all cases are low and never exceed 1%. To explore possible non-linearities in the data, Figure A16 presents a series of binned scatter plots of the alternative-to-risky share against various covariates, after controlling for time fixed effects. The plots show that the poor model fits in columns (1) to (4) of Table A11 are not driven by our choice of a linear model.

The panel regression in Equation (23) weighs all years in our sample equally. Yet, for many policy questions, it is also useful to understand variation in the current portfolio structure of public pensions, especially because the aggregate portfolio appears to have reached a new steady state in recent years (Figure 1b). To this end, columns (5) to (8) of Table A11 run regression (23) only using data from 2021, the last year in which BEA-adjusted funding ratios are available. Standard errors for these columns are clustered by state for regressions run at the system level and are robust for those run at the state level.

The results in columns (5) to (8) echo those in columns (1) to (4). The point estimates continue to be statistically insignificant for all covariates and the model fits are also poor. The point estimates are also largely comparable in sign and magnitude with the exception of the hurdle rate (liability discount rate) in column (7). The negative sign in that column runs counter to intuition, as it suggests that pensions with higher hurdle rates are more likely to have lower alternative-to-risky shares in 2021. Figures A17a and A17b show that non-linearities are unlikely to bias these results.

F.1.5 Portfolio Constraints

In Section 4.1, we test whether binding portfolio constraints can explain why some pensions have a higher alternative-to-risky share than others. We measure the tightness of constraints as deviations l_{pt} of actual risky shares from target ones. Figure A18a plots the raw value of l_{pt} through time, both in aggregate and in the cross-section. In aggregate, the average value of l_{pt} during our sample equals -0.9, implying that the risky share of U.S. pensions tends to be below target. However, the negative in-sample mean is heavily influenced by the fall in risky asset prices during the 2008-2009 Global Financial Crisis. To strip out the effect of market fluctuations, we therefore first regress l_{pt} on contemporaneous 1-year returns. Figure A18b then visualizes the cross-section of the residuals in each year, as well as their aggregate value. The mean of the residualized aggregate series is close to zero (-0.0) and the time-series does not deviate from zero for long periods of time. These patterns weigh against the idea that U.S. pensions are constrained in aggregate.

F.2 Spending-Based Constraints

In Section 4.3 we analyze how the alternative-to-risky share responded to the Covid-19 pandemic, which plausibly reduced current and expected spending needs through its impact on mortality rates. The test is motivated by the idea that a pension's effective risk aversion is, at least in part, shaped by its spending obligations (e.g., Campbell and Sigalov, 2022). We find little evidence of a strong link between expected spending needs and the alternative-to-risky share. In this subsection, we present several complementary

pieces of evidence that reinforce the conclusion that spending needs play a limited role in explaining variation in alternative use.

F.2.1 Forecasting Spending and Spending Constraints

We start by testing whether the alternative-to-risky share predicts future spending (or spending-related constraints) across pensions. To motivate this test, suppose that, in reduced form, each pension's alternative-to-risky share depends linearly on its expected future spending:

$$\Omega_{p,t} = a + b\mathbb{E}_t[Spending_{p,T}] \tag{24}$$

where $Spending_{p,T}$ measures spending constraints at some later date $T > t$. Under rational expectations, $Spending_{p,T} = \mathbb{E}_t[Spending_{p,T}] + \varepsilon_{p,T}$ for an unpredictable noise term $\varepsilon_{p,T}$. This relation implies that equation (24) can be tested using forecasting regressions of the following form:

$$Spending_{p,T} = \alpha + \beta\Omega_{p,t} + \eta_{p,t},$$

where $\alpha = -\frac{a}{b}$, $\beta = 1/b$, and $\eta_{p,t} = \varepsilon_{p,t}$. Intuitively, if investment in alternatives today is driven by anticipated spending, then the alternative-to-risky share should forecast realized spending. A similar logic applies to anticipated cash-flow or spending constraints.

We regress the alternative-to-risky share in 2011 (half the sample), $\Omega_{p,2011}$, to forecast a measure of spending constraints in the second half of the sample. We consider two measures of spending constraints. Our first measure of the spending constraint is the contribution-to-payout ratio $CPR_{p,2012-2021}$ averaged over the second half of our sample (2012-2021). A lower contribution-to-payout ratio indicates greater reliance on investment income or asset sales to fund pension payouts, rather than incoming contributions from plan sponsors and current employees. This measure is analogous to the concept of "endowment dependence" in higher education, where spending policies exceed what can be funded through tuition revenue and donations. Our second measure that we forecast is the fraction of years between 2012 and 2021 in which plan sponsors failed to make their Actuarially Required Contribution (ARC). The ARC is the actuarially determined amount necessary to sustain the pension under GASB standards. Consistent failure to meet ARC payments is widely recognized as a signal of fiscal stress (Novy-Marx and Rauh, 2014).

Table A12a presents the results. Generally speaking, the alternative-to-risky share has very little forecasting power for future spending or spending constraints. In column (1), there is a small, yet statistically insignificant, correlation between $\Omega_{p,2011}$ and $CPR_{p,2012-2021}$. In column (3), the correlation between the alternative-to-risky share and future ARC payment delinquency is effectively zero. Columns (2) and (4) replace the alternative-to-risky share in 2011 with its change between 2002 and 2011 as the key regressor. These specifications test whether pensions that increased their alternative-to-risky share in the 2000s did so in anticipation of future financial stress. Once again, we find no evidence of such a link. We also obtain similar results when using different start years for measuring $\Omega_{p,year}$, and when using long-run changes in our funding stress measures as the outcome variable, rather than their average levels between 2012 and 2021.

F.2.2 Current Spending

A related hypothesis is that plan sponsors like the illiquidity of alternatives because it constrains current and future spending.[33] To test this proposition, we run the following set of panel regressions:

$$y_{p,t} = \lambda_t + \beta CPR_{p,t} + \varepsilon_{p,t},$$

where $y_{p,t}$ is the share of all alternatives, private equity (PE), or real assets (RA) relative to risky assets for pension p in year t. The first three columns of Table A12b show results, where columns (2) and (3) focus specifically on PE and RA in due to their inherent illiquidity. $CPR_{p,t}$ is the contemporaneous contribution-to-payout ratio of pension p. Low values of $CPR_{p,t}$ indicate that current spending is high, since there is more reliance on investment returns or asset sales to meet payout obligations. This measure is closely connected to the notion of "endowment dependence" for universities (Campbell et al., 2024). Across all three specifications, there is no economically or statistically significant relationship between alternative use and current spending.

Columns (4)-(6) of Table A12b show a complimentary regression of long-run changes in portfolio composition on contemporaneous changes in spending:

$$\Delta y_p = a + b\Delta CPR_p + \eta_p,$$

where long-run changes are computed from 2002 to 2021. This regression tests whether pensions that decreased their spending (i.e., increases in ΔCPR_p) were also those that shifted to alternatives. Across all three columns, we continue to find little evidence of a link between spending and alternative use.

F.3 The Risky Share

For completeness, we now repeat our analysis of the alternative-to-risky share for the overall risky share. We first investigate variation in long-run changes in the risky share and then study factors associated with its level in the cross-section of pensions.

F.3.1 Long-Run Changes

To start, we analyze long-run changes in the risky share using the following linear regression:

$$\Delta\omega_{risky,p} = a + \beta\Delta X_p + \varepsilon_p$$

where $\Delta\omega_{risky,p}$ is the change in the risky share for pension p. The explanatory variable ΔX_p is the contemporaneous change in the variables considered in Section 4.2. In some cases, X_p is available only at the state level, in which case the regression is based on state-level changes. All changes are measured between 2002 and 2021. Standard errors are clustered by state and fiscal year for regressions run at the system level and are otherwise robust. Table A13 summarizes the estimation results.

[33] We thank our editor, Tarun Ramadorai, for making this suggestion.

Column (1) of the table investigates the relationship between changes in the risky share and changes in GASB 25 funding ratios. The point estimate is close to zero and not statistically different from zero at conventional confidence levels. Column (2) shows that the weakness of this relationship is not an artifact of overinflated GASB 25 funding ratios (Brown and Wilcox, 2009; Novy-Marx and Rauh, 2011). The regression in column (2) is run at the state level and uses BEA-adjusted funding ratios, which use more appropriate liability discount rates based on AAA-rated corporate bond yield curves. The negative point estimate does indicate that pensions who became more underfunded did increase their risky share, but its magnitude is small and not statistically discernible from zero.

In column (3), we evaluate the view that increases in the risky share are driven by high nominal hurdle rates. Consistent with this mechanism, there is a positive and statistically significant relationship between changes in the risky share and hurdle rates in the cross-section of pensions. The point estimate implies that moving from the 10th to 90th percentile of changes in hurdle rates is associated with an increase in the risky share of 7.3 pp. For reference, over this period, the 10th percentile pension decreased its discount rate by 0.3 pp whereas the 90th percentile decreased it by 1.6 pp.

In column (4), we instead proxy for risk-seeking motives using the fraction of retirees. The motivation for this proxy comes from Andonov et al. (2017), who argue that U.S. pensions have an increased incentive to take on risk as their members retire because of how liabilities are discounted. Consistent with Andonov et al. (2017), there is a positive and statistically significant relationship between changes in the risky share and changes in the fraction of retirees. However, the relationship is fairly weak in economic terms, as the point estimate implies that moving from the 10th to 90th percentile of changes in the fraction of retired members predicts an increase in the risky share of 3.3 pp. Within the regression sample, the actual difference between the 10th and 90th percentile of changes in the risky share is 29 pp.[34]

Columns (5) to (8) instead predict changes in the risky share between 2002 and 2020 with the characteristics in 2002, such as the level of funding. Columns (5) and (6) find no evidence that the most underfunded pensions at the turn of the century subsequently increased their risky share. Column (7) indicates that those with lower hurdle rates in 2002 were actually more likely to increase their risky share, inconsistent with the idea that high nominal hurdle rates induce pension risk-taking in the cross-section, but this relationship is not significant. Column (8) shows that pensions with a lower initial share of retirees were also more likely to increase their risky share. In both cases, though, the point estimates are also not economically meaningful.

Another noteworthy feature of this analysis is that the R^2's are uniformly low across all specifications in Table A13, with all well below 10%. Figure A19 shows further these poor model fits are not driven by non-linearities in the data. These results suggest that funding- and accounting-based explanations can only partially explain why U.S. pensions have taken on more risk since the turn of the century. One possibility is that these factors were more important in the 1980s and 1990s, which saw a much larger increase in the risky share compared to the one experienced after 2000 (see Section 2.2.1).

[34] Our results are somewhat weaker than those found in Andonov et al. (2017), who use a different regression specification and study U.S. pensions between 1990 and 2012. Part of this difference may be driven by the fact that U.S. pensions dramatically increased their risky share during the 1990s (see Section 2.2.1), but have slowed the pace of increase in the risky share since the 2000s.

F.3.2 Discussion

Overall the results in this subsection suggest that cross-pension variation in the risky share is weakly correlated with measures of funding, hurdle rates, and fund maturity. These results are somewhat surprising given the commonly held view that pensions are reaching-for-yield because the decline in interest rates has worsened their funding or made it harder to clear nominal hurdle rates.

To understand these results better, it is helpful to consider the precise mechanism through which underfunding would lead to reach-for-yield behavior in the context of U.S. public pensions. The premise for this idea is that public pensions and corporate plan sponsors have similar investment incentives. Corporate incentives to reach-for-yield due to underfunding are discussed and studied extensively in Rauh (2008). In the corporate setting, a plan sponsor (i.e., the shareholders) is essentially a debtor to pension beneficiaries. Consequently, the agency theory introduced by Jensen and Meckling (1976) suggests that it is optimal for the sponsor to take more risk when the pension becomes underfunded. Sharpe (1976) and Treynor (1977) argue that this moral hazard or risk-shifting problem is exacerbated in the United States by the existence of the federal Pension Benefit Guaranty Corporation (PBGC), which acts as an insurer for corporate pensions. In effect, the PBGC sells the plan sponsor a put option on the value of the pension assets. However, because the premium paid to the PBGC does not depend on asset risk, the plan sponsor can maximize the net value of its put option by increasing the risk of pension assets.

Rauh (2008) also points out several forces that may offset the aforementioned risk-shifting incentives. For one, in the United States, corporate sponsors of underfunded pensions must make contributions according to legally specified formulas that are generally increasing in the size of the funding gap (Rauh (2006)). In the presence of external financing costs, these mandatory requirements may in turn divert resources from valuable investment opportunities (Froot, Scharfstein, and Stein (1993)), thereby incentivizing plan sponsors to decrease asset risk as a pension becomes underfunded. These types of risk-management considerations may also arise if there are dead-weight costs to bankruptcy or if taxes are a convex function of earnings (Mayers and Smith (1982); Smith and Stulz (1985)).

In the context of public DB pensions, tax payers are the most natural analogue to corporate plan sponsors or equity holders. However, unlike corporate equity, tax payers are not protected by limited liability in the event of pension default. For example, if a state-sponsored pension were to default on its obligations, the state must either raise taxes, reduce other public spending, or cut pension benefits, all of which are costly. Taxes are distortionary for the usual textbook reasons (e.g., through their effect on labor markets) and therefore act as a form of costly external finance.[35] Cutting pension benefits may also distort labor markets and likely comes with political costs, especially when age is positively correlated with voter turnout. This line of argument suggests that risk-management may outweigh risk-shifting incentives for public plan sponsors.

On the other hand, risk-shifting motives may dominate if public plan sponsors believe that they will be bailed out by the federal government in the event of default (Biggs, 2021). In this case, the federal government creates investment incentives for public pensions much like the PBGC does for corporate pensions.

[35] Reductions in public investment may also reduce welfare for policies that have positive multiplier effects (e.g., Hendren and Sprung-Keyser, 2020).

Another important consideration is the alignment between public plan sponsors and pension managers. In the corporate setting, the contracting environment is arguably designed to align the incentives of shareholders and managers (e.g., equity-linked compensation), though these factors are likely weaker for public pension mangers. Indeed, to the extent that pension managers have compensation contracts that are convex in the assets of the pension, their incentive to risk-shift will be larger compared to tax payers. Understanding the primary forces driving the overall risky share for U.S. public pensions is an important and interesting avenue for future research.

G Aggregate Implications

G.1 Risk-Seeking and Portfolio Constraints

In Section 5.1 we simulate how a decline in risk aversion coupled with a binding portfolio constraint on fixed income impacts pension behavior in the model of Section 2.3. To do so, we first assume that the aggregate pension portfolio was unconstrained in 2001 but then a minimum constraint on fixed income becomes binding by 2020. This means that the observed fixed income share in 2020 is the minimum required share ω_f^{min}. We then draw a random set of beliefs from the following uniform distributions: (i) expected excess returns on public equities, $\mu_E \sim U(0.02, 0.07)$; (ii) the variance of excess equity returns, $\sigma_E^2 \sim U(0.02, 0.09)$; (iii) risk-adjusted expected returns to alternatives relative to equities, $\alpha \sim U(0, 0.05)$; (iv) the CAPM- beta of alternatives relative to public equities, $\beta \sim U(0, 1.5)$. Realized excess returns on alternatives are therefore given by:

$$r_A - r_f = \alpha + \beta(r_E - r_f) + \eta_A,$$

where σ_η^2 is the idiosyncratic variance of alternatives. Together, these belief parameters fully define the expected excess return and variance-covariance matrix that determine optimal asset allocation. For each parameter draw, we set σ_η^2 to match the aggregate alternative-to-risky share Ω in 2001 according to Equation (1) and the relationship $\Omega := \omega_A/(\omega_A + \omega_E)$. We further restrict the set of beliefs S^* to imply positive idiosyncratic variance $\sigma_\varepsilon^2 > 0$, $\gamma_1 \geq 1$, and $\sigma_A^2 \leq 0.25$. We then draw random belief sets until $S^* = 100{,}000$. Figure A20 shows the distribution of belief parameters in the set S^*.

Given this set of admissible initial beliefs, we then assume the portfolio becomes constrained in 2021 and solve for the new risk aversion γ_{2021} needed to generate the observed $\Delta\Omega$ in the data. Panels A and B in Table 9 summarize the simulation approach. In all cases, we impose a constraint that alternatives cannot be shorted, meaning $\omega_A \geq 0$. The headline result is that in the vast majority of cases, a decline in risk aversion cannot match the observed increase in the alternative-to-risky share. The basic intuition is that by revealed preference, the low alternative-to-risky share in 2001 means that pensions would tilt towards equities over alternatives when facing risk constraints.

Of the 100,000 simulations, 484 are such that a decline in risk aversion generates an increase in the alternative-to-risky share, but not enough to match the data before risk aversion hits its lower bound of 1. For roughly half of these simulations, the alternative-to-risky share does not exceed 30%, quite far from the actual alternative-to-risky share of 39% in 2020.

Pane C of Table 9 does show, however, that a decline in risk aversion can generate the aggregate rise in the alternative-to-risky share in 0.4% of simulations (379 out of 100,000). Figure A21a indicates that the implied average reduction in risk aversion, $\gamma_{2001} - \gamma_{2021}$, is about 4.51 for these cases. For this subset of simulations, we calculate the shadow cost of the binding risk constraint. To do so, define $M = \mathbb{E}[r_p] + \frac{1}{2}(1-\gamma)\mathbb{V}[r_p]$ as the function that investors maximize. The shadow cost of the portfolio constraint is then given by:

$$ShadowCost = M_{counterfactual} - M_{portfolio}, \tag{25}$$

where $M_{portfolio}$ is the utility of the constrained portfolio in 2021 and $M_{counterfactual}$ is the utility that pensions would receive in the absence of the risk constraint. The shadow cost can therefore be interpreted as the fee that pensions would be willing to pay to relax the constraint, increase their risky share, and invest in their perceived tangency portfolio. The unit of this shadow cost is the same as the log return. Figure A21b presents the distribution of the shadow cost across the 379 simulations in which risk-seeking motives can generate the observed $\Delta\Omega$. The average shadow cost across these simulations is 732 basis points, which we view as implausibly large given that asset hurdle rates for most public pensions are around 700 basis points. This is especially true given the simulations assume a fairly high minimum fixed income constraint of roughly 24%, the aggregate level in 2021.

G.2 Simulating a Pure Belief Shift

G.2.1 Full Simulation Results

In Section 5.2, we use the mean-variance model to simulate how much pension beliefs about alpha would need to shift in order to fully explain the aggregate rise in the alternative-to-risky share. In each simulation, we first shift beliefs about α and then solve for the risk aversion γ_{2021} needed to match the risky share in 2021. Interestingly, absent a change in risk aversion, the shift in α can also generate the bulk (55%) of the observed rise in the *risky* share as well. To demonstrate this more clearly, Figure A22b shows both the initial distribution of risk aversion γ_{2001} needed to match the risky share in 2001 and the new γ_{2021} distribution needed to match the 2021 risky share. The first takeaway from the plot is that the two distributions are very close to each other. Thus, after allowing for shifts in beliefs, only small changes in risk aversion are needed to match the observed increase in the aggregate risky share.

To understand the reason behind this result, notice that the risky share also depends on α in the model:

$$\frac{\partial \omega_E + \omega_A}{\partial \alpha} = \frac{1}{\gamma}\frac{1}{\sigma_\eta^2}(1-\beta)$$

This is because the optimal holding of the risky portfolio depends on its Sharpe ratio, which is impacted by α. If $\beta < 1$, an increase in α will also lead to an increase in the overall risky share. Figure A20 shows that this condition holds for virtually all of our simulations. Moreover, the median consultant's perceived beta from CMAs has never exceeded 0.75 over the last two decades. Because the risky share meaningfully rises with α for most simulations, risk aversion does not need to change by much to match the observed increase in the risky share.

The second thing to note from Figure A22b is that risk aversion must actually increase slightly to match the aggregate risky-asset share in 2021 This is because the rise in α induces a rise in the simulated risky share that is larger than the observed one. A decomposition within each simulation indicates that on average, 55% of the increase in the risky share can be generated by beliefs, with the remaining 45% coming from risk aversion.

Figure A23 conducts a complimentary exercise in which we increase the perceived alpha $\Delta\alpha$ by a fixed amount and compute the distribution of the implied change in the alternative-to-risky share $\Delta\Omega(\Delta\alpha)$ across simulations. This allows us to assess how the alternative-to-risky share responds to varying increases in perceived alpha, including those derived from estimates in the academic literature (e.g., Korteweg, 2019). The horizontal dashed line in the plot represents the actual aggregate change in the alternative-to-risky share from 2001 to 2021. The main takeaway from the plot is that, depending on initial beliefs in 2001, an increase in perceived alpha ranging from $\Delta\alpha =$ 20-70 bps can reasonably account for the full rise in the aggregate alternative-to-risky share.

G.2.2 Incorporating Consultant Effects

To incorporate the causal impact of consultants into each simulation, we first express the pension sector's total perceived alpha α_p as:

$$\alpha_p = \zeta\,\alpha_c + \xi_p, \tag{26}$$

where α_c is the consultant sector's perceived alpha and ξ_p is the pension sector's private belief about alpha. ζ is the causal impact of consultant beliefs on pension beliefs. To infer ζ in a given simulation, we leverage the fact that:

$$\begin{aligned}\frac{\partial\Omega}{\partial\alpha_c} &= \frac{\partial\Omega}{\partial\alpha_p}\times\frac{\partial\alpha_p}{\partial\alpha_c}\\ &= \Lambda\times\zeta,\end{aligned} \tag{27}$$

where $\Lambda \equiv \partial\Omega/\partial\alpha_p$ comes from the model. Equation (3) shows that Λ depends on beliefs and risk aversion, meaning it can be computed given the initial parameters governing each simulation.

The left-hand side of the expression in (27) is the elasticity of the alternative-to-risky share to consultant beliefs, which we denote as b_1 in the main text. This parameter is pinned down by our IV strategy in Section 3.1 and is roughly $b_1 \approx 4$. Thus, in each simulation, $\zeta = b_1/\Lambda$. The distribution of ζ has a mean of 0.138 and a median of 0.075 in all simulations

Armed with a ζ, it is straightforward to compute the implied change in private beliefs needed to match the alternative-to-risky share. First, we compute the total required change $\Delta\alpha_p$. Then, we subtract out $\zeta\Delta\alpha_c$, where α_c is taken directly from Figure 4a ($\Delta\alpha_c \approx 60$ bps). The remaining component $\Delta\xi_p$ is the required change in pension's private beliefs to match the shift in risky portfolio weights from 2001 to 2021. Figure A24 shows the distribution of $\Delta\xi_p$ from the simulations. On average, private beliefs add about 60 bp to pension-perceived alpha, and the consultant's beliefs add about 10 bp ($\approx 0.138\times 60$ bp), totaling a 70 bp increase in alpha.

G.2.3 Peer Multipliers

In Section 3.3, we documented the existence of peer effects: pensions are more likely to allocate risky investments to alternatives when other nearby pensions do, even after controlling for the effect of shared consultants. How do the peer effects we document in the cross-section impact the aggregate alternative-to-risky share in equilibrium? To make progress on this question, notice that the peer effects regression in Equation (8) can be written in vector form as:

$$\begin{aligned}\boldsymbol{\Omega}_t &= \beta_z \boldsymbol{D}\boldsymbol{\Omega}_t + \tilde{\mathbf{X}}_t\boldsymbol{\Gamma} + \boldsymbol{\varepsilon}_t \\ &= (\boldsymbol{I} - \beta_z \boldsymbol{D})^{-1}\tilde{\mathbf{X}}_t\boldsymbol{\Gamma} + (\boldsymbol{I} - \beta_z \boldsymbol{D})^{-1}\boldsymbol{\varepsilon}_t\end{aligned}$$

where $\boldsymbol{\Omega}_t$ is a $P \times 1$ vector of alternative-to-risky shares and $\boldsymbol{I}$ is the $P \times P$ identity matrix. $\boldsymbol{D}$ is a $P \times P$ matrix for which element $D_{i,j} = d_{ij}^{-1} \times \left(\sum_k d_{i,k}^{-1}\right)^{-1}$ is the inverse distance between pension i and j, normalized so that the rows of D sum to 1 and $d_{i,i}^{-1} = 0$. $\tilde{\mathbf{X}}_t$ is a $P \times K$ matrix of controls and fixed effects used in the regression and $\boldsymbol{\Gamma}$ is the associated $K \times 1$ vector of regression coefficients. β_z captures the strength of peer effects and we find it to be roughly $\beta_z \approx 0.4$ for pensions that have little incentive to herd as in Scharfstein and Stein (1990). We therefore interpret it as measuring the strength with which beliefs about alternatives propagate through the peer network defined by $\boldsymbol{D}$. $\boldsymbol{\varepsilon}_t$ is a vector of idiosyncratic shocks to the alternative-to-risky share that can also be interpreted as being driven by beliefs.

The Leontief inverse matrix $\boldsymbol{\Psi} = (\boldsymbol{I} - \beta_z \boldsymbol{D})^{-1}$ is the key object describing how the peer network impacts the alternative-to-risky share in equilibrium. For a given idiosyncratic shock vector $\boldsymbol{\varepsilon}_t$, the equilibrium impact on the alternative-to-risky share of all pensions equals $\boldsymbol{\Psi}\boldsymbol{\varepsilon}_t$.[36] Similarly, the aggregate impact of a given set of idiosyncratic belief shocks is given by $\boldsymbol{s}_t'\boldsymbol{\Psi}\boldsymbol{\varepsilon}_t$, where s_t is a $P \times 1$ vector of weights based on each pension's AUM. We define the peer multiplier for a given shock vector as:

$$M(\boldsymbol{\Psi}, \boldsymbol{\varepsilon}_t) = \frac{\boldsymbol{s}_t'\boldsymbol{\Psi}\boldsymbol{\varepsilon}_t}{\boldsymbol{s}_t'\boldsymbol{\varepsilon}_t}.$$

The numerator in $M(\boldsymbol{\Psi}, \boldsymbol{\varepsilon}_t)$ is the aggregate impact of shocks after accounting for peer effects, whereas the denominator is the aggregate impact of shocks ignoring the peer network. $M(\boldsymbol{\Psi}, \boldsymbol{\varepsilon}_t)$ is thus a measure of how much the peer network amplifies shocks in aggregate, hence why we call it a peer multiplier.

To develop a sense of the magnitude of $M(\boldsymbol{\Psi}, \boldsymbol{\varepsilon}_t)$, we study an isolated shock to the alternative-to-risky share of CalPERS, the largest U.S. pension by assets under management. That is, we set the entry for CalPERS in $\boldsymbol{\varepsilon}_t$ equal to one and set all other values to zero. The resulting peer multiplier $M(\boldsymbol{\Psi}, \boldsymbol{\varepsilon}_t)$ for CalPERS equals 1.23. Roughly speaking, this number implies that peer effects account for 19% ($= (1.23 - 1)/1.23$) of the aggregate rise in the alternative-to-risky share.

[36] This analysis implicitly assumes that the supply of alternatives is not fixed, so that an increase in the alternative-to-risky share of one pension need not require a reduction by another.

H Other Explanations

H.1 Agency-Based Explanations

H.1.1 Reporting Benefits

In this section, we further discuss whether the return-smoothing properties of alternative investments can explain their adoption among U.S. public pensions.

Ex-Post Performance To start, we analyze the realized volatility of pensions in different asset classes to evaluate the scope for the return-smoothing mechanism. Specifically, for each alternative asset class $a \in \{\text{PE, RE, HF}\}$, we retain observations where pensions' reported returns in the PPD data are simultaneously available for both a and public equities. For each pension p, we then compute the annual volatility of alternatives and public equities, using identical sample periods to ensure consistency in the comparison.

Table A14a reports the median volatility by asset class and shows that indeed alternative asset classes exhibit lower realized volatility than public equities. The volatility difference is relatively modest between private and public equities, but is more pronounced for hedge funds and real estate. This result leaves room for an interpretation based on reporting benefits. Under this view, pensions may not actually believe that alternatives offer superior risk-adjusted returns, but invest in them anyways because they believe doing so will artificially enhance the reported Sharpe ratio and CAPM alpha of their portfolio.

This proposition is directly testable. For each pension p that has complete data from 2002 to 2021, we calculate the average alternative-to-risky share, denoted by $\bar{\Omega}_p$. Based on the computed $\bar{\Omega}_p$, we sort pensions and assign them to quartiles, $q(p)$, with the $q = 4$ corresponding to the top quartile of $\bar{\Omega}_p$. Figure A12d plots the average Sharpe ratio across pensions in each quartile. The interesting result is that pensions with lower alternative-to-risky shares exhibit slightly larger Sharpe ratios than those with higher shares. Figure A12b further shows that average net-of-fee return is also similar across quartiles, implying that portfolio volatility must also be similar.

At first glance, this result seem at odds with Table A14a, which shows that alternatives tend to have lower volatility than public equities. However, these findings are not in conflict for a few reasons. First, the divergence in public equity exposure between low- and high-alternative users became more pronounced following the Global Financial Crisis (GFC)—precisely when public equities began a prolonged period of strong performance. Pensions that shifted more aggressively into alternatives during this time therefore missed out on much of the post-GFC equity bull market. Second, the correlations between public equities, fixed income, and alternatives are higher for pensions with greater alternative exposure, limiting the diversification benefits of combining the three. Third, the fixed income and equity holdings of low-alternative users tend to have lower volatility. Table A14b, which is based on sample averages for low and high alternative users, shows how these latter two factors lead to very similar total portfolio volatilities across the two groups.

Of course, this analysis is based on ex-post realized returns, and it is possible that the pensions who shifted aggressively into alternatives for return-smoothing purposes simply experienced bad luck ex-post. Still, we think it is notable that if return smoothing was indeed the motive, the strategy has not led to obvious

measured outperformance at the portfolio level.

Accounting Reasons to Smooth As we discuss in the main text, return-smoothing incentives may explain the aggregate rise of alternatives. However, our results put fairly strong restrictions on their ability to explain the cross-section of pension behavior. More specifically, to rationalize the data, return-smoothing motives would need to vary meaningfully across pensions in ways orthogonal to observable characteristics, such as size, funding status, discount rates, contribution requirements, spending rates, beneficiary composition, board structure (Andonov et al., 2018), home bias (Hochberg and Rauh, 2012), and state-level trends effects. Moreover, given the wide dispersion in alternative adoption (Figure 3b), any such incentive would also need to have strengthened over time for some pensions but not others.

GASB accounting rules provide a possible microfoundation for such as a return-smoothing incentive. Before 2014, pensions could smooth market asset fluctuations over a chosen window when calculating actuarial assets, directly affecting reported funding ratios and required contributions. After 2014, GASB required market valuations, removing this smoothing option.[37] Most pensions (~60%) used a five-year window, while others use no smoothing or up to a ten-year window. These "asset smoothing" policies were typically set decades before the rise of alternatives and were rarely updated during our sample period. From 2001 to 2014, over half of pensions made no changes to their smoothing windows, and 75% changed them at most once. We therefore treat smoothing windows as largely exogenous to alternative allocations.

In the cross-section, we can use this fact to test whether asset smoothing policies impact the alternative-to-risky share. Specifically, we run the following panel regression:

$$\Omega_{p,t} = \lambda_t + \beta SW_{p,t} + \varepsilon_{p,t},$$

where $\Omega_{p,t}$ is pension p's alternative-to-risky share in year t, and $SW_{p,t}$ is its smoothing window. Standard errors are clustered by state and time. Since pensions with shorter smoothing windows presumably have larger incentives to smooth their returns using alternatives, it seems natural to expect that $\beta > 0$. However, prior to 2014, we estimate a small and statistically insignificant slope of $\beta = 0.003$ ($t = 0.440$).

To alleviate concerns about the joint determination of $\Omega_{p,t}$ and $SW_{p,t}$, we can also instrument $SW_{p,t}$ with its initial level in 2001, assuming of course that smoothing windows in 2001 were not driven by anticipated changes in alternative use. The first-stage F-statistic is well over 70, confirming the relevance of the instrument. The IV estimate is $\beta = -0.000$ ($t = -0.013$). Thus, there is no evidence that variation in smoothing windows materially influenced alternative investments in the cross section.

Real Assets vs. Private Equity Taking a step back, we recognize the inherent challenge of determining whether stated beliefs reflect one's true convictions or are instead driven by an underlying agency friction. Nevertheless, several pieces of evidence help to at least to rule in a meaningful role for genuine beliefs. For one, as we document in Section 2.2.4, pensions differ substantially in how they allocate within alternative

[37] If pensions prefer stable funding ratios and ARCs for budget-planning purposes, this shift should have increased incentives to use alternatives for return smoothing after 2014. However, aggregate evidence from Figure 2c indicates the opposite: the largest increases in the alternative-to-risky share occurred prior to 2012, when the new GASB rules were first announced, and growth slowed markedly thereafter.

asset classes, even those that are similarly illiquid. This degree of within-alternative variation is hard to reconcile with a generic return-smoothing motive and is more naturally explained by differences in beliefs.

More directly, this heterogeneity lines up closely with consultant beliefs. We establish this fact using the following panel regression:

$$\Omega_{a,p,t} = f_{p,t} + \beta V_{c(p),a,t} + \varepsilon_{p,a,t}, \tag{28}$$

where $\Omega_{a,p,t}$ is the target weight (relative to risky assets) of pension p in alternative type a at time t and $V_{c(p),a,t,}$ is the contemporaneous alpha of a reported by its consultant $c(p)$.[38] $f_{p,t}$ is a pension-by-time fixed effect. The fixed effect is particularly useful because it absorbs any time-varying pension characteristics, including risk aversion and a general desire to invest in unmarked assets, thereby isolating the link between consultant beliefs and portfolio choice. Standard errors in the regression are double-clustered by consultant and pension.

Figure A25 visualizes Regression (28) using a binned scatter plot of $\Omega_{a,p,t}$ against $V_{c(p),a,t}$, after absorbing a pension-by-time fixed effect. The estimated coefficient from Regression (28) equals 1.92 (t =6.83) and indicates that pensions tend to allocate more toward private equity or real assets when their consultant reports it to have a high alpha.[39] The within-R^2 further shows that 15% of within-pension variation in the composition of private-market investments (real assets versus private equity) is explained by consultant beliefs.

H.1.2 Board Composition

Andonov et al. (2018) document that the presence of state officials on pension boards is correlated with poor performance of private equity investments. They attribute this result to agency frictions, namely that state officials support private equity investments that benefit them politically or personally. Motivated by this evidence, we now test whether board composition is correlated with alternative use and adoption in the cross-section. The mechanism we have in mind is that boards with higher state representation have higher private incentives to conceal risk via alternatives, as this allows pension management and officials to avoid public scrutiny (Dyck et al., 2018). The first panel regression we run is:

$$\Omega_{p,t} = \alpha_t + \Gamma B_{p,t} + \beta X_{pt} + \varepsilon_{pt},$$

where $\Omega_{p,t}$ is pension p's target alternative-to-risky share, $B_{p,t}$ is one of several measures of board composition, and $X_{p,t}$ is a vector of controls. α_t is a time fixed effect. The data we use for board composition is taken from Aleksandar Andonov's website, which update the data from Andonov et al. (2018) through 2020. Pensions in these data are identified by their Preqin ID, which we map by hand to the pension systems in our data. Standard errors are clustered by time and pension system, as board composition is relatively static through time.

[38] Because beliefs about infrastructure not consistently available during our sample, we proxy for them using beliefs about real estate. For the same reason, we proxy for private debt beliefs using private equity. Recall that the definition of private equity in the PPD includes private equity and private debt, and our definition of real assets includes real estate and commodities.

[39] A p-value of 0.001 is obtained when using a wild bootstrap to test if $\beta = 0$ (Cameron et al., 2008).

Table A15 summarizes the results. Column (1) shows estimates when including no additional controls and column (2) shows results when adding controls for GASB 25 funding, each system's asset hurdle rate (liability discount rate), log assets, required contributions scaled by payroll, total spending scaled by payroll, and overall board size. The punchline is that board composition has little to no correlation with the alternative-to-risky share in the cross-section, as all of the reported point estimates are economically small and statistically insignificant. Moreover, the low within-R^2 in both specifications indicates that most of the cross-sectional variation in the alternative-to-risky share is orthogonal to board composition.

H.1.3 Home Bias

A related agency-based explanation for the rise of alternatives stems from Hochberg and Rauh (2012), who show that U.S. public pensions exhibit home bias when investing in private equity. They interpret this home bias as reflecting an agency problem, mainly because it tends to be associated with poor performance. Here, we test whether a similar mechanism can account for cross-state variation in the rise of alternatives. We do so via the following state-level regression:

$$\Delta\Omega_{s,2002\to2021} = a + b\text{Alt Activity}_s + \varepsilon_s, \tag{29}$$

where $\Delta\Omega_{s,2002\to2021}$ is the change in state s's target alternative-to-risky share from 2002 to 2021 andAlt Activity_s is one of two measures capturing the extent to which alternative assets are domiciled in state s. The first measure is the fraction of U.S.-based private capital funds domiciled in state s from vintages 2000 to 2022. The second measure is the simply the fraction of AUM domiciled in each state over the same period, as opposed to total counts. To construct these variables, we downloaded the "Fund Details" file from the WRDS version of Preqin and merged it with the "Managers Detail" file. These datasets contain information on fund size, vintage, and the state in which each general partner (GP) is domiciled, primarily for private equity and venture capital funds. We restrict the sample to funds managed by U.S.-based GPs with vintages between 2000 and 2022. The slightly wider vintage window (relative to the one used for constructing $\Delta\Omega_{s,2002\to2021}$) allows for the fact that Preqin may define vintage years slightly differently than how pension funds define fiscal years. For each state s, we then compute the fraction of total funds raised or total AUM domiciled in that state.

Table A16 summarizes the results. Both columns of the table show a weak relationship between states with more alternative assets domiciled and the increase in the alternative-to-risky share. The point estimates are small, statistically insignificant, and the R^2 are low. Thus, we conclude that home bias is unlikely to explain cross-state variation in the rise of alternatives.

H.2 Supply-Side Considerations

In Section 6.1, we explore whether a desire to hold the global market portfolio (GMP) can explain the aggregate rise of alternatives. The GMP used in this analysis is sourced from State Street.[40] To facilitate a direct comparison to the PPD data, we adjust the weights in the GMP in a few ways. First, we con-

[40] https://www.ssga.com/us/en/institutional/insights/global-market-portfolio-2024

solidate convertibles, emerging-market debt, government bonds, high-yield bonds, investment-grade credit, inflation-linked bonds, and securitized instruments into a single category called fixed income. Second, we reallocate the market capitalization of real estate investment trusts (REITs) from real estate and to public equities. The market capitalization of REITs is sourced from Nareit, a U.S. industry group that includes REITs and other real estate companies.[41] Because Nareit's data covers only the U.S. market, the adjusted weight of real estate in the GMP is likely overstated, while the size of public equities is likely understated. Third, we introduce a new category in the GMP to reflect hedge funds. Conceptually, it is not entirely clear that hedge funds should be treated as a distinct asset class, as they primarily invest in public equities and fixed income markets. Nevertheless, we treat hedge funds as a separate category to align with the PPD classification. To do so, we calculate the AUM of equity-focused hedge funds using data from Hedge Fund Research (HFR) and subtract this amount from the total market capitalization of public equities in the GMP. We repeat the same process for fixed-income hedge funds. We classify hedge funds as equity-focused based on whether their broad strategy in HFR is listed as "Equity", "Equity Hedge", and "Fund of Funds", while fixed-income hedge funds encompass all other strategies. This classification likely understates the true size of public equities and overstates the global alternative-to-risky share, as fund-of-funds may not be entirely equity-focused in practice.

The analysis in the main text shows that, in aggregate, the risky portfolio of U.S. pensions has become less diversified with the rise of alternatives. A similar story holds in the cross-section of U.S. pensions. In 2002, 59% of U.S. public pensions were underweight alternatives. By 2021, this fraction had declined sharply to 16%, indicating that the vast majority of pensions became overweight alternatives. Here, we classify a pension as being underweight alternatives if its target alternative-to-risky share is at least 1 pp less than the contemporaneous alternative-to-risky share in the GMP. Consistent with the aggregate trends discussed above, the degree of this overweighting is substantial. Among pensions that were overweight alternatives in 2021, the average deviation of their target alternative-to-risky share from the GMP was 19 pp. Figure A26b further illustrates this point by plotting the percentage of U.S. pensions for which the alternative-to-risky share exceeds the GMP by at least 10 pp ("heavily overweight"). In 2002, only about 10% of pensions were heavily overweight alternatives, compared to over 60% in 2021.

Additionally, there is little evidence to suggest that diversification motives can explain why some pensions have shifted more aggressively to alternatives more than others over the last twenty years. To more formally test this hypothesis, we estimate the following cross-sectional regression:

$$\Delta\Omega_{p,2002\to2021} = \beta_0 + \beta_1 D_{p,20202} + \varepsilon_s, \tag{30}$$

where $\Delta\Omega_{p,2002\to2021}$ is the change in p's target alternative-to-risky share between 2002 and 2021 and $D_{p,2002}$ is the distance between p's risky portfolio and that of the GMP in 2002. If poorly diversified pensions in 2002 increased their alternative-to-risky shares to better align with the GMP, we should observe $\beta_1 > 0$. Table A17 shows that the opposite is true empirically. In column (1), where $D_{p,2002}$ is defined based on each pension's target portfolio shares, the estimated β_1 is negative and statistically insignificant. In column

[41] See https://www.reit.com/nareit for more details.

(2), we obtain a similar result when constructing $D_{p,2002}$ based on actual pension portfolio shares. These findings suggest that, if anything, better diversified pensions in 2002 were more likely to increase their alternative-to-risky shares.

In the last two columns of Table A17, we replace the level of $D_{p,2002}$ in regression (30) with its change between 2002 and 2021, denoted $\Delta D_{p,2002\to2021}$. Column (3) measures D using target shares, while Column (4) uses actual shares. In both cases, the estimated regression coefficient is positive and statistically significant. Thus, the pensions that increased their alternative-to-risky share the most between 2002 and 2021 were also the ones that became the *least* diversified. Moreover, the relatively large R^2 values in both regressions suggest that declines in the degree of diversification across pensions are largely explained by increases in their alternative-to-risky shares.

I Derivations

I.1 CRRA preference with log-normal distribution

The investors's objective function is

$$\max \mathrm{E}\left[u\right] = \mathrm{E}\left[\frac{W^{1-\gamma}}{1-\gamma}\right] \Longleftrightarrow \max \frac{W_0^{1-\gamma}}{1-\gamma}\mathrm{E}[R_p^{1-\gamma}] \Longleftrightarrow \max \mathrm{E}[R_p^{1-\gamma}].$$

Taking the log of the last term,

$$\begin{aligned} &\max \log \mathrm{E}[R_p^{1-\gamma}] = (1-\gamma)r_p + \frac{1}{2}(1-\gamma)^2\sigma_p^2 \\ \Longrightarrow &\max r_p + \frac{1}{2}(1-\gamma)\sigma_p^2 \end{aligned}$$

I.2 Solving the optimization problem with minimum fixed income constraint

We have

$$\begin{aligned} &\max_{\mathbf{w}_r} \mathrm{E}[r_p] + \frac{1}{2}(1-\gamma)\mathrm{Var}\left[r_p\right], \\ s.t. \quad &\mathrm{E}[r_p] = r_f + \mathbf{w}_r'\mu + \frac{1}{2}\mathbf{w}_r'\boldsymbol{\sigma}^{\mathbf{2}} - \frac{1}{2}\mathbf{w}_r'\boldsymbol{\Sigma}\mathbf{w}_r, \\ &\mathrm{Var}\left[r_p\right] = \mathbf{w}_r'\boldsymbol{\Sigma}\mathbf{w}_r, \\ &\omega_f = 1 - \omega_A - \omega_E \geq \omega_f^{min} \geq 0, \end{aligned}$$

where

$$\mu = \begin{bmatrix} \mu_A \\ \mu_E \end{bmatrix}, \boldsymbol{\sigma}^{\mathbf{2}} = \begin{bmatrix} \sigma_A^2 \\ \sigma_E^2 \end{bmatrix}, \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_A^2 & \sigma_{AE} \\ \sigma_{AE} & \sigma_E^2 \end{bmatrix}.$$

Note that μ_i denotes the mean ***excess*** return on asset i. Re-arranging the last term in the constraints we get

$$\omega_A + \omega_E - 1 \leq -\omega_f^{min}.$$

The Lagrangian is

$$\mathscr{L} = \mathrm{E}[r_p] + \frac{1}{2}(1-\gamma)\mathrm{Var}[r_p] - \lambda(\omega_f^{min} + \omega_A + \omega_E - 1).$$

We consider the following two cases:

1. $\lambda < 0$. The constraint is not binding and we can drop the constraint in the Lagrangian. The FOC is

$$\nabla\mathscr{L}(\mathbf{w}_r) = (\frac{\partial\mathscr{L}}{\partial\omega_A}, \frac{\partial\mathscr{L}}{\partial\omega_E}) = \mu + \frac{1}{2}\boldsymbol{\sigma^2} - \boldsymbol{\Sigma}\mathbf{w}_r + (1-\gamma)\boldsymbol{\Sigma}\mathbf{w}_r = 0.$$

   With some algebra we obtain

$$\mathbf{w}_r = \frac{1}{\gamma}\boldsymbol{\Sigma}^{-1}(\mu + \frac{1}{2}\boldsymbol{\sigma^2}), \quad \text{where } \boldsymbol{\Sigma}^{-1} = \frac{1}{\sigma_A^2\sigma_E^2 - \sigma_{AE}^2}\begin{bmatrix} \sigma_E^2 & -\sigma_{AE} \\ -\sigma_{AE} & \sigma_A^2 \end{bmatrix}.$$

   This is equivalent to

$$\omega_A = \frac{1}{\gamma}\frac{\sigma_E^2(\mu_A + \frac{1}{2}\sigma_A^2) - \sigma_{AE}(\mu_E + \frac{1}{2}\sigma_E^2)}{\sigma_A^2\sigma_E^2 - \sigma_{AE}^2},$$
$$\omega_E = \frac{1}{\gamma}\frac{\sigma_A^2(\mu_E + \frac{1}{2}\sigma_E^2) - \sigma_{AE}(\mu_A + \frac{1}{2}\sigma_A^2)}{\sigma_A^2\sigma_E^2 - \sigma_{AE}^2}.$$

   The optimal share of riskless asset is

$$\begin{aligned}\omega_f &= 1 - \omega_A - \omega_E \\ &= 1 - \frac{1}{\gamma}\frac{(\sigma_A^2 - \sigma_{AE})(\mu_E + \frac{1}{2}\sigma_E^2) + (\sigma_E^2 - \sigma_{AE})(\mu_A + \frac{1}{2}\sigma_A^2)}{\sigma_A^2\sigma_E^2 - \sigma_{AE}^2}\end{aligned}$$

   The allocation of the risky portfolio to alternatives, Ω, and equities, $1-\Omega$, equals

$$\Omega = \frac{\omega_A}{1-\omega_f}, \quad 1-\Omega = \frac{\omega_E}{1-\omega_f}.$$

   It's easy to see that when the constraint is not binding, the risky portfolio is the tangency portfolio and Ω does not depend on the risk aversion γ.

2. $\lambda \geq 0$. The constraint is binding and $\omega_f = \omega_f^{min}$. The optimization problem becomes

$$\begin{aligned} &\max_{w_A} \mathrm{E}[r_p] + \frac{1}{2}(1-\gamma)\mathrm{Var}[r_p], \\ s.t. \quad &\mathrm{E}[r_p] = r_f + \mathbf{w}_r'\mu + \frac{1}{2}\mathbf{w}_r'\boldsymbol{\sigma}^{\mathbf{2}} - \frac{1}{2}\mathbf{w}_r'\boldsymbol{\Sigma}\mathbf{w}_r, \\ &\mathrm{Var}[r_p] = \mathbf{w}_r'\boldsymbol{\Sigma}\mathbf{w}_r, \\ &\mathbf{w}_r = \begin{bmatrix} \omega_A \\ 1-\omega_f^{min}-\omega_A \end{bmatrix}, \end{aligned}$$

where

$$\mu = \begin{bmatrix} \mu_A \\ \mu_E \end{bmatrix}, \boldsymbol{\sigma}^{\mathbf{2}} = \begin{bmatrix} \sigma_A^2 \\ \sigma_E^2 \end{bmatrix}, \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_A^2 & \sigma_{AE} \\ \sigma_{AE} & \sigma_E^2 \end{bmatrix}.$$

Substituting into the Lagrangian we obtain

$$\mathscr{L} = r_f + \underbrace{\mathbf{w}_r'\mu + \frac{1}{2}\mathbf{w}_r'\boldsymbol{\sigma}^{\mathbf{2}}}_{①} - \underbrace{\frac{\gamma}{2}\mathbf{w}_r'\boldsymbol{\Sigma}\mathbf{w}_r}_{②}.$$

Expanding the matrix forms in ① and ② we get

$$\begin{aligned} ① &= \omega_A(\mu_A + \frac{1}{2}\sigma_A^2) + (1-\omega_f^{min}-w_A)(\mu_E + \frac{1}{2}\sigma_E^2), \\ ② &= \frac{\gamma}{2}\left[\omega_A^2\sigma_A^2 + 2\omega_A(1-\omega_f^{min}-w_A)\sigma_{AE} + (1-\omega_f^{min}-\omega_A)^2\sigma_E^2\right]. \end{aligned}$$

Taking the FOC with respect to w_A,

$$\frac{\partial\mathscr{L}}{\partial\omega_A} = (\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2) - \frac{\gamma}{2}\left[2w_A\sigma_A^2 + 2\sigma_{AE}(1-\omega_f^{min}) - 4\omega_A\sigma_{AE} - 2(1-\omega_f^{min}-\omega_A)\sigma_E^2\right] = 0$$

$$\begin{aligned} &\implies 0 = (\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2) - \gamma\left[\omega_A\sigma_A^2 + \sigma_{AE}(1-\omega_f^{min}) - 2\omega_A\sigma_{AE} - (1-\omega_f^{min}-\omega_A)\sigma_E^2\right] \\ &\implies \omega_A = \frac{1}{\gamma}\underbrace{\frac{(\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2)}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}}_{K} + (1-\omega_f^{min})\underbrace{\frac{\sigma_E^2 - \sigma_{AE}}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}}_{C}. \end{aligned}$$

Hence, the expression for the alternative asset share is

$$\omega_A = \frac{1}{\gamma}K + (1-\omega_f^{min})C,$$

$$\text{where} \quad K = \frac{(\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2)}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}, \quad C = \frac{\sigma_E^2 - \sigma_{AE}}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}.$$

The public equity asset share is

$$\begin{aligned}\omega_E &= 1-\omega_f^{min}-\omega_A \\ \implies \omega_E &= -\frac{1}{\gamma}K+(1-\omega_f^{min})(1-C).\end{aligned}$$

The risky asset portfolio allocation is as of follows:

$$\begin{aligned}\text{Alternatives: } \Omega &= \frac{\omega_A}{1-\omega_f^{min}} = \frac{1}{\gamma}\frac{1}{1-\omega_f^{min}}K+C, \\ \text{Equities: } 1-\Omega &= \frac{\omega_E}{1-\omega_f^{min}} = -\frac{1}{\gamma}\frac{1}{1-\omega_f^{min}}K+(1-C).\end{aligned}$$

What is the minimum level of risk aversion such that the riskless share constraint is not binding?

$$\begin{aligned}&\omega_f^{min} < \omega_f \\ &\omega_f^{min} < 1-\frac{1}{\gamma}\frac{(\sigma_A^2-\sigma_{AE})(\mu_E+\frac{1}{2}\sigma_E^2)+(\sigma_E^2-\sigma_{AE})(\mu_A+\frac{1}{2}\sigma_A^2)}{\sigma_A^2\sigma_E^2-\sigma_{AE}^2} \\ &\implies \gamma > \gamma^{min} = \frac{1}{1-\omega_f^{min}}\cdot\frac{(\sigma_A^2-\sigma_{AE})(\mu_E+\frac{1}{2}\sigma_E^2)+(\sigma_E^2-\sigma_{AE})(\mu_A+\frac{1}{2}\sigma_A^2)}{\sigma_A^2\sigma_E^2-\sigma_{AE}^2}\end{aligned}$$

In vector form, it can be expressed as

$$\gamma^{min} = \frac{1}{1-\omega_f^{min}}\cdot(\mu+\frac{1}{2}\boldsymbol{\sigma^2})^T\boldsymbol{\Sigma}^{-1}\iota.$$

I.3 Using CAPM notation in solution with minimum fixed income constraint

Another approach is to express μ_A, σ_A^2, σ_{AE} as functions of α and β, where α, β are the constant and coefficient terms of the regression of log alternative returns on log public equity returns. We have

$$r_A - r_f = \alpha+\beta(r_E-r_f)+\eta_A$$

$$\begin{aligned}\implies \mu_A &= \alpha+\beta\mu_E \\ \sigma_A^2 &= \beta^2\sigma_E^2+\sigma_\eta^2 \\ \sigma_{AE} &= \beta\sigma_E^2\end{aligned}$$

1. Constraint is not binding. Optimal asset allocation can be expressed as follows:

$$\begin{aligned}\omega_A &= \frac{1}{\gamma}\left(\frac{\alpha\sigma_E^2+\frac{1}{2}(\beta-1)\beta\sigma_E^4+\frac{1}{2}\sigma_\eta^2\sigma_E^2}{\sigma_\eta^2\sigma_E^2}\right) \\ \implies \omega_A &= \frac{1}{\gamma}\left(\frac{\alpha}{\sigma_\eta^2}+\frac{1}{2}(\beta-1)\beta\frac{\sigma_E^2}{\sigma_\eta^2}+\frac{1}{2}\right)\end{aligned}$$

$$\omega_E = \frac{1}{\gamma}\left(\frac{\mu_E\sigma_\eta^2 - \alpha\beta\sigma_E^2 + \frac{1}{2}(1-\beta)\sigma_E^2(\beta^2\sigma_E^2+\sigma_\eta^2)}{\sigma_\eta^2\sigma_E^2}\right)$$
$$\implies \omega_E = \frac{1}{\gamma}\left(\frac{\mu_E}{\sigma_E^2} - \frac{\alpha\beta}{\sigma_\eta^2} + \frac{1}{2}(1-\beta)(\beta^2\frac{\sigma_E^2}{\sigma_\eta^2}+1)\right)$$

$$1-\omega_f = \omega_A + \omega_E$$
$$\implies 1-\omega_f = \frac{1}{\gamma}\left(\frac{\mu_E(\sigma_\eta^2/\sigma_E^2)+\alpha(1-\beta)}{\sigma_\eta^2} - \frac{1}{2}\beta(1-\beta)^2\frac{\sigma_E^2}{\sigma_\eta^2} + 1 - \frac{\beta}{2}\right)$$

2. Constraint is binding. Optimal asset allocation can be expressed as follows:

$$\omega_A = \frac{1}{\gamma}K + (1-\omega_f^{min})C, \quad \omega_A^* = \frac{1}{\gamma}\cdot\frac{1}{1-\omega_f^{min}}K + C$$

$$K = \frac{(\mu_A+\frac{1}{2}\sigma_A^2)-(\mu_E+\frac{1}{2}\sigma_E^2)}{\sigma_A^2-2\sigma_{AE}+\sigma_E^2}$$
$$= \frac{(\alpha+(\beta-1)\mu_E)/\sigma_E^2}{(\beta-1)^2+\sigma_\eta^2/\sigma_E^2} + \frac{1}{2}\frac{(\beta^2-1)+\sigma_\eta^2/\sigma_E^2}{(\beta-1)^2+\sigma_\eta^2/\sigma_E^2}$$
$$\implies = \frac{(\alpha+(\beta-1)\mu_E)/\sigma_E^2}{(\beta-1)^2+\sigma_\eta^2/\sigma_E^2} + \left(\frac{\beta-1}{(\beta-1)^2+\sigma_\eta^2/\sigma_E^2} + \frac{1}{2}\right)$$

$$C = \frac{\sigma_E^2-\sigma_{AE}}{\sigma_A^2-2\sigma_{AE}+\sigma_E^2}$$
$$\implies = \frac{1-\beta}{(\beta-1)^2+\sigma_\eta^2/\sigma_E^2}$$

What is the minimum level of risk aversion such that the riskless share constraint is not binding?

$$\gamma > \gamma^{min} = \frac{1}{1-\omega_f^{min}}\cdot\frac{(\sigma_A^2-\sigma_{AE})(\mu_E+\frac{1}{2}\sigma_E^2)+(\sigma_E^2-\sigma_{AE})(\mu_A+\frac{1}{2}\sigma_A^2)}{\sigma_A^2\sigma_E^2-\sigma_{AE}^2}$$
$$\implies = \frac{1}{1-\omega_f^{min}}\cdot\frac{\mu_E\sigma_\eta^2+\alpha(1-\beta)\sigma_E^2-\frac{1}{2}\beta(1-\beta)^2\sigma_E^2+(1-\frac{\beta}{2})\sigma_\eta^2\sigma_E^2}{\sigma_\eta^2\sigma_E^2}$$
$$\implies = \frac{1}{1-\omega_f^{min}}\cdot\left(\frac{\mu_E}{\sigma_E^2}+\frac{\alpha(1-\beta)}{\sigma_\eta^2}-\frac{1}{2}\frac{\beta(1-\beta)^2}{\sigma_\eta^2}+(1-\frac{\beta}{2})\right)$$

References

Anagol, S., V. Balasubramaniam, and T. Ramadorai (2021). Learning from noise: Evidence from india's ipo lotteries. *Journal of Financial Economics 140*(3), 965–986.

Andonov, A., R. M. Bauer, and K. M. Cremers (2017). Pension fund asset allocation and liability discount rates. *The Review of Financial Studies 30*(8), 2555–2595.

Andonov, A., Y. V. Hochberg, and J. D. Rauh (2018). Political Representation and Governance: Evidence from the Investment Decisions of Public Pension Funds. *The Journal of Finance*.

Andonov, A. and J. D. Rauh (2022). The Return Expectations of Public Pension Funds. *The Review of Financial Studies 35*(8), 3777–3822.

Andrews, I., J. H. Stock, and L. Sun (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics 11*(Volume 11, 2019), 727–753.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy 126*(6), 2224–2276.

Bailey, M., D. Johnston, T. Kuchler, J. Stroebel, and A. Wong (2022). Peer effects in product adoption. *American Economic Journal: Applied Economics 14*(3), 488–526.

Begenau, J. and E. N. Siriwardane (2022). Fee variation in private equity. *Journal of Finance*.

Biggs, A. (2021). *Prelude to a State Pension Bailout*. The Wall Street Journal.

Brown, G., C. Lundblad, and W. Volckmann (2025). Risk-adjusted performance of private funds: What do we know? White paper, Institute for Private Capital.

Brown, J. R. and D. W. Wilcox (2009). Discounting state and local pension liabilities. *The American Economic Review 99*(2), 538–542.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2009). Fight or flight? portfolio rebalancing by individual investors. *The Quarterly Journal of Economics 124*(1), 301–348.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics 90*(3), 414–427.

Campbell, J. Y. and R. Sigalov (2022). Portfolio choice with sustainable spending: A model of reaching for yield. *Journal of Financial Economics 143*(1), 188–206.

Campbell, J. Y., J. C. Stein, and A. A. Wu (2024, May). Economic budgeting for endowment-dependent universities. Working Paper 32506, National Bureau of Economic Research.

Couts, S. J., A. S. Goncalves, and A. Rossi (2020). Unsmoothing returns of illiquid funds. *SSRN Working paper*.

Cremers, K. J. M. and A. Petajisto (2009, 08). How active is your fund manager? a new measure that predicts performance. *The Review of Financial Studies 22*(9), 3329–3365.

Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics 7*(2), 197–226.

Dyck, A., P. Manoel, and A. Morse (2018). Outrage and compensation in public pension performance. Working Paper.

Eaton, T. V. and J. R. Nofsinger (2004). The effect of financial constraints and political pressure on the management of public pension plans. *Journal of Accounting and Public Policy 23*(3), 161–189.

Egan, M., G. Matvos, and A. Seru (2019). The market for financial adviser misconduct. *Journal of Political Economy 127*(1), 233–295.

Foerster, S., J. T. Linnaimaa, B. T. Melzer, and A. Previtero (2017). Retail financial advice: Does one size fit all? *The Journal of Finance 72*(4), 1441–1482.

Froot, K. A., D. S. Scharfstein, and J. C. Stein (1993). Risk management: Coordinating corporate investment and financing policies. *The Journal of Finance 48*(5), 1629–1658.

Gervais, S. and T. Odean (2001). Learning to be overconfident. *The review of financial studies 14*(1), 1–27.

Gupta, A. and S. Van Nieuwerburgh (2021). Valuing private equity investments strip by strip. *The Journal of Finance 76*(6), 3255–3307.

Harris, R. S., T. Jenkinson, and S. N. Kaplan (2014). Private equity performance: What do we know? *The Journal of Finance 69*(5), 1851–1882.

Hendren, N. and B. Sprung-Keyser (2020, 03). A Unified Welfare Analysis of Government Policies*. *The Quarterly Journal of Economics 135*(3), 1209–1318.

Hochberg, Y. V. and J. D. Rauh (2012, 12). Local Overweighting and Underperformance: Evidence from Limited Partner Private Equity Investments. *The Review of Financial Studies 26*(2), 403–451.

Jensen, M. C. and W. H. Meckling (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of financial economics 3*(4), 305–360.

Jurek, J. W. and E. Stafford (2015). The cost of capital for alternative investments. *The Journal of Finance 70*(5), 2185–2226.

Kaplan, S. N. and A. Schoar (2005). Private equity performance: Returns, persistence, and capital flows. *The Journal of Finance 60*(4), 1791–1823.

Korteweg, A. (2019). Risk adjustment in private equity returns. *Annual Review of Financial Economics 11*, pp. 131–152.

Korteweg, A. (2023). Chapter 8 - risk and return in private equity. In B. E. Eckbo, G. M. Phillips, and M. Sorensen (Eds.), *Private Equity and Entrepreneurial Finance*, Volume 1 of *Handbook of the Economics of Corporate Finance*, pp. 343–418. North-Holland.

Korteweg, A. and S. Nagel (2016). Risk-adjusting the returns to venture capital. *The Journal of Finance 71*(3), 1437–1470.

Korteweg, A. and S. Nagel (2024, 10). Risk-adjusted returns of private equity funds: A new approach. *The Review of Financial Studies*, hhae067.

Lian, C., Y. Ma, and C. Wang (2018, 09). Low Interest Rates and Risk-Taking: Evidence from Individual Investment Decisions. *The Review of Financial Studies 32*(6), 2107–2148.

Linnainmaa, J. T. (2011). Why do (some) households trade so much? *The Review of Financial Studies 24*(5), 1630–1666.

Lu, L., M. Pritsker, A. Zlate, K. Anadu, and J. Bohn (2019). Reach for yield by us public pension funds. *FRB Boston Risk and Policy Analysis Unit Paper No. RPA*, 19–2.

Mayers, D. and C. W. Smith (1982). On the corporate demand for insurance. *The Journal of Business 55*(2), 281–296.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, 867–887.

Mitchell, O., E. Hustead, and W. S. P. R. Council (2001). *Pensions in the Public Sector*. Pension Research Council Publications. University of Pennsylvania Press, Incorporated.

Mittal, V. (2024). Desperate capital breeds productivity loss: Evidence from public pension investments in private equity. *SSRN Working paper*.

Mohan, N. and T. Zhang (2014). An analysis of risk-taking behavior for public defined benefit pension plans. *Journal of Banking and Finance 40*, 403 – 419.

NACUBO (2021). 2021 nacubo-tiaa study of endowments. Accessed: 2025-01-31.

NASRA (2020). Public retirement system rebalancing policies. https://www.nasra.org/files/Compiled%20Resources/rebalancing%20policies.pdf. Accessed: 2024-01-15.

Novy-Marx, R. and J. Rauh (2009). The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives 23*(4), 191–210.

Novy-Marx, R. and J. Rauh (2011). Public pension promises: how big are they and what are they worth? *The Journal of Finance 66*(4), 1211–1249.

Novy-Marx, R. and J. Rauh (2014, February). The revenue demands of public employee pension promises. *American Economic Journal: Economic Policy 6*(1), 193–229.

Pennacchi, G. and M. Rastad (2011). Portfolio allocation for public pension funds. *Journal of Pension Economics & Finance 10*(2), 221–245.

Rauh, J. D. (2006). Investment and financing constraints: Evidence from the funding of corporate pension plans. *The Journal of Finance 61*(1), 33–71.

Rauh, J. D. (2008, 06). Risk Shifting versus Risk Management: Investment Policy in Corporate Pension Plans. *The Review of Financial Studies 22*(7), 2687–2733.

Scharfstein, D. S. and J. C. Stein (1990). Herd behavior and investment. *Amercian Economic Review 80*(Jun.), 465–479.

Sharpe, W. F. (1976). Corporate pension funding policy. *Journal of Financial Economics 3*(3), 183–193.

Smith, C. W. and R. M. Stulz (1985). The determinants of firms' hedging policies. *The Journal of Financial and Quantitative Analysis 20*(4), 391–405.

Stafford, E. (2022). Replicating private equity with value investing, homemade leverage, and hold-to-maturity accounting. *The Review of Financial Studies 35*(1), 299–342.

Staiger, D. and J. H. Stock (1997). Instrumental variables regression with weak instruments. *Econometrica 65*(3), 557–586.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *The Review of Economic Studies 25*(2), 65–86.

Treynor, J. L. (1977). The principles of corporate pension finance. *The Journal of Finance 32*(2), 627–638.

Zorn, P. (1997). Survey of state and local government employee retirement systems. Chicago: Public Pension Coordinating Council, c/o Government Finance Officers Association.

Figure A1: Aggregate High-Yield Share

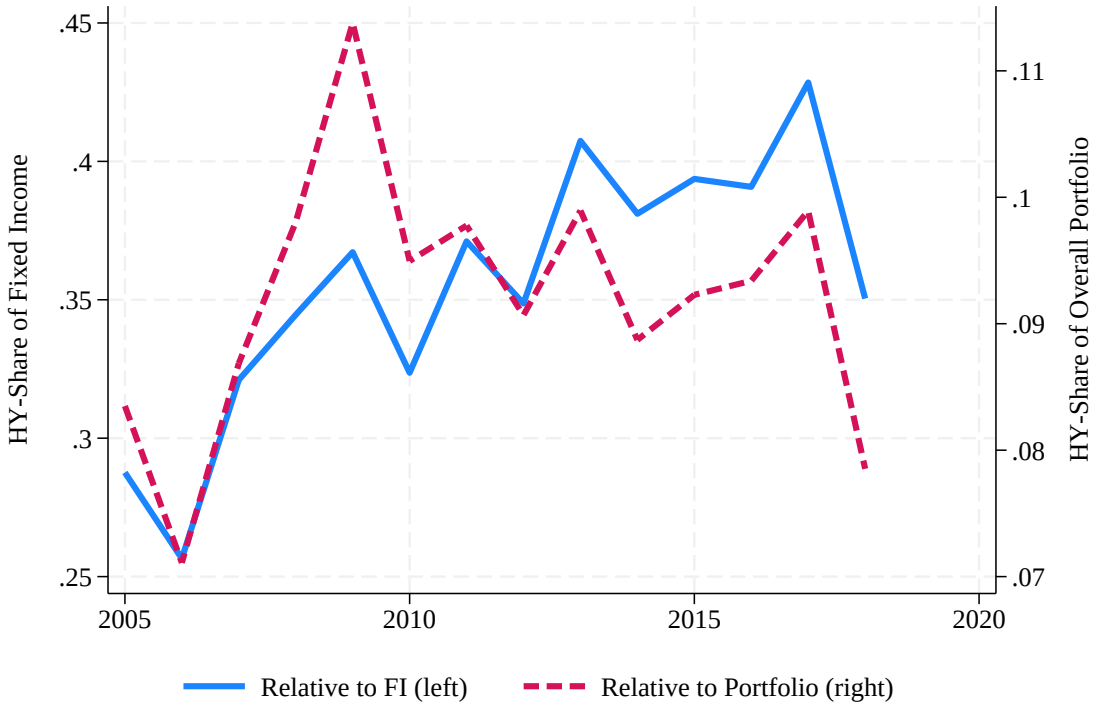

Notes: This plot shows the aggregate U.S. pension share of high-yield debt relative to fixed income holdings (left axis) and to the whole portfolio (right axis). See Internet Appendix A.1.3 for complete details on the data underlying the plot.

Figure A2: Example of PENDAT Survey Question about Target Shares

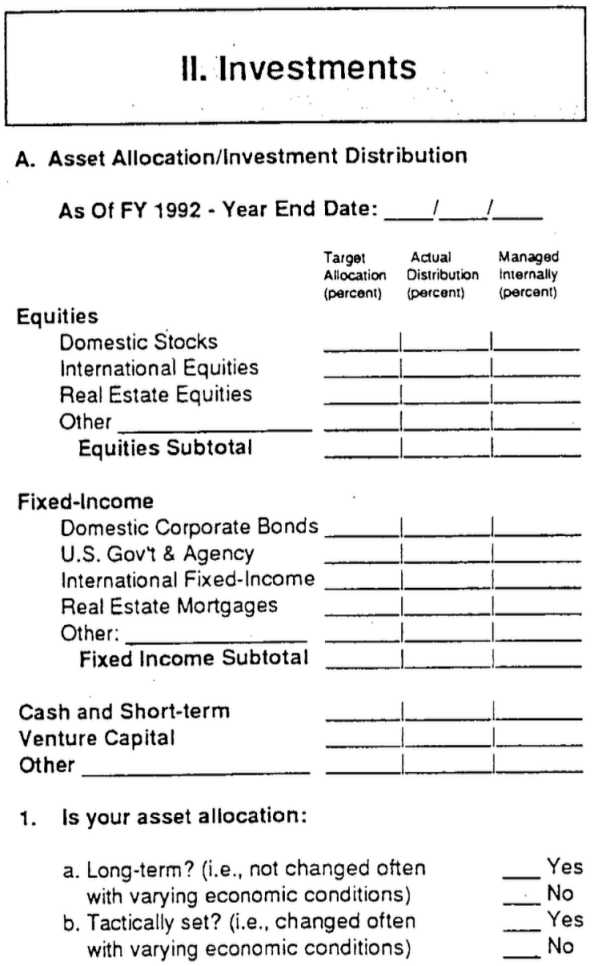

II. Investments

A. Asset Allocation/Investment Distribution

As Of FY 1992 - Year End Date: ____/____/____

	Target Allocation (percent)	Actual Distribution (percent)	Managed Internally (percent)
Equities			
Domestic Stocks			
International Equities			
Real Estate Equities			
Other _____			
Equities Subtotal			
Fixed-Income			
Domestic Corporate Bonds			
U.S. Gov't & Agency			
International Fixed-Income			
Real Estate Mortgages			
Other: _____			
Fixed Income Subtotal			
Cash and Short-term			
Venture Capital			
Other _____			

1. Is your asset allocation:

a. Long-term? (i.e., not changed often with varying economic conditions) ___ Yes ___ No

b. Tactically set? (i.e., changed often with varying economic conditions) ___ Yes ___ No

Notes: This figure provides an example of a PENDAT survey question about target asset allocation.

Figure A3: PENDAT Data Coverage

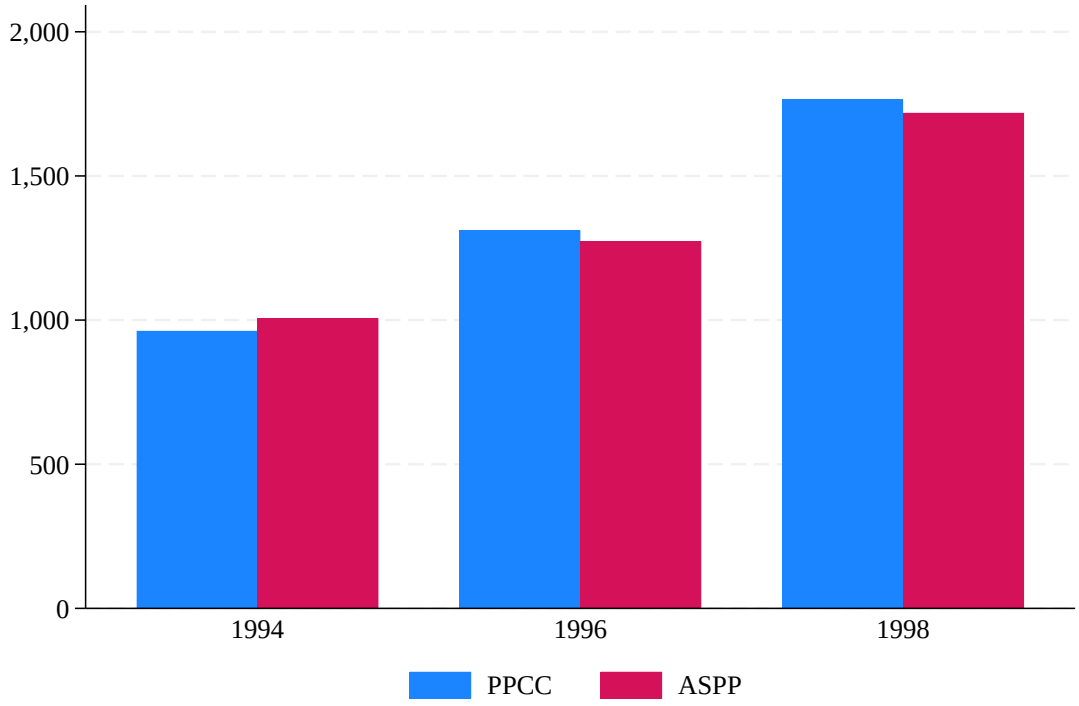

Notes: This plot shows the aggregate amount of assets in FY1994, 1996, and 1998 within the PENDAT data. For comparison, it also shows the aggregate amount of U.S. pension assets listed in the ASPP. See Internet Appendix A.4 for more details on the PENDAT data.

Figure A4: Aggregate Decomposition of Fixed Income Allocation

(a) Treasuries, Agency, and Corporate Debt (ASPP)

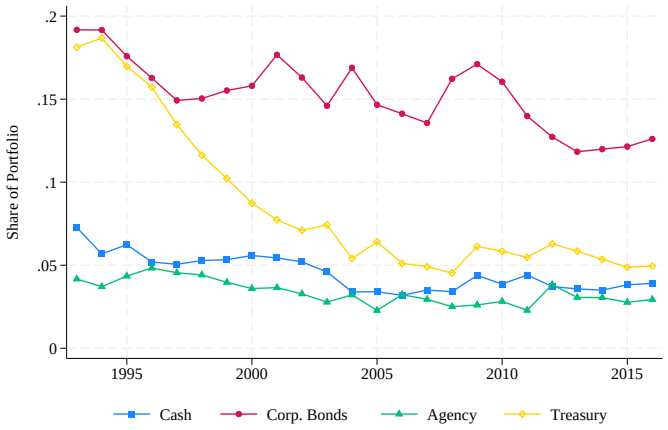

(b) U.S Government Sponsored vs Corporate Debt (QSPP)

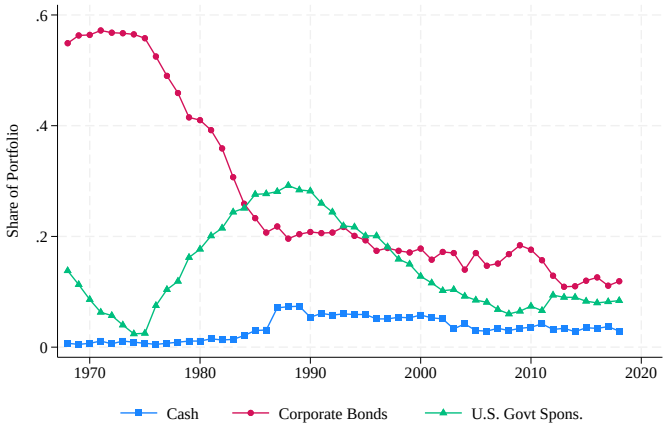

Notes: Panel (a) uses the U.S. Census Bureau's Annual Survey of Public Pensions (ASPP) to decompose the fixed income share into cash, corporate bonds, U.S. Treasuries, and Agency debt. Panel (b) plots a similar decomposition based on the QSPP but over a longer time horizon. See Section A.2 for complete details.

Figure A5: The Risky Share in the Cross Section

(a) Cross-Sectional Distribution

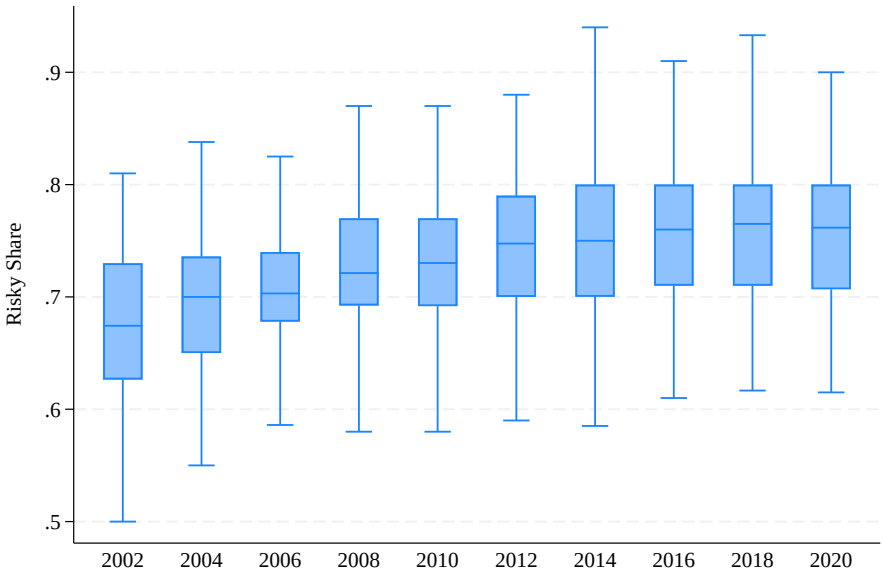

(b) Change in Risky Share Between 2002 and 2021

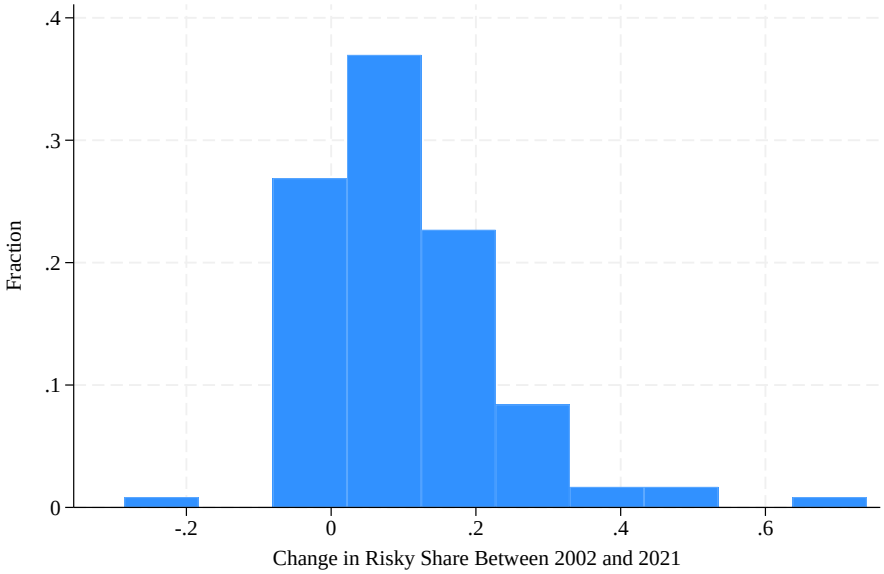

Notes: This figure visualizes the distribution of the risky share in the cross-section (panel a) and the distribution of the change in the risky share across U.S. pension systems between 2002 and 2021. The risky share is defined as all holdings outside of fixed income and cash. All data are from the PPD. Portfolio weights are based on target shares. Even years are shown in panel (a) for readability.

Figure A6: System-Level Share of Alternatives

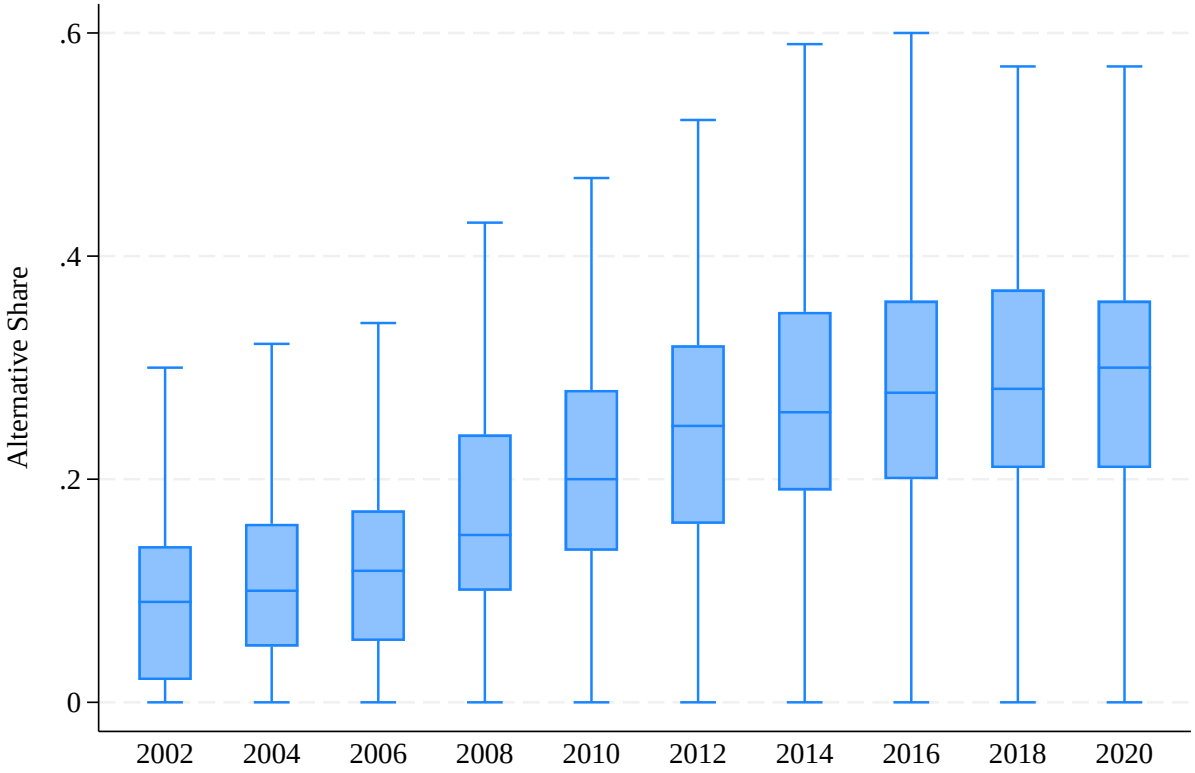

Notes: This figure depicts the distribution of the raw target alternative share across pension systems through time. Each box plot summarizes the distribution for the corresponding year on the *x*-axis. Only even years are plotted to make the graph more readable. All data are from the PPD.

Figure A7: Consultant IV Estimates under Bayesian Learning

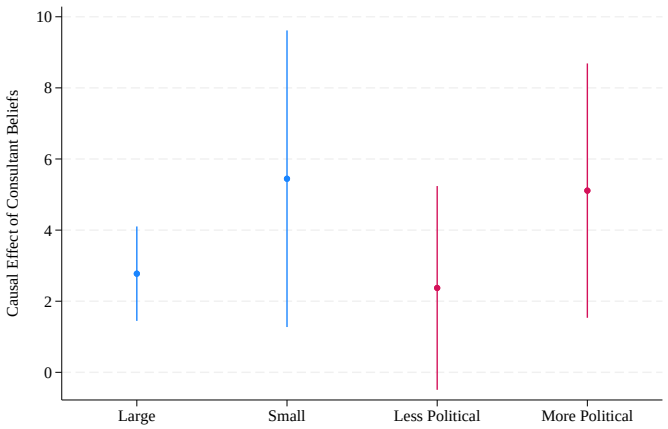

Notes: This figure reports the coefficient from an instrumental variables regression of pension p's target alternative-to-risky share in year t on the reported α of alternatives by pension p'*s*consultant in year t, interacted with measures of relative consultant signal precision. The specification follows the Bayesian learning framework outlined in Section D.1 and Eq. 15. The measures of relative consultant signal precision s_p are the pension's size and the degree of board politicization. Large (small) denotes pensions above (below) the median by AUM each year. More (less) political splits pensions into above (below) the median in terms of fraction of political board members as in Andonov et al. (2018). Pension controls include the level of pension funding, (log) size, hurdle rate, the ratio of required actuarial contributions too payroll, the ratio of administrative expenses to payroll, the cash share, annual return of the fund, and the residualized difference between each pension's target and actual risky share. Confidence intervals are based on clustered standard errors at the consultant and year levels.

Figure A8: Decomposing Changes in Consultant-Perceived Beta

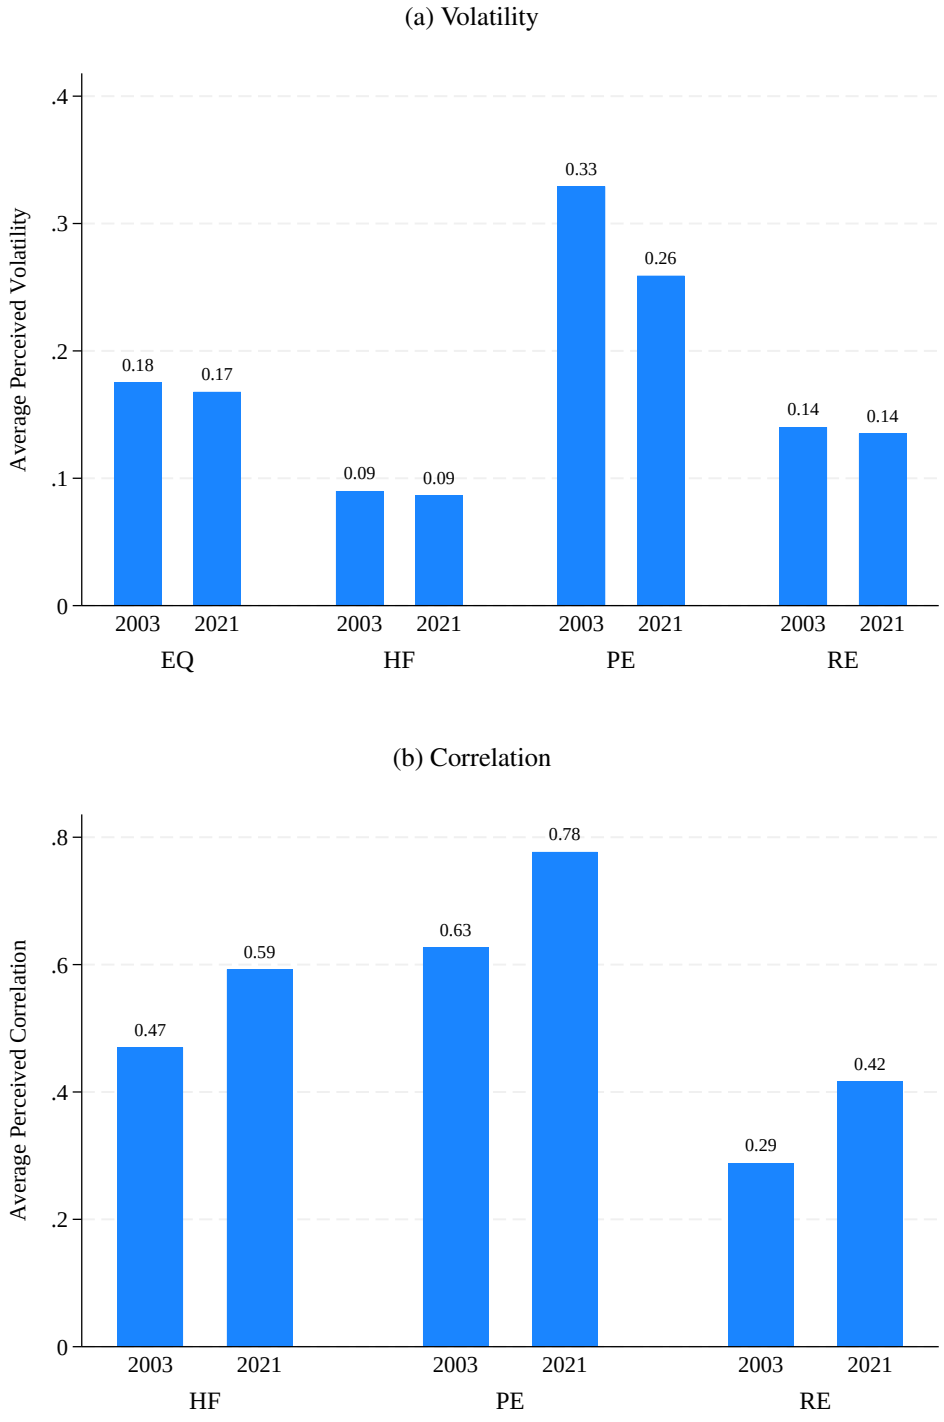

Notes: Panel (a) of the plot shows the average consultant-perceived volatility of different asset classes through time. Panel (b) repeats the exercise for correlation with respect to large-cap U.S. equities. Within a year and asset class, averages are taken over consultants that have at least ten years of data. The data for the plot are based on hand-collected capital market assumption (CMA) reports that were obtained under confidentiality agreements. See Section 3.1.2 for details.

Figure A9: U.S. Public vs Private-Sector Clients by Consultant

(a) Alternative-to-Risky Share

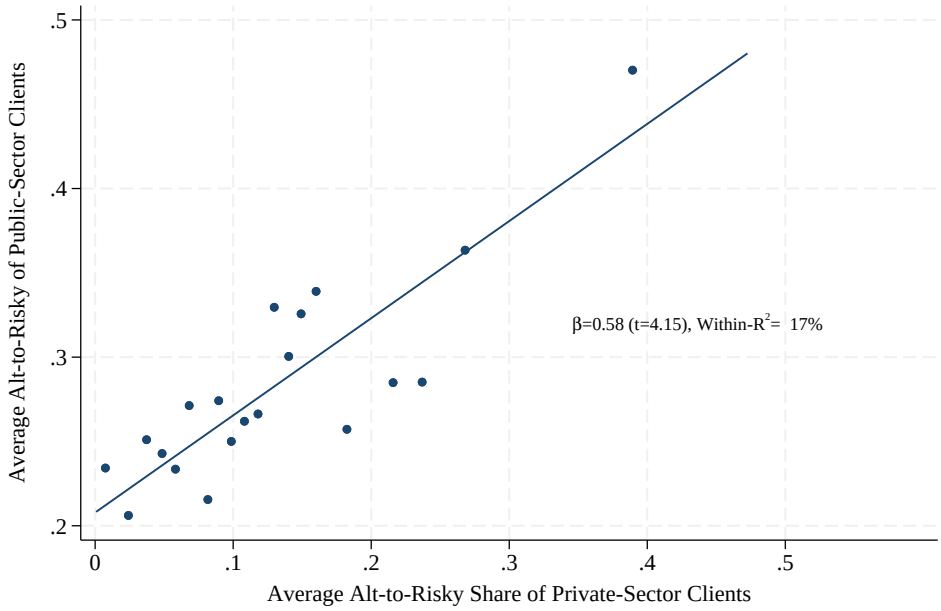

(b) Risky Share

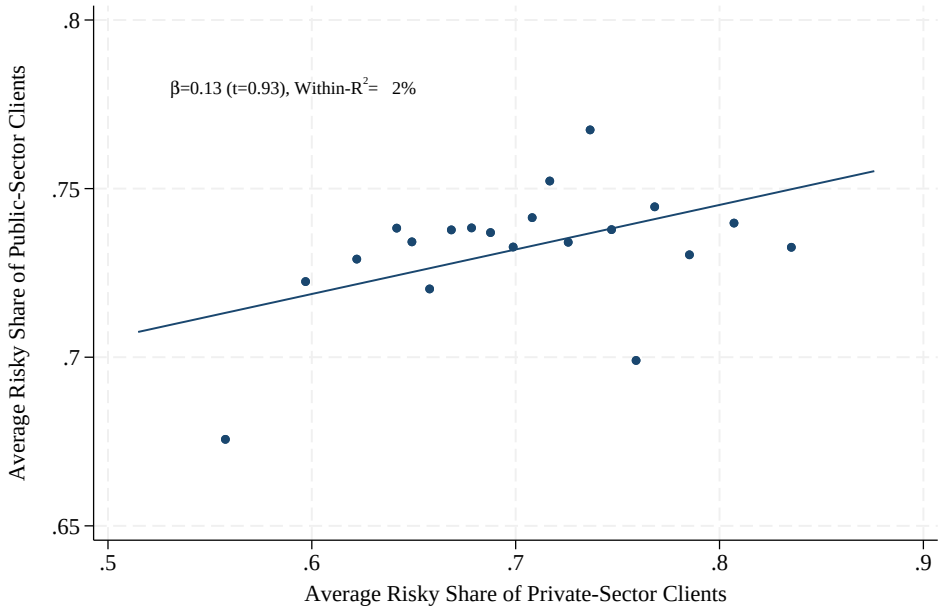

Notes: For each consultant c in year t, we compute the average alternative-to-risky and risky share of U.S. public pension clients, denoted by $\bar{u}^a_{c,t}$ and $\bar{u}^r_{c,t}$, respectively. In addition, we compute the average alternative-to-risky and risky share of U.S. private-sector clients, denoted by $\bar{p}^a_{c,t}$ and $\bar{p}^r_{c,t}$, respectively. Private-sector clients include defined-benefit corporate pensions, unions, and endowments. Panel (a) is a binscatter of $\bar{u}^a_{c,t}$ against $\bar{p}^a_{c,t}$ and panel (b) is a binscatter of $\bar{u}^r_{c,t}$ against $\bar{p}^r_{c,t}$ Both binscatters absorb a year fixed effect. The plot also shows the associated linear regression line with standard errors clustered by consultant. Consultants must have more than five private-sector clients to be included in the plot. Public sector allocations are based on the PPD and private-sector allocations are based on S&P's Money Market Directory. See Internet Appendix A for more details on the data.

Figure A10: What Explains Variation in Specific Alternative Use?

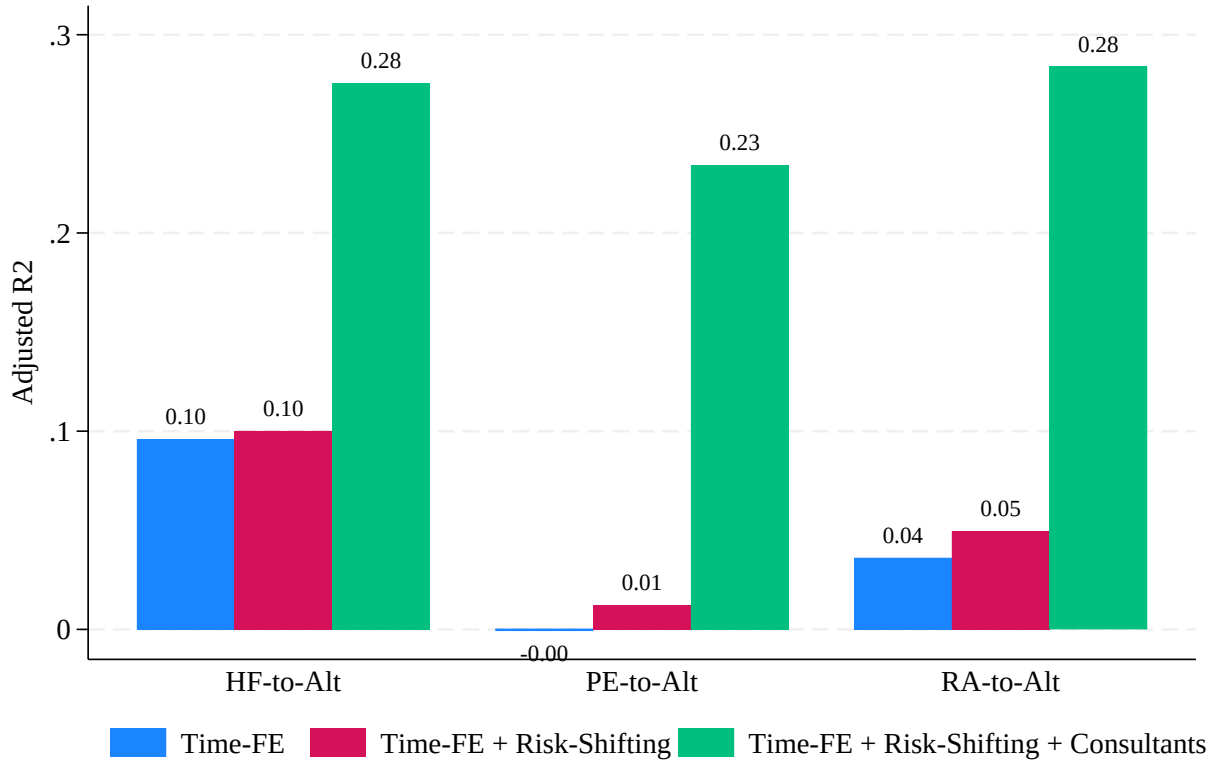

Notes: This figure summarizes variance decompositions based on the following regression:

$$y_{p,t} = \alpha_t + \lambda_c + \theta RS_{p,t} + \varepsilon_{p,t},$$

where $y_{p,t}$ is the target share of alternatives invested in either private equity (PE), real assets (RA), or hedge funds (HF) for pension p in year t. α_t is a time fixed effect, λ_c is a consultant fixed effect, and $RS_{p,t}$ are a set of proxies for risk-seeking motives. $RS_{p,t}$ includes GASB 25 funding level, liability discount rates, the fraction of retired members in a given year, the ratio of required contributions to payroll, the ratio of total spending to payroll, the cash share, the difference between each pension's target and actual risky share, and realized 1-year portfolio returns. The blue, red, and green bars respectively show the adjusted R^2 from regressions that includes only the time fixed effect (α_t), the time fixed effect and risk-seeking proxies (α_t and $RS_{p,t}$), and all covariates.

Figure A11: Historical Performance of Alternatives vs Consultant Beliefs

(a) CAPM Alpha

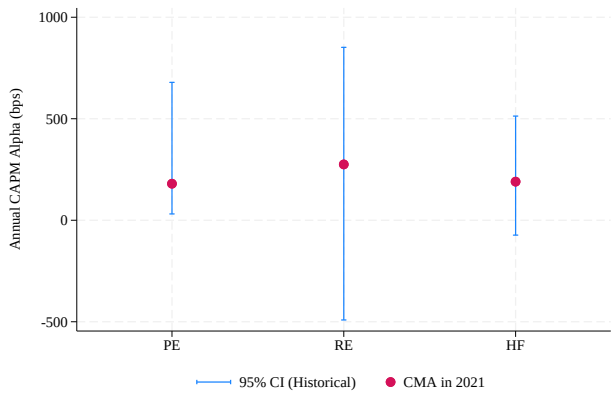

(b) CAPM Beta

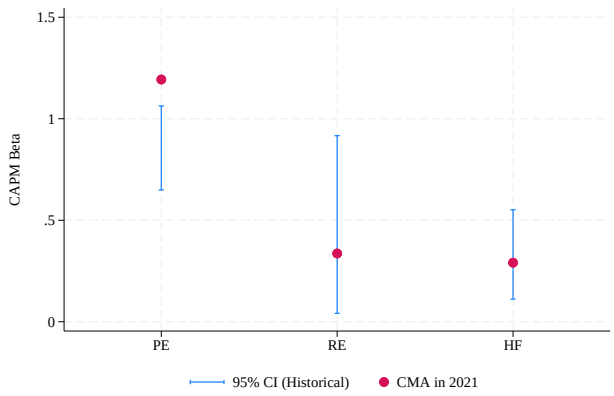

(c) Total Return

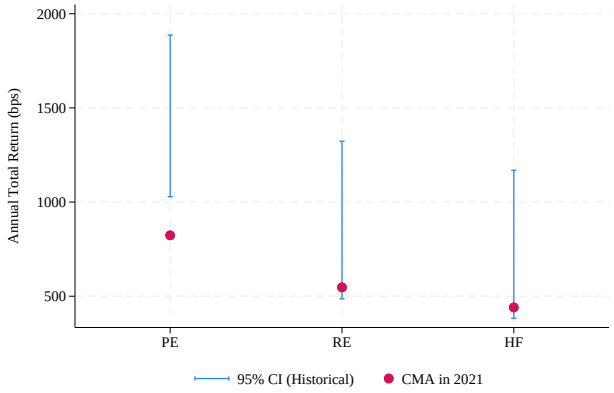

Notes: This plot compares the annual historical performance of U.S. private equity, U.S. real estate, and hedge funds to consultant beliefs in 2021. Historical returns to PE, RE, and HF are based on the Cambridge Associates U.S. Private Equity Index (1983-2020), the Cambridge Associates U.S. Real Estate Index (1986-2020), and the Credit Suisse Hedge Fund Index (1994-2020). Each panel shows 95% on confidence bands for a given performance metric based on historical returns, along with expected future value of the metric based on the average consultant's beliefs in 2021. The performance metrics in panels (a), (b), and (c) are annual CAPM alpha (bps), CAPM beta, and annual total return (bps), respectively. Historical CAPM alphas are the constant from a regression of excess returns on the contemporaneous and lagged returns of the CRSP Value-Weighted Index. Historical CAPM betas are the sum of the slope coefficients from the same regression. Excess returns are relative to the annual return on 1-month Treasury bills.

Figure A12: Performance Differences Across Alternative Users

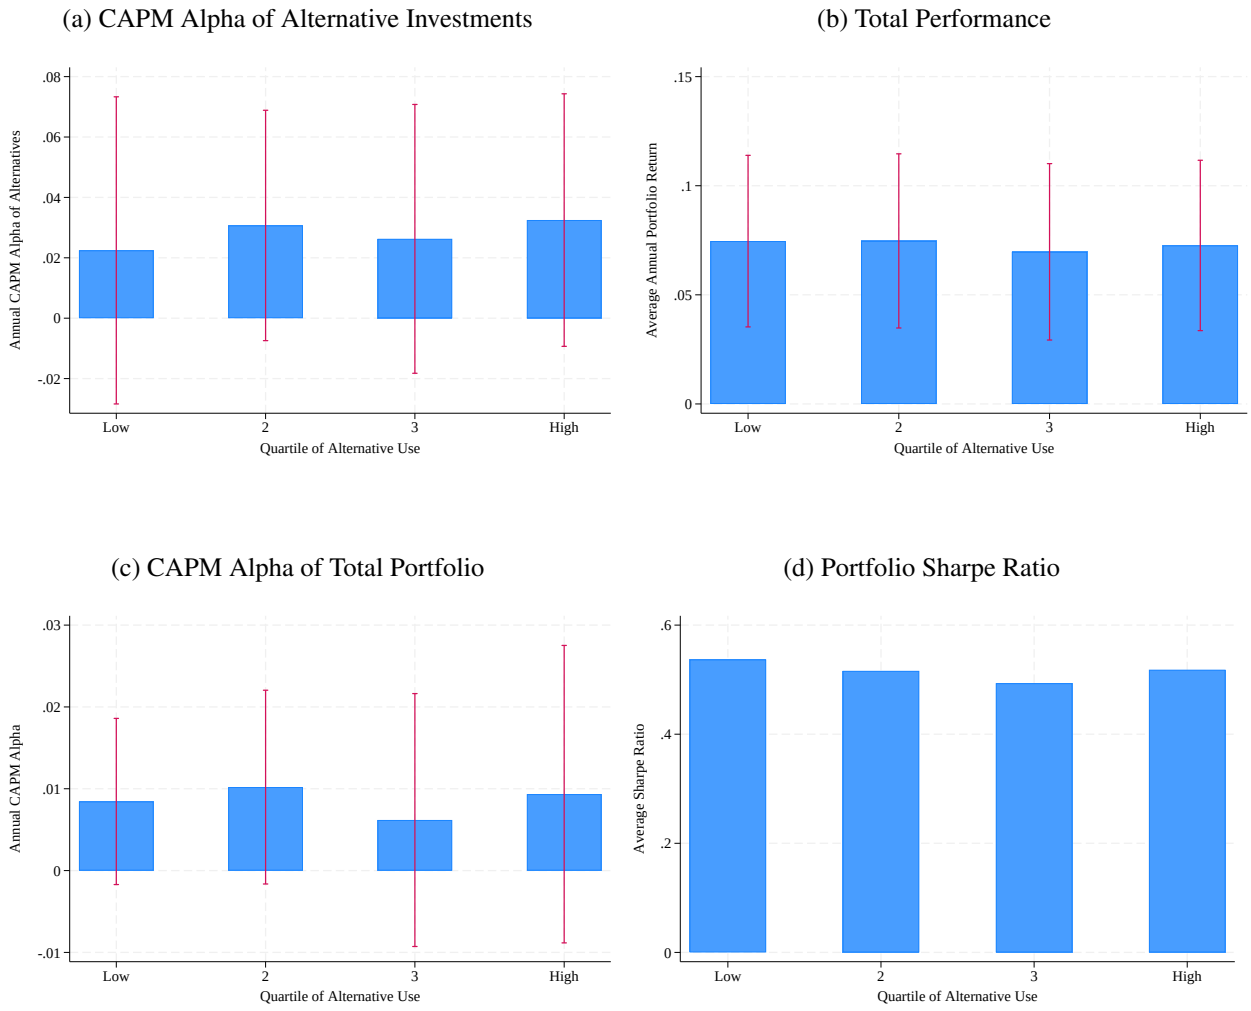

Notes: This panel summarizes differences in returns across quartiles of alternative use. For each pension p that has complete data from 2002 to 2021, we calculate the average alternative-to-risky share, denoted by $\bar{\Omega}_p$. Based on the computed $\bar{\Omega}_p$, pensions are sorted and assigned to quartiles, $q(p)$, with the $q = 4$ corresponding to the top quartile of $\bar{\Omega}_p$. Panel (a) shows the average CAPM alpha of alternative investments by quartile q. Each bar is based on a regression of alternative returns on two lags of the CRSP Value-weighted index, both in excess of the annual return on one-month Treasury bills. The 95% confidence bands are based on standard errors that are clustered by time. Panel (b) shows the average total excess portfolio returns, constructed based on beginning-of-year assets and net investment income as in Andonov and Rauh (2022). Panel (c) repeats the analysis from panel (a) but for total excess portfolio returns. Panel (d) shows the average Sharpe ratio of total portfolio returns across pensions in each quartile.

Figure A13: U.S. Public Pension Funding Status

(a) National Funding Ratios

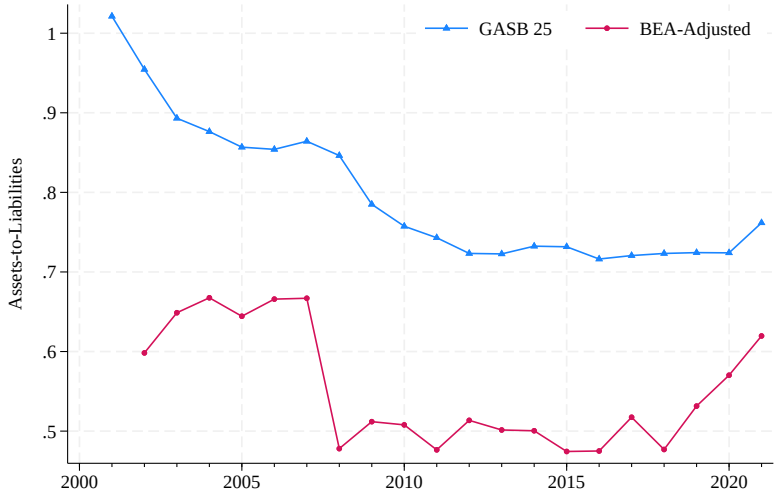

(b) Cross-Sectional Funding Ratios

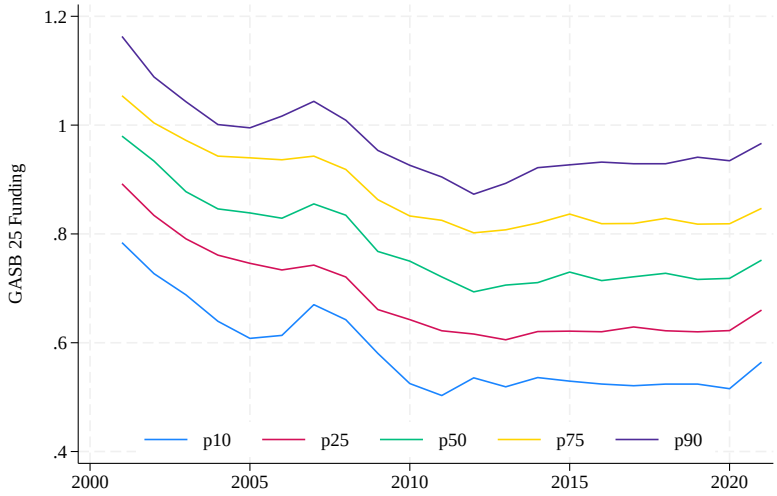

Notes: Panel (a) of this figure plots the time-series of the national pension funding ratios, defined as pension assets divided by pension liabilities. The line with triangular markers is the GASB 25 funding ratio. Under GASB 25 accounting standards, assets are smoothed over a trailing window and liabilities equal the present value of future benefits discounted using each plan's assumed long-run investment return (or hurdle rate). The line with the circular markers is the national funding ratio reported in the Enhanced Financial Accounts (EFA) of the United States published by the U.S. Federal Reserve. Pension liabilities in the EFA are computed by discounting future promises using AAA-rated corporate bond yield curves. Panel (b) plots the percentiles of GASB 25 funding in the cross section of pensions.

Figure A14: Funding-Based Explanations and Changes in the Alternative-to-Risky Share

(a) GASB 25 Funding

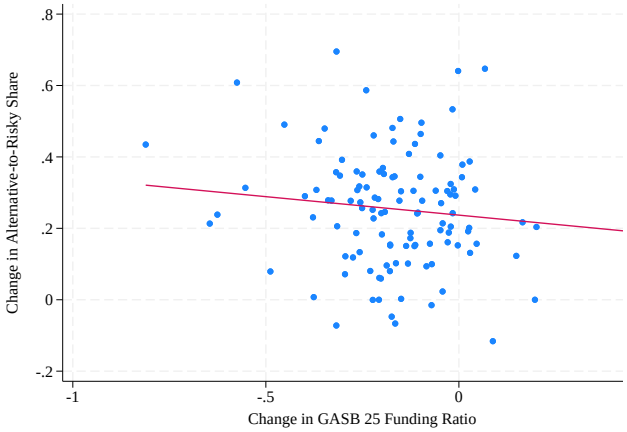

(b) Liability Discount Rate

(c) Fraction of Retired Members

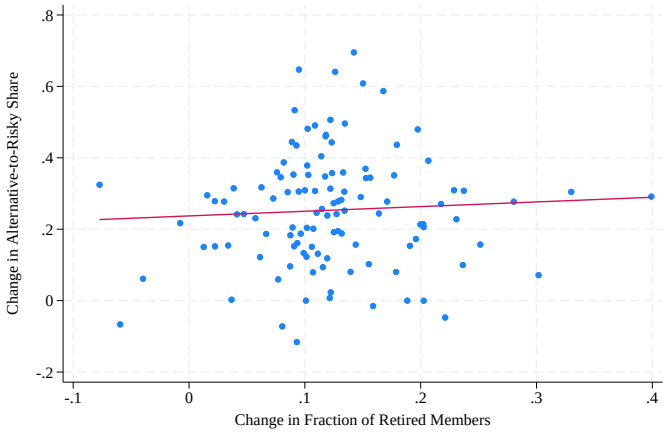

Notes: The figure contains scatter plots of changes in the target alternative-to-risky share from 2002 to 2021 against contemporaneous changes in GASB 25 funding (panel a), hurdle rates or, equivalently, liability discount rates (panel b), and the fraction of retired members (panel c). Data are from the PPD and the unit of observation in each plot is at the pension-system level.

Figure A15: Alternative Adoption, Annual Required Contributions and Size

(a) Annual Required Contributions and Alternative-to-Risky Share

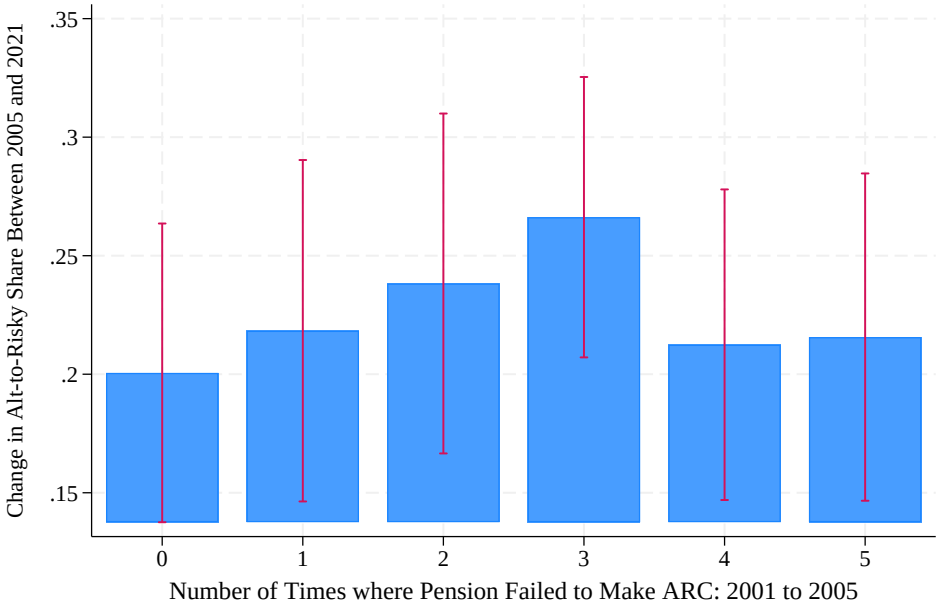

(b) Size and Alternative-to-Risky Share

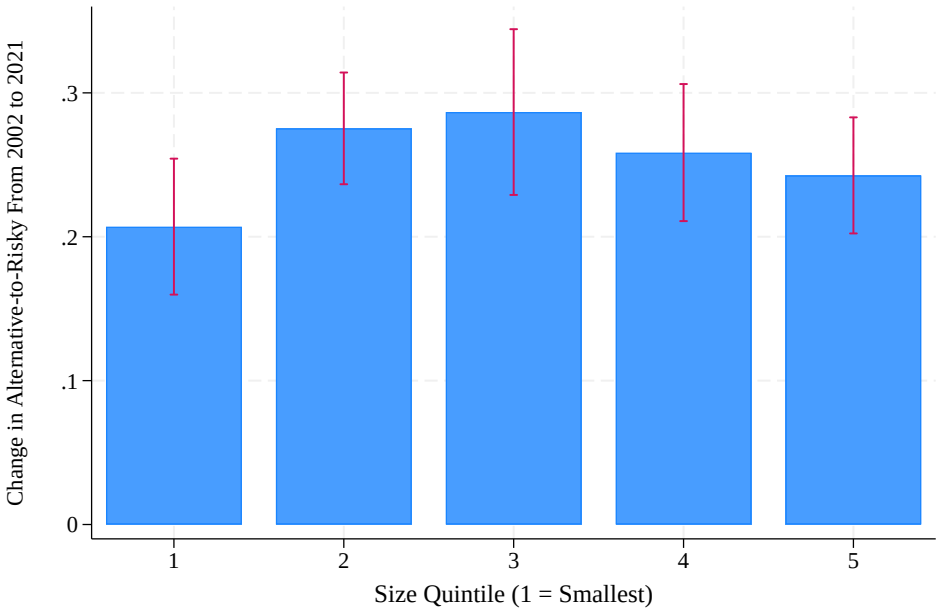

Notes: In panel (a), we first count the number of times a pension failed to make a required actuarial contribution between 2001 and 2005. We then sort pensions by count and compute the average change in the target alternative-to-risky share. In panel (b), we sort pensions into quintiles based on the assets under management in 2002. We then compute the average change in the target alternative-to-risky share within each quintile and plot it in the graph. The orange lines in both plots are 95% standard error bands. Data are from the PPD.

Figure A16: Variation in the Level of the Alternative-to-Risky Share

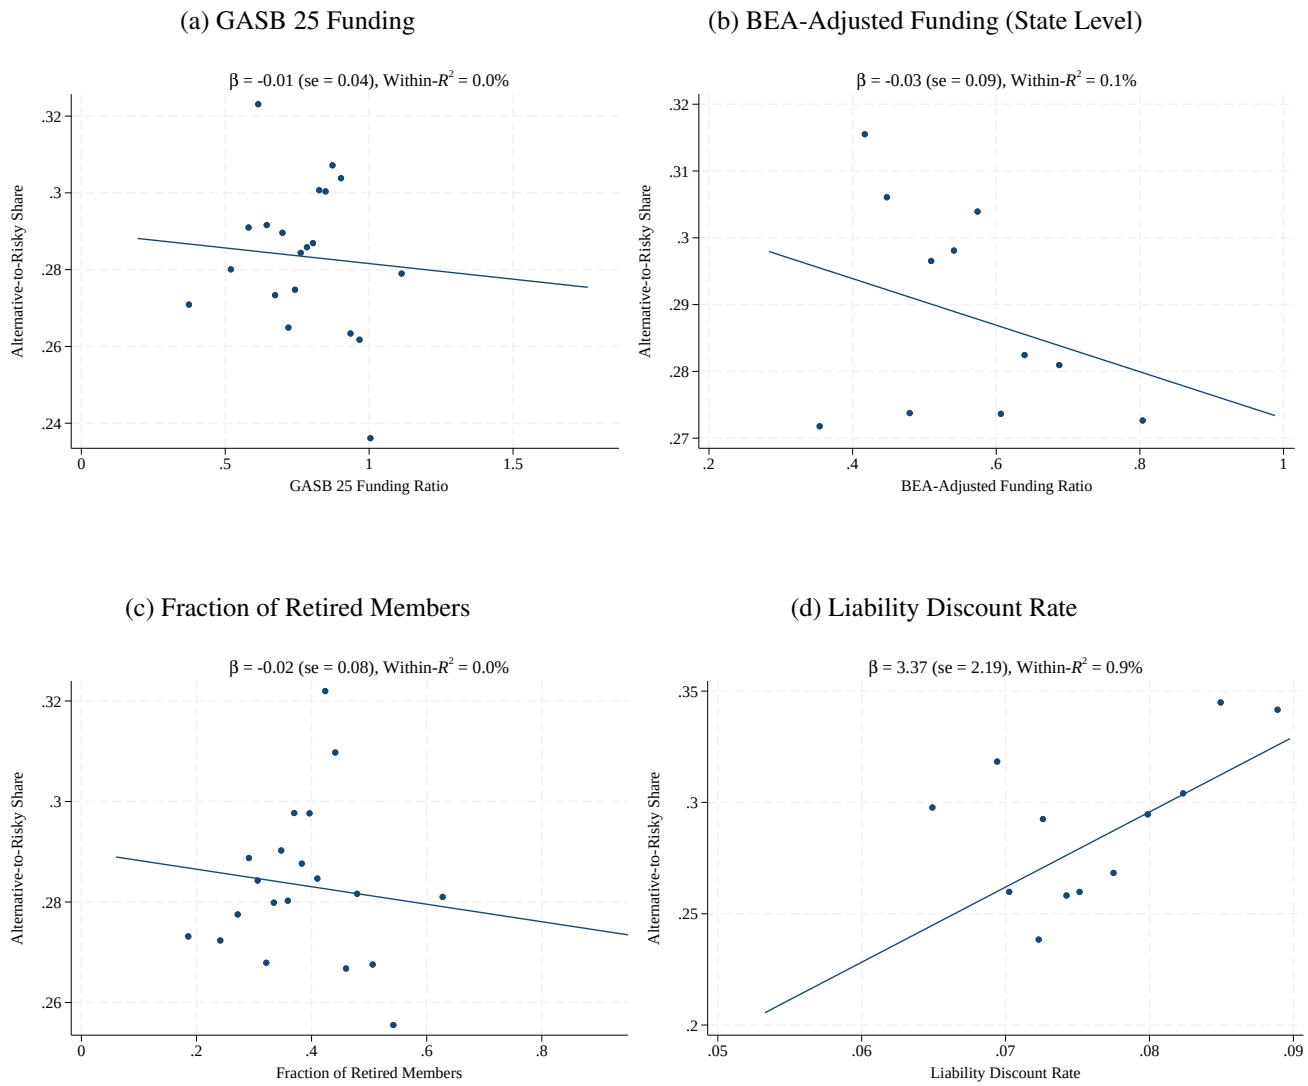

Notes: This figure shows a series of binscatter plots of each pension's target alternative-to-risky share against its: (panel a) GASB 25 funding ratio; (panel b) BEA-adjusted funding ratio (state level); (panel c) the fraction of retired members; and (panel d) hurdle rate (liability discount rate). Each plot controls for a time fixed effect. All data are based on the PPD.

Figure A17: Alternative-to-risky Share, Funding Ratio, and Hurdle Rate

(a) Alternative-to-Risky Share vs BEA Funding

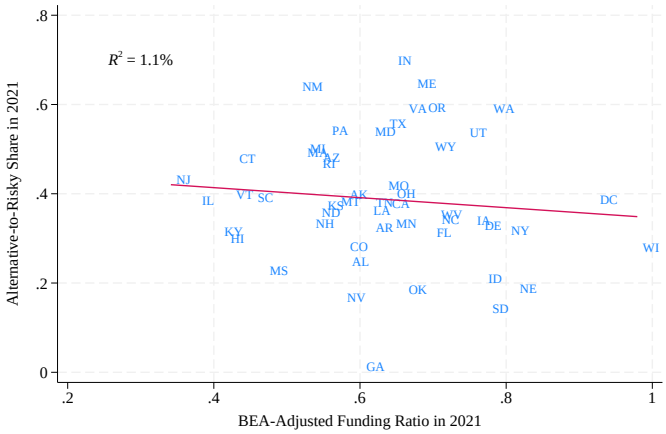

(b) Alternative-to-Risky Share vs Hurdle Rate

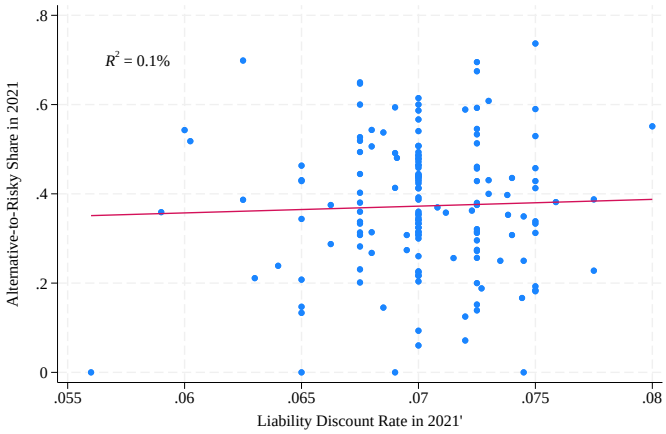

Notes: This figure shows 2020 target alternative-to-risky shares against: (a) BEA-adjusted funding ratios at the state level; and (b) hurdle rates (liability discount rate) at the pension level. All portfolio data are from PPD.

Figure A18: Deviations of Actual from Target Weights

(a) Raw Values

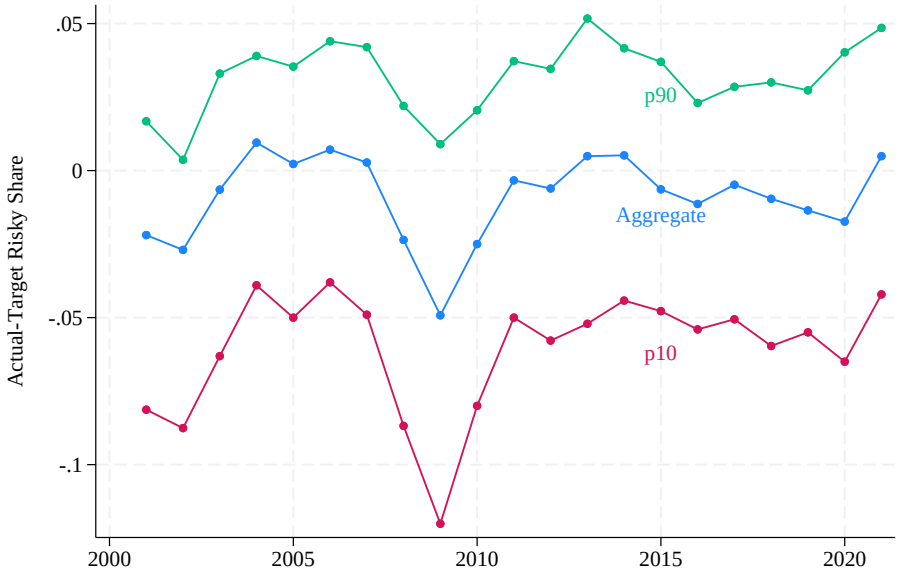

(b) Residualized to 1-Year Returns

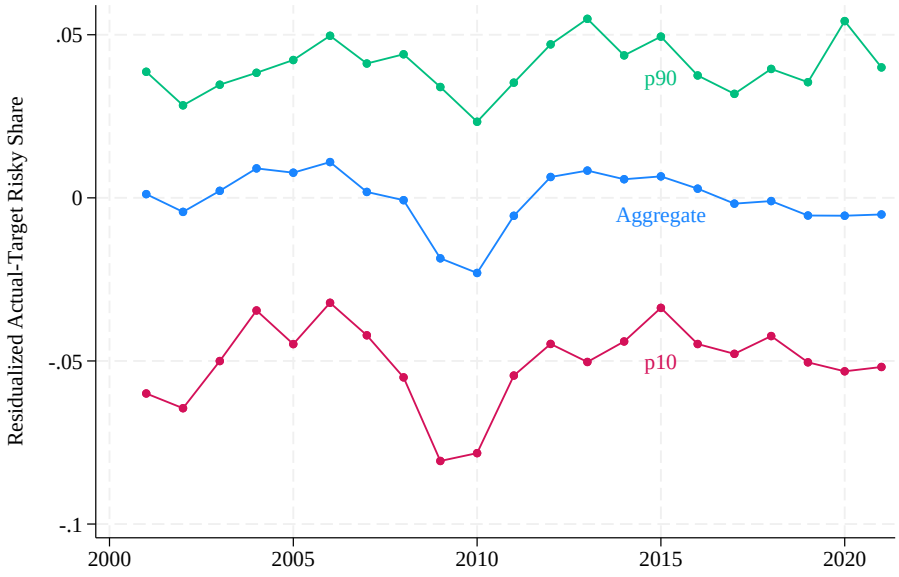

Notes: Panel (a) of the figure shows the aggregate, 10th, and 90th percentiles of l_{pt}, defined as the difference between the actual and target risky share for pension p in year t. Panel (b) plots the same series, after residualizing l_{pt} to the annual return of each pension. See Section 4.1 for details.

Figure A19: Funding-Based Explanations and Changes in the Risky Share

(a) GASB 25 Funding

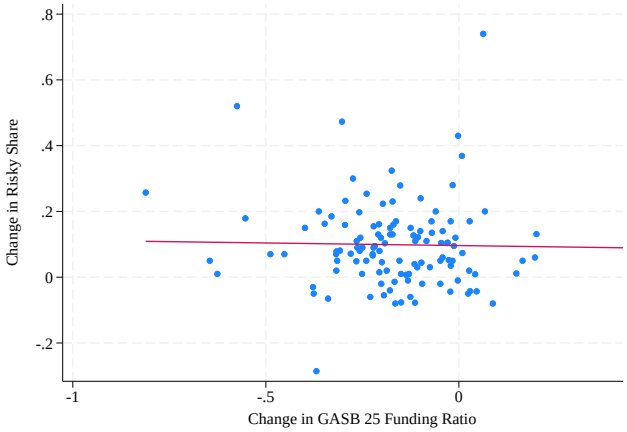

(b) Liability Discount Rate

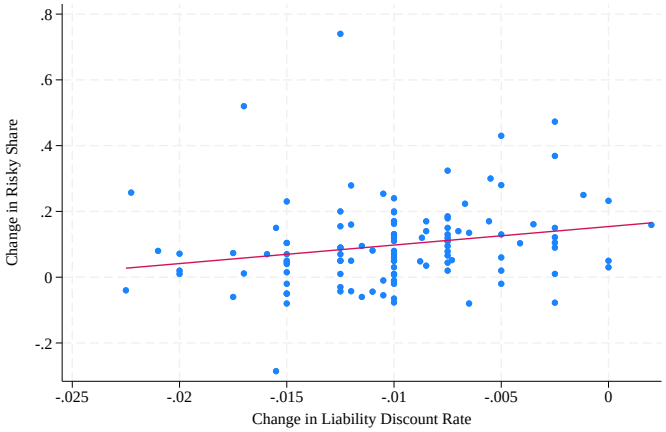

(c) Fraction of Retired Members

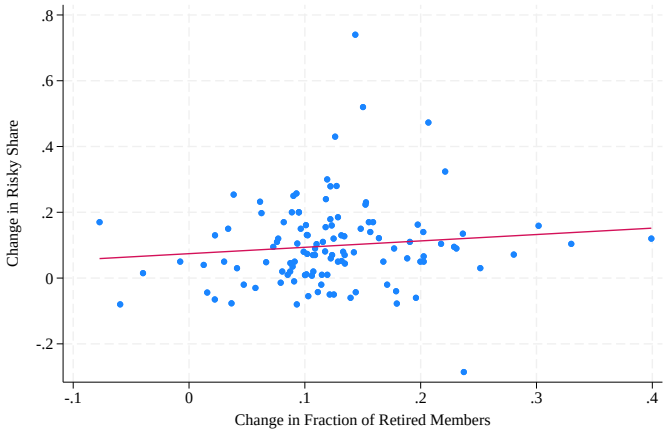

Notes: The figure contains scatter plots of changes in the target risky share from 2002 to 2021 against contemporaneous changes in GASB 25 funding (panel a), hurdle rates or, equivalently, liability discount rates (panel b), and the fraction of retired members (panel c). Data are from the PPD.

Figure A20: Distribution of admissible beliefs

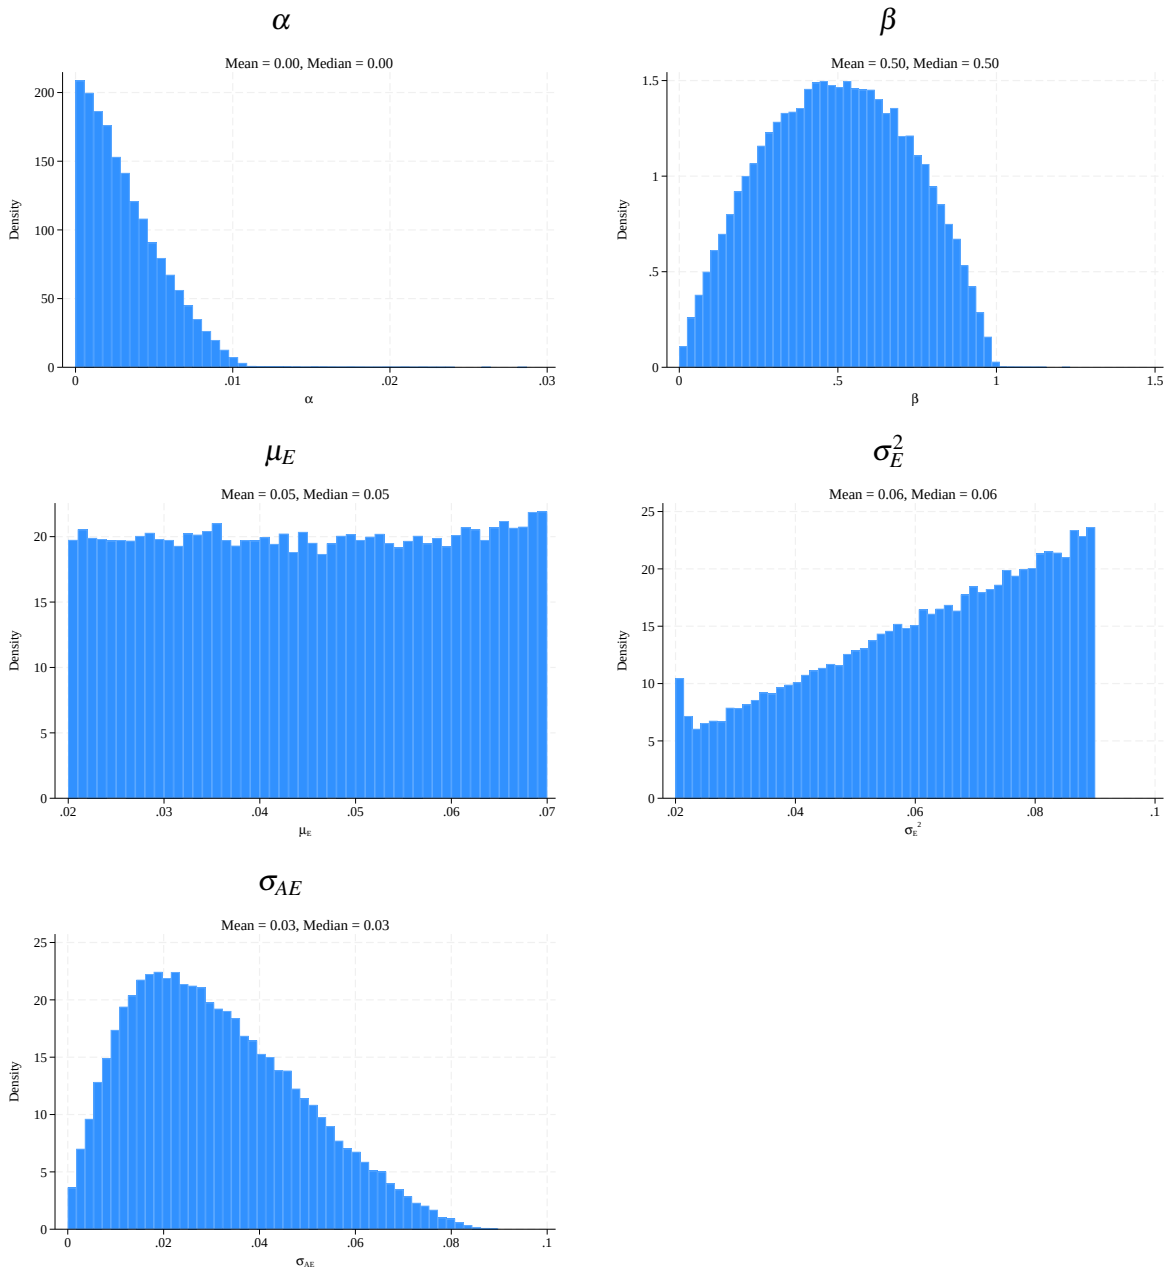

Notes: This figure plots the distributions of admissible beliefs when simulating the portfolio choice model of Section 2.3. See Section 5.1 of the main paper for details on the simulation.

Figure A21: Simulating a Decline in Risk Aversion Plus Binding Portfolio Constraints

(a) Implied Change in Risk Aversion

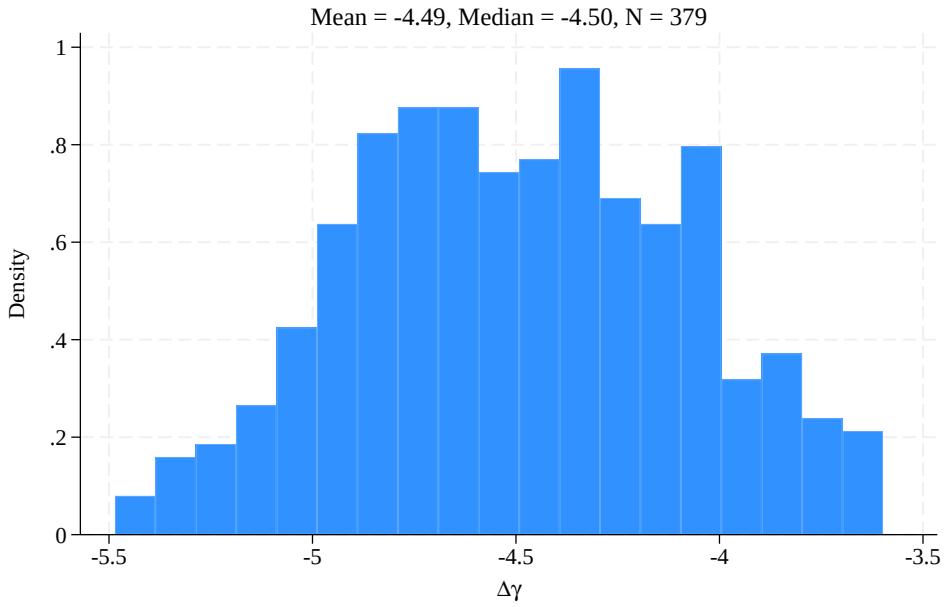

(b) Implied Shadow Cost

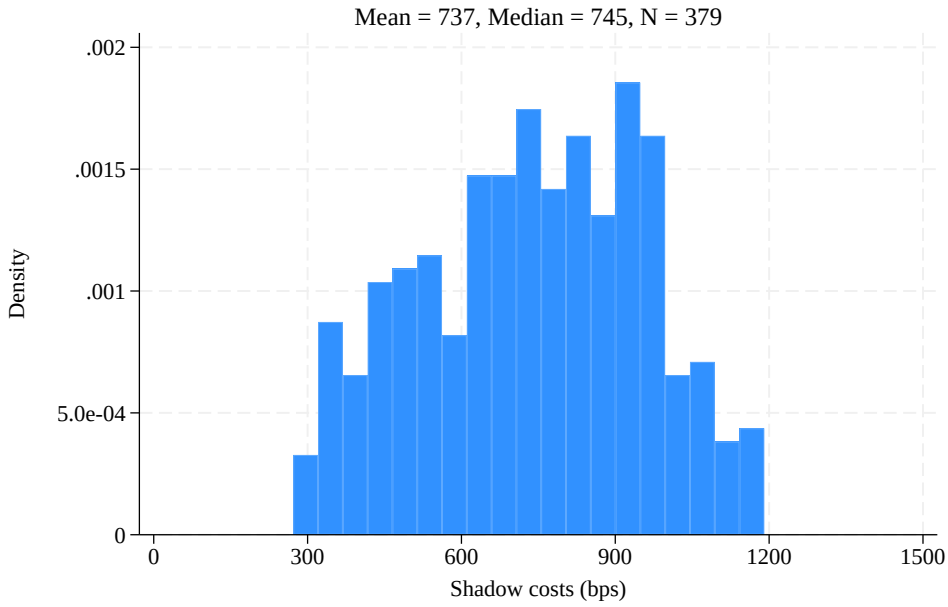

Notes: This figure shows simulation results based on the model in Section 2.3, in which pensions face a minimum constraint on fixed income investment. Panel (a) of the figure shows how the change in risk aversion required to match portfolio weights in 2001 and 2021 varies across simulations of initial beliefs in 2001. Panel (b) shows the distribution of the implied shadow cost of the portfolio constraint, expressed in units of returns. The simulations in this figure target aggregate portfolio weights, assume that the aggregate pension portfolio faced a minimum constraint on fixed income that became binding in 2021, and hold initial beliefs in 2001 constant. Out of 100,000 initial simulations that match the initial portfolio weights with reasonable initial beliefs in 2001, the plots in this figure show the 379 simulations in which a decline in risk aversion would explain the aggregate portfolio shifts. See Section 5.1 for details.

Figure A22: Simulated Change in Beliefs About α

(a) Required Increase in Alpha

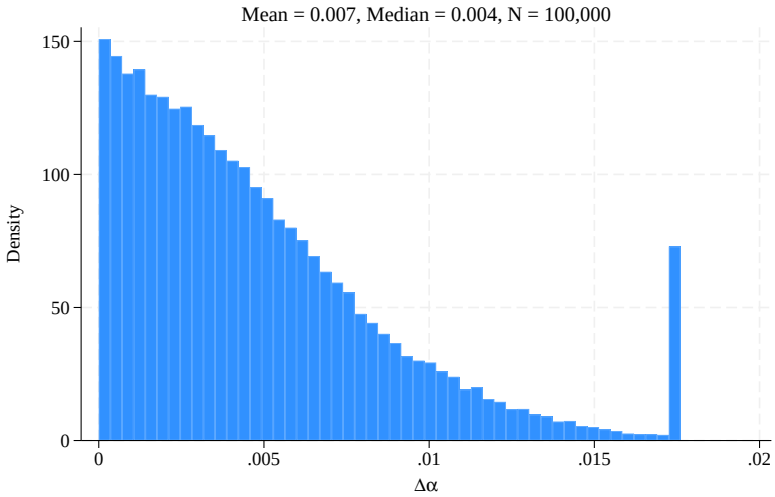

(b) Required Change in Risk Aversion

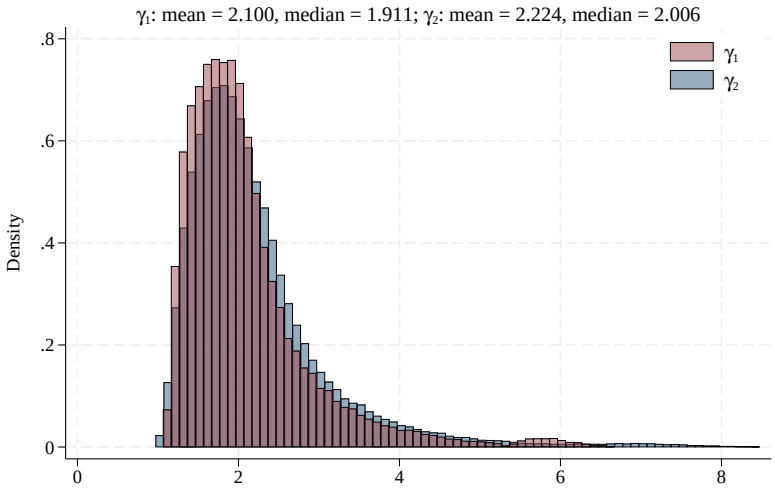

Notes: This figure shows belief change simulation results based on the model in Section 2.3. For the same initial sets of beliefs as used in Section 5.1 that match the initial portfolio weights in 2001, Figure A22a shows the change in CAPM-α required to match the change in the alternative-to-risky share from 2001 to 2021, $\Delta\omega_A^*$. Given the change in α from Figure A22a to match $\Delta\omega_A^*$, Figure A22b shows the distribution of γ_1 needed to match the 2001 risky-asset share and the distribution of γ_2 needed to match the 2021 risky-asset share. See Sections 5.1 and Internet Appendix G.2 for details.

Figure A23: Changes in the Alternative-to-Risky Share vs Changes in Perceived Alpha

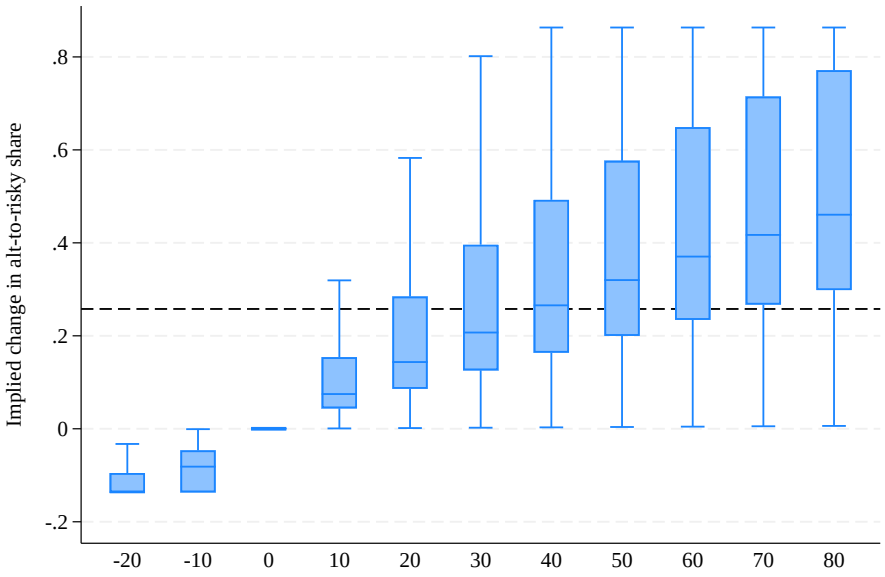

Notes: This plot shows how the aggregate alternative-to-risky share Ω_t depends on changes in the perceived alpha of alternatives within the simulated model from Section 2.3. For the same initial sets of beliefs as used in Section 5.1 that match the initial portfolio weights in 2001, we increase the perceived alpha by a fixed amount of $\Delta\alpha$ in each simulation. Each box plot in the figure depicts how the distribution of the implied change in the alternative-to-risky share $\Delta\Omega(\Delta\alpha)$ varies across simulations. See Sections 5.1 and Internet Appendix G.2 for details.

Figure A24: Implied Change in Pension's Private Beliefs from Simulations

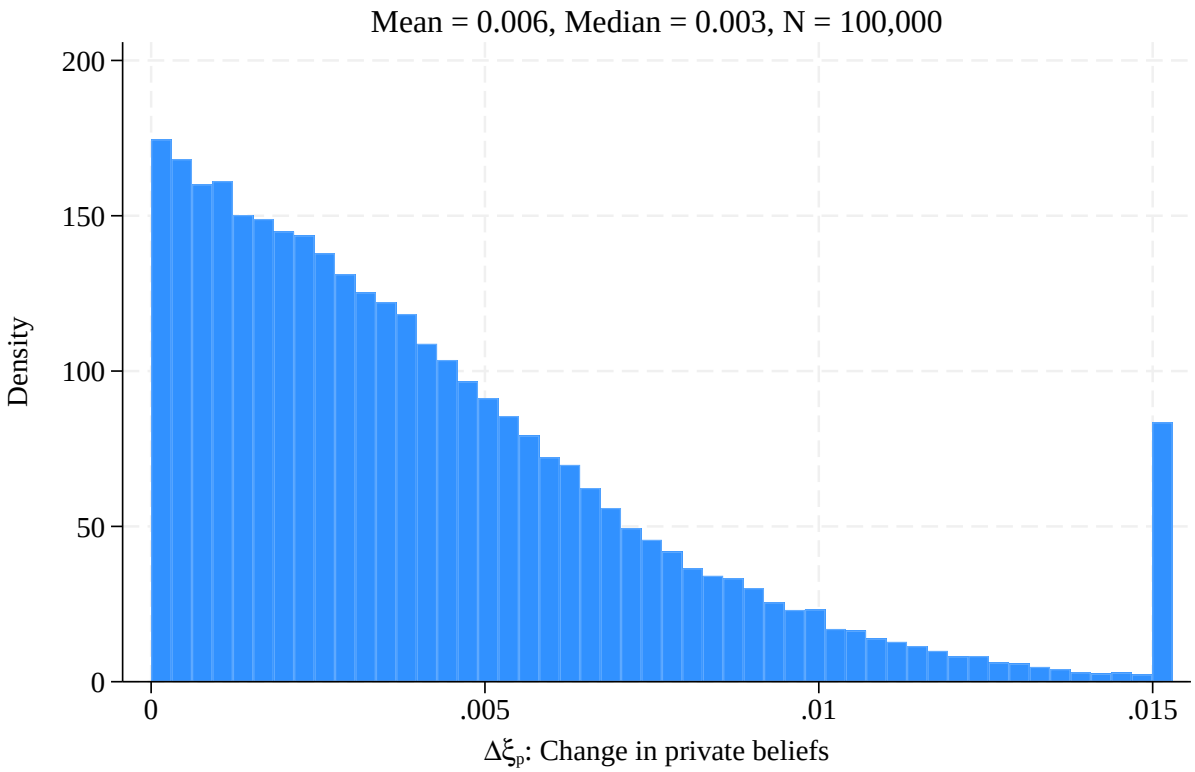

Notes: The figure plots the distribution of the changes in pensions' private beliefs $\Delta\xi_p$ from the model $\Delta\alpha_p = \zeta\Delta\alpha_c + \Delta\xi_p$, where $\Delta\alpha_c$ is the observed change in consultant-perceived alpha in Figure 4a and $\Delta\alpha_p$ is the required change in pension alpha in Figure A22a that matches the change in the alternative-to-risky share from 2001 to 2021in the aggregate portfolio. We construct ζ using the derivative of alt-to-risky share with respect to pension alpha $\frac{\partial\Omega}{\partial\alpha_p}$ from the simulations, as well as the estimated elasticity of the alternative-to-risky share to consultant beliefs from the IV strategy in Section 3.1. The results are based on 100,000 initial simulations that match the initial portfolio weights with reasonable initial beliefs in 2001. See Section G.2.2 for details. For illustration purposes, we winsorize data at the 97.5th percentile in the figure.

Figure A25: Consultant Beliefs and the Choice Between Opaque Assets

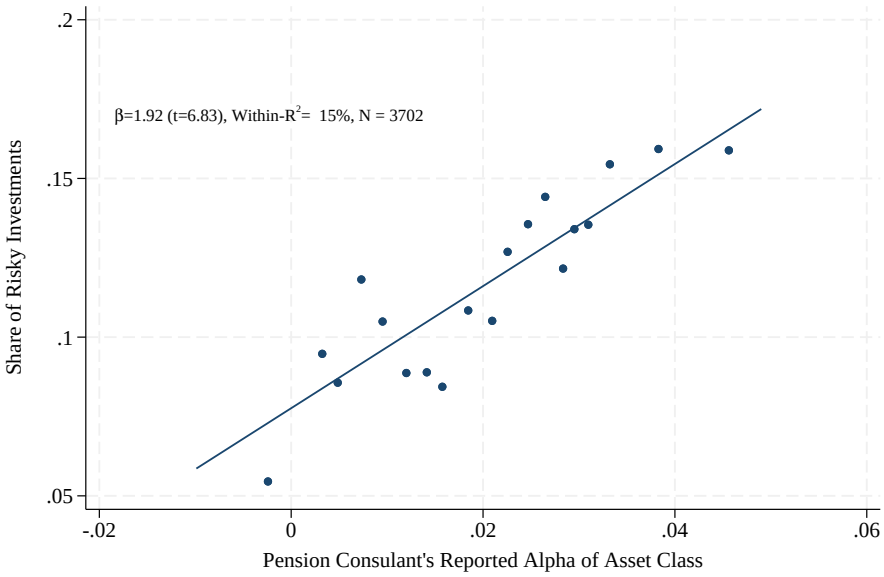

Notes: This plot shows a binscatter of each pension's allocation of risky investments to real assets or private equity against its consultant's reported alpha in each subcategory, after controlling for a pension-by-time fixed effect. The reported t-statistic is based on standard errors that are clustered by consultant and pension. See Section H.1.1 for details.

Figure A26: Diversification Motives

(a) Asset Class Breakdown

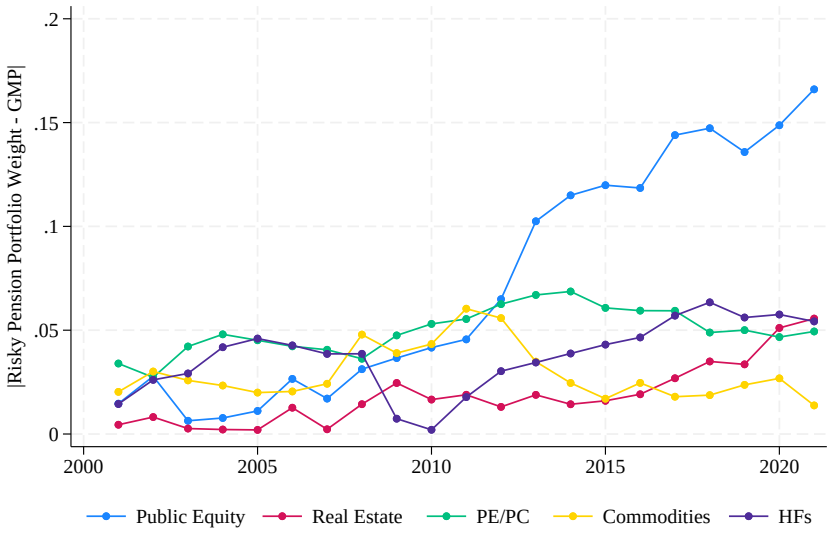

(b) The Cross-Section

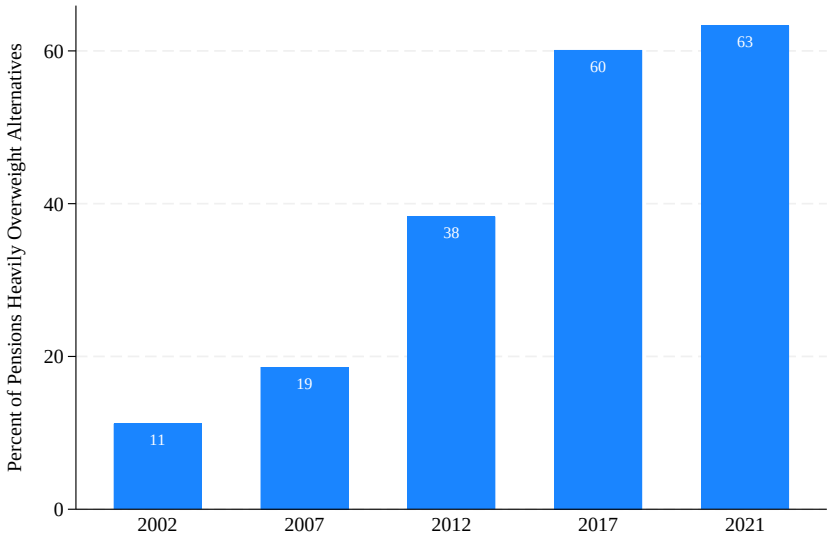

Notes: Panel (a) plots the absolute deviation between target and GMP portfolio shares for different subcomponents of the risky portfolio, with all shares measured relative to the risky portfolio. Panel (b) plots the percent of U.S. public pensions that are heavily overweight alternatives within their risky portfolios through time. In a given year, a pension is considered heavily overweight alternatives if its target alternative-to-risky share is at least 10 pp higher than the alternative-to-risky share in the GMP. Data for the GMP is taken from State Street, with adjustments made for REITs and hedge funds. See Section 6.1 and Internet Appendix H.2 for complete details of the data.

Table A1: Transition Rates for Quartiles of Portfolio Weights

(a) Risky Share

	2002				
2021	**1**	**2**	**3**	**4**	**Total**
	%	%	%	%	%
1	36	21	27	15	100
2	29	11	25	36	100
3	13	30	33	23	100
4	29	32	18	21	100
Total	27	24	26	24	100

(b) Alternative-to-Risky Share

	2002				
2021	**1**	**2**	**3**	**4**	**Total**
	%	%	%	%	%
1	33	30	23	13	100
2	30	20	23	27	100
3	17	33	27	23	100
4	21	14	29	36	100
Total	25	25	25	25	100

Notes: This table shows transition rates between quintiles of the share of the risky share (panel a) and the share of alternatives relative to risky investment (panel b) between 2002 and 2021. For each year y, we divide pensions into quartiles $q(y)$ based on their risky or alternatives-to-risky share. Quartile 1 contains the lowest set of shares and quartile 4 contains the highest. We then compute the percent of pensions in $q(y = 2021)$ who were in quintile $q(y = 2002)$. Rows in the table are normalized to sum to 1. Risky investments includes everything outside of fixed income and cash. Data are from the PPD and portfolio weights are target shares.

Table A2: Consultant Beliefs and Pension Allocations - Robustness

	Alternative-to-Risky Share$_{p,t}$			
	(1)	(2)	(3)	(4)
$\alpha_{c(p),t}$	3.26**	3.31**	3.01**	3.77**
	(2.85)	(3.04)	(3.80)	(4.72)
Estimation Type	OLS	OLS	IV	IV
Controls	No	Yes	Yes	Yes
Fixed-Effects	Time	Time	Time	Time
$p_{wild}(\beta = 0)$	0.01	0.01	0.02	0.01
IV Location Date			t	2005
First-Stage F			83.44	19.90
β_{IV}/β_{OLS}			0.91	1.18
Within-R^2	0.08	0.12		
N	1,946	1,932	1,932	1,617

Notes: This table summarizes linear regressions of pension p's target alternative-to-risky share in year t on the reported alpha of alternatives of pension p's consultant in year t. Pension controls include the level of pension funding, (log) size, hurdle rate, the ratio of required actuarial contributions too payroll, the ratio of administrative expenses to payroll, the cash share, annual return of the fund, and the residualized difference between each pension's target and actual risky share. Standard errors are clustered at the consultant and year levels, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. Missing target shares are imputed based on active changes in portfolio composition following Calvet et al. (2009). See Sections A.4 and C for more details.

Table A3: Consultants and the Composition of Risky Investments - Robustness

			Fixed Effects							
	Dep. Variable:	Controls	Time	State-Time	Consultant	F	p	Adj. R^2	C	N
(1)	Alternatives	x	x					0.31		3,004
(2)	Alternatives	x	x		x	18.72	0.00	0.48	54	2,999
(3)	Alternatives	x		x				0.49		2,609
(4)	Alternatives	x		x	x	14.09	0.00	0.62	50	2,603
(5)	Private Equity/Credit	x	x					0.19		3,004
(6)	Private Equity/Credit	x	x		x	12.28	0.00	0.34	54	2,999
(7)	Private Equity/Credit	x		x				0.32		2,609
(8)	Private Equity/Credit	x		x	x	10.40	0.00	0.45	50	2,603
(9)	Hedge Funds	x	x					0.15		3,004
(10)	Hedge Funds	x	x		x	12.11	0.00	0.29	54	2,999
(11)	Hedge Funds	x		x				0.29		2,609
(12)	Hedge Funds	x		x	x	9.45	0.00	0.42	50	2,603
(13)	Real Assets	x	x					0.15		3,004
(14)	Real Assets	x	x		x	15.70	0.00	0.34	54	2,999
(15)	Real Assets	x		x				0.28		2,609
(16)	Real Assets	x		x	x	14.54	0.00	0.47	50	2,603

Notes: This table shows fixed effects regressions of the following form:

$$\Omega^*_{p(c),t} = \lambda_t + \lambda_c + \Gamma X_{p,t} + \varepsilon_{p(c),t}$$

where $\Omega^*_{p(c),t}$ is one of several target asset allocations for pension system p matched with consultant c in fiscal year t. λ_t is a year or year-by-state fixed effect, $X_{p,t}$ is a vector of characteristics of pension p in year t, and λ_c is a consultant fixed effect. Control variables include (log) size, GASB 25 funding, asset hurdle rates, the ratio of required actuarial contributions to payroll, and the ratio of administrative expenses to payroll. Some regressions also include a pension fixed effect. The listed F-statistic is the result of testing the null hypothesis that the consultant effects λ_c are equal to each other. C is the number of consultant effects that are included in the test. Sample sizes differ across models because we drop singleton groups for any included fixed effects. All allocations are relative to the overall target share of risky investments. Missing target shares are imputed based on active changes in portfolio composition following Calvet et al. (2009). See Sections A.4 and C for more details.

Table A4: Experience in the 1990s - Robustness

	Δ Alt-to-Risky$_{2002\to2021}$			Δ Risky$_{2002\to2021}$
	(1)	(2)	(3)	(4)
Δ Equity$_{1996\to2002}$	0.53**	0.89**	0.96**	0.15
	(3.18)	(3.71)	(3.40)	(0.98)
Δ Equity$_{1990\to2002}$		-0.25	-0.28	0.15
		(-1.10)	(-0.99)	(0.96)
Fraction of Restricted Assets$_{1998}$			0.05	0.09
			(0.74)	(1.34)
Financial Sophistication$_{1998}$			0.01	0.02
			(0.07)	(0.63)
Political Board$_{2002}$			-0.01	0.05
			(-0.21)	(0.83)
Controls	No	No	Yes	Yes
R^2	0.17	0.21	0.27	0.45
N	47	46	46	46

Notes: This table shows OLS regressions of the following form:

$$\Delta\Omega_{s,2002\to2021} = c + \beta_1 \Delta\text{Equity}_{s,1996\to2002} + \beta_2 \Delta\text{Equity}_{s,1990\to2002} + \Gamma X_s + \varepsilon_s.$$

In columns (1)-(2) the dependent variable is the state-level change in the target alternative-to-risky share and in column (3) it is the change in the target risky share, both computed using data from the PPD between 2002 and 2021. All regressions include the change in the target share of public equities between 1996 and 2002 ($\Delta\text{Equity}_{s,1996\to2002}$). Column (1) shows univariate regression results. Column (2) controls for the change in the target share of public equities between 1990 and 2002 ($\Delta\text{Equity}_{s,1990\to2002}$). Columns (2)-(3) also control for the fraction of assets that are restricted by each state's constitution as of 1998, the change in hurdle rates, the change in the fraction of retired members, the level of BEA-adjusted funding in 2002, and the log of assets under management in 2002, with all changes measured between 2002 and 2021. Investment restrictions from 1996 are used when they are not available in the 1998 PENDAT survey. When constructing changes in target shares between t and T, we substitute active changes in portfolio shares, following Calvet et al. (2009), for states where target shares are missing for more than 50% of assets in either t or T. Robust standard errors are used in all regressions, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A5: Peer Effects - Robustness

	Alternative-to-Risky Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Peers' Alt-to-Risky Share	0.71**	0.57**	0.71**	0.75**		
	(3.16)	(2.93)	(3.13)	(3.34)		
× Established-CIO		0.05				
		(0.35)				
× Well-Funded			-0.05			
			(-0.42)			
× High-Performing				-0.20**		
				(-2.30)		
Lagged Peers' Alt-to-Risky Share					0.74**	
					(3.12)	
Non-Pension Peers' Alt-to-Risky Share						0.44**
						(2.30)
FE	(c,y,d)	(c,y,d)	(c,y,d)	(c,y,d)	(c,y,d)	(c,y)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.15	0.11	0.16	0.16	0.15	0.02
Total R^2	0.68	0.60	0.68	0.68	0.67	0.44
N	2,009	875	2,009	2,009	1,890	2,393

Notes: This table summarizes linear regressions of pension p's target alternative-to-risky share in year t on the target alternative-to-risky share of its peers. In columns (1)-(5), for a given pension, the target alternative-to-risky share of its peers is defined as a weighted average of all other pensions' target shares of alternatives, where weights are defined as the inverse of the distances between the given pension and each of its peers, respectively. In column (6), the alternative-to-risky share of each pension's peers is defined as the average (actual) alternative-to-risky share of unions, private-sector pensions, and endowments in its state. Columns (1)-(5) include consultant-by-time-by-census-division fixed effects and column (6) includes consultant-by-year fixed effects. All regressions include controls for funding ratio, asset hurdle rate, (log) size, required actuarial contributions relative to payroll, and administrative expenses relative to payroll. Columns (1) to (4) show the baseline relationship between the alternative-to-risky share for pension p in year t and its peers' share, as well as the interaction with several characteristics. In all cases, the regression includes both the characteristic and its interaction with peer target share of alternatives. The indicator variable "Established-CIO" in column (2) equals one if the pension has had the same CIO for at least five years. The regression in column (2) is estimated only using data after 2005. The indicator for "Well-Funded" in column (3) equals one if the pension is in the top quartile of funding for the entire sample. The indicator for "High-Performing" in column (4) equals one if the pension is in the top quartile of annual net-of-fee performance in a given year. Column (5) shows the regression of the alternative-to-risky share on the lagged peer target alternative share. Standard errors are clustered at the year and state levels, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. Missing target shares are imputed based on active changes in portfolio composition following Calvet et al. (2009). See Sections A.4 and C for more details.

Table A6: Risk-Seeking Explanations for the Rise in Alternatives - Robustness

	ΔAlternative-to-Risky Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ GASB 25 Funding Ratio	-0.10				-0.08		
	(-1.13)				(-0.92)		
Δ BEA-Adjusted Funding Ratio		-0.10					
		(-0.45)					
Δ Liability Discount Rate			-1.50			-1.60	
			(-0.44)			(-0.41)	
Δ Fraction of Retired Members				0.14			0.11
				(0.75)			(0.60)
Aggregation	System	State	System	System	System	System	System
Constraint Control	No	No	No	No	Yes	Yes	Yes
Total R^2	0.01	0.00	0.00	0.00	0.02	0.02	0.02
N	117	47	116	114	115	114	112

Notes: This table shows regressions of changes in the target alternative-to-risky share on contemporaneous changes in several covariates. All changes are computed between 2002 and 2021. When changes are not available, we instead use active changes in the alternative-to-risky share based on Calvet et al. (2009). See Sections A.4 and C for more details. Risky investments are any holdings outside of fixed income and cash. Alternatives encompass investments in hedge funds, private equity, private debt, and real assets. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate borrowing rates. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled "Aggregation" specifies whether the regression is run at the system or state level. Active changes in the alternative-to-risky share are used for states in which changes in target shares are missing for over 50% of assets. The row "Constraint Control" indicates whether controls are included for the average difference between each pension's actual and its target risky share between 2002 and 2020 (when both are available), as well as the change in their target cash share over the same period. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. See Section 2.1 for how we filter the data and Section 4.1 for more details on how we define the tightness of risk constraints.

Table A7: What Has Driven the Rise in Consultant-Perceived α?

	Expected Returns (R)	Public Equity RP (P)	Diversification (D)
All Alternatives	98	31	-38
PE	51	40	-9
RE	239	25	-117
HF	70	18	-27

Notes: This table decomposes the change in average consultant-perceived alpha between 2003 and 2021 into a portion coming from an increase in expected excess returns (Expected Returns), a decline in expected excess returns to public equities (Public Equity RP), and a decline in beta (Diversification). Negative values in the "Diversification" column indicate that the average beta with respect to public equities has increased, thereby contributing negatively to the increase in average alpha. The row listed as "All Alternatives" is based on an equal-weighted average of PE, RE, and HFs. See Internet Appendix D.2 for the derivation of the decomposition. Consultants must have at least ten years of data to be included in the analysis.

Table A8: Consultants as the Source of Peer Effects?

	Alternative-to-Risky Share		
	(1)	(2)	(3)
Peers' Alt-to-Risky Share	0.71**	0.63**	0.86**
	(3.30)	(2.32)	(3.84)
Peer Set	All Pensions	All Pensions	Different Consultant
FE	(c,y,d)	(y,d)	(c,y,d)
Controls	Yes	Yes	Yes
Within-R^2	0.15	0.12	0.15
Total R^2	0.68	0.45	0.68
N	2,039	2,039	2,039

Notes: This table summarizes linear regressions of pension p's target alternative-to-risky share in year t on the target alternative-to-risky share of its peers. In columns (1) and (2), target alternative-to-risky share of its peers is defined as a weighted average of all other pensions' target shares of alternatives, where weights are defined as the inverse of the distances between the given pension and each of its peers, respectively. In column (3), it is instead computed by averaging over pensions that do not share a consultant with p, with weights for the averaging again determined by inverse distance. These differences are reflected in the row labeled "Peer Set" in the table. Columns (1) and (3) include consultant-by-time-by-census-division fixed effects, whereas column (2) includes only a census-division-by-year fixed effect. Standard errors are clustered at the year and state levels, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A9: Risk-Seeking Explanations for the Rise in Alternatives - Post-Crisis Analysis

	ΔAlternative-to-Risky Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ GASB 25 Funding Ratio	0.16*					
	(1.93)					
Δ BEA-Adjusted Funding Ratio		-0.25				
		(-1.35)				
Δ Liability Discount Rate			2.92			
			(0.98)			
Δ Fraction of Retired Members				-0.30		
				(-1.09)		
GASB 25 Funding Ratio$_{2010}$					0.02	
					(0.36)	
BEA-Adjusted Funding Ratio$_{2010}$						0.04
						(0.29)
Aggregation	System	State	System	System	System	State
Total R^2	0.02	0.03	0.01	0.01	0.00	0.00
N	141	51	138	138	141	51

Notes: This table shows regressions of changes in the target alternative-to-risky share on contemporaneous changes in several covariates. All changes are computed between 2010 and 2021. Risky investments are any holdings outside of fixed income and cash. Alternatives encompass investments in hedge funds, private equity, private debt, and real assets. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate borrowing rates. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled "Aggregation" specifies whether the regression is run at the system or state level. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A10: Predicting Changes in the Alternative-to-Risky Share

	ΔAlternative-to-Risky Share (2002-2020)			
	(1)	(2)	(3)	(4)
GASB 25 Funding Ratio in 2002	0.14*			
	(1.81)			
BEA-Adjusted Funding Ratio in 2002		-0.17		
		(-0.76)		
Liability Discount Rate in 2002			2.39	
			(0.57)	
Fraction of Retired Members in 2002				-0.10
				(-0.64)
Aggregation	System	State	System	System
Total R^2	0.02	0.01	0.00	0.00
N	118	47	117	118

Notes: This table shows regressions of changes in the target alternative-to-risky share between 2002 and 2021 on the level of several covariates in 2002. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate yield curves. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's asset return target. The row labeled aggregation specifies whether the regression is run at the system or state level. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. See Section 2.1 for detailed asset class definitions and how we filter the data.

Table A11: Variation in the Level of the Alternative-to-Risky Share

	Alternative-to-Risky Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GASB 25 Funding Ratio	-0.01				-0.01			
	(-0.18)				(-0.17)			
BEA-Adjusted Funding Ratio		-0.03				-0.11		
		(-0.40)				(-0.99)		
Liability Discount Rate			3.37				1.52	
			(1.54)				(0.32)	
Fraction of Retired Members				-0.02				-0.07
				(-0.22)				(-0.52)
Aggregation	System	State	System	System	System	State	System	System
Time-FE	Yes	Yes	Yes	Yes	No	No	No	No
Sample	Full	Full	Full	Full	2020	2020	2020	2020
Within-R^2	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.00
Total-R^2	0.31	0.39	0.32	0.32	0.00	0.01	0.00	0.00
N	3,122	1,008	3,091	3,105	158	51	157	155

Notes: This table shows panel regressions of the target alternative-to-risky share on several covariates. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate yield curves. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled aggregation specifies whether the regression is run at the system or state level. The first four columns are a panel regression that includes a time fixed effect. The last four columns run the regression using only 2021 data, the last year in which BEA-adjusted fund ratios are available. Standard errors are clustered by state and time. *t*-statistics are listed below point estimates. ** indicates a *p*-value of 0.05 and * indicates a *p*-value of 0.10.

Table A12: Spending Constraints

(a) Predictive Regressions

	$\text{FMRC}_{p,2012-2021}$		$\text{CPR}_{p,2012-2021}$	
	(1)	(2)	(3)	(4)
$\Omega_{p,2011}$	0.01		0.21	
	(0.04)		(1.43)	
$\Delta\Omega_{p,2002\to2011}$		-0.15		0.19
		(-0.55)		(1.37)
R^2	0.00	0.00	0.02	0.02
N	152	120	152	120

(b) Current Spending

	$\Omega_{p,t}$			$\Delta\Omega_p$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{CPR}_{p,t}$	0.02	0.00	0.01			
	(1.32)	(0.27)	(0.99)			
$\Delta\ \text{CPR}_p$				0.00	-0.00	-0.01
				(0.09)	(-0.04)	(-0.69)
Alternative Type	All	PE	RA	All	PE	RA
Total-R^2	0.32	0.10	0.14	0.00	0.00	0.00
Within-R^2	0.00	0.00	0.00			
N	3,122	3,122	3,122	118	118	118

Notes: This table analyzes the link between spending and alternative use. Panel (a) shows a series of predictive regressions at the pension level. The outcome variable is either the average contribution-to-payout ratio from 2012 to 2021 ($\Delta CR_{2012\to2022}$) or the fraction of years between 2012 and 2021 in which a plan sponsor fails to make its actuarial required contribution (ARC) under GASB standards. The main regressors in panel (a) are either the alternative-to-risky share in 2011 ($\Omega_{p,2011}$) or its change between 2002 and 2011 ($\Delta\Omega_{p,2002\to2011}$). Panel (b) explores the relationship between current spending and alternative use. The first three columns show panel regressions of the alternative-to-risky, PE-to-risky, and RA-to-risky shares, respectively, on the contemporaneous ratio of contributions to beneficiary payouts ($CPR_{p,t}$). These regressions include a time fixed effect. The last three columns show regressions of long-run changes of the alternative-to-risky, PE-to-risky, and RA-to-risky shares, respectively, on contemporaneous changes in the contributions to beneficiary payouts (ΔCPR_p). Long-run changes are computed from 2002 to 2021. Standard errors in the last three specifications are clustered by state. Standard errors in panel (a) are clustered by state. In panel (b), standard errors in the first three columns are clustered by time and state. In the last three columns, they are clustered by state. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A13: Funding-Related Measures and Changes in the Risky Share

	Δ Risky Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ GASB 25 Funding Ratio	-0.02							
	(-0.18)							
Δ BEA-Adjusted Funding Ratio		-0.25						
		(-1.10)						
Δ Liability Discount Rate			5.63**					
			(2.30)					
Δ Fraction of Retired Members				0.19				
				(1.14)				
GASB 25 Funding Ratio in 2002					0.07			
					(0.85)			
BEA-Adjusted Funding Ratio in 2002						-0.08		
						(-0.46)		
Liability Discount Rate in 2002							-5.43	
							(-1.30)	
Fraction of Retired Members in 2002								-0.32*
								(-1.76)
Aggregation	System	State	System	System	System	State	System	System
Total R^2	0.00	0.06	0.04	0.01	0.01	0.00	0.02	0.05
N	119	47	118	116	119	47	118	119

Notes: The first four columns of the table show regressions of changes in the target risky share on contemporaneous changes in several covariates, whereas the last four columns show regressions on the 2002 level of the covariates. Risky assets are defined as everything outside of fixed income and cash. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate borrowing rates. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled aggregation specifies whether the regression is run at the system or state level. All changes are computed between 2002 and 2021. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A14: Realized Pension Performance in Different Asset Classes

(a) Realized Pension Performance by Alternative Type

	Alternative Asset			Public Equities		
	Excess Return	Volatility	SR	Excess Return	Volatility	SR
Private Equity	11.6	15.9	0.74	8.4	17.3	0.50
Real Estate	7.4	11.6	0.65	8.2	17.7	0.45
Hedge Funds	4.5	7.6	0.62	10.6	15.7	0.68

(b) Portfolio Volatility of Low vs High-Alternative Users

	Alternative Use	
Parameter	Low	High
Equity Weight	0.59	0.43
Alternative Weight	0.11	0.33
Fixed Income Weight	0.30	0.24
Equity Volatility	0.17	0.18
Alternative Volatility	0.12	0.11
Fixed Income Volatility	0.04	0.05
Alt/Equity Correlation	0.55	0.68
Equity/FI Correlation	–0.09	0.11
Alt/FI Correlation	–0.11	0.05
Portfolio Volatility	**10.76%**	**10.73%**

Notes: Panel (a) shows the average annual excess return (%), volatility (%), and Sharpe ratio (SR) of different asset classes based on U.S. public pension returns from 2001 to 2021. For each alternative asset class a, we retain observations where pension returns are available for both a and public equities. Within each pension, we then compute the average excess return, volatility, and Sharpe ratio of a and public equities using the same overlapping sample period. The table then reports the median of each of these statistics across pensions for both a and public equities. Panel (b) computes the total portfolio volatility of two hypothetical portfolios, one with a low alternative share and one with a low share. The parameters in the example are chosen to roughly match their average sample counterparts in the data. Low (high) alternative users refer to pensions in the lowest (highest) quartile of average alternative-to-risky share over the sample period.

Table A15: Board Composition and the Alternative-to-Risky Share

	Alternative-to-Risky Share	
	(1)	(2)
State-appointed	0.19	0.18
	(1.67)	(1.65)
State-exofficio	0.04	0.03
	(0.68)	(0.50)
Participant-elected	0.05	0.03
	(0.98)	(0.60)
Public-appointed	0.07	0.07
	(1.07)	(1.00)
Year-FE	Yes	Yes
Controls	No	Yes
Within-R^2	0.06	0.07
Total R^2	0.40	0.40
N	2,516	2,502

Notes: This table shows panel regressions of the target alternative-to-risky share on board composition. State-appointed, State-exofficio, Participant-elected, and Public-appointed respectively measure the percentage of: (i) appointed board members that are state officials; (ii) ex officio board members who are state officials; (iii) board members who are elected by plan participants; and (iv) board members that are appointed by the general public. We also control for the percentage representation of State-elected, Participant-exofficio, Public-exofficio and Public-elected board members. The omitted category is Participant-appointed. Column (1) has no additional controls and column (2) adds controls for GASB 25 funding, liability discount rate, each system's asset hurdle rate (liability discount rate), log assets, required contributions scaled by payroll, total spending scaled by payroll, and total board size. Board composition is based on Andonov et al. (2018) and is only available from 2001 to 2020. All regressions include a year fixed effect and are run at the system level. Standard errors are clustered by pension system and time. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A16: Home Bias

	Δ Alt-to-Risky$_{2002\rightarrow2021}$	
	(1)	(2)
Fraction of PE/VC Funds Domiciled in State	-0.03	
	(-0.16)	
Fraction of PE/VC AUM Domiciled in State		-0.03
		(-0.23)
Constant	0.26**	0.26**
	(9.40)	(9.61)
R^2	0.00	0.00
N	47	47

Notes: This table shows state-level regressions of changes in the target alternative-to-risky between 2002 and 2021 on two state-level measures of alternative activity. The first is the fraction of private-capital funds domiciled in each state with vintages between 2000 and 2022. The second is the fraction of private-capital AUM domiciled in each state with vintages between 2000 and 2022. Both measures are based on data from Preqin, which is sourced from WRDS, and focuses mainly on private equity and venture capital funds. All changes are computed between 2010 and 2021. t-statistics based on robust standard errors are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table A17: Diversification Benefits in the Cross-Section

	ΔAlternative-to-Risky Share$_{p,2002\to2021}$			
	(1)	(2)	(3)	(4)
$D_{p,2002}$ (Target)	-1.57			
	(-0.86)			
$D_{p,2002}$ (Actual)		-1.46		
		(-0.82)		
$\Delta D_{p,2002\to2021}$ (Target)			2.41**	
			(5.27)	
$\Delta D_{p,2002\to2021}$ (Actual)				2.35**
				(9.07)
Constant	0.27**	0.27**	0.18**	0.19**
	(10.74)	(10.56)	(12.02)	(13.72)
Total R^2	0.01	0.01	0.44	0.39
N	118	118	118	118

Notes: This table shows OLS regressions for the form:

$$\Delta\Omega_{p,2002\to2021} = \beta_0 + \beta_1 X_p + \varepsilon_s,$$

where $\Delta\Omega_{p,2002\to2021}$ is the change in p's target alternative-to-risky share between 2002 and 2021 and X_p is one of several pension-specific measures of portfolio diversification. In column (1), X_p measures the distance between pension p's target risky portfolio and that of the global market portfolio (GMP) in 2002, with distance defined using the "Active Share" metric from Cremers and Petajisto (2009). Column (2) instead measures 2002 distance using actual risky portfolio weights. In column (3), X_p is the change in the distance between each pension's target risky portfolio and the GMP between 2002 and 2021. In column (4), X_p is the change in the distance between each pension's actual risky portfolio and the GMP between 2002 and 2021. The GMP is defined using an adjusted version of State Street's Global Market Portfolio (see Section 6.1 and Internet Appendix H.2 for more details). The risky portfolio is defined as all investments outside of fixed income. t-statistics based on standard errors that are clustered by state are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.