Internet Appendix:

The Rise of Alternatives

Juliane Begenau[*] Pauline Liang[†] Emil Siriwardane[‡]

Abstract

In this Internet Appendix, we provide: (i) details of all data used in our analysis, as well as additional facts about the U.S. pension sector; (ii) additional evidence for our analysis of beliefs and risk-seeking motives; (iii) the full set of results from our simulation of a mean-variance model of the aggregate pension portfolio; (iv) evidence on agency-based explanations for the rise of alternatives, and (v) all derivations for the mean-variance model 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.

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 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 75% of pension assets in the U.S..

As mentioned in Section 4.1.2, 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 4.1.2. 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 $\mu_{Geometric} = \mu_{Arithmetic} - 0.5\sigma^2$. 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 4.2, 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, with each wave asking 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). The 1995 file of the PENDAT data contains survey data from the 1994 fiscal year, as well as some time-series data for fiscal years between 1990 and 1994. We specifically extract data on the share of equities held in FY1990 by each reporting pension and aggregate it to the state level based on 1990 AUM. We aggregate to the state level to avoid fuzzy matching between pension systems in the PENDAT and PPD datasets. Based on user guides of the PENDAT data, 1990 portfolio shares reflect actual weights, not target. Across all PENDAT data, the user guides indicate that values of -9 means missing and -1 means not applicable. We set both instances to missing.

To construct 1996 equity shares, we use the 1997 PENDAT file, which clearly indicates that it contains data for FY1996. Here, both target and actual shares are available. For our baseline analysis, we use actual shares because they are much more populated than target shares (roughly 8% missing vs 30%).

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, as this is the last wave containing answers to either.

Figure A2 provides 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.

A.5 S&P Money Market Directory

In Section 4.3, we compute the average alternative-to-risky and risky share for each consultant's U.S. private-sector clients (endowments, corporate pensions, and union-sponsored plans). The data for this 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.[1] 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 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. After matching plans to consultants, we collapse portfolio shares to the consultant-year level by taking simple averages. We require that a consultant has at least two private-sector clients in each year.

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

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 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, miscellaneous, and unquoted/private equities.[2] 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.[3] 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.

[2] The miscellaneous share is a small part of the portfolio. It never exceeds 6% and has been around 1% since 2016.
[3] See https://www.nacubo.org/Research/2022/Public-NTSE-Tables.

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 A3a 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.

Figure A3b 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.[4] 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 A3b also further confirms that the decline in fixed income since

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

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 A4a 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 A4a 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 A4b 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).[5] 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 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.[6]

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 15 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

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

[6] 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.

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.

Figure A5 provides a geographical sense of risk-taking by aggregating risky shares to the state level. The graph is broken out into four subfigures, one for each Census region in the United States. Blue squares in the plot depict the risky share of each state in 2002 and red circles depict the 2021 share. The average change in the risky share for states in the Midwest, Northeast, South, and West was 4 pp, 5 pp, 18 pp, and 11 pp, respectively. The risky share in each of these regions as of 2021 was 73%, 78%, 78%, and 78%, respectively.

B.2 Heterogeneity in the Composition of Risky Investments

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.[7] Similarly, the Indiana Public Retirement System has targeted roughly 70% of its risky investments in alternatives since 2012.[8]

Figure A6 shows how the alternative-to-risky share varies across states in both 2002 (blue squares) and 2021 (red circles). Over this period, states in the Midwest, Northeast, South, and West increased their alternative-to-risky shares by 20 pp, 29 pp, 26 pp, and 27 pp, respectively. Their respective alternative-to-risky shares in 2021 were 37%, 47%, 34%, and 42%.

For completeness, Figure A7 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

[7] 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.

[8] 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.

share between the 10th and 90th percentile pension was only 18 pp, roughly half the analogous spread in 2021.

C Belief-Based Explanations: Additional Evidence

C.1 Statistical model of consultant effects

In Section 4.1.2 of the main text, we use a simple statistical model of the alternative-to-risky share to motivate our empirical analysis of consultants. There, we assume that the alternative-to-risky share for pension p can be written as:

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

where $\alpha_{c(p)}$ is its consultant's belief about alpha and ξ_p is a pension-specific belief. ν_p reflects is any non-belief preferences for alternatives, such as those arising from agency-based or risk-seeking incentives. We now show how this equation can be derived from the simple portfolio choice model outlined in Section 3. In particular, we make use of a first-order Taylor approximation of the alternative-to-risky ratio, $\omega_A^*(\alpha_p, \beta_p)$, around known values of $\alpha_0 = 0$ and $\beta_0 = 1$ (the CAPM benchmark). To avoid notational clutter, we drop the subscripts p in our subsequent derivations.

To derive the first-order Taylor expansion of $\omega_A^*(\alpha, \beta)$ around $(\alpha_0, \beta_0) = (0, 1)$, we will start with the expression:

$$\omega_A^*(\alpha, \beta) = \frac{\omega_A(\alpha, \beta)}{\omega_A(\alpha, \beta) + \omega_E(\alpha, \beta)},$$

where

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

and

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

We need to evaluate the function and its first-order partial derivatives at the point $(\alpha_0, \beta_0) = (0, 1)$.

First, let's evaluate ω_A and ω_E at $(\alpha_0, \beta_0) = (0, 1)$:

$$\omega_A(0,1) = \frac{1}{\gamma}\left(0+0+\frac{1}{2}\right) = \frac{1}{2\gamma},$$

$$\omega_E(0,1) = \frac{1}{\gamma}\left(\frac{\mu_E}{\sigma_E^2}+0+0\right) = \frac{\mu_E}{\gamma\sigma_E^2}.$$

So, the value of $\omega_A^*(\alpha,\beta)$ at $(\alpha_0,\beta_0) = (0,1)$ is:

$$\omega_A^*(0,1) = \frac{\frac{1}{2\gamma}}{\frac{1}{2\gamma}+\frac{\mu_E}{\gamma\sigma_E^2}} = \frac{\frac{1}{2}}{\frac{1}{2}+\frac{\mu_E}{\sigma_E^2}} = \frac{1}{1+2\frac{\mu_E}{\sigma_E^2}}.$$

Next, we calculate the partial derivatives of $\omega_A^*(\alpha,\beta)$ with respect to α and β. The partial derivative with respect to α:

$$\frac{\partial\omega_A^*}{\partial\alpha} = \frac{(\omega_A(\alpha,\beta)+\omega_E(\alpha,\beta))\frac{\partial\omega_A(\alpha,\beta)}{\partial\alpha} - \omega_A(\alpha,\beta)\left(\frac{\partial\omega_A(\alpha,\beta)}{\partial\alpha}+\frac{\partial\omega_E(\alpha,\beta)}{\partial\alpha}\right)}{(\omega_A(\alpha,\beta)+\omega_E(\alpha,\beta))^2},$$

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

Substituting these into the expression for $\frac{\partial\omega_A^*}{\partial\alpha}$ and evaluating at $(\alpha_0,\beta_0) = (0,1)$:

$$\left.\frac{\partial\omega_A^*}{\partial\alpha}\right|_{(0,1)} = \frac{\left(\frac{1}{2\gamma}+\frac{\mu_E}{\gamma\sigma_E^2}\right)\frac{1}{\gamma\sigma_\eta^2} - \frac{1}{2\gamma}\left(\frac{1}{\gamma\sigma_\eta^2}-\frac{1}{\gamma\sigma_\eta^2}\right)}{\left(\frac{1}{2\gamma}+\frac{\mu_E}{\gamma\sigma_E^2}\right)^2},$$

which can be simplified to

$$\left.\frac{\partial\omega_A^*}{\partial\alpha}\right|_{(0,1)} = \frac{1}{\sigma_\eta^2\left(\frac{1}{2}+\frac{\mu_E}{\sigma_E^2}\right)}$$

The partial derivative with respect to β is

$$\frac{\partial\omega_A^*}{\partial\beta} = \frac{(\omega_A(\alpha,\beta)+\omega_E(\alpha,\beta))\frac{\partial\omega_A(\alpha,\beta)}{\partial\beta} - \omega_A(\alpha,\beta)\left(\frac{\partial\omega_A(\alpha,\beta)}{\partial\beta}+\frac{\partial\omega_E(\alpha,\beta)}{\partial\beta}\right)}{(\omega_A(\alpha,\beta)+\omega_E(\alpha,\beta))^2},$$

where $\frac{\partial\omega_A(\alpha,\beta)}{\partial\beta} = \frac{1}{\gamma}\left(\beta\frac{\sigma_E^2}{\sigma_\eta^2}-\frac{1}{2}\frac{\sigma_E^2}{\sigma_\eta^2}\right)$, and $\frac{\partial\omega_E(\alpha,\beta)}{\partial\beta} = \frac{1}{\gamma}\left(-\frac{\alpha}{\sigma_\eta^2}+\frac{1}{2}(1-\beta)\left(2\beta\frac{\sigma_E^2}{\sigma_\eta^2}\right)-\frac{1}{2}\left(\beta^2\frac{\sigma_E^2}{\sigma_\eta^2}+1\right)\right)$. The two partial derivatives evaluated at $(\alpha_0,\beta_0) = (0,1)$ are

$$\left.\frac{\partial\omega_A(\alpha,\beta)}{\partial\beta}\right|_{(0,1)} = \frac{1}{2\gamma}\frac{\sigma_E^2}{\sigma_\eta^2}, \quad \left.\frac{\partial\omega_E(\alpha,\beta)}{\partial\beta}\right|_{(0,1)} = -\frac{1}{2\gamma}\left(\frac{\sigma_E^2}{\sigma_\eta^2}+1\right).$$

Substituting these into the expression for $\frac{\partial \omega_A^*}{\partial \beta}$ and evaluating at $(\alpha_0, \beta_0) = (0,1)$ yields:

$$\left.\frac{\partial \omega_A^*}{\partial \beta}\right|_{(0,1)} = \frac{\left(\frac{1}{2\gamma} + \frac{\mu_E}{\gamma\sigma_E^2}\right)\left(\frac{1}{2\gamma}\frac{\sigma_E^2}{\sigma_\eta^2}\right) + \left(-\frac{1}{2\gamma}\left(\frac{1}{2\gamma}\frac{\sigma_E^2}{\sigma_\eta^2} - \frac{1}{2\gamma}\frac{\sigma_E^2}{\sigma_\eta^2} - \frac{1}{2\gamma}\right)\right)}{\left(\frac{1}{2\gamma} + \frac{\mu_E}{\gamma\sigma_E^2}\right)^2},$$

$$\left.\frac{\partial \omega_A^*}{\partial \beta}\right|_{(0,1)} = \frac{\sigma_E^2 + 2\mu_E + \sigma_\eta^2}{4\sigma_\eta^2\left(\frac{1}{2} + \frac{\mu_E}{\sigma_E^2}\right)^2}.$$

As a result, the first-order Taylor expansion of $\omega_A^*(\alpha, \beta)$ around $(\alpha_0, \beta_0) = (0,1)$ is:

$$\omega_A^*(\alpha, \beta) \approx \omega_A^*(0,1) + \left.\frac{\partial \omega_A^*}{\partial \alpha}\right|_{(0,1)} \alpha + \left.\frac{\partial \omega_A^*}{\partial \beta}\right|_{(0,1)} (\beta - 1).$$

Substituting the values yields

$$\omega_A^*(\alpha, \beta) \approx \frac{1}{1 + 2\frac{\mu_E}{\sigma_E^2}} + \frac{1}{\sigma_\eta^2\left(\frac{1}{2} + \frac{\mu_E}{\sigma_E^2}\right)}\alpha + \frac{\sigma_E^2 + 2\mu_E + \sigma_\eta^2}{4\sigma_\eta^2\left(\frac{1}{2} + \frac{\mu_E}{\sigma_E^2}\right)^2}(\beta - 1).$$

Assuming no change in beliefs about equity performance and β, we can write

$$\omega_A^*(\alpha) \approx \text{constant} + \pi\alpha \tag{1}$$

with the constant being $\frac{1}{1+2\frac{\mu_E}{\sigma_E^2}} + \frac{\sigma_E^2+2\mu_E+\sigma_\eta^2}{4\sigma_\eta^2\left(\frac{1}{2}+\frac{\mu_E}{\sigma_E^2}\right)^2}(\beta - 1)$ and $\pi = \left(\sigma_\eta^2\left(\frac{1}{2} + \frac{\mu_E}{\sigma_E^2}\right)\right)^{-1}$.

Next, assume each pension's total belief about alpha α can be decomposed as:

$$\alpha = \lambda\alpha_c + (1-\lambda)\xi + \nu, \tag{2}$$

where α_c is the belief of the pension's consultant, ξ is a pension-specific belief, and ν is a non-belief motive for wanting alternatives. ν should be thought of as being positive when a pension knowingly inflates the belief about α that it plugs into the mean-variance model because it wants to invest in alternatives for agency or risk-seeking reasons. λ reflects the relative weight that the pension puts on its consultant's belief about alpha. Plugging Equation (2) in (1) yields the linear expression used in Equation (6) in the main text, with $\theta = \pi\lambda$ and $\kappa = \pi(1-\lambda)$.

C.2 Variance Decompositions

In this subsection, we use the statistical framework from Section 4.1.2 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.

C.2.1 Simple Case

Assuming ν_p and ξ_p are orthogonal, Equation (6) implies that the total cross-sectional variance of ω_p^* is given by:

$$\mathbb{V}(\omega_p^*) = \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + \mathbb{V}(\pi\nu_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \pi\nu_p), \tag{3}$$

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:

$$B = \frac{\mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p)}{\mathbb{V}(\omega_p^*)}. \tag{4}$$

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 (3) and (4) 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:

$$\begin{aligned}\mathbb{V}(\lambda_c) &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\mathbb{E}_c[\kappa\xi_p]) + 2\mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]) \\ &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) - \mathbb{E}[\mathbb{V}_c(\kappa\xi_p)] + 2\mathbb{C}(\theta\alpha_{c(p)}, \mathbb{E}_c[\kappa\xi_p]) \\ &= \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) - \mathbb{E}[\mathbb{V}_c(\kappa\xi_p)] + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p) \\ &\leq \mathbb{V}(\theta\alpha_{c(p)}) + \mathbb{V}(\kappa\xi_p) + 2\mathbb{C}(\theta\alpha_{c(p)}, \kappa\xi_p), \end{aligned} \tag{5}$$

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 (5) 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, $\mathbb{E}[\mathbb{V}_c(\kappa\xi_p)]$.

C.2.2 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 C.2.1 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, ν_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

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

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

$$\begin{aligned}\mathbb{V}_{X_p}(\lambda_c) &= \mathbb{V}_{X_p}(\theta\alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa\xi_p) - \mathbb{E}_{X_p}[\mathbb{V}_{c,X_p}(\kappa\xi_p)] + 2\mathbb{C}_{X_p}(\theta\alpha_{c(p)},\kappa\xi_p)\\ &\leq \mathbb{V}_{X_p}(\theta\alpha_{c(p)}) + \mathbb{V}_{X_p}(\kappa\xi_p) + 2\mathbb{C}_{X_p}(\theta\alpha_{c(p)},\kappa\xi_p) \qquad (6)\end{aligned}$$

Dividing both side of Equation (6) 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.

C.2.3 Selection on Non-Beliefs: Evidence from Private-Sector Pensions

The variance decompositions in Section 4.1.3 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 4.1.2. 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 3, including size and pension funding.

Figure A8a 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.[9] To compute $\bar{p}^a_{c,t}$, we use data on 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 5 pp ($t = 5.06$) increase in the alternative-to-risky share of its public-sector clients.

Figure A8b 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 A8 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.

[9] The resulting plot 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 3 before collapsing to the consultant-year level.

C.3 Decomposing Changes in Perceived Alpha

Figure 4a in the main text shows that the median consultant's perceived alpha of alternatives relative to public equities has steadily risen since the early 2000s. Recall that alpha is defined as:

$$\alpha = \mu_A - \beta_A \mu_E,$$

where μ_A is the expected excess return of alternatives, β_A is the beta of alternatives with respect to public equities, and μ_E is the expected excess return on public equities. For a given consultant, this implies that the change in alpha over any two periods can be decomposed as:

$$\Delta\alpha = \Delta\mu_A - \Delta(\beta_A \mu_E).$$

Thus, the increase in perceived alpha must come from a combination of an increase in the expected excess return to alternatives, a decline in β_A (improved diversification), or a decline in the expected excess return of public equities.

Table A2 provides a sense of how each has contributed to the increase in perceived alpha. We focus on consultants' perceived beliefs on private equity and real estate since there are no data on hedge funds in 2001 in our sample. In each year, we compute α, β_A, μ_A, μ_E, and the product term $(\beta_A \mu_E)$ as averages across consultants. Panel A presents the average consultant beliefs in 2001 and 2021. The expected return of alternatives increases by 81 bps, whereas the risk premium $(\beta_A \mu_E)$ increases by 33 bps. Next, we compute the implied changes in α assuming that there is no change in one of the components. Panel B presents the decomposition results. The increase in the expected return of alternatives (μ_A) contributes to 168% of the rise in consultant-perceived alpha. The combined effect of beta and expected return of public equity reduces consultant-perceived alpha rather than increasing it.

The fact that the rise in consultant alpha has been driven entirely by expected asset returns is inconsistent with the idea that return-smoothing motives have driven the rise of alternatives. In the simplest version of the return-smoothing story, pensions want to invest in alternatives because they conceal risk. To the extent that consultants cater to or even feed this view, it seems natural to expect that beliefs about risk would drive the rise in consultants' reported alpha. However, Panel A shows that there has been no meaningful change in their perceived diversification benefits (via lower β). Moreover, Panel A also indicates that consultants have consistently viewed private markets as riskier than public markets in terms of total volatility. It is of course possible that consultants cater to return-smoothing motives by reporting a high expected return of alternatives, but we believe such motives would more likely appear in risk measures.

D Evaluating Risk-Seeking Explanations: Additional Evidence

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

Section 5.2 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.

D.1.1 Long-Run Changes: Model Misspecification

One potential issue with Table 6 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 A10 presents a series of scatter plots that show the raw data underneath the regressions in Table 6. 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.

D.1.2 Long-Run Changes: Additional Hypotheses

Table A3 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 6 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 A11a. 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 A12a. 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 A12b also shows that size is not a particularly strong predictor of the average level of the alternative-to-risky share.

D.1.3 Panel Analysis of Levels

Our analysis thus far has focused on understanding changes in the alternative-to-risky share. However, the model in Section 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:

$$\omega^*_{A,p,t} = \alpha_t + \beta X_{pt} + \varepsilon_{pt} \tag{7}$$

where $\omega^*_{A,p,t}$ is the alternative-to-risky share of pension p in fiscal year t.[10] 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 6, measured at time t. Standard errors are double-clustered by state and year.

Table A4 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 A13

[10]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.

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 A4 are not driven by our choice of a linear model.

The panel regression in Equation (7) 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 A4 run regression (7) 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 A14a and A14b show that non-linearities are unlikely to bias these results.

D.1.4 Portfolio Constraints

In Section 5.1, 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 A15a 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 A15b 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 (-0.1) 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.

D.2 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.

D.2.1 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 5.2. 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 A5 summarizes the estimation results.

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 7.5 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 5.7 pp. Within the regression sample, the actual difference between the 10th and 90th percentile of changes in the risky share is 30 pp.[11]

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 A5, with all well below 10%. Figure A16 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).

D.2.2 Analysis of Levels

Next, we analyze the level of the risky share using the same panel regression setup in Equation (7). To assess economic magnitudes in this panel setting, we proceed in three steps. First, we regress $\omega_{risky,t}$ on a time fixed effect and compute the standard deviation of the residual, denoted as $\sigma_{CS}(\omega_{risky,t})$. We repeat the same process for each explanatory variable X_{pt} and denote the standard deviation of the resulting residual $\sigma_{CS}(X_{pt})$. Finally, we compare the predicted impact of a one-standard deviation increase within the cross-

[11] 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.

section of pensions, $\hat{\beta}\sigma_{CS}(X_{pt})$, to $\sigma_{CS}(\omega_{risky,t})$. In what follows, $\hat{\beta}\sigma_{CS}(X_{pt})$ is what we refer to as the impact of a one-standard deviation change in X_{pt}. Both $\omega_{risky,t}$ and X_{pt} are purged of common time-series variation because our goal is to study cross-pension behavior.

Column (1) of Table A6 examines the link between the level of risk-taking in the cross-section and GASB 25 funding ratios. Consistent with reaching-for-yield due to lack of funding, the negative point estimate indicates that underfunded pensions have higher risky shares, though the estimate is only marginally significant. The second column of the table reruns the panel regression at the state level and uses BEA-adjusted funding ratios. The point estimate in this case becomes more negative but is also only marginally significant. It implies that a one-standard deviation decline in BEA-funding is associated with a 1.1 pp increase in the risky share. This is relatively modest, as the standard deviation of the risky share across states – and after removing common time trends – is 6.5 pp. In addition, the within-R^2 of 3.1% in column (2) indicates that most of the cross-state variation in the risky share is not explained by variation in BEA-adjusted funding levels.

The third column of Table A6 tests whether pensions with relatively higher hurdle rates are also more likely to invest in risky assets. The positive and statistically significant point estimate is consistent with the idea that high hurdle rates induces reach-for-yield behavior. It implies that a one-standard deviation increase in the hurdle rate is associated with a 1.7 pp increase in the risky share, or about a third of the actual cross-sectional standard deviation. The within-R^2 in this specification is relatively low (4.3%), suggesting that the large majority of cross-pension variation in the risky-share is not related to variation in hurdle rates.

Column (4) in Table A6 tests whether pensions that have more retired members tend to have higher risky shares. The positive point estimate is once again consistent with the logic in Andonov, Bauer, and Cremers (2017), but is not statistically different from zero and modest in economic magnitude. For example, the point estimate implies that a pension with all of its members retired would have a 5 pp higher risky share compared to one with no retired members.

Columns (5) to (8) repeat these regressions using only data from 2021 and largely echo the full panel results. For example, the point estimate on BEA-adjusted funding in column (6) is comparable to the one in column (2). Thus, there is some evidence that states with more underfunded pensions in 2021 also have higher risky shares.

Perhaps more importantly, the R^2's are fairly low across all specifications in Table A6 and we confirm in Figures A14c, A14d, and A17 that this is not due to any obvious model misspecification issues.

D.2.3 Analysis of Trends

The consistent theme of our results thus far is that the mechanisms through which reach-for-yield is thought to arise within public pensions have little cross-sectional support. This raises doubts about such mechanisms to fully account for the national trends in portfolio composition since the 2000s. To probe this claim further, we independently sort pensions based on their average GASB 25 funding ratio and hurdle rate during our sample. Pensions in the top quartile of average funding and the bottom quartile of average hurdle rate are labeled as having a low incentive to reach-for-yield. Conversely, pensions in the bottom quartile of average funding and top quartile of average hurdle rate are labeled as having a high incentive to reach-for-yield.

We then compare whether there are differential time-series trends in the risky share of the two groups. To understand the logic of this test, consider the view that U.S. pensions have shifted their portfolios because the decline in global interest rates put a strain on their funding. Under this view, it should also be the case that underfunded pensions rebalanced more aggressively than well-funded ones. A similar logic holds for pensions with the highest hurdle rates, since it became harder for these pensions to meet return targets as interest rates fell. Thus, if the fall in rates influenced aggregate pension behavior through funding or hurdle rates, the portfolios of pensions with low and high incentives to reach-for-yield should follow different time-series trends. If not, then it seems less likely that either channel can explain the trend towards riskier assets at the national level.

Figure A18a starts by plotting how the GASB 25 funding ratios of the low- and high-incentive group have evolved since 2001. The low-incentive group has been nearly fully funded throughout the sample, whereas the high-incentive group started at roughly 80% funded and fell to below 60% funding by 2020. Despite this divergence, Figure A18b indicates that both groups have increased their risky share at roughly the same pace since the early 2000s. The same holds true when plotting the evolution of the alternative-to-risky share in Figure A18c. The fact that the two groups share a time-series trend cuts against the idea that falling interest rates have caused pensions to reach-for-yield through their interaction with funding or hurdle rates. However, we want to emphasize that falling rates could still have played a role in the portfolio shifts that have occurred since the 2000s. For example, one of the primary arguments of our paper is that beliefs about risk and return could have also changed with the interest rate environment.

D.2.4 Discussion

Overall the results in this subsection suggest that cross-pension variation in the risky share – both in long-run changes and levels – 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.[12] 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. 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.

E Simulation of Mean-Variance Model

E.1 Risk-Seeking and Portfolio Constraints

In Section 6.1 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 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

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

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_A^* in 2001 according to Equations (2) to (3) and the relationship $\omega_A^* := \omega_A/(\omega_A + \omega_E)$. 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 A19 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 $\Delta\omega_A^*$ in the data. Panels A and B in Table 7 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, 501 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. Figure A21 shows the cumulative distribution function for the resulting alternative-to-risky share in these simulations. 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 7 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 (419 out of 100,000). Figure A20a 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:

$$ShadowCost = M_{counterfactual} - M_{portfolio}, \tag{8}$$

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 A20b presents the distribution of the shadow cost across the 419 simulations in which risk-seeking motives can generate the observed $\Delta\omega_A$. 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.

E.2 Simulating a Pure Belief Shift

E.2.1 Full Simulation Results

In Section 6.2, 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 (47%) 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 A19 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, 47% of the increase in the risky share can be generated by beliefs, with the remaining 53% coming from risk aversion.

E.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:

$$\alpha_p = \zeta \alpha_c + \nu_p, \tag{9}$$

where α_c is the consultant sector's perceived alpha and ν_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:

$$\begin{aligned} \frac{\partial \omega_A^*}{\partial \alpha_c} &= \frac{\partial \omega_A^*}{\partial \alpha_p} \times \frac{\partial \alpha_p}{\partial \alpha_c} \\ &= \Lambda \times \zeta, \end{aligned} \tag{10}$$

where $\Lambda \equiv \partial \omega_A^* / \partial \alpha_p$ comes from the model. Equation (4) 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 (10) is the elasticity of the alternative-to-risky share to consultant beliefs, which we denote as ϕ in the main text. This parameter is pinned down by our IV strategy in Section 4.1 and is roughly $\phi \approx 4$. Thus, in each simulation, $\zeta = \phi / \Lambda$. Figure A23 shows the distribution of ζ from the 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 ($\Delta\alpha_c \approx 70$ bps). The remaining component $\Delta\nu_p$ 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 $\Delta\nu_p$ from the simulations. On average, private beliefs add 60 bp to pension-perceived alpha, and the consultant's beliefs add 10 bp, totaling a 70 bp increase in alpha.

E.2.3 Peer Multipliers

In Section 4.3, 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 (10) can be written in vector form as:

$$\begin{aligned}\boldsymbol{\omega}^*_{A,t} &= \beta_z \boldsymbol{D}\boldsymbol{\omega}^*_{A,t} + \tilde{\mathbf{X}}_t\boldsymbol{\Gamma} + \boldsymbol{\varepsilon}_t \\ &= (\boldsymbol{I} - \beta_z\boldsymbol{D})^{-1}\tilde{\mathbf{X}}_t\boldsymbol{\Gamma} + (\boldsymbol{I} - \beta_z\boldsymbol{D})^{-1}\boldsymbol{\varepsilon}_t\end{aligned}$$

where $\boldsymbol{\omega}^*_{A,t}$ 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 $\boldsymbol{\Psi}\boldsymbol{\varepsilon}_t$.[13] 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

[13] 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.

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.

F Agency-Based Explanations

F.1 Real Assets vs. Private Equity

In Section 7.1.1, we explore the idea that pensions have moved to alternatives because of a desire to conceal risk. While we cannot rule out this possibility out fully, one way to at least rule in the importance of beliefs is to study the choice between real asset and private equity, as both asset classes offer similar abilities to conceal risk. We implement this idea via the following panel regression:

$$\omega^*_{a,p,t} = f_{p,t} + \beta V_{c(p),a,t} + \varepsilon_{p,a,t}, \tag{11}$$

where $\omega^*_{a,p,t}$ is the 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 $c(p)$. $f_{p,t}$ 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 (11) 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 (11) equals 1.92 (t =6.43) and indicates that pensions tend to allocate more toward private equity or real assets when their consultant reports it to have a high alpha.[14] The within-R^2 further shows that 16% of within-pension variation in the composition of private-market investments (real assets versus private equity) is explained by consultant beliefs.

[14] A p-value of 0.003 is obtained when using a wild bootstrap to test if $\beta = 0$ (Cameron et al., 2008).

F.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 in the cross-section.[15] The basic idea 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 specific panel regression we run is:

$$\omega^*_{A,p,t} = \alpha_t + \Gamma B_{p,t} + \beta X_{pt} + \varepsilon_{pt},$$

where $\omega^*_{A,p,t}$ is pension p's 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. Following Andonov et al. (2018), $B_{p,t}$ contains the fraction of each pension's board that belongs to one of eight categories: state-appointed, state-exofficio, state-elected, participant-exofficio, participant-elected, public-appointed, public-exofficio, and public-elected. As in Andonov et al. (2018), we omit the category of participant-appointed from the regression and only report coefficient estimates for the state-appointed, state-exofficio, participant-elected, and public-appointed categories. Standard errors are clustered by time and pension system, as board composition is relatively static through time.

Table A7 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, and total spending scaled by payroll. 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.

[15] As in Begenau and Siriwardane (2022), we extend the board composition from Andonov et al. (2018) through 2020. We thank Josh Rauh for sharing board composition data with us.

G Derivations

G.1 CRRA preference with log-normal distribution

The investors's objective function is

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

Taking the log of the last term,

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

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

We have

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

where

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

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 is ω_A^* and ω_E^* where

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

It's easy to see that when the constraint is not binding, the risky portfolio is the tangency portfolio and neither ω_A^* nor ω_E^* depend on the risk aversion γ.

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

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

where

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

Substituting into the Lagrangian we obtain

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

Expanding the matrix forms in ① and ② we get

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

Taking the FOC with respect to w_A,

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

$$\Longrightarrow 0 = (\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2) - \gamma\Big[\omega_A\sigma_A^2 + \sigma_{AE}(1-\omega_f^{min}) - 2\omega_A\sigma_{AE} - (1-\omega_f^{min}-\omega_A)\sigma_E^2\Big]$$
$$\Longrightarrow \omega_A = \frac{1}{\gamma}\underbrace{\frac{(\mu_A + \frac{1}{2}\sigma_A^2) - (\mu_E + \frac{1}{2}\sigma_E^2)}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}}_{K} + (1-\omega_f^{min})\underbrace{\frac{\sigma_E^2 - \sigma_{AE}}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2}}_{C}.$$

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

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

The risky asset portfolio allocation is as of follows:

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

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

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

In vector form, it can be expressed as

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

G.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:

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

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

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

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

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

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

$$
\begin{aligned}
C &= \frac{\sigma_E^2 - \sigma_{AE}}{\sigma_A^2 - 2\sigma_{AE} + \sigma_E^2} \\
\Longrightarrow &= \frac{1 - \beta}{(\beta-1)^2 + \sigma_\eta^2/\sigma_E^2}
\end{aligned}
$$

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

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

References

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*.

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

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.

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

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.

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

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.

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.

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.

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.

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

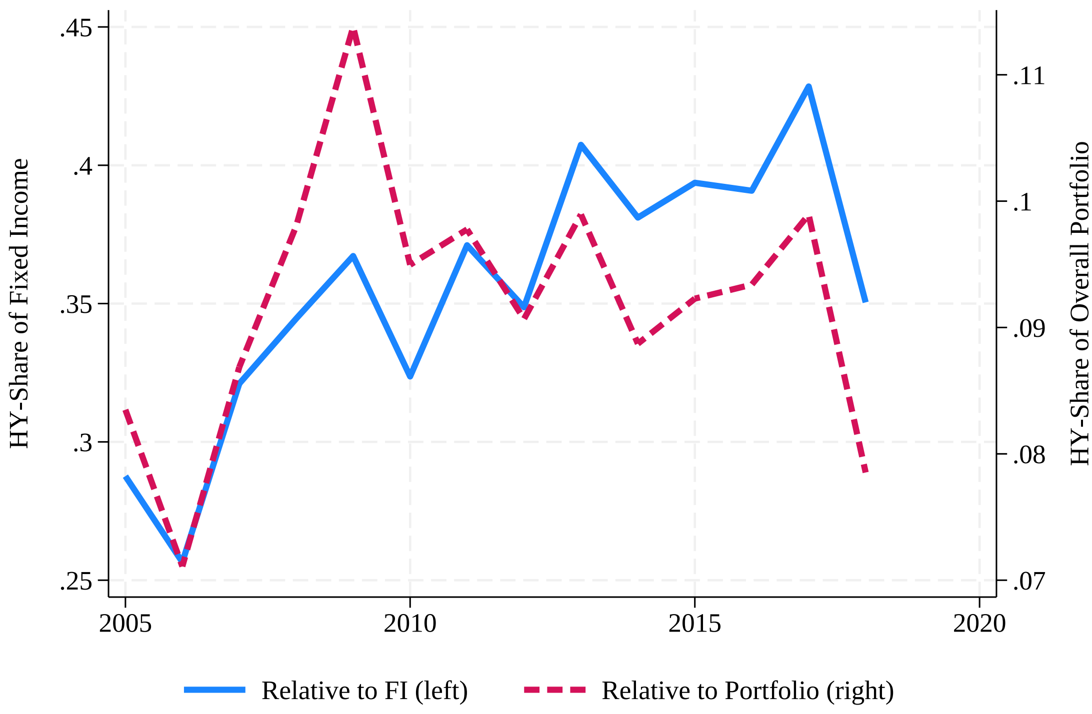

Notes: This plot shows the aggregate U.S. pension share of high-yield debt relative to fixed income holdings (left axis) and to the whole portfolio (right axis). See Internet Appendix A.1.3 for complete details on the data underlying the plot.

Figure A2: PENDAT Data Coverage

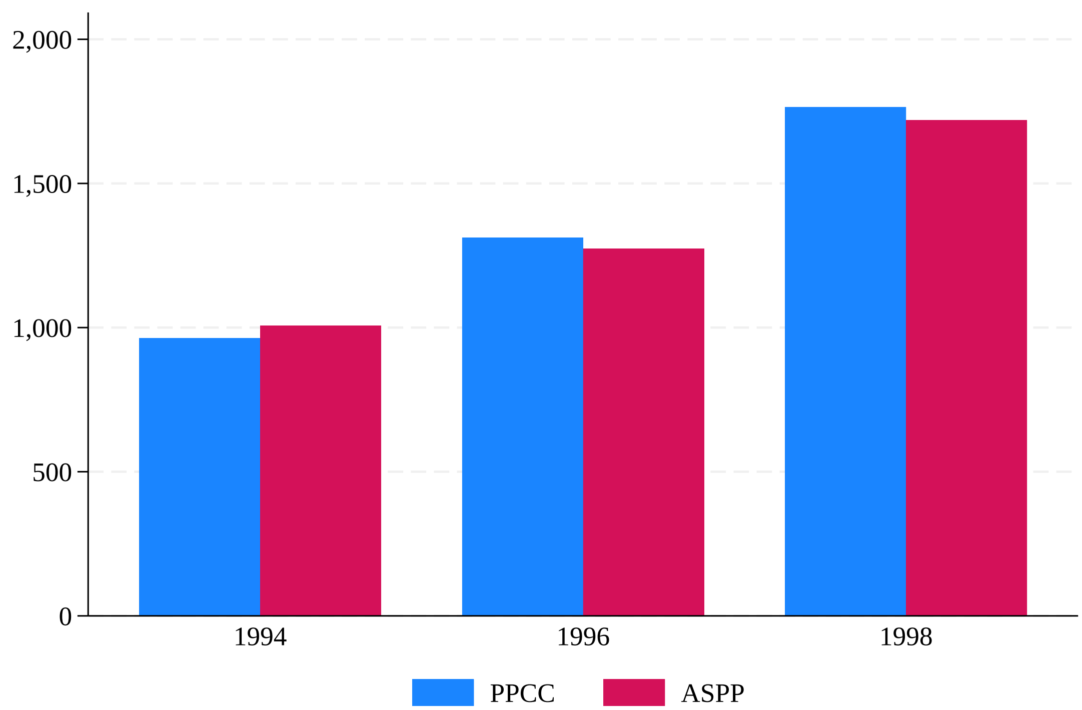

Notes: This plot shows the aggregate amount of assets in FY1994, 1996, and 1998 within the PENDAT data. For comparison, it also shows the aggregate amount of U.S. pension assets listed in the ASPP. See Internet Appendix A.4 for more details on the PENDAT data.

Figure A3: Aggregate Decomposition of Fixed Income Allocation

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

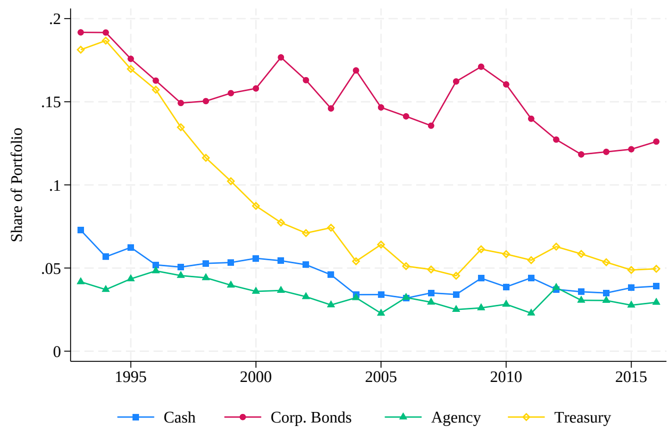

(b) U.S Government Sponsored vs Corporate Debt (QSPP)

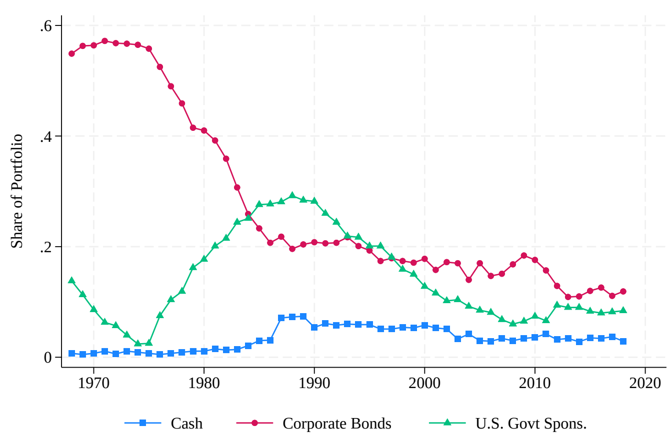

Notes: Panel (a) uses the U.S. Census Bureau's Annual Survey of Public Pensions (ASPP) to decompose the fixed income share into cash, corporate bonds, U.S. Treasuries, and Agency debt. Panel (b) plots a similar decomposition based on the QSPP but over a longer time horizon. See Section A.2 for complete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

Figure A4: The Risky Share in the Cross Section

(a) Cross-Sectional Distribution

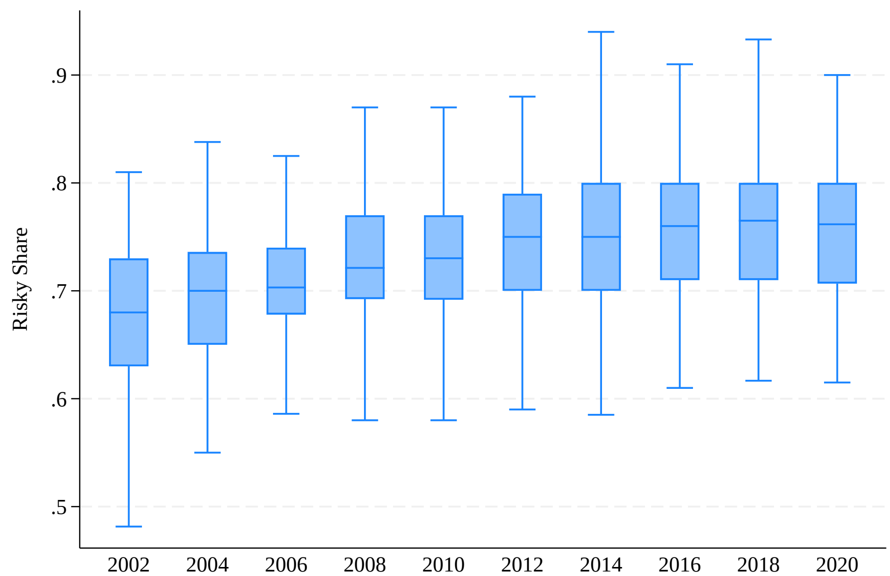

(b) Change in Risky Share Between 2002 and 2021

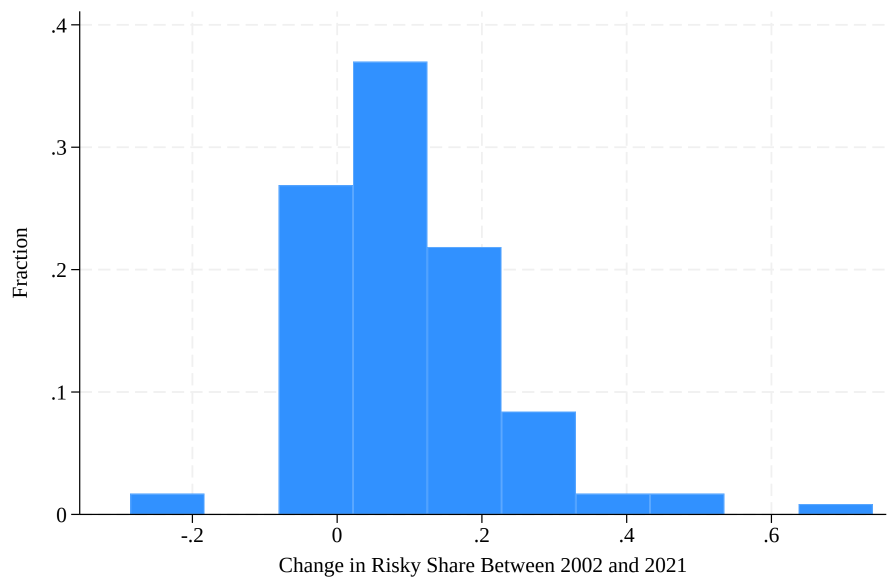

Notes: This figure visualizes the distribution of the risky share in the cross-section (panel a) and the distribution of the change in the risky share across U.S. pension systems between 2002 and 2021. The risky share is defined as all holdings outside of fixed income and cash. All data are from the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

Figure A5: Risky Share by Geography, 2002 and 2021

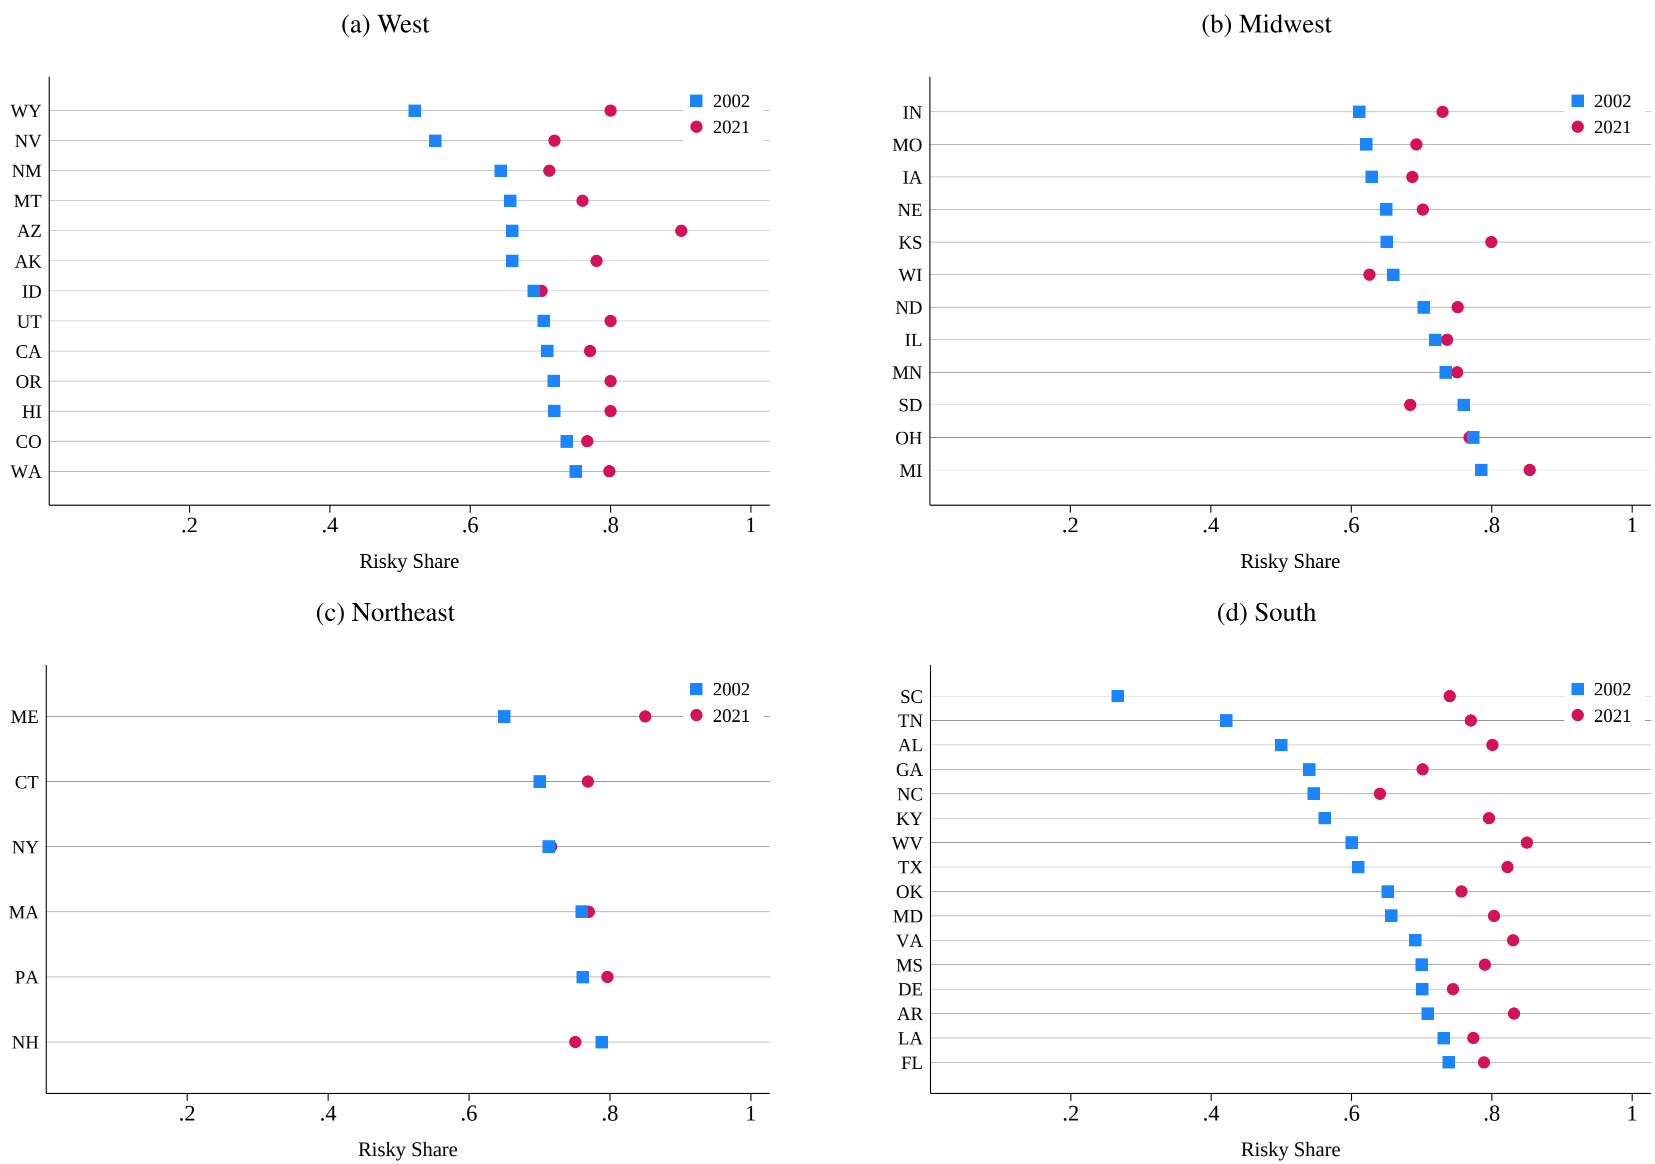

Notes: This plot shows state-level target risky shares for 2002 (blue squares) and 2021 (orange circles). Each panel of the plot corresponds to a different U.S. Census region. All data are from the PPD. Risky investments are defined as any holding outside of fixed income and cash.

Figure A6: Alternative-to-Risky Share by Geography, 2002 and 2021

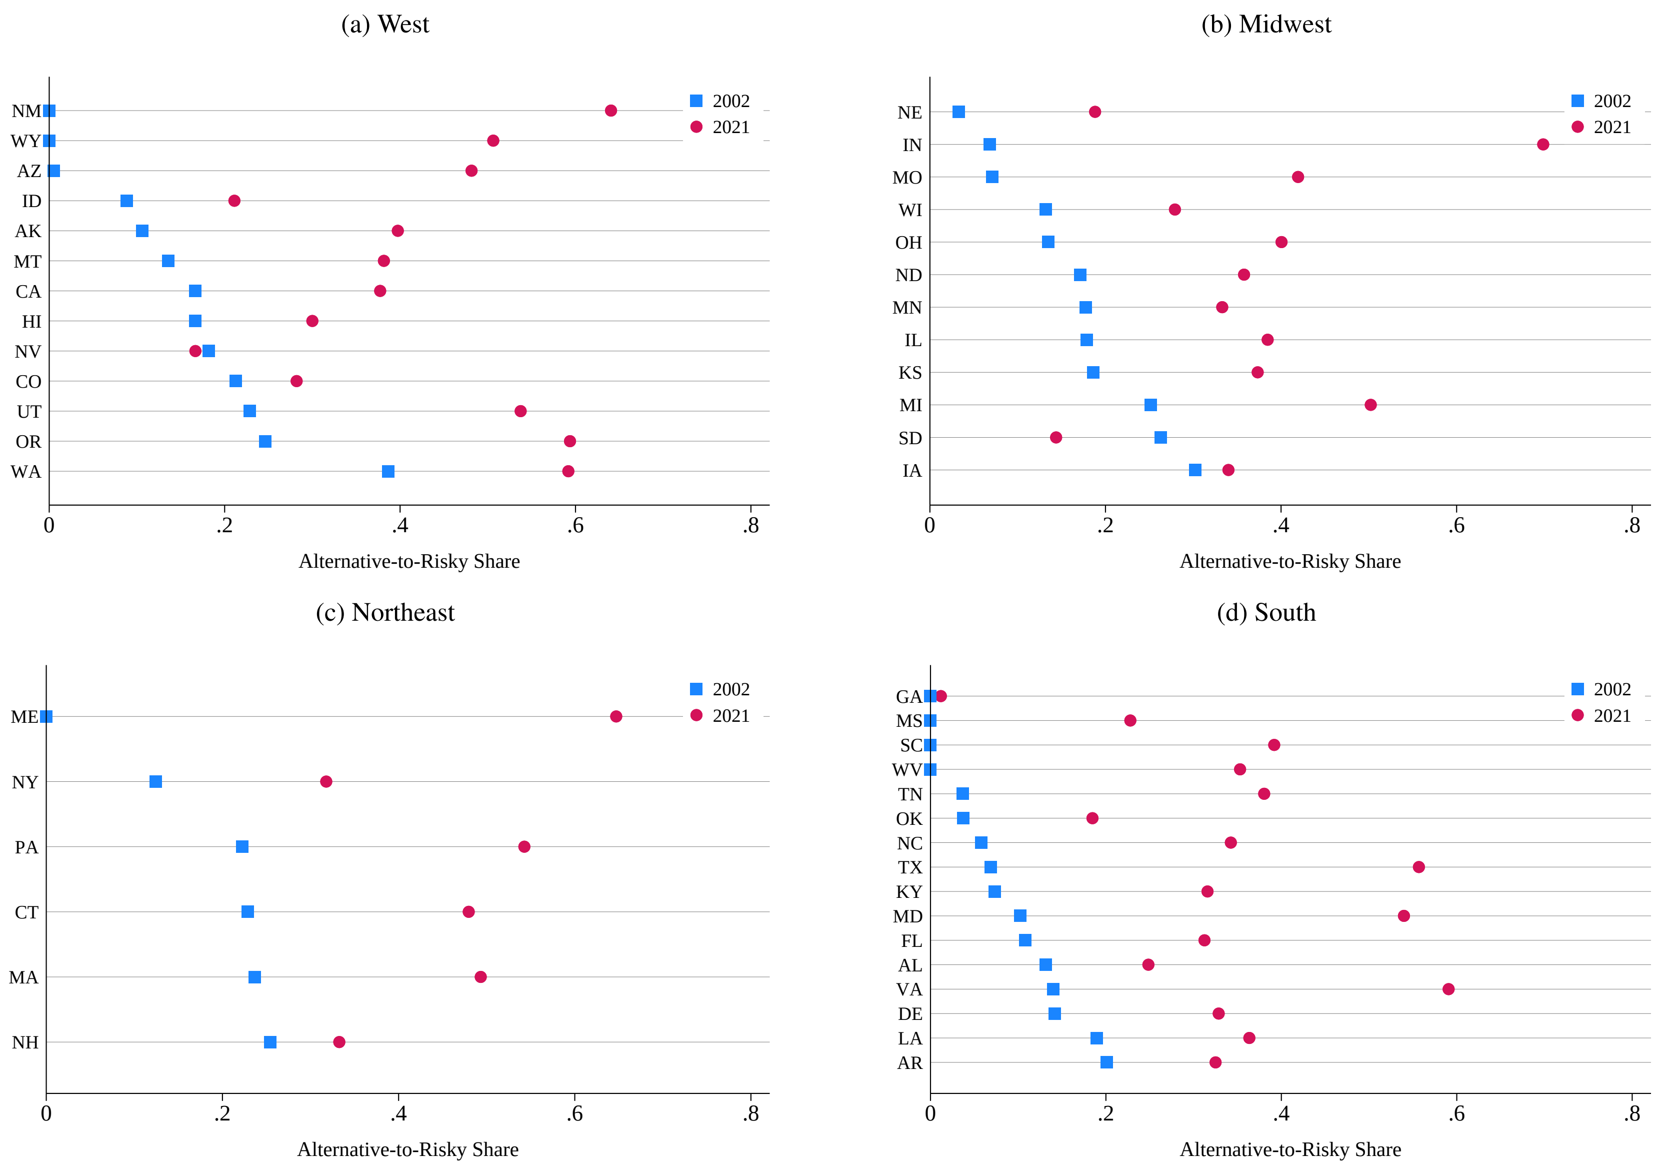

Notes: This plot shows state-level target alternative-to-risky shares for 2002 (blue squares) and 2021 (orange circles). Each panel of the plot corresponds to a different U.S. Census region. All data are from the PPD. Risky investments are defined as any holding outside of fixed income and cash.

Figure A7: System-Level Share of Alternatives

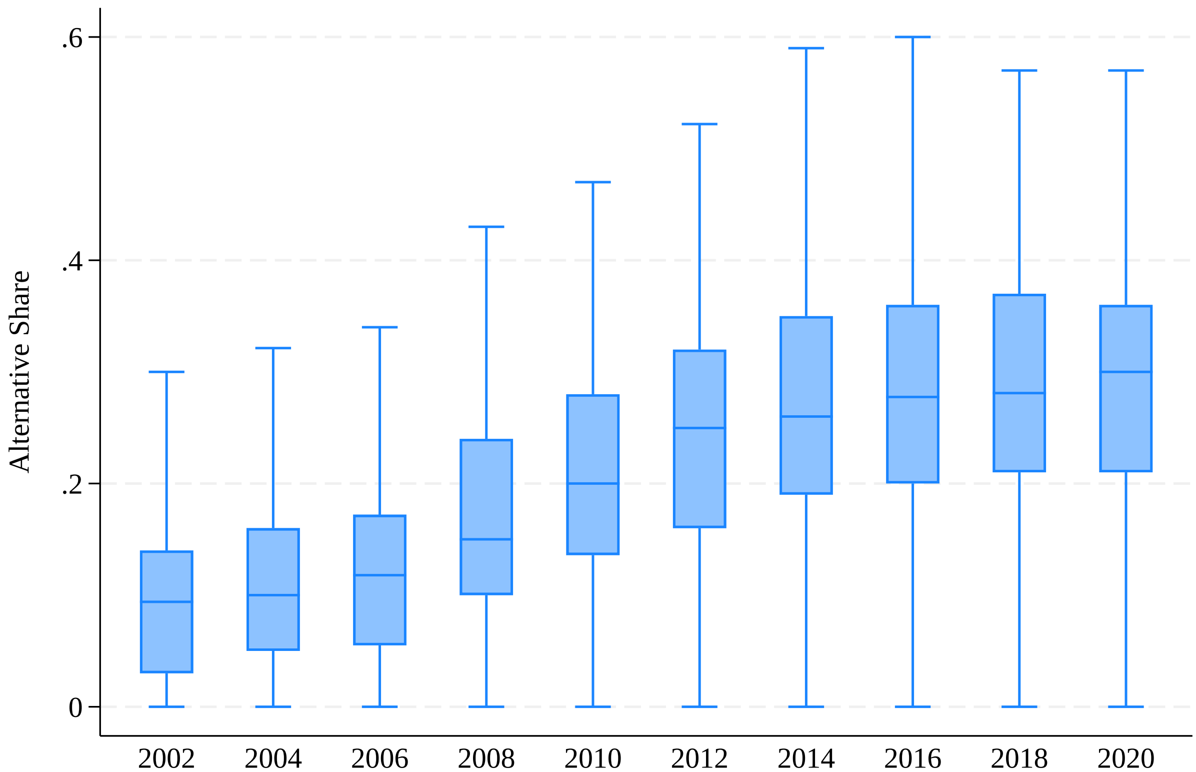

Notes: This figure depicts the distribution of the raw target alternative share across pension systems through time. Each box plot summarizes the distribution for the corresponding year on the *x*-axis. Only even years are plotted to make the graph more readable. All data are from the PPD.

Figure A8: U.S. Public vs Private-Sector Clients by Consultant

(a) Alternative-to-Risky Share

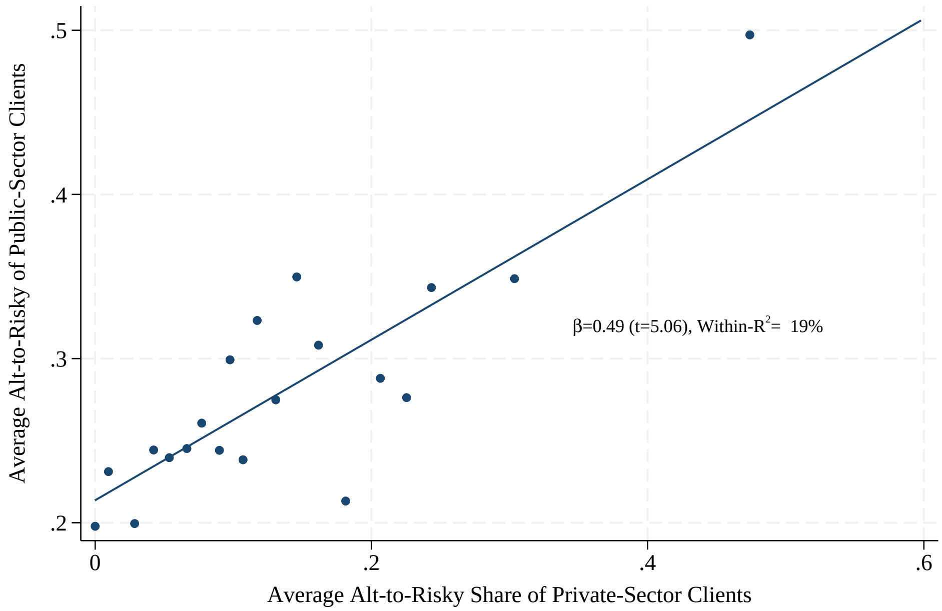

(b) Risky Share

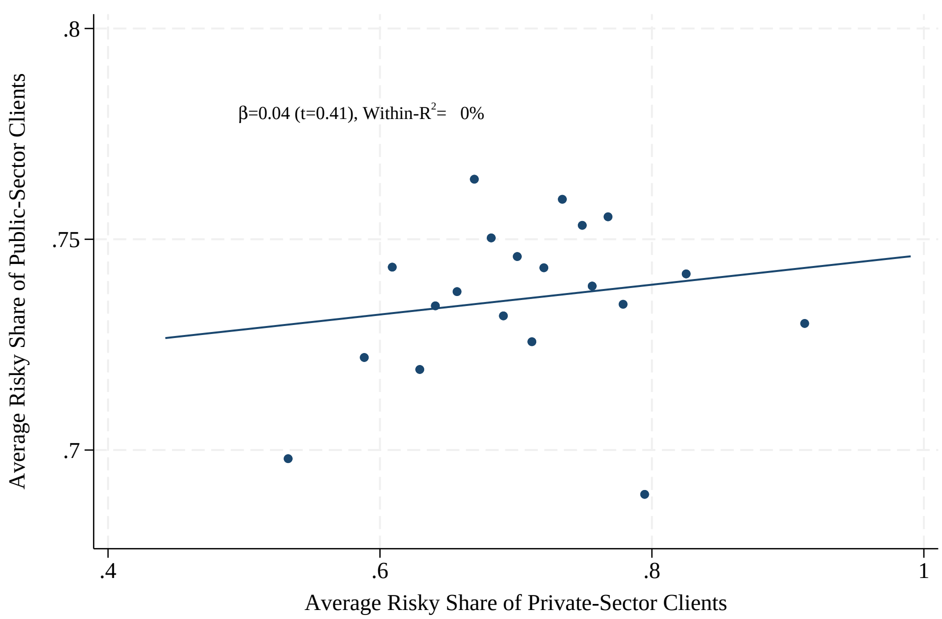

Notes: For each consultant c in year t, we compute the average alternative-to-risky and risky share of U.S. public pension clients, denoted by $\bar{u}^a_{c,t}$ and $\bar{u}^r_{c,t}$, respectively. In addition, we compute the average alternative-to-risky and risky share of U.S. private-sector clients, 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 at least two 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 A9: U.S. Public Pension Funding Status

(a) National Funding Ratios

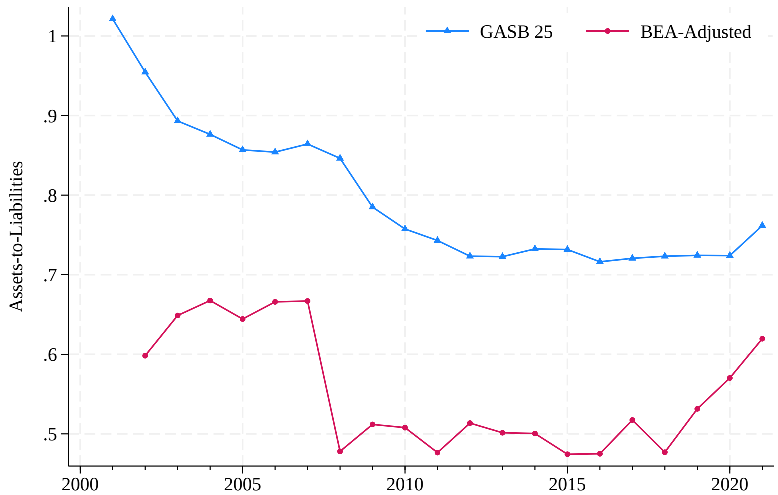

(b) Cross-Sectional Funding Ratios

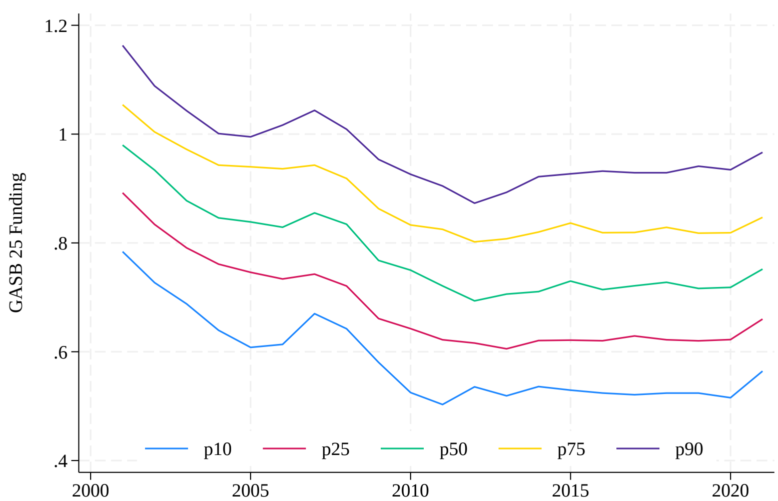

Notes: Panel (a) of this figure plots the time-series of the national pension funding ratios, defined as pension assets divided by pension liabilities. The line with triangular markers is the GASB 25 funding ratio. Under GASB 25 accounting standards, assets are smoothed over a trailing window and liabilities 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 A10: Funding-Based Explanations and Changes in the Alternative-to-Risky Share

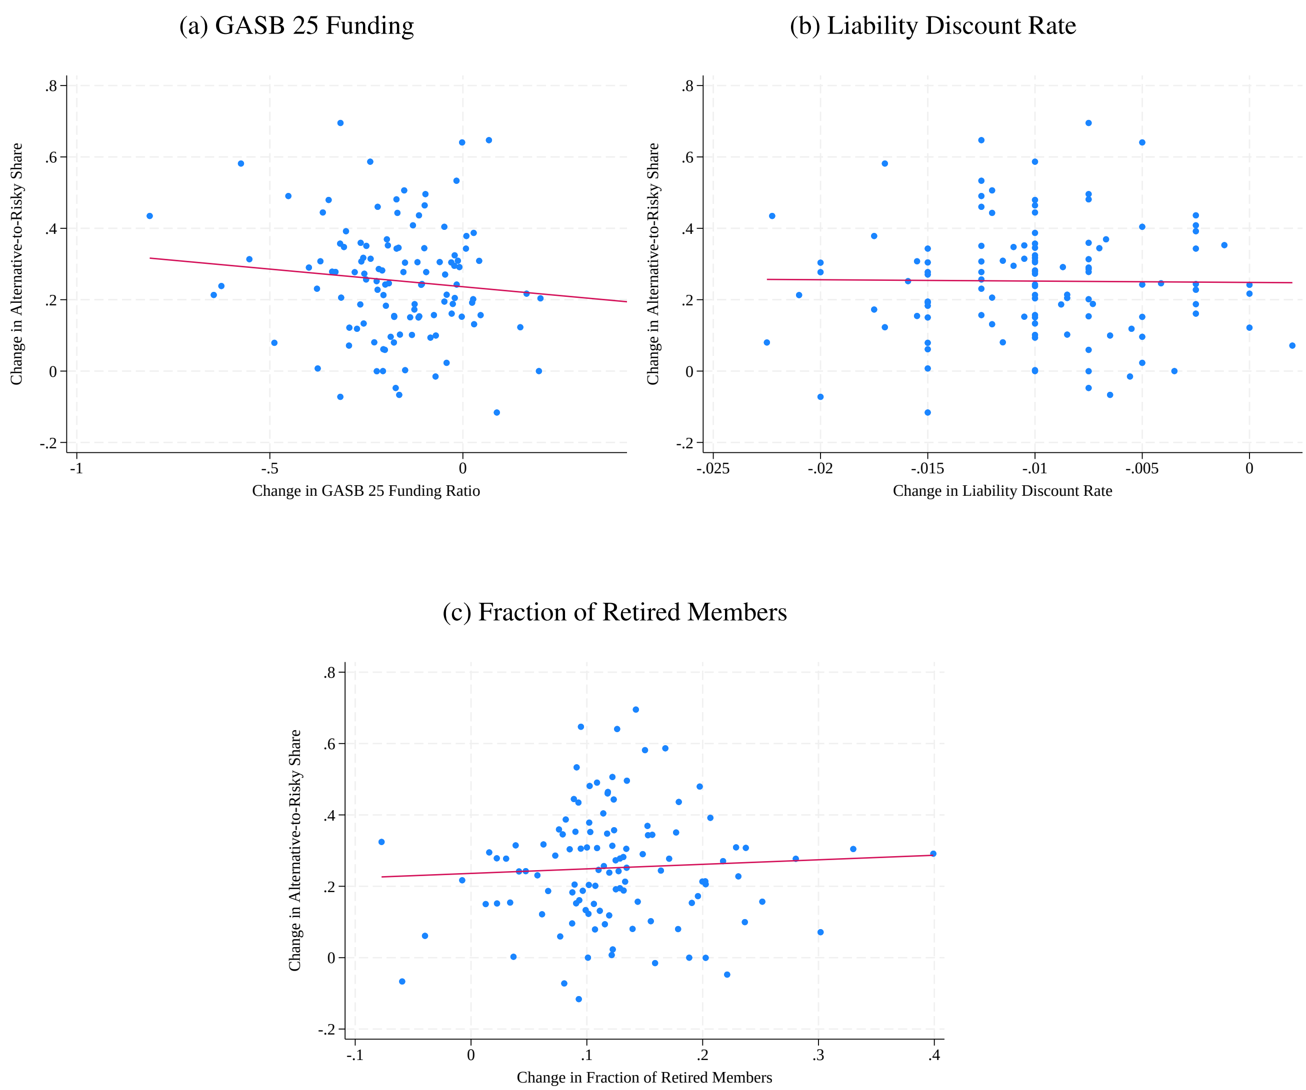

Notes: The figure contains scatter plots of changes in the target alternative-to-risky share from 2002 to 2021 against contemporaneous changes in GASB 25 funding (panel a), hurdle rates or, equivalently, liability discount rates (panel b), and the fraction of retired members (panel c). Data are from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

Figure A11: Annual Required Contributions and Investment Behavior

(a) Alternative-to-Risky Share

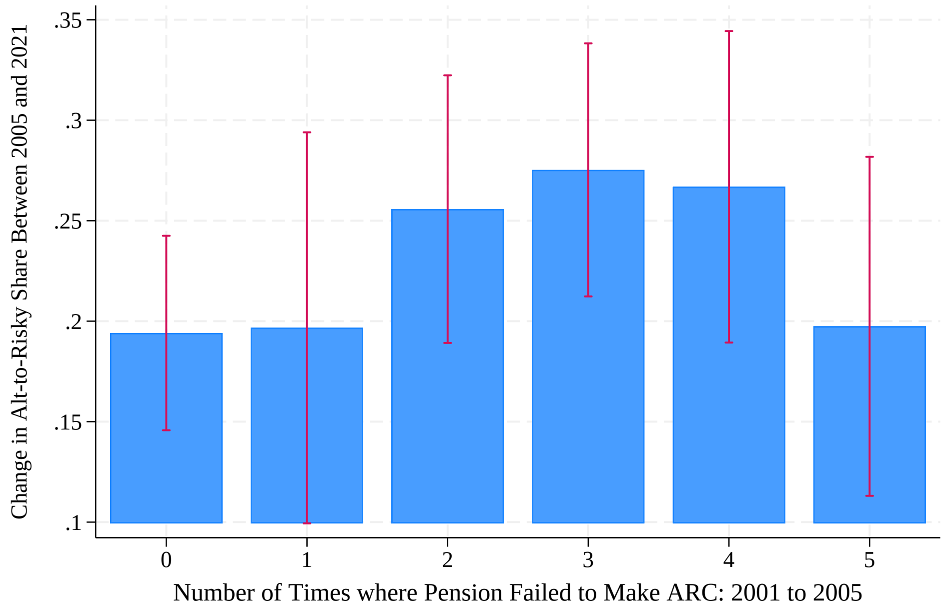

(b) Risky Share

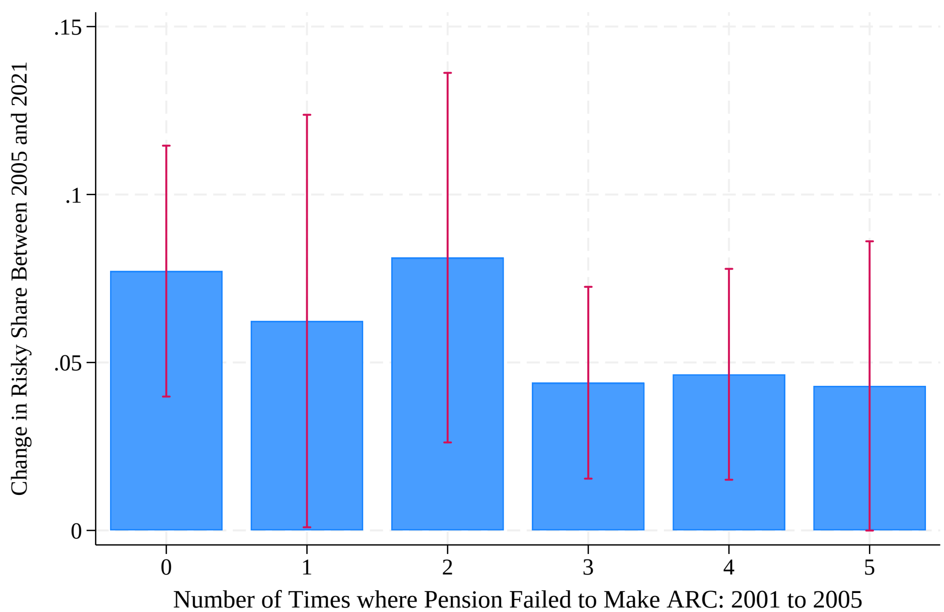

Notes: In this figure, we first count the number of times a pension failed to make a required actuarial contribution between 2001 and 2005. We then sort pensions by count and compute the average change in the target alternative-to-risky share (panel a) and the overall risky share (panel b) within each group. The orange lines in both plots are 95% standard error bands. Data are from the PPD.

Figure A12: Alternative Adoption and Use versus Size

(a) Long-Run Changes in the Alternative-to-Risky Share

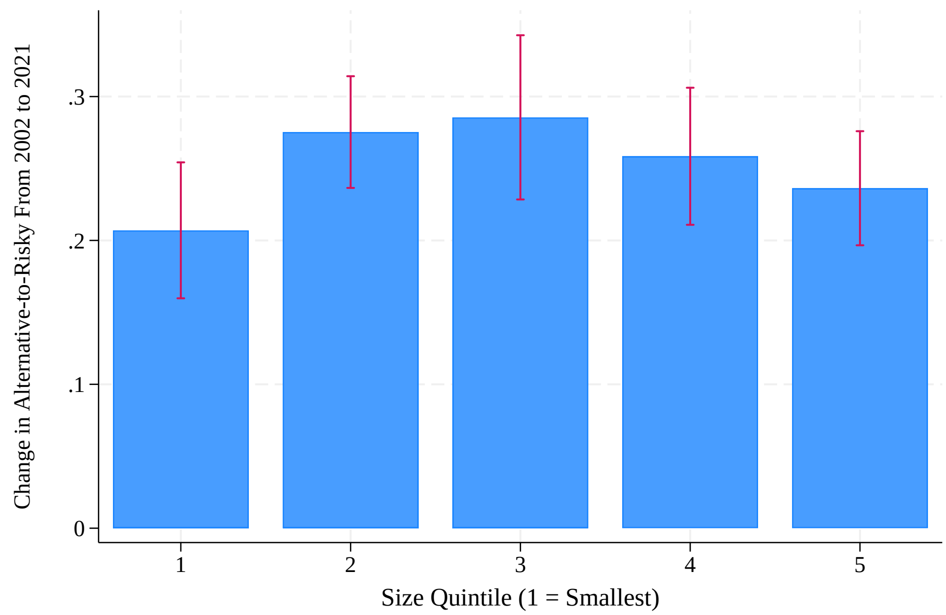

(b) The Average Level of the Alternative-to-Risky Share

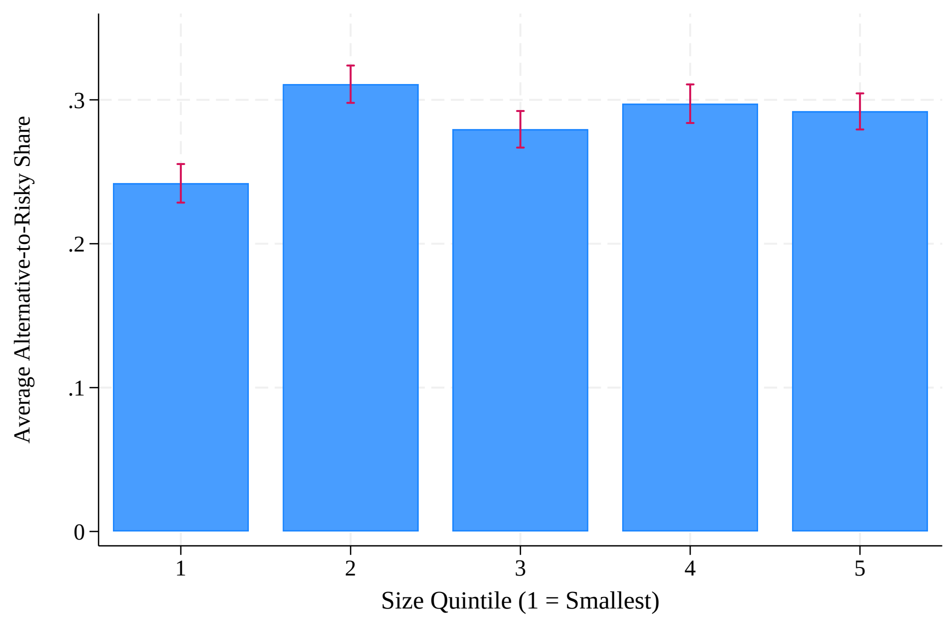

Notes: In panel (a), we sort pensions into quintiles based on the assets under management in 2002. We then compute the average change in the target alternative-to-risky share within each quintile and plot it in the graph. In panel (b), we sort pensions into quintiles based in size in each year and plot the average alternative-to-risky share for each quintile. The orange lines in both plots are 95% standard error bands. Data are from the PPD.

Figure A13: Variation in the Level of the Alternative-to-Risky Share

(a) GASB 25 Funding

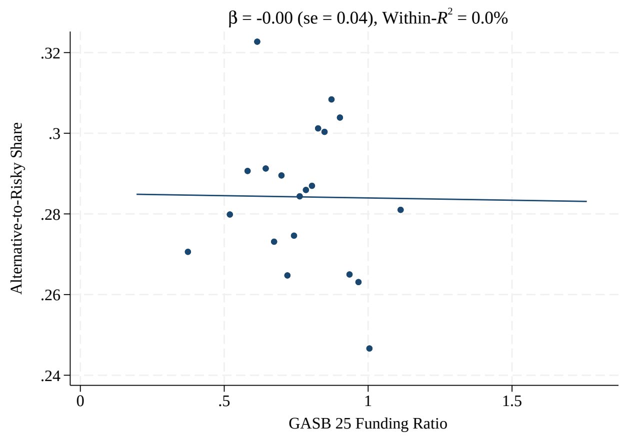

(b) BEA-Adjusted Funding (State Level)

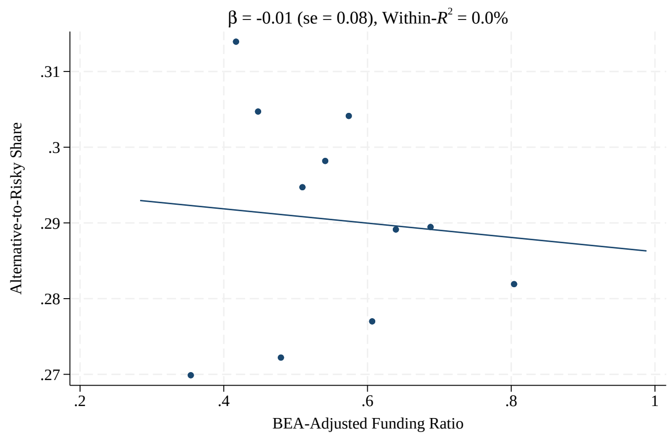

(c) Fraction of Retired Members

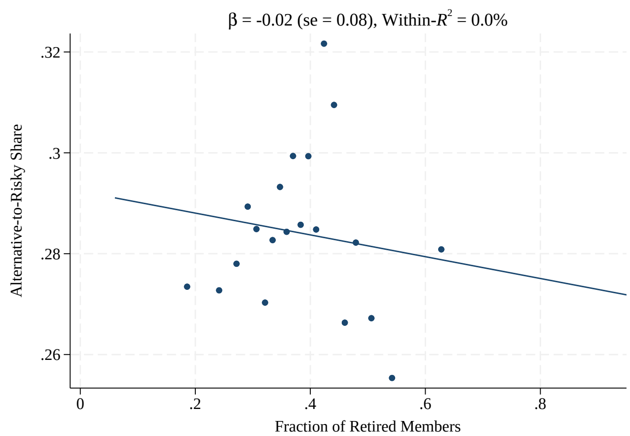

(d) Liability Discount Rate

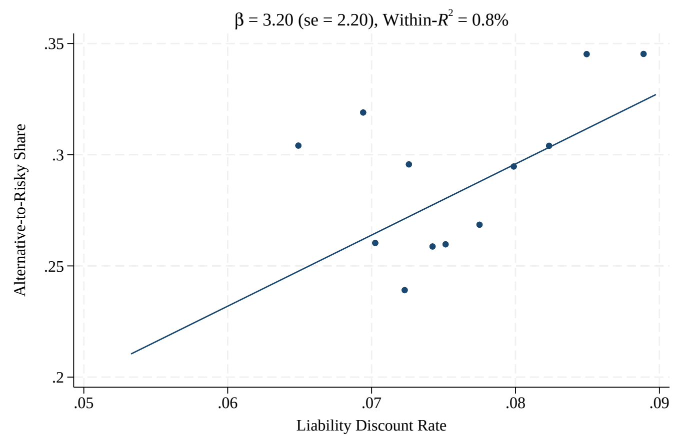

Notes: This figure shows a series of binscatter plots of each pension's target alternative-to-risky share against its: (panel a) GASB 25 funding ratio; (panel b) BEA-adjusted funding ratio (state level); (panel c) the fraction of retired members; and (panel d) hurdle rate (liability discount rate). Each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

Figure A14: Portfolio Composition in 2021

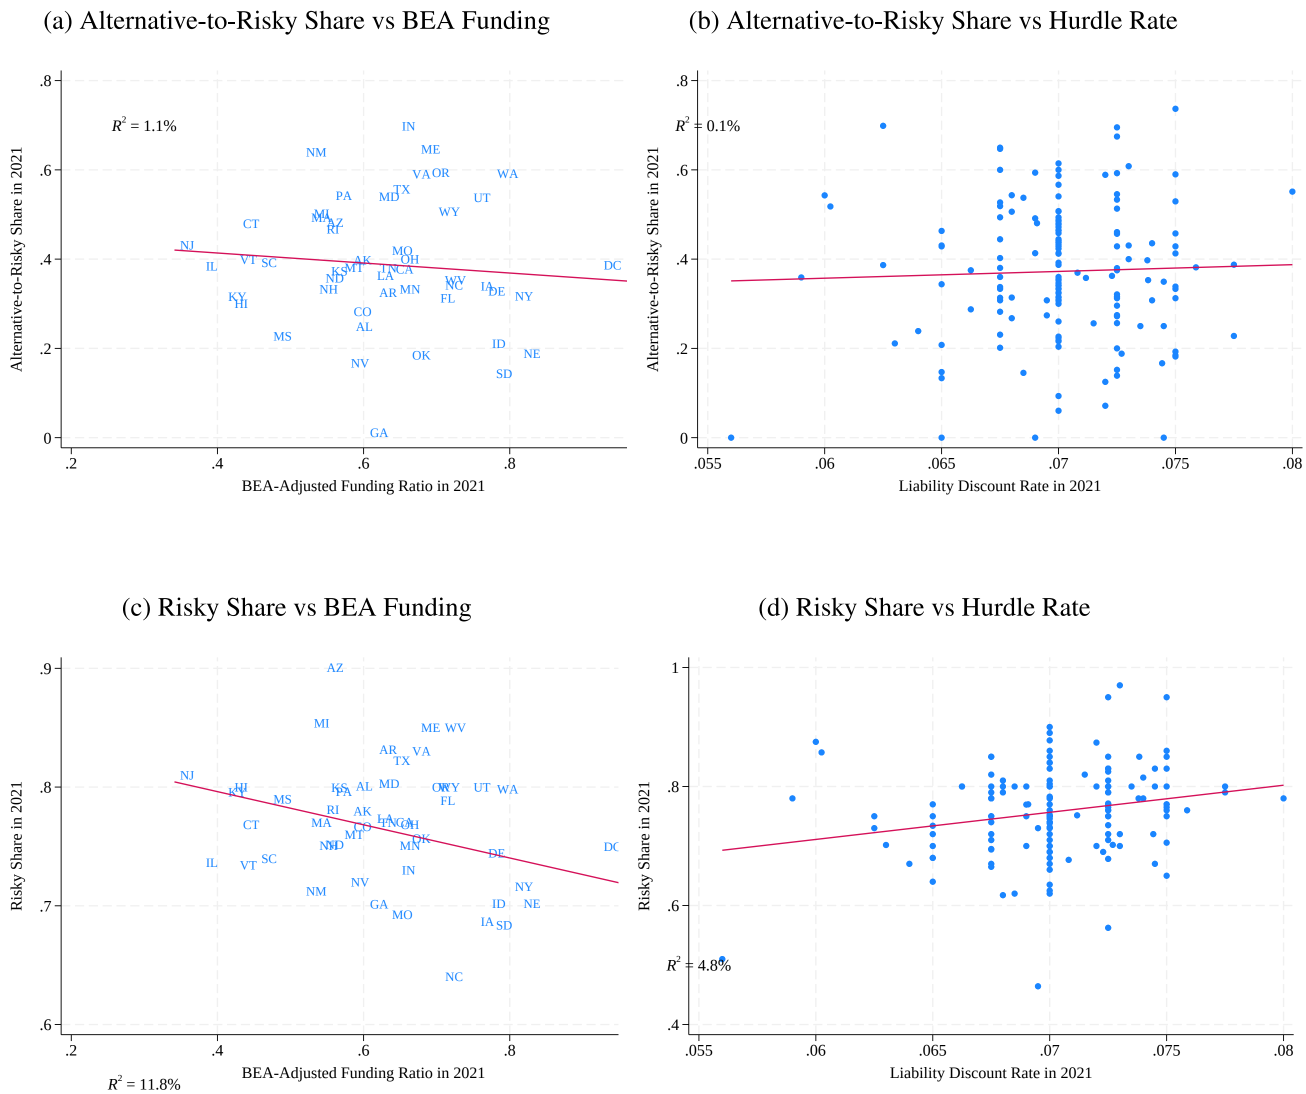

Notes: The top two panels respectively show 2020 target alternative-to-risky shares against: (a) BEA-adjusted funding ratios at the state level; and (b) hurdle rates (liability discount rate) at the pension level. The bottom two panels repeat the plots but use the risky share instead. All portfolio data are from PPD.

Figure A15: Deviations of Actual from Target Weights

(a) Raw Values

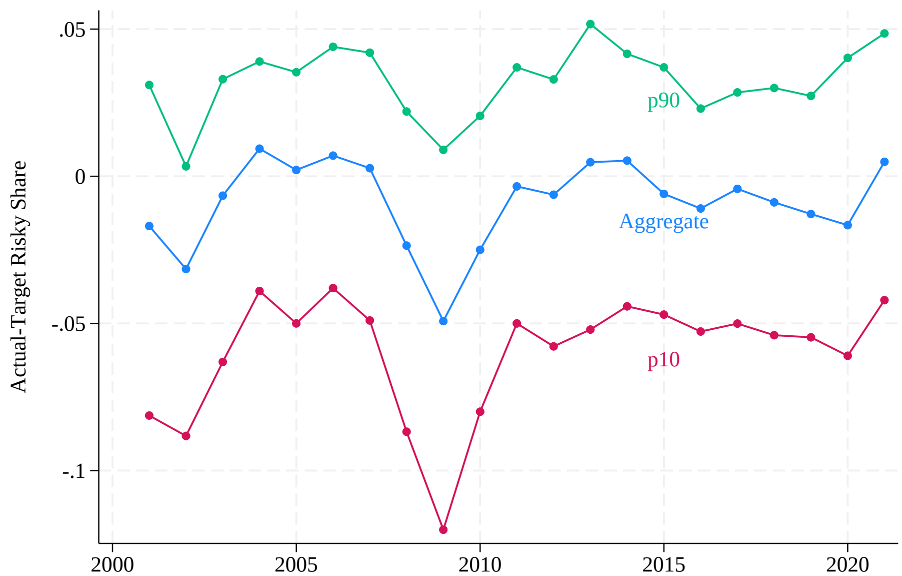

(b) Residualized to 1-Year Returns

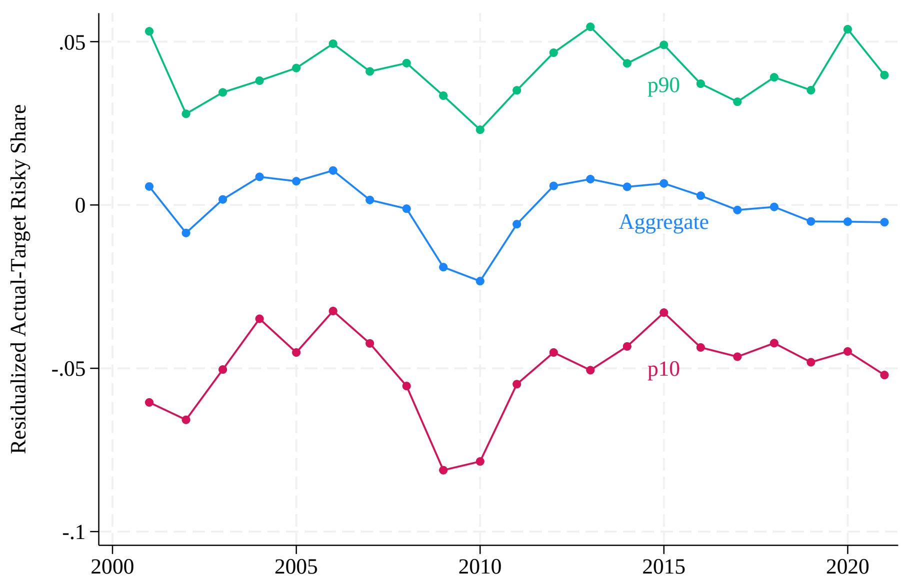

Notes: Panel (a) of the figure shows the aggregate, 10th, and 90th percentiles of l_{pt}, defined as the difference between the actual and target risky share for pension p in year t. Panel (b) plots the same series, after residualizing l_{pt} to the annual return of each pension. See Section 5.1 for details.

Figure A16: Funding-Based Explanations and Changes in the Risky Share

(a) GASB 25 Funding

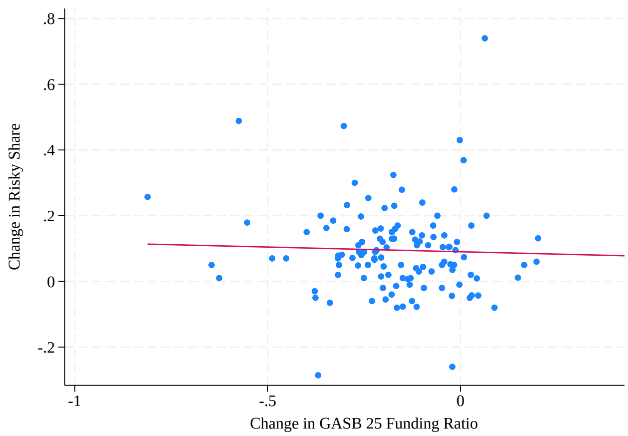

(b) Liability Discount Rate

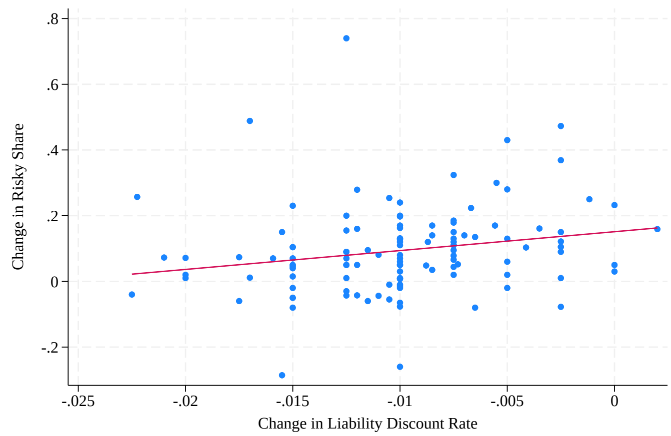

(c) Fraction of Retired Members

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 A17: Variation in the Level of the Risky Share

(a) GASB 25 Funding

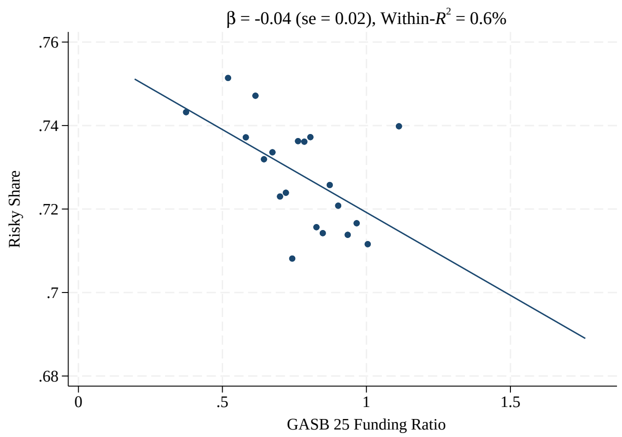

(b) BEA-Adjusted Funding (State Level)

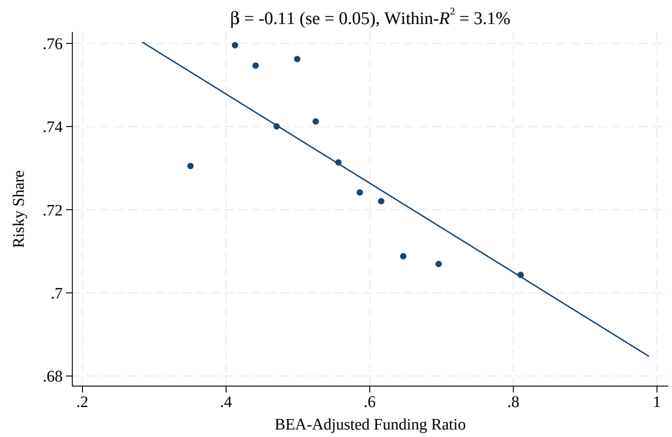

(c) Fraction of Retired Members

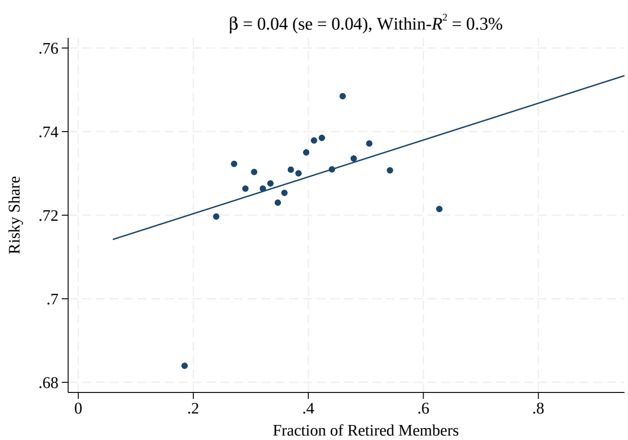

(d) Liability Discount Rate

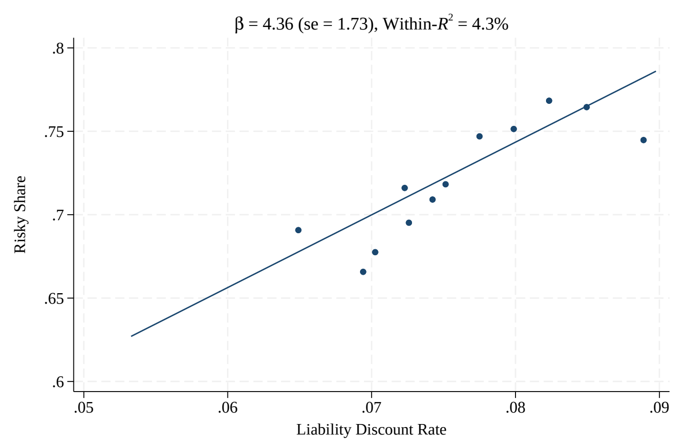

Notes: This figure shows a series of binscatter plots of each pension's target risky share against its: (a) GASB 25 funding ratio; (b) BEA-adjusted funding ratio (state level); (c) the fraction of retired members; and (d) hurdle rate (liability discount rate). The risky share is defined as all investments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

Figure A18: Portfolio Trends, Reach-for-Yield Incentives, and Falling Interest Rates

(a) GASB 25 Funding

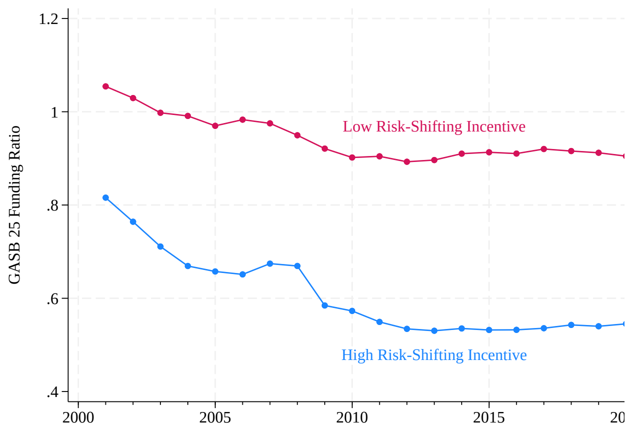

(b) Target Risky Share

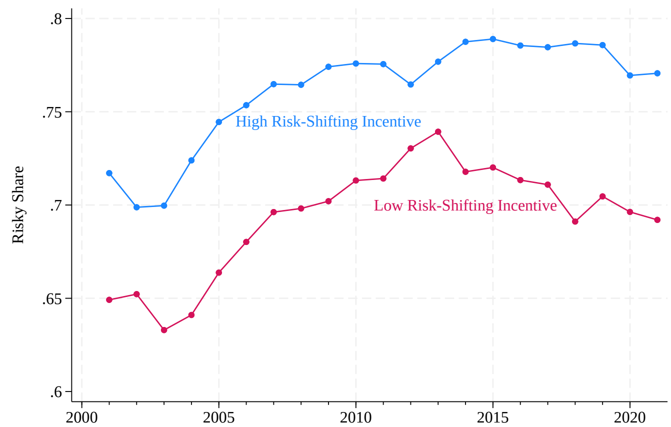

(c) Target Alternative-to-Risky Share

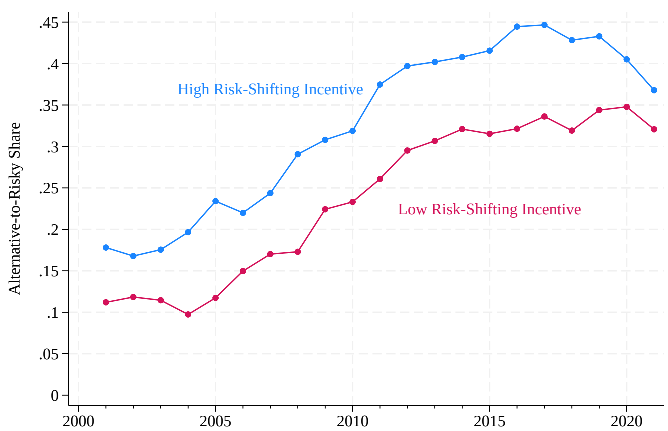

Notes: Panel (a) of the plot shows the average funding ratio through time for pensions with high and low incentives to reach-for-yield as interest rates fall. High-incentive pensions are those in the bottom quartile of average funding and the top quartile of average hurdle rate, where both averages are taken over the full sample. Low-incentive pensions are the opposite. Panel (b) plots the average risky share for both groups through time and Panel (c) plots the average target alternative-to-risky share.

Figure A19: Distribution of admissible beliefs

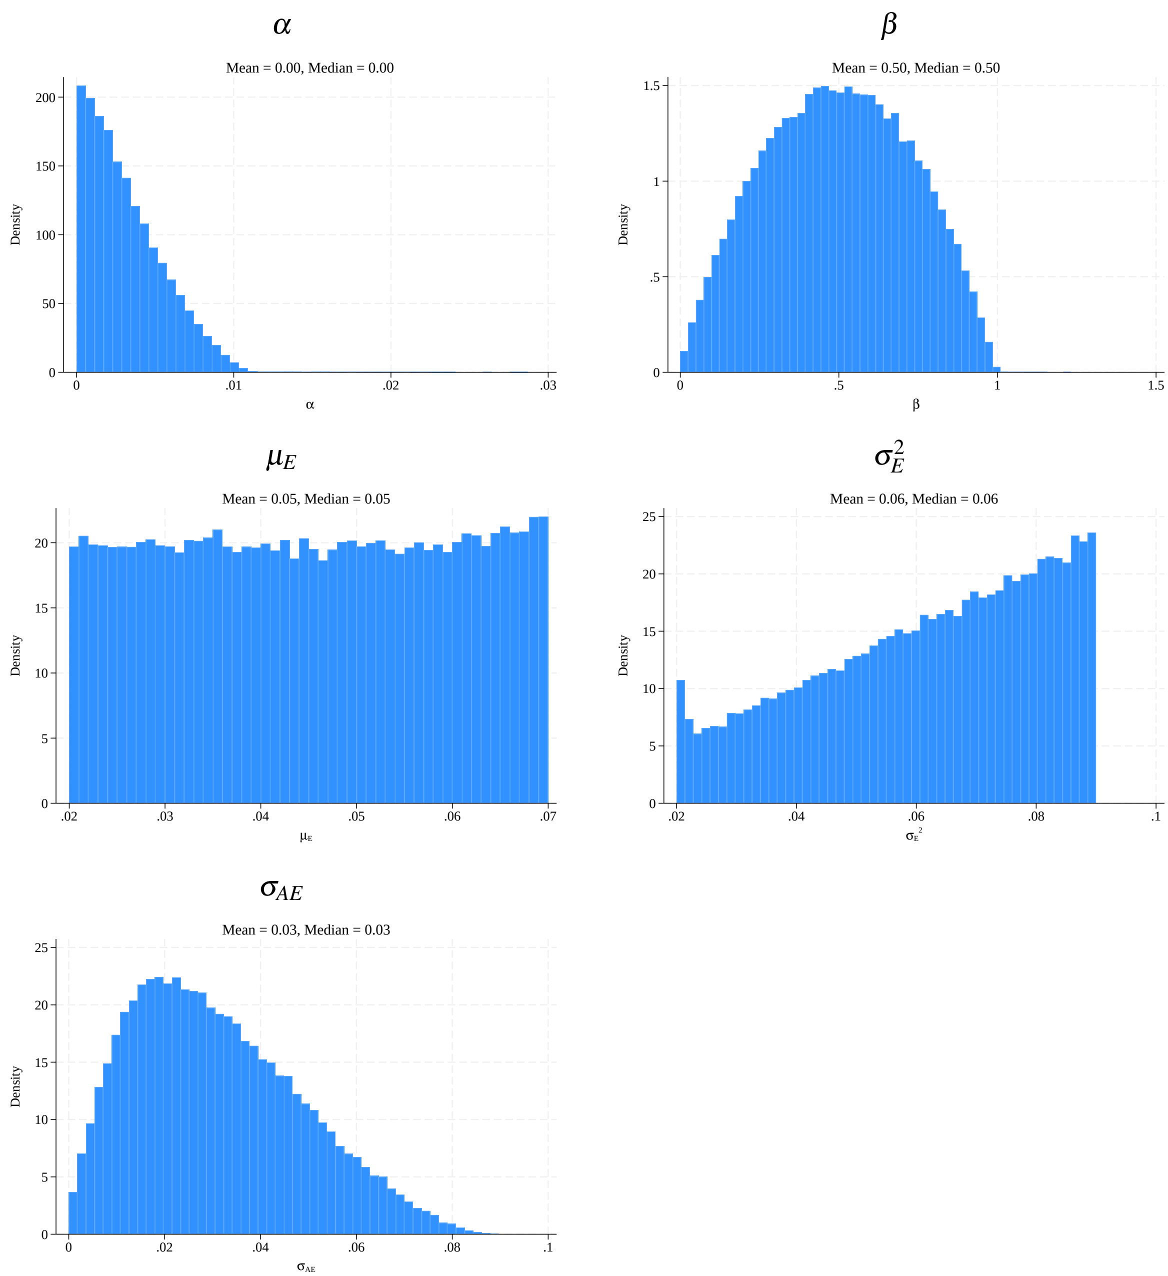

Notes: This figure plots the distributions of admissible beliefs when simulating the portfolio choice model of Section 3. See Section 6.1 of the main paper for details on the simulation.

Figure A20: Simulating a Decline in Risk Aversion Plus Binding Portfolio Constraints

(a) Implied Change in Risk Aversion

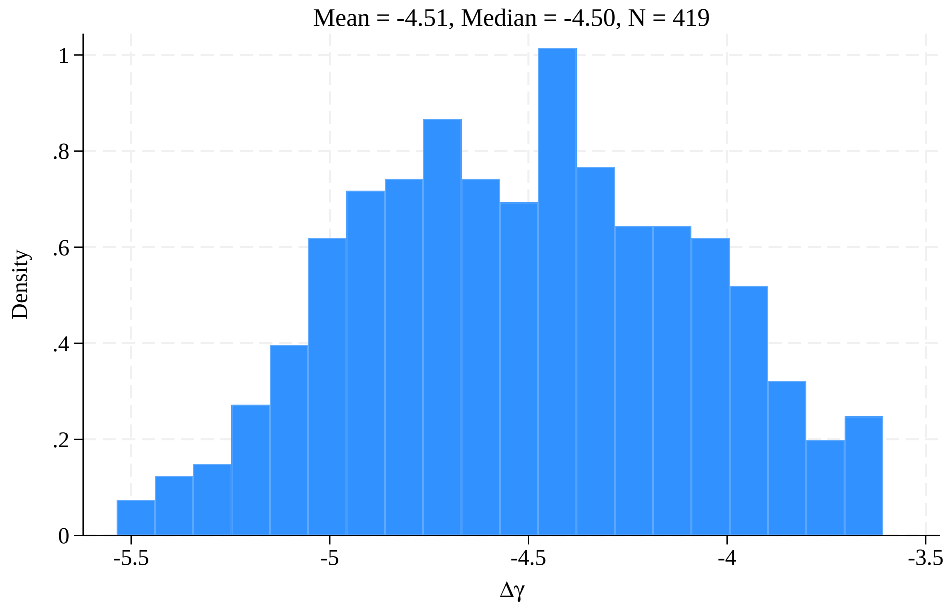

(b) Implied Shadow Cost

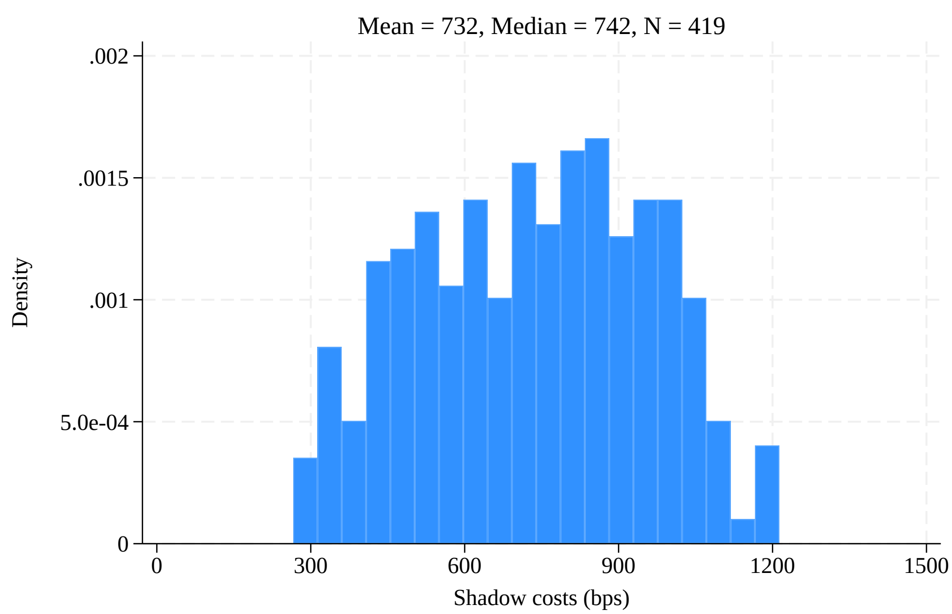

Notes: This figure shows simulation results based on the model in Section 3, in which pensions face a minimum constraint on fixed income investment. Panel (a) of the figure shows how the change in risk aversion required to match portfolio weights in 2001 and 2021 varies across simulations of initial 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 419 simulations in which a decline in risk aversion would explain the aggregate portfolio shifts. See Section 6.1 for details.

Figure A21: Simulated Alternative-to-Risky Share when Risk Aversion Goes to 1

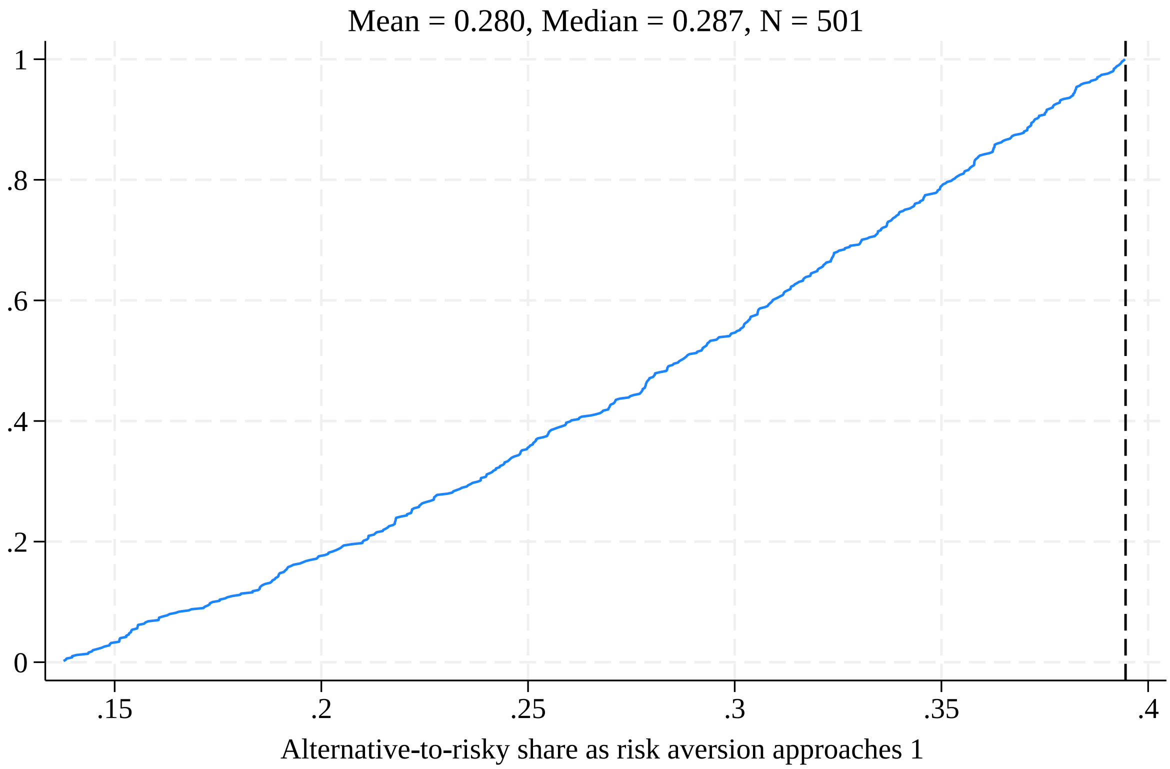

Notes: This figure shows results based on the simulation used in Section 6.1. In particular, the plot shows the cumulative distribution function for the 2021 alternative-to-risky share for the cases in which a decline in risk aversion plus a binding constraint is able to generate an increase in the alternative-to-risky share, but not enough to match the data before risk aversion hits its lower bound of 1. This corresponds to 501 out of 100,000 simulations.The black dotted line shows the true alternative-to-risky share as of 2021. See Section 6.1 for complete details.

Figure A22: Simulated Change in Beliefs About α

(a) Required Increase in Alpha

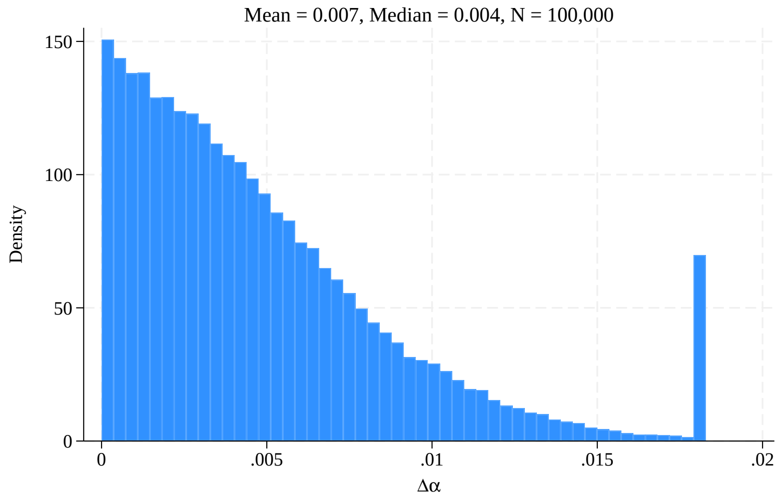

(b) Required Change in Risky Share

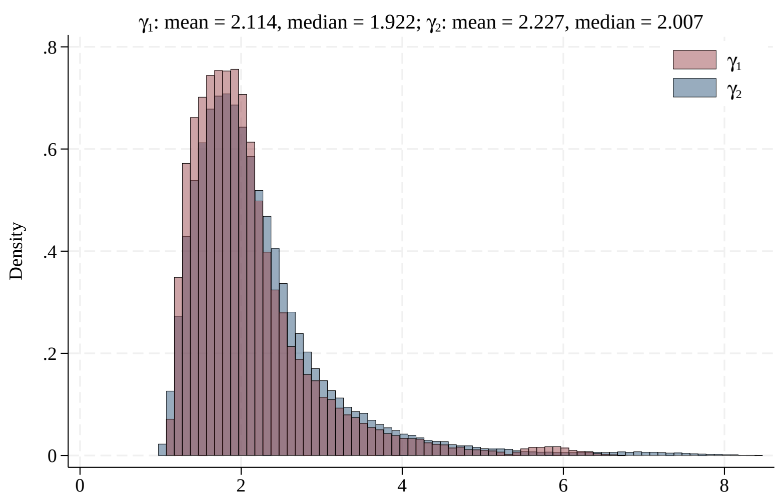

Notes: This figure shows belief change simulation results based on the model in Section 3. For the same initial sets of beliefs as used in Section 6.1 that match the initial portfolio weights in 2001, Figure A22a shows the change in CAPM-α required to match the change in the alternative-to-risky share 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 6.1 and Internet Appendix E.2 for details.

Figure A23: Implied Causal Impact of Consultant Beliefs on Pension Beliefs

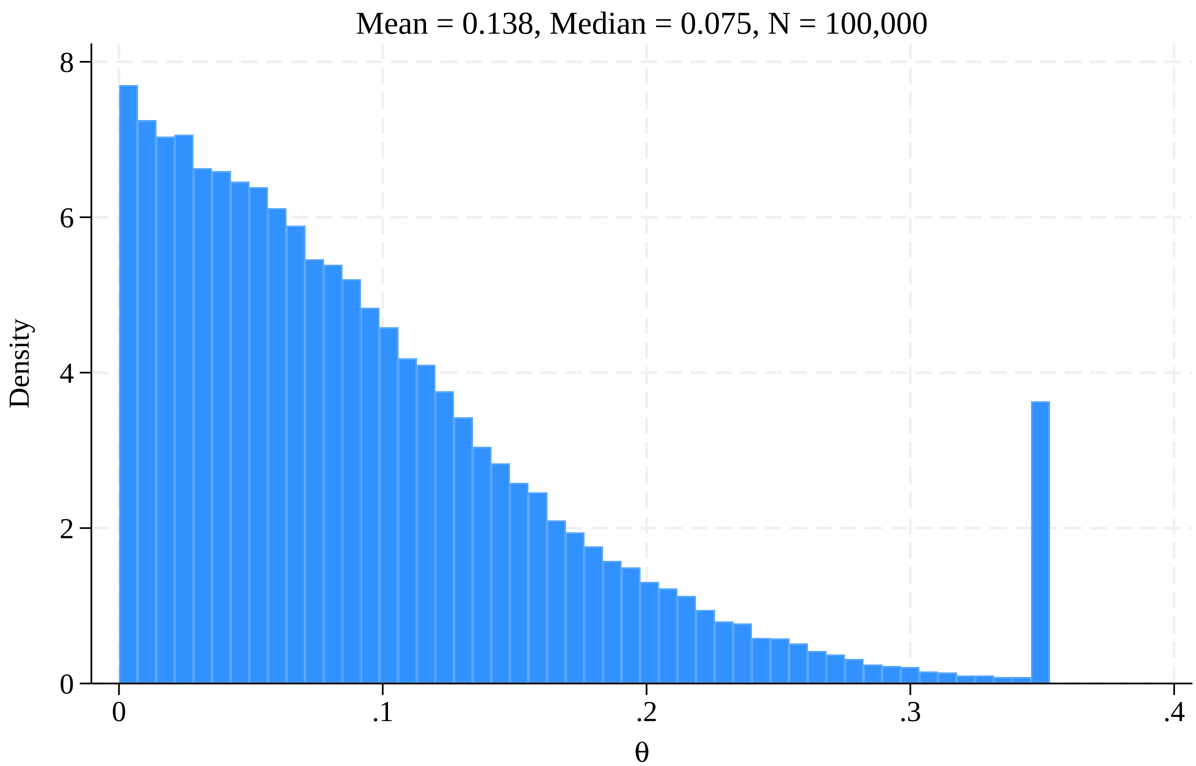

Notes: The figure plots the distribution of the elasticity of the pension-perceived alpha to consultant-perceived alpha from the expression $\Delta\alpha_p = \theta\Delta\alpha_c + \Delta\nu_p$. We construct $\theta = \phi/\Lambda$ using the derivative of alt-to-risky share with respect to pension alpha (Λ) from the simulations, as well as the estimated elasticity of the alternative-to-risky share to consultant beliefs ($\phi \approx 4$) from the IV strategy in Section 4.1. The results are based on 100,000 initial simulations that match the initial portfolio weights with reasonable initial beliefs in 2001. See Section E.2.2 for details. For illustration purposes, we winsorize data at the 97.5th percentile in the figure.

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

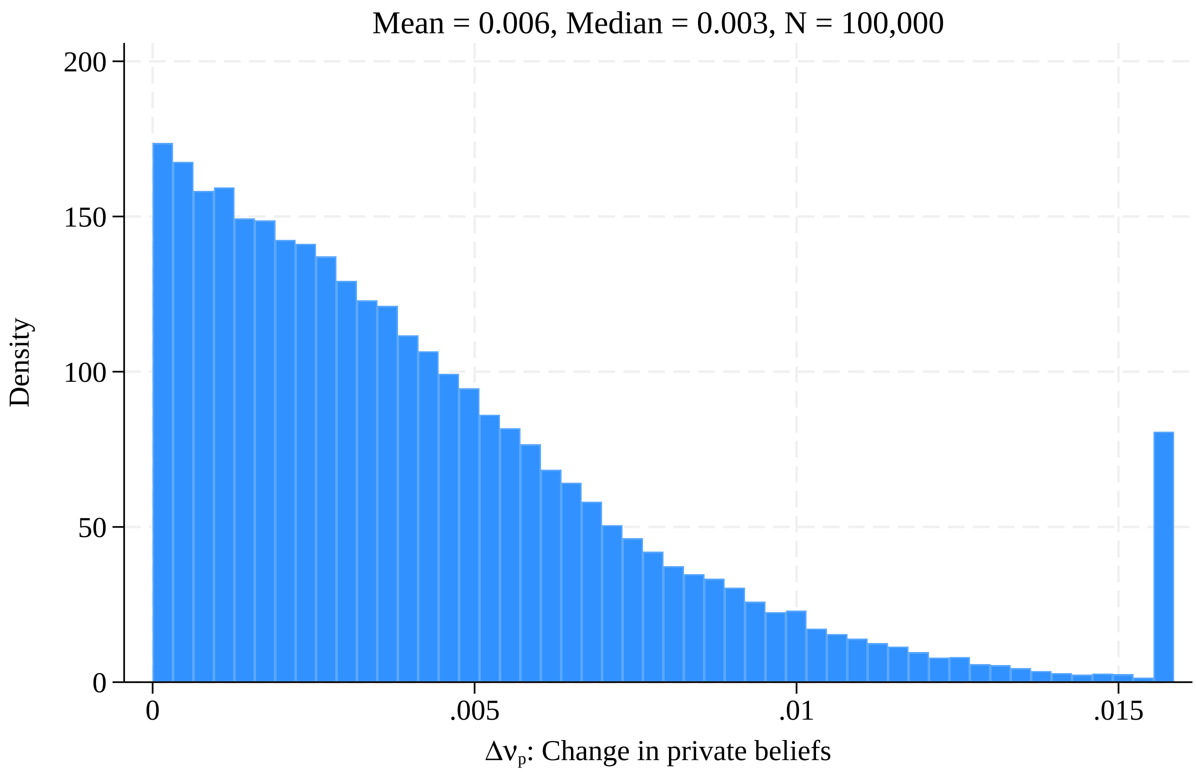

Notes: The figure plots the distribution of the changes in pensions' private beliefs $\Delta\nu_p$ from the model $\Delta\alpha_p = \zeta\Delta\alpha_c + \Delta\nu_p$, where $\Delta\alpha_c$ is the observed change in consultant-perceived alpha in Figure 4a and $\Delta\alpha_p$ is the required change 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 $\frac{\partial\omega_A^*}{\partial\alpha_p}$ from the simulations, as well as the estimated elasticity of the alternative-to-risky share to consultant beliefs from the IV strategy in Section 4.1. The results are based on 100,000 initial simulations that match the initial portfolio weights with reasonable initial beliefs in 2001. See Section E.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

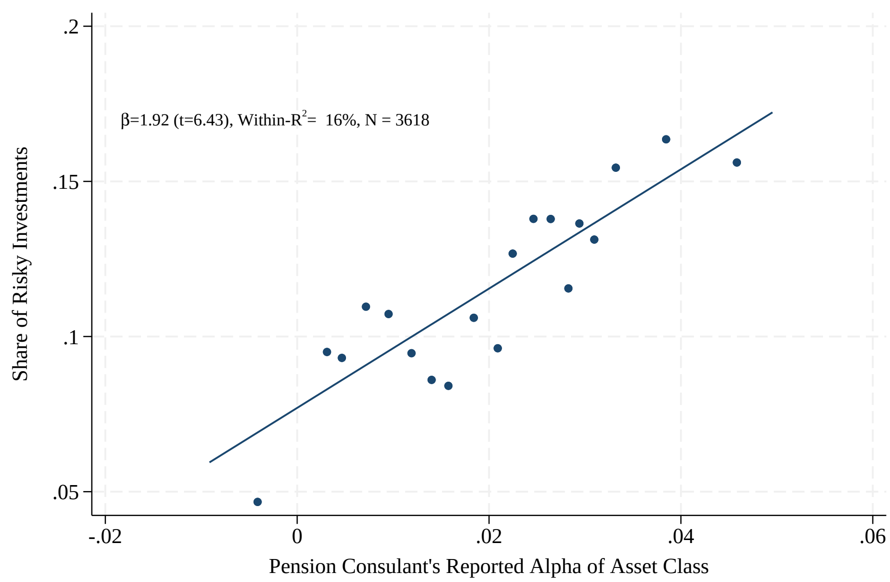

Notes: This plot shows a binscatter of each pension's allocation of risky investments to real assets or private equity against its consultant's reported alpha in each subcategory, after controlling for a pension-by-time fixed effect. The reported *t*-statistic is based on standard errors that are clustered by consultant and pension. See Section F.1 for details.

Table A1: Transition Rates for Quartiles of Portfolio Weights

(a) Risky Share

	2002				
2021	**1**	**2**	**3**	**4**	**Total**
	%	%	%	%	%
1	33	27	24	15	100
2	25	11	25	39	100
3	13	30	33	23	100
4	29	32	18	21	100
Total	25	25	25	24	100

(b) Alternative-to-Risky Share

	2002				
2021	**1**	**2**	**3**	**4**	**Total**
	%	%	%	%	%
1	37	30	20	13	100
2	27	20	27	27	100
3	17	37	23	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 Belief Decomposition

Panel A: Consultant Beliefs in 2001 and 2021

	2001	2021	Changes
α	0.0185	0.0233	0.0048
μ_A	0.0492	0.0573	0.0081
$\beta \cdot \mu_E$	0.0307	0.0340	0.0033
μ_E	0.0427	0.0443	0.0016
β	0.7204	0.7575	0.0371
$Vol(r_A)$	0.2311	0.1949	-0.0363
$Vol(r_E)$	0.1737	0.1681	-0.0056
$Corr(r_A, r_E)$	0.4500	0.6022	0.1522

Panel B: Decomposition

	Implied $\Delta\alpha$	Contribution
Actual	0.0048	100%
No changes in		
μ_A	-0.0033	168.4%
$\beta \cdot \mu_E$	0.0081	-68.4%
β	0.0069	-42.1%
μ_E	0.0064	-32.9%
Interaction effects		6.6%

Panel A of this table summarizes consultant's beliefs in 2001and 2021 from the CMA data. Panel B presents the implied changes in consultant perceived alpha assuming that there is no change in one of the components in the expression $\Delta\alpha = \Delta\mu_A - \Delta(\beta_A \mu_E)$.

Table A3: Predicting Changes in the Alternative-to-Risky Share

	ΔAlternative-to-Risky Share (2002-2020)			
	(1)	(2)	(3)	(4)
GASB 25 Funding Ratio in 2002	0.12			
	(1.65)			
BEA-Adjusted Funding Ratio in 2002		-0.18		
		(-0.81)		
Liability Discount Rate in 2002			2.26	
			(0.55)	
Fraction of Retired Members in 2002				-0.11
				(-0.66)
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 A4: 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	-0.00				-0.01			
	(-0.03)				(-0.17)			
BEA-Adjusted Funding Ratio		-0.01				-0.11		
		(-0.11)				(-0.99)		
Liability Discount Rate			3.20				1.52	
			(1.45)				(0.32)	
Fraction of Retired Members				-0.02				-0.07
				(-0.28)				(-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	0.41	0.32	0.32	0.00	0.01	0.00	0.00
N	3,121	1,008	3,090	3,104	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 A5: 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	-0.03							
	(-0.32)							
Δ BEA-Adjusted Funding Ratio		-0.25						
		(-1.10)						
Δ Liability Discount Rate			5.74**					
			(2.40)					
Δ Fraction of Retired Members				0.34**				
				(2.30)				
GASB 25 Funding Ratio in 2002					0.05			
					(0.67)			
BEA-Adjusted Funding Ratio in 2002						-0.08		
						(-0.46)		
Liability Discount Rate in 2002							-5.53	
							(-1.37)	
Fraction of Retired Members in 2002								-0.40**
								(-2.27)
Aggregation	System	State	System	System	System	State	System	System
Total R^2	0.00	0.06	0.04	0.03	0.00	0.00	0.02	0.08
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 A6: Variation in the Level of the Risky Share

	Risky Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GASB 25 Funding Ratio	-0.04				-0.06**			
	(-1.69)				(-2.02)			
BEA-Adjusted Funding Ratio		-0.11**				-0.14**		
		(-2.11)				(-2.46)		
Liability Discount Rate			4.36**				4.55**	
			(2.53)				(2.38)	
Fraction of Retired Members				0.04				-0.07
				(1.05)				(-0.99)
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	0.01	0.03	0.04	0.00	0.02	0.12	0.05	0.01
Total-R^2	0.14	0.22	0.17	0.14	0.02	0.12	0.05	0.01
N	3,128	1,008	3,097	3,111	158	51	157	155

Notes: This table shows panel regressions of the target risky share on several covariates. The risky share is defined as the share all investments 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 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 A7: Board Composition and the Alternative-to-Risky Share

	Alternative-to-Risky Share	
	(1)	(2)
State-appointed	0.11	0.10
	(0.68)	(0.67)
State-exofficio	0.01	0.03
	(0.07)	(0.29)
Participant-elected	0.01	0.00
	(0.22)	(0.04)
Public-appointed	0.05	0.08
	(0.62)	(0.96)
Year-FE	Yes	Yes
Controls	No	Yes
Within-R^2	0.06	0.08
Total R^2	0.45	0.46
N	1,861	1,843

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, and total spending scaled by payroll. 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.