The Rise of Alternatives*

Juliane Begenau† Pauline Liang‡ Emil Siriwardane§

January 2024

Abstract

Since the 2000s, U.S. public pensions have rotated out of public equities and into alternative assets like private equity and hedge funds. We explore several explanations for these trends, focusing on those implied by standard portfolio theory. Our evidence suggests that the rise of alternatives has been fueled by an increase in their perceived risk-adjusted returns relative to public equities. Pension beliefs are shaped by investment consultants, experience in the 1990s, and peers. Explanations rooted in risk-seeking motives, such as those driven by pension underfunding or low interest rates, have weaker empirical support.

*We are grateful to Aleksandar Andonov, John Y. Campbell, Josh Coval, Mark Egan, Daniel Green, Sam Hanson, Ben Hebert, Christian Heyerdahl-Larsen (discussant), Victoria Ivashina, Arvind Krishnamurthy, Hanno Lustig, Yueran Ma (discussant), Josh Rauh, Monika Piazzesi, Martin Schneider, Amit Seru, Erik Stafford, Jeremy Stein, Adi Sunderam, Boris Vallee, Luis Viceira, and participants at the Bank of Canada, Q-Group Fall 2023 Seminar, NBER AP SI 2023, Norges Bank Workshop "Women in Central Banking", MIT Sloan, Dartmouth, Stockholm School of Economics, Harvard, Stanford, and the University of Utah for their comments and discussions. We also thank the Kroner Center for Financial Research for their research support. This article represents the views of the authors and not the Kroner Center for Financial Research. Jiayi Cheng and Aaron Leung provided excellent research assistance. First draft is from November 2021. Earlier versions circulated under the titles "Investing with Style: the effect of consultants and peers on public pension portfolio composition" and "What determines public pension investment decisions?".

†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

1 Introduction

Over the last few decades, public pensions in the United States have fundamentally changed the way they take risk. At the turn of the century, risky investments – defined as any holding outside of fixed income and cash – were mostly in the form of public equities. Alternative assets like private equity, real estate, and hedge funds accounted for just 14% of risky investments in 2001 but grew to 39% by 2021.[1] These national trends also mask considerable heterogeneity across pensions. For example, the alternative-to-risky share for pensions in states like Maine, New Mexico, Indiana, Wyoming, and Texas has increased by an average of 58 percentage points since 2001, whereas it has hardly changed for pensions in South Dakota, Nevada, Georgia, Iowa, and Colorado.

In this paper, we investigate the key factors driving the adoption of alternatives, focusing on two explanations that are implied by modern portfolio theory (Markowitz (1952); Merton (1969); Campbell and Viceira (2002)). The first is that pensions have updated their beliefs about the risk-return properties of alternatives – specifically, they have become optimistic about the so-called alpha of alternatives relative to public equities, some more so than others.[2] The second is that pensions want to take more risk but are constrained from increasing their overall risky share, forcing them to contort the composition of their risky investments towards alternatives. The evidence we present is more consistent with the first explanation and less so with the second, leading us to conclude that beliefs have played a central role in the rise of alternatives.

The first half of the paper evaluates belief-based mechanisms. We start by analyzing the role of investment consultants, who provide advice on portfolio construction for almost all U.S. public pensions (Andonov, Bonetti, and Stefanescu, 2023). Following Couts, Goncalves, and Loudis (2023), the beliefs of investment consultants are extracted from capital market assumption (CMAs) reports that are published by many of the major investment consultancies. These data directly confirm that consultant-reported beliefs about the alpha of alternatives relative to public equities have risen steadily through time, increasing by about 68 basis points since 2001. Simulations show that

[1] The overall risky share went from 68% to 76% over the same period.

[2] Alpha is defined as $\alpha = \mu_A - \beta\mu_E$, where μ_A and μ_E are the expected excess returns of alternatives and public equities, respectively. β is the covariance between the two asset classes divided by the variance of public equities.

this change is large enough to generate the entire observed increase in the aggregate alternative-to-risky share for the myopic long-run investor modeled in Campbell and Viceira (2002). In the cross-section, there is also a strong and positive relation between a consultant's reported alpha and the alternative-to-risky share of its U.S. public pension clients, even after controlling for state-by-time fixed effects, pension funding, size, and several other attributes. Within the class of alternatives, consultant beliefs further explain why some pensions invest more in private equity relative to real assets. Since both asset classes have similar levels of liquidity, focusing on this choice helps isolate the link between beliefs and alternative investments, independent of a pension's inclination to invest in illiquid and opaque assets (Couts, Goncalves, and Rossi (2020); Stafford (2022)).

More generally, consultant identity is far more correlated with alternative use compared to a wide range of attributes like pension funding and size. These "consultant effects" are economically meaningful: clients of the 5th percentile consultant have an average alternative-to-risky share of 8%, whereas clients of the 95th percentile consultant have an average share of 51%. Additional tests suggest that consultants exert some causal influence on the beliefs of their clients, though matching between pensions and consultants based on shared beliefs is also possible. Either way, belief heterogeneity appears important for why the identity and stated beliefs of consultants correlate so strongly with alternative use in the cross-section of U.S. pensions.

Drawing on the literature linking belief formation to experience (Malmendier and Nagel (2016); Andonov and Rauh (2021)), we then show that pensions who experienced relatively poor performance during the 1990s were more inclined to move towards alternatives in the 2000s. Experience in the 1990s alone accounts for one-fifth of the cross-pension variation in changes in the alternative-to-risky share from 2002 to 2021. Our interpretation is that pensions who were late to invest in equities during the dot-com bubble of the late 1990s experienced lower returns, leading them to then perceive alternatives as more favorable than public equities. Notably, experience in the 1990s retains its predictive power even after controlling for funding levels in 2002, cutting against the notion that poor performance worsened funding frictions and instead supporting a belief-based channel. However, missing return data from the 1990s for some pensions limits the strength of this conclusion.

Motivated by household finance research on belief formation and social networks (Bailey et al. (2018, 2022)), we also test whether pensions' beliefs are influenced by their peers, where peers are based on distance. Consistent with this idea, pensions allocate similarly to their peers when choosing between alternatives and public equities, and the strength of this effect is large compared to the impact of funding, size, and other pension attributes. To account for consultant effects and common unobserved economic shocks (Angrist (2014)), peer effects are identified using variation within the same consultant, year, and census division. There are at least two possible interpretations of these findings. The first is that pensions learn about the risk-reward trade-off of alternatives from pensions who are geographically close, perhaps because their investment staff share information or attend local investment conferences. Another interpretation is based on the model of Scharfstein and Stein (1990), whereby pensions herd with their nearby peers to avoid public backlash for contrarian behavior. In reality, both channels are likely present in the data, however we rule in the first by showing that pensions with low incentives to herd still allocate similarly to their peers.

The second half of the paper explores whether risk-seeking motives can explain the rise of alternatives. We consider several proxies for pension risk-seeking motives, including those related to underfunding (Pennacchi and Rastad (2011); Mohan and Zhang (2014); Lu et al. (2019)), pension accounting and fund aging (Andonov, Bauer, and Cremers (2017)), and nominal return targeting. There is a weak and inconsistent link between these proxies and the alternative-to-risky share in the cross-section of pensions, both in long-run changes and levels. For example, a simple cross-sectional regression of changes in the alternative-to-risky share between 2002 and 2020 on changes in pension funding yields an R^2 of 1%. These broad conclusions are robust to several regression specifications and additional measures of risk-seeking incentives, including initial funding levels, funding ratios based on market discount rates (Novy-Marx and Rauh (2011)), and whether plan sponsors have failed to make actuarial required contributions.

In canonical portfolio choice models, greater risk-seeking motives will only affect the composition of risky investments (i.e., the alternative-to-risky share) if investors face a maximum constraint on their risky share. We explore this mechanism using two different approaches, both of which cast doubt on the ability of binding portfolio constraints to explain the rise of alternatives. First, we develop a measure of constraint tightness based on deviations of actual from target risky shares and show that it has little predictive power for the alternative-to-risky share in the cross-section of pensions. Second, we simulate a standard portfolio choice model (Merton (1969); Campbell and Viceira (2002)) with binding portfolio constraints, which allows us to quantify the strength of this potential channel without needing proxies for constraint slack or pension risk aversion. In each simulation, we draw beliefs about the expected risk and return of public equities and alternatives from a wide distribution and determine the risk aversion needed to match the aggregate pension portfolio in 2001. Holding beliefs fixed, we then lower risk aversion and assume that pensions are constrained from raising their risky share above its observed level in 2021. In virtually all simulations (99.6%), it is not possible to lower risk aversion enough to match the observed change in the alternative-to-risky share. This is because the low alternative-to-risky share in 2001 implies that pensions would shift toward public equities and not alternatives when their portfolio constraints bind. To be clear, the main point of this exercise is not to rule out mechanisms that involve risk-seeking or binding portfolio constraints. Rather, the model highlights that neither of these channels can generate the observed shift in the composition of risky investments without a simultaneous shift in beliefs.

In the final part of the paper, we consider the idea that agency frictions have driven the rise of alternatives. Our analysis puts restrictions on the nature of any such friction. For one, it would need to cause a wide range of institutions that differ in governance, regulation, and geography to rotate out of public equities and into alternatives, but not concurrently increase their overall risky share. This follows from the fact that U.S. public pensions, U.S. endowments, U.S. corporate pensions, and U.K. corporate pensions have all increased their alternative-to-risky shares since the 2000s, even though their risky shares have diverged sharply. Moreover, any agency friction must vary in the cross-section of public pensions and be orthogonal to a host of attributes including funding and size. Finally, it must lead to matching with consultants in a way that lines up with the alpha they report in CMAs, even within asset classes that provide a similar scope for hiding risk like private equity and real assets (Stafford (2022)).

Supply-side factors may also have contributed to the rise of alternatives. For example, investor access to privately held firms via private equity limited partnerships has improved over time. Indeed, as we show in Section 6.2, the supply of alternatives has expanded from 2% of all global risky assets in 2000 to 8% in 2020. While these types of supply-side factors are clearly relevant for understanding aggregate pension behavior, they cannot explain the large cross-sectional variation in alternative adoption.

This paper contributes to both public economics and finance, starting with research on how interest rates influence investor behavior. Much of this literature has focused on why investors appear to take more risk as interest rates fall (Borio and Zhu (2012)), with common explanations centering around agency frictions (Becker and Ivashina (2015)), institutional constraints (Campbell and Sigalov (2022)), and behavioral biases (Lian, Ma, and Wang (2018)). In the context of U.S. public pensions, underfunding and accounting distortions have been used to explain why the risky share has increased from 68% to 76% since the turn of the century (Pennacchi and Rastad (2011); Mohan and Zhang (2014); Andonov, Bauer, and Cremers (2017); Lu, Pritsker, Zlate, Anadu, and Bohn (2019)). In contrast, we ask why pensions have shifted the composition of their riskier assets away from public equities and toward alternative assets over this period. This compositional shift is large: since 2001, for every dollar that has flowed out of fixed income, $2.60 has moved into alternatives and $1.60 has flowed *out* of public equities.[3] Consistent with intuition from the two-fund separation theorem (Tobin (1958)), we argue that shifting beliefs are necessary to understand the adoption of alternatives.

An important precursor to our paper is Ivashina and Lerner (2018), who document that private and public sector pension funds around the globe have all increased their overall portfolio share of alternatives since 2008. We extend their research by studying the evolution of the alternative-to-risky share in conjunction with the risky share. This analysis reveals that the rise of alternatives

[3] Even under the assumption that hedge funds only invest in the stock market, the net decline in the portfolio share of public equities is 7.4 pp, as opposed to 12.8 pp without netting (see Section 2.2.2).

has been driven by a shift in the composition of risky investments, as opposed to a broader global expansion into all types of risky assets. For example, the alternative-to-risky shares for U.S. and U.K. corporate pensions have both increased since the early 2000s, yet their overall risky shares have both decreased. Our paper suggests that beliefs can help to jointly explain these trends.

Finally, our research connects to a growing literature that studies the relationship between beliefs and heterogeneity in household investment behavior (Pástor (2000); Leombroni et al. (2020); Beutel and Weber (2022); Giglio et al. (2023)). We highlight this link in the context of U.S. public pensions. In particular, our findings imply that disagreement about the alpha of alternatives may explain why dispersion in the alternative-to-risky share has risen so sharply since the 2000s. It is natural to expect disagreement in our setting given the challenges associated with measuring the risks of alternatives (Stafford (2022); Couts, Goncalves, and Rossi (2020); Korteweg (2019)).[4] This discussion also raises a related question of how beliefs are formed. Consistent with prior work, we show that consultants, peers, and past experience all play a role (Foerster et al. (2017); Bailey et al. (2022); Bordalo, Gennaioli, and Shleifer (2022)).

The subsequent sections of the paper are structured as follows. Section 2 provides an overview of the data utilized in our analysis and presents several facts about U.S. public pension investment behavior that serve as the foundation for our study. Section 3 uses the portfolio choice model of Campbell and Viceira (2002) to highlight possible explanations for these facts, most notably the increase in the alternative-to-risky share. Section 4 explores the role of beliefs for understanding the rise of alternatives and Section 5 focuses on risk-seeking channels. Section 6 discusses agency-based and supply-side explanations, then concludes. Additional details and results are available in an internet appendix.

[4] A growing body of work evaluates the investment performance of private equity (e.g., Kaplan and Schoar (2005); Phalippou and Gottschalg (2008); Harris, Jenkinson, and Kaplan (2014); Korteweg and Nagel (2016); Gupta and Van Nieuwerburgh (2019); Stafford (2022); Korteweg, Panageas, and Systla (2023)).

2 Data and Motivating Facts

The data for this study comes primarily from the Public Plans Data (PPD) that is maintained by the Center for Retirement Research (CRR) at Boston College. This section describes these data and their coverage of the broader U.S. defined benefit (DB) pension system. We then document that the aggregate alternative-to-risky share has risen sharply since the early 2000s, as has dispersion in the cross-section of U.S. public pensions.

2.1 Sample description and variable definitions

We obtain annual information on individual pension plans using the PPD. These data are based on comprehensive annual financial reports (CAFRs) that are filed annually by each DB public pension in the United States. The exact content and format of CAFRs vary across pensions and years, but all contain data on various plan characteristics like assets under management, portfolio composition, expected asset returns, actuarial value of liabilities, contribution rates, and information on beneficiaries. Throughout this paper, we refer to the expected rate of return on pension assets as an "asset hurdle rate" or "liability discount rate" because pension liabilities are discounted using the expected rate of return on pension assets under GASB 25 accounting standards. We supplement the PPD using hand-collected data on the identity of investment consultants (see Section 4.1.2). Information on chief investment officers and their salaries are taken directly from Lu et al. (2023).

Data in the PPD are reported at the plan level, though in many cases the assets of multiple plans are pooled and managed by pension "systems." For example, the board of the Colorado Pension Public Employees' Retirement Association invests and manages the pension assets of Colorado state employees, local school districts, the state's judicial system, and many local municipalities. Because our focus is on asset allocation decisions, our main unit of analysis is therefore at the system level. We map individual plans to larger pension systems based on hand-collected information and data from the Center for Retirement Research, and then aggregate plans to the system level accordingly.[5]

[5] Our mapping takes into account a handful of pension mergers that have occurred during our sample period and is

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. Unless otherwise noted, target weights are used in our subsequent analysis of the PPD data because these are chosen by pensions, whereas actual weights would also reflect market fluctuations.

available upon request.

Table 1 presents summary statistics for our sample of U.S. public pension funds, broken out over four equally-spaced time periods. We discuss these summary statistics and the PPD sample in greater detail in Internet Appendix A.1, where we also validate it against data from the U.S. Census Bureau. From 2006 onwards, the PPD covers over 90% of total U.S public pension assets. The table further shows that the PPD coverage is also fairly large in comparison to the broader U.S. pension system (roughly one-quarter), defined as the total amount of assets held by all private and public sector pension funds in Table L.117 of the Financial Accounts of the United States.

2.2 Trends in Portfolio Composition

We now present some basic facts about the investment behavior of U.S. public pensions, starting with the overall risky share and then turning to the composition of risky investments.

2.2.1 The Risky Share in the Aggregate and Cross-Section

Though our main focus is the composition of risky investments, the portfolio choice theory we outline later suggests that the risky share is informative for interpreting movements in the alternative-to-risky share. Recall that risky investments are defined as all holdings outside of cash and fixed income. In principle, high-yield debt securities should be included in risky investments, but granular data on the credit rating of fixed income investments are limited.[6] Consistent with prior work on public pensions (e.g., Andonov, Bauer, and Cremers (2017)), Figure 1a shows that the risky share has risen steadily since the turn of the century, albeit at a relatively modest pace. From 2001 to 2021, it went from 68% to 76%.

Figure 1b provides a longer-run perspective using the U.S. Census Bureau's Quarterly Survey of Public Pensions (QSPP), which covers the largest 100 pensions in the country starting in 1968. The first thing to notice is that risky shares from the QSPP and PDD largely overlap when both are available, reinforcing the quality of the PPD data.[7] Second, and more strikingly, the 8 pp increase

[6]The PPD does have a "PensionCreditRating" dataset that contains fixed income holdings by credit rating for a subset of pensions from 2004 to 2018. We discuss these data and their limitations in Internet Appendix A.1.2.

[7]In the plot, target weights are used from PPD whereas QSPP reflects actual weights.

in the risky share since the 2000s is relatively small compared to the 35 pp increase that occurred between 1970 and 2000. The risky share increased by nearly 17 pp in the 1990s alone, most of which was driven by a rotation out of U.S. Treasuries (see Internet Appendix B.1.1). Figure 1b also suggests that the risky share for U.S. public pensions may have reached a new steady state of just under 80% in the last five years.

Figure 1c 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. In Internet Appendix B.1.2, we show that this plot masks some turnover in which pensions take the most or least risk. We also discuss outliers in the risky share and provide a geographical sense of the risky share at the state level.

2.2.2 The Composition of the Aggregate Risky Portfolio

While the increase in the risky share since the 2000s has been relatively modest, there has been a considerable change in the *way* that U.S. pensions take risk. To illustrate this more clearly, Figure 2a starts by plotting the raw portfolio weights for fixed income, public equities, and alternatives based on PPD data. The first thing that stands out from the plot is the rise of alternatives. From 2001 to 2021, the share of alternatives in the national portfolio increased from 9% to 30%, mirroring a broader trend by pensions around the world (Ivashina and Lerner, 2018; Betermier et al., 2021). At the same time, the share of public equities fell from 59% to 46%. These flows imply that for every dollar that has shifted out of fixed income since 2001, $2.60 has moved into alternatives and $1.60 has flown *out* of public equities.

The preceding decomposition may overstate how much pension capital has moved out of public equities if hedge funds ultimately invest in the stock market. Figure 2b sheds some light on this issue by breaking out alternatives into subcategories for both 2001 and 2020. Hedge fund exposure has indeed increased by 5.4 pp during this period, though this change is not large enough to offset the contemporaneous decline of 12.8 pp decline in the public equity share. Thus, even if hedge funds are fully invested in public equities, public pensions have still decreased their exposure to the stock market by 7.4 pp.

Figure 2b further shows that all forms of alternatives have risen since the turn of the century. In 2001, the respective shares of real assets, private equity and credit, and hedge funds were 4%, 4%, and 0%. In 2021, their respective shares were 12%, 10%, and 6%.

As a simple way to summarize how the nature of risk-taking has changed, Figure 2c plots the evolution of the alternative-to-risky share (as opposed to raw shares). This object is a useful summary statistic for our purposes because it depends solely on beliefs in models where the two-fund separation theorem of Tobin (1958) holds (see Section 3). From 2001 to 2021, the alternative-to-risky share rose from 14% to 39%. Thus, U.S. public pensions are increasingly using alternatives over public equities to take investment risk.

2.2.3 Heterogeneity in the Composition of Risky Investments

We now analyze how the composition of risky investments varies across public pensions, which we summarize using the alternative-to-risky share $\omega^*_{A,pt}$ for pension p in year t. Figure 3a shows the distribution of $\omega^*_{A,pt}$ for each even year since 2001. There is a striking degree of cross-sectional variation in $\omega^*_{A,pt}$: in 2021, the alternative-to-risky shares for the 10th and 90th percentile pensions were 18% and 57%, respectively. This dispersion has also widened considerably over time. In 2001, the spread between the 10th and 90th percentile pensions was 26 pp.

Figure 3b shows the distribution of changes in $\omega^*_{A,pt}$ across pensions systems from 2002 to 2021. Echoing the widening dispersion in Figure 3a, the shift into alternatives has varied strongly across pensions: the 25th pension increased its alternative-to-risky share by 15 pp whereas the 75th percentile increased by 35 pp. Table A1b of the Internet Appendix shows that these changes have resulted in some amount of turnover in terms of the pensions with the highest alternative-to-risky share. For example, 21% of the pensions who were in the top quartile of $\omega^*_{A,pt}$ in 2021 were in the bottom quartile in 2002. At the same time, 13% of pensions who were in the bottom quartile of $\omega^*_{A,pt}$ in 2021 were also in the top quartile in 2002.

In Internet Appendix Section B.2, we discuss specific outlier pensions in terms of alternative usage, provide state-level summaries of the alternative-to-risky share, and repeat our cross-sectional analysis on the portfolio share of alternatives (as opposed to the alternative-to-risky share). Overall, the results in this section reveal large heterogeneity in the use and adoption of alternatives across U.S. public pensions.

3 Conceptual Framework

We now use standard portfolio theory to develop testable hypotheses about the potential causes of the sharp rise in the level and dispersion of the alternative-to-risky share. Details and derivations are in Internet Appendix E. In the tradition of Markowitz (1952) and Merton (1969), we study a power-utility investor who lives two periods and begins with wealth W_0. As shown in Campbell and Viceira (2002), the optimal portfolio for this investor is equivalent to that of a myopic, long-lived investor who ignores intertemporal hedging motives (Merton, 1973).[8] The investable universe consists of three assets: a riskless asset, public equities, and alternatives. The log return on each is denoted by r_f, r_E, and r_A, respectively. Returns on public equities and alternatives are assumed to be jointly log-normal, with the perceived mean μ and variance-covariance matrix Σ of log excess returns given by:

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

[8] It seems plausible that U.S. public pensions ignore or respond weakly to the intertemporal hedging motives of Merton (1973), even if this behavior may not be optimal in the long run. The reason why is that consultant and pension annual reports rarely make mention of large-scale market-timing strategies, but often mention standard elements of mean-variance optimization.

In what follows, it will be useful to express the elements of μ and Σ as parameters from a CAPM-style regression of alternatives on public equities:

$$r_A - r_f = \alpha + \beta(r_E - r_f) + \eta_A, \tag{1}$$

where η_A is an idiosyncratic shock to alternatives that has a volatility of σ_η. Note that this implies $\mu_A = \alpha + \beta\mu_E$, $\sigma_A^2 = \beta^2\sigma_E^2 + \sigma_\eta^2$, and $\sigma_{AE} = \beta\sigma_E^2$.

The investor chooses portfolio weights $\omega = [\omega_f, \omega_A, \omega_E]'$ in each asset class to maximize $\mathrm{E}\left[(1-\gamma)^{-1}(W_0 R_p)^{1-\gamma}\right]$, where R_p is the simple return of its portfolio. Using the results in Chapter 2.1.3 of Campbell and Viceira (2002), we show in Internet Appendix E.2-E.3 that the optimal portfolio weights are as follows:

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

$$\omega_E = \frac{1}{\gamma}\left(\frac{\mu_E}{\sigma_E^2} - \frac{\alpha\beta}{\sigma_\eta^2} + \frac{1}{2}(1-\beta)(\beta^2\frac{\sigma_E^2}{\sigma_\eta^2} + 1)\right), \tag{3}$$

$$\omega_f = 1 - \omega_A - \omega_E.$$

Here, we have expressed the portfolio weights in terms of the CAPM-regression coefficients in Equation (1). The optimal alternative-to-risky share is defined as $\omega_A^* := \omega_A/(\omega_A + \omega_E)$.

Belief-Based Explanation. The optimal portfolio shares in Equations (2) and (3) highlight how beliefs could explain the rise of alternatives. More formally, they imply that the alternative-to-risky share ω_A^* is increasing in the "alpha" (α) of alternatives relative to public equities:

$$\frac{\partial\omega_A^*}{\partial\alpha} = \frac{(\beta\omega_A + \omega_E)}{\gamma\sigma_\eta^2(\omega_A + \omega_E)^2} > 0, \tag{4}$$

if $\beta\omega_A + \omega_E > 0$. This condition is trivially satisfied in practice. Intuitively, the investor tilts the composition of its risky investments towards alternatives when its perceived alpha of alternatives rises. This result also means that cross-sectional dispersion in the alternative-to-risky share is modulated by the degree of disagreement about alpha.

Risk-Based Explanation. Equations (2) and (3) further show that a version of the Tobin (1958) separation theorem holds in this model, as the optimal alternative-to-risky share does not depend on risk aversion γ.[9] However, it is well known that this separation result breaks if there is a constraint on the maximum risky share, or equivalently, the minimum amount of riskless investment, $\omega_f \geq \omega_f^{min}$. When the constraint is binding ($\omega_f = \omega_f^{min}$), the optimal portfolio shares are instead given by:

$$\begin{aligned} \omega_A &= \frac{1}{\gamma}K + (1-\omega_f^{min})C, \\ \omega_E &= -\frac{1}{\gamma}K + (1-\omega_f^{min})(1-C), \end{aligned} \tag{5}$$

where K and C are functions of beliefs (see Internet Appendix E.3 for more details). In this case, it is straightforward to show that a change in risk aversion γ affects the composition of the risk-portfolio ω_A^* as follows:

$$\frac{\partial \omega_A^*}{\partial \gamma} = -\frac{1}{\gamma^2}\frac{1}{1-\omega_f^{min}}K.$$

Thus, for some initial beliefs ($K > 0$), the alternative-to-risky share can increase if investors want to take on more risk but are constrained from doing so. In practice, an increased desire to take risk could be driven by factors that are outside of our model, such as underfunding or low interest rates.

In summary, standard portfolio theory suggests two potential explanations for the rise of alternatives. The first is that pensions have become more optimistic about the alpha of alternatives relative to public equities. The second is that pensions have become more risk-seeking but are constrained from increasing their risky share. The main objective of the paper is to gauge the relative empirical strength of both explanations, though these channels are not mutually exclusive.

[9] Though risk aversion γ appears in Equation (4), it is straightforward to show that it actually cancels out because Equations (2) and (3) show that ω_A and ω_E are both functions of γ.

4 Belief-Based Explanations

In this section, we examine whether the rise of alternatives can be explained by a change in beliefs about their risk-return properties. We start by documenting that the perceived alpha of alternatives reported by pension investment consultants has indeed risen since the 2000s. Moreover, the magnitude of the increase appears large enough to explain the observed aggregate change in the alternative-to-risky share, at least from the perspective of the portfolio choice model from Section 3. Additionally, the stated beliefs and identity of pension consultants correlate strongly with alternative use in the cross-section. We then examine two classic factors that influence beliefs: past experiences and peer behavior. Finally, we show that many types of institutions have shifted toward alternatives, not just U.S. public pensions. Together, these findings support the idea that beliefs have played a key role in the rise of alternatives.

4.1 Evidence from Pension Consultants

4.1.1 Directly Reported Beliefs

Most public pension investment boards rely on an outside consultant to provide advice on portfolio construction. The nature of advice provided by consultants can range from broad asset allocation decisions to specific portfolio manager selection. Our focus is on broader asset allocation decisions. For this reason, we hand-collect data on the identity of the general consultant used by each pension for each year in the PPD data. In some cases, this information can be found easily on publicly available comprehensive annual reports. When it is not available from publicly available sources, we file Freedom of Information Act requests (FOIAs) with the pensions directly, asking for the "identity of the consultant that primarily advises on broad asset allocation decisions (e.g., percent in public equities, fixed income, private equity, etc.)." These collection efforts produce a fairly high coverage rate in our sample, as we have general consultant information for over 98% of the system-year observations in the PPD data. After accounting for mergers and name changes, there are 57 distinct consultants that have multiple public pension clients in a given year.[10]

After matching consultants with pension systems, we extract consultant beliefs using annual reports on Capital Market Assumptions (CMAs) that are published by most major consultancies. A typical CMA contains beliefs about expected returns, volatilities, and correlations of different asset classes (Couts et al., 2023). Depending on the consultant, these can range from five to thirty-year forecasts. The consultants for whom we have CMAs manage the majority of U.S. pension assets (~75%) but are somewhat limited in total count. We have CMAs from 14 consultants. Additional details on our CMA data and how we process it can be found in Internet Appendix A.6. We recognize that beliefs reported in CMAs may not reflect true beliefs and could instead be shaped by agency frictions. For example, consultants may strategically report that private equity has a high alpha to attract business from pensions that want to justify their allocation to private equity. However, determining whether CMAs reflect the true beliefs of consultants is outside of the purview of this study and we assume that they do for the remainder of the paper.

Figure 4a plots the median consultant's reported alpha for each year since 2001.[11] Alternatives include real estate, private equity, and hedge funds, and alpha is computed by averaging across the asset classes that are available in a given CMA. The perceived alpha for the median consultant has risen steadily since the early-2000s, going from 158 basis points in 2001 to 226 bps in 2021. In Internet Appendix D.1, we decompose this increase into two main factors. First, the expected excess return of alternatives has risen, contributing roughly 32% to the total change in perceived alpha. Second, the perceived diversification benefits of alternatives have risen, as measured by a decline in their perceived beta with respect to public equities.

The portfolio theory outlined in Section 3 suggests that an increase in perceived alpha is one way to explain the rise of alternatives. To get a back-of-the-envelope sense of magnitudes, let the portfolio shares of alternatives and public equities equal their aggregate values in 2001 and assume the beta of alternatives is the average reported by consultants in CMAs (see Figure A19b). In the

[10] We assume that no consultant was used when we could not find information for a pension in a given year.

[11] Alpha is computed using the expected excess return of each asset class over cash and the variance-covariance matrix in the CMAs.

language of the model, this means that $\omega_A = 0.09$, $\omega_E = 0.59$, and $\beta = 0.6$. Further assume that the idiosyncratic volatility of alternatives equals $\sigma_\eta = 13.6\%$ and the volatility of public equities equals $\sigma_E = 17\%$, roughly its long-run historical average. This implies that the volatility of alternatives also equals $\sigma_A = 17\%$. For a risk aversion of $\gamma = 2$, Equation (4) implies that the increase in perceived alpha in Figure 4a translates to an increase in the alternative-to-risky share of 26 pp, matching the actual change shown in Figure 2c.

In Internet Appendix D.2, we enrich this back-of-the-envelope calculation by simulating changes in alpha within the model. For each simulation, we draw an initial set of random beliefs and solve for the required value change in the risk aversion parameter to match the aggregate pension portfolio in 2001. We then solve for the implied change in α needed to match the observed change in the alternative-to-risky share. The average required increase in α across simulations is 70 bps, strikingly close to the increase of 68 bps shown in Figure 4a.

The preceding analysis of the aggregate alternative-to-risky share implicitly assumes that consultant beliefs are relevant for understanding pension portfolios. Next, we confirm that this is indeed the case by studying the cross-sectional relationship between a consultant's beliefs and the composition of its clients' risky portfolio. Figure 4b depicts this relationship using a binscatter plot of each pension's alternative-to-risky share in year t against its consultant's reported belief about alpha in the same year.[12] Importantly, state-by-year fixed effects are included in the plot to absorb any state-level trends driving alternative adoption, such as regulation, governance, or agency-based factors. The plot also includes controls for the level of pension funding, (log) size, and the expected return on each pension's assets (i.e., its hurdle rate), the ratio of required actuarial contributions to payroll, and the ratio of administrative expenses to payroll. These controls are designed to capture risk-seeking motives and their logic is discussed further in Section 5.1. Figure 4b clearly shows a strong and positive relationship between each pension's alternative-to-risky share and its consultant's reported alpha, and a linear regression yields a coefficient of 3.00 ($t = 3.35$).[13] In Internet

[12] Throughout this paper, we follow Cattaneo et al. (2023) when constructing binscatter plots.

[13] Standard errors in the plot are double-clustered by consultant and pension. Because the number of consultant clusters is relatively small, we also test the null that the regression coefficient on reported alpha equals zero using the wild bootstrap procedure in Cameron et al. (2008). The p-value in this case equals 0.037.

Appendix C.1.3, we show that the controls in the regression add almost no explanatory power to the regression, implying that the within-R^2 of 12% is attributable almost entirely to variation in consultant-reported α.

Two additional tests further highlight the importance of consultant beliefs for explaining variation in the composition of risky investments. First, controlling for each pension's overall risky share in Figure 4b has a small impact on the linear regression coefficient, as it falls only slightly to 2.68 ($t = 2.97$). Second, we test whether beliefs explain the use of different types of alternatives using the following panel regression:

$$\omega^*_{a,p,t} = f_{p,t} + \beta V_{c(p),a,t} + \varepsilon_{p,a,t}, \tag{6}$$

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. Standard errors in the regression are double-clustered by consultant and pension.

Figure 4c visualizes Regression (6) 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 types of alternatives considered in the plot are real assets versus private equity. The choice between these two is useful to study because they are similar in many dimensions, most notably liquidity and leverage, making it more likely that the identifying variation in the regression is driven by beliefs and not agency frictions (see Section 6.1). The estimated coefficient from Regression (6) equals 1.92 (t =6.43) and indicates that pensions tend to allocate more toward a specific type of alternative 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.

Overall, these findings lend direct support to a belief-driven explanation for the rise of alternatives. We now develop this argument further by measuring the extent to which consultant identity

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

can explain variation in alternative use by public pensions.

4.1.2 Consultant Effects

One limitation with our CMA data is that it is available for 62% of the pension system-year pairs in our sample. This makes it harder to develop a general sense of how of much consultant beliefs matter for portfolio composition in the cross-section of U.S. pensions. As one way to make progress on this question, we build on Bertrand and Schoar (2003)'s study of corporate managers and use a series of fixed-effects regressions to gauge whether pension portfolio structure varies systematically with consultants. Specifically, we estimate the following regression:

$$\omega^*_{A,p(c),t} = \lambda_t + \sum_k \beta^k_t \times X^k_{pt} + \lambda_c + \varepsilon_{pct}, \tag{7}$$

where $\omega^*_{A,p(c),t}$ is the alternative-to-risky share for pension p matched with consultant c in fiscal year t. λ_t is a time fixed effect and X^k_{pt} is a set of pension attributes whose coefficients are allowed to vary through time. The set of attributes that are included are (log) size, GASB 25 funding, asset hurdle rates, the ratio of required actuarial contributions to payroll, and the ratio of administrative expenses to payroll. Given the results from Section 4.1.1, it is natural to interpret the consultant fixed effects λ_c as standing in for each consultant's average belief about the alpha of alternatives.

The first row of Table 2 presents the adjusted R^2 obtained from estimating the model with only time fixed effects and the second provides an additional benchmark by adding pension attributes. Consistent with results in later sections (Section 5), the inclusion of pension attributes only adds 1 pp of explanatory power to the regression model, even when their coefficients are allowed to vary through time. In contrast, the third row shows that consultant fixed effects add a considerable amount of explanatory power, as the adjusted R^2 rises by 17 pp after their inclusion. The F-statistic and associated p-value reflect this large incremental R^2 and indicate that a test of equal consultant effects is easily rejected by conventional statistical standards.

Though the F-test points to the existence of consultant effects, it is less informative about magnitudes. As a way to assess economic size, Figure 5 plots the distribution of the consultant fixed effects after applying a Bayesian shrinkage to account for sampling error (Casella (1992), Eqs. 7.11 and 7.13). The average alternative-to-risky share is added back to fixed effects to facilitate their interpretation. The plot illustrates that consultant effects are an economically important source of cross-pension variation in portfolio choice. Clients of the 5th percentile consultant have an average alternative-to-risky share of 8%, whereas clients of the 95th percentile consultant have an average share of 51%. It is worth emphasizing that these estimates are orthogonal to pension attributes, most notably size and funding.

Rows (5) to (16) of Table 2 repeat this exercise by testing whether consultant effects exist for different types of alternatives, not just the alternative-to-risky share. Rows (7), (11), and (15) indicate that this is indeed the case as there are also large consultant effects for private equity, hedge fund, and real asset investments. The private equity effects have a correlation of 8% and 36% with the real asset and hedge fund effects, respectively. The real asset and hedge fund effects are -2% correlated. These relatively low correlations indicate that consultants whose clients invest heavily in one type of alternative may not necessarily invest heavily in others. We interpret these results as consultants disagreeing on the alpha of different alternatives and advising their clients accordingly.

4.1.3 Interpretation

The results in Sections 4.1.1 and 4.1.2 indicate that the identity and stated beliefs of investment consultants are strongly correlated with alternative use in the cross-section of public pensions. These results can be interpreted in at least three ways. First, consultants may exert a direct, causal influence on the beliefs – and hence the portfolio composition – of their clients. Second, pensions may choose consultants based on shared beliefs about the alpha of alternatives. Third, pensions may prefer alternatives for risk-seeking or agency-based reasons and choose consultants who justify these preferences.

To better understand which of these three mechanisms is operant in the data, we now show that public- and private-sector pensions that share a consultant have correlated alternative-to-risky shares, yet the same is not true of their 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 2, including size and pension funding.

Figure 6a 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.[15] 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 Internet Appendix A.3. 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 6b 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 6 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. We revisit possible agency-based interpretations later in Section 6.1, after presenting all of our other results.

[15] 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 2 before collapsing to the consultant-year level.

This discussion still leaves open the question of whether consultant effects arise due to a causal channel or selection by pensions based on beliefs. One way to rule in causality is to test if the consultant effects in Table 2 survive the inclusion of pension and state-by-time fixed effects. This is a valid test of causality if conditional on the controls in Table 2, each pension's idiosyncratic preference for alternatives can be fully decomposed into a component that does not vary through time and one that varies through time, but only at the state-level (e.g., state-specific changes in regulation or agency frictions). Rows (4), (8), (12), and (16) of Table 2 confirm that consultant effects are still present in this more stringent specification, suggesting that consultants may have some causal impact on portfolio construction.

Taking a step back, we want to emphasize that our goal is not to settle whether consultant effects arise due to selection or causality, as both are likely to occur in reality. Regardless of whether consultants have a causal effect on clients, the evidence in this subsection makes a broader point, namely that beliefs are likely an important driver of why cross-pension variation in the composition of risky investments loads so heavily on consultants. 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.

4.2 Experience

Our analysis thus far has highlighted the importance of beliefs about alternative assets' risk-return properties for understanding cross-pension differences in alternative investment intensity, with a specific focus on investment consultants. There is a growing body of research that shows beliefs are also often shaped by experience. For example, Malmendier and Nagel (2016) show how individual inflation expectations are strongly influenced by the amount of realized inflation experienced during one's lifetime. Andonov and Rauh (2021) show public pension return expectations are influenced by past returns. Motivated by this work, we now provide suggestive evidence that investment experience during the 1990s shaped the subsequent evolution of beliefs about the risk-reward benefits of alternatives.

Figure 8 plots the change in the alternative-to-risky share for each pension between 2002 and 2021 ($\Delta\omega^*_{A,p}$) against its geometric average return from 1992-2002 (e_p). The sample size for the plot is lower than our main analysis sample because many pensions do not report 10-year returns on their annual reports and we generally do not have annual reports prior to the 2000s. With that said, the 39 pension systems in the plot are fairly large and cover 51% of PPD assets in 2002. The figure clearly shows that the pensions who had low performance during the 1990s were also more likely to shift to alternatives after the 2000s. Experience alone can explain 18.5% of the cross-sectional variation in the change in alternatives. As we show later in Section 5.1, this R^2 is three-to-four times larger than what is obtained when using various proxies for changes in pension risk aversion, including pension underfunding.

Our preferred interpretation of Figure 8 is as follows: many U.S. pensions rapidly increased their overall risky share during the 1990s (see Figure 1b), primarily through an increase in public equities.[16] Public markets during the 1990s famously went through a large boom and bust cycle that culminated with the bursting of the dot-com bubble in 2000. Pensions who were late to shift into equities were therefore more exposed to the bursting of the bubble relative to pensions who invested earlier, resulting in lower relative returns for the period between 1992 and 2002. The poor experience then caused these pensions to view alternatives as more favorable relative to public equities on a risk-return basis, thereby explaining why they shifted more aggressively towards alternatives after the 2000s.

A different interpretation of Figure 8 is that poor performance in the 1990s led to a decline in pension funding and thus a desire to take more risk (e.g., reaching-for-yield). As shown in the model from Section 3, such an increase could cause a shift to alternatives if there are constraints on the maximum risky share. However, inconsistent with this hypothesis, the level of GASB 25 funding in 2002 does not predict the change in the alternative-to-risky share from 2002 to 2021 (see Internet Appendix C.1.2). Moreover, we obtain a similarly strong relationship between $\Delta\omega^*_{A,p}$

[16] According to the ASPP, virtually every dollar that left fixed income during the 1990s went into public equities.

and e_p even after controlling for the level of GASB 25 funding in 2002. For instance, the point estimate of a linear regression of $\Delta\omega^*_{A,p}$ on e_p moves from -5.9 (t = -2.9) to -6.5 (t = -3.0) after controlling for the level of 2002 funding, whereas the coefficient on funding is not statistically different from zero.[17]

Overall, Figure 8 indicate that pension experience during the 1990s played a meaningful role in the rise of alternatives through its impact on beliefs. With that said, we view the evidence in Figure 8 as more suggestive in nature given the lack of complete data.

4.3 Peer Effects

Finally, we investigate whether pension portfolio decisions can be explained by the behavior of peers. This line of inquiry is motivated by recent research in household finance showing that beliefs about asset prices and product selection are shaped by social networks (Bailey et al. (2018, 2022)). We define peer networks in our context based on geographical distance. More formally, for each pension p in year t, we define the target alternative-to-risky share of peers as follows:

$$n_{pt} \equiv \sum_{k \neq p} \delta_{pk} \times \omega^*_{A,k,t},$$

with weight $\delta_{pk} = d^{-1}_{pkt} \times \left(\sum_{j \neq p} d^{-1}_{p,j,t}\right)^{-1}$. The target share of alternatives for pension k at time t is $\omega^*_{A,k,t}$. d_{pkt} is the distance (in kilometers) between the headquarters of pension p and pension k in year t. We measure distance using the 5-digit zip codes of each pension in our sample.[18] The weights δ_{pk} are therefore based on the inverse distance between pension systems.

We then measure the degree of peer effects using variants of the following panel regression:

$$\omega^*_{A,p,t} = \lambda_{cdt} + \beta_z n_{pt} + \theta' \mathbf{X}_{pt} + \varepsilon_{pt}, \tag{8}$$

[17] We show in Internet Appendix D.4 that poor 1990s performance also predicts an increase in the risky share, albeit weaker than the alternative-to-risky share. See the appendix for a complete discussion.

[18] For pension systems located in the same zip code, we assume they are 1.6 kilometers (1 mile) apart.

where λ_{cdt} is a consultant-by-time-by-census-division fixed effect and $\mathbf{X}_{pt}$ is a vector of pension observables that includes each pension's GASB 25 funding ratio, asset hurdle rate, (log) size, required actuarial contribution relative to payroll, and administrative expenses relative to payroll. The fixed effect λ_{cdt} is included for two reasons. First, if peers are more likely to choose the same consultant, then β_z will overstate the impact of peers because it will also reflect consultant effects (Section 4.1.2). Second, a common issue in the peer effects literature is separating the impact of peers from common shocks that may affect portfolio choice (Angrist (2014)). For example, pensions who are close in distance may allocate similarly because they experience similar local economic conditions. The inclusion of the fixed effect controls for both potential confounding factors and means that β_z is identified from variation within the same consultant, census division, and year. Standard errors in the regression are clustered by state and time.

Column (1) of Table 3 reports our baseline estimate of the regression specification in (8). The point estimate is estimated with statistical precision and indicates a fairly large pass-through of peer portfolio choices. On average, a 10 pp increase in the alternative-to-risky share of a pension's peers is associated with a 6.9 pp increase in its own alternative-to-risky share. As we show later in Section 5.1, this elasticity further suggests that peer behavior explains the composition of risky investments much better than a pension's own funding level, asset hurdle rate, or other proxies for risk aversion.

There are at least two possible interpretations of this finding. The first is that pensions learn about the risk-reward tradeoff of alternatives from pensions who are geographically close. There are several ways this type of learning could occur. For example, it seems plausible that the investment staff and chief investment officers (CIOs) of nearby pensions are more likely to interact with each other, perhaps by attending the same investment conferences or workshops. Another interpretation is based on the herding model of Scharfstein and Stein (1990), whereby pensions herd with their nearby peers to avoid public backlash for contrarian behavior. While both channels are likely present in the data, columns (2) to (4) provide suggestive evidence for the first channel by testing whether peer effects are still present in subsets of the data where herding incentives are likely weaker.

In column (2), we develop a test based on the idea that CIOs with more job security have less of an incentive to herd. We measure job security by first computing CIO tenure for each pension-year observation in our sample. CIO identities are taken from Lu et al. (2023) and are only available for a subset of pensions from 2001 onward. This means that we cannot perfectly measure tenure for CIOs who began prior to 2001. As one way to get around this issue, we create an indicator variable based on whether CIO tenure is at least five years and only estimate Regression (8) using data after 2005. The cutoff of five is roughly the median tenure in the post-2005 data. This indicator variable should capture CIOs who are relatively "established", even if we cannot perfectly measure their tenure. When estimating Regression (8), we then interact the indicator variable with the alternative-to-risky share of each pension's peer and include it (along with the indicator itself) in the regression. The positive interaction term in column (2) suggests that, if anything, peer effects are stronger for established CIOs, however the estimate is not statistically different from zero.

In column (3), we instead create an indicator variable for whether a pension is well-funded, defined as being in the top quartile of GASB 25 funding over our full sample. In practice, this means we only include pensions whose GASB 25 funding is above 88%. The assumption in this test is that well-funded pensions are under less public scrutiny and therefore have less incentive to herd. The negative interaction term in column (3) is consistent with this assumption, but nonetheless implies that well-funded pensions are responsive to their peers.

Finally, in column (4), we create an indicator variable based on whether a pension is in top quintile of overall performance. Similar to column (3), the idea is that high-performing pensions have less negative public scrutiny and should therefore be less inclined to herd as in Scharfstein and Stein (1990). The negative and statistically significant interaction term in column (3) suggests that this may be the case, but much like our previous results, we continue to find an economically and statistically meaningful peer pass-through coefficient for this subgroup of pensions.

For robustness, in column (5) we replace the contemporaneous alternative-to-risky share of peers with its lagged value. Though imperfect, this is one way to allay fears that the regression coefficients are driven by reverse causality. The point estimate in column (5) is similar in magnitude to that in column (1) and remains statistically significant. We therefore interpret the findings in Table 3 as evidence that peers likely have some influence on pension beliefs about the risk and return properties of alternatives. These results also support our broader argument that beliefs have played a large role in the rise of alternatives, especially in light of the fixed effects and controls used in Table 3.

4.4 Evidence from Other Institutions

We conclude our analysis by exploring the idea that the perceived alpha has increased for all types of investors, not just public pensions. We do so by documenting aggregate trends in the portfolio composition of DB pensions sponsored by U.S. corporations and unions, U.K. corporations, as well as U.S. endowments. Data for U.S. corporate pensions is based on the corporate pension funding study by Milliman, which contains asset allocation data for the top 100 U.S. corporate DB pensions by assets. These pensions held $1.8 trillion of assets as of 2021. Asset allocation data for U.K. corporate DBs come from the U.K. Pension Protection Authority and U.S. endowment information comes from the National Association of College and University Business Officers (NACUBO). Internet Appendix A contains more background on these sources and their aggregate coverage.

Figure 7a plots the alternative-to-risky share for each institution through time, revealing a clear and common upward trend for all institutions.[19] In contrast, Figure 7b shows that there is no such common trend in the overall risky share. Despite starting at relatively similar levels in the mid-2000s, the risky share for corporate pensions in the U.K. and U.S. has sharply declined, whereas it has slightly increased for U.S. endowments and public pensions. U.K. pensions provide the most striking example of these diverging trends. From 2004 to 2021, their alternative-to-risky share

[19] Gabaix, Koijen, Mainardi, Oh, and Yogo (2022) document a similar rise in alternatives among ultra-rich households.

more than tripled from 11% to 50% while their risky share more than halved from 69% to 31%.

These findings are closely related to the work of Ivashina and Lerner (2018), who document that public and private sector DB pensions have both increased their overall share of alternatives since 2008. We extend their results by showing that this trend cannot be explained by a common desire of all institutions to take on more risk and has instead largely occurred through a change in the *composition* of risky investments. Beliefs about asset returns are a natural explanation for this fact, especially from the perspective of canonical portfolio choice models. For instance, it seems natural to expect that investors have updated their beliefs about the risk-reward properties of alternatives as the asset class has matured. It is also conceivable that the structural forces behind the decline in interest rates have simultaneously impacted the perceived benefits of alternatives relative to public equities. These types of belief-based explanations are appealing because they can account for why institutions that vary widely in governance structure, regulation, funding, and economic function (e.g,. endowments vs. DB pensions) have all shifted the composition of their risky investments – but not their overall risky share – in the same way. We discuss alternative interpretations of the data in Section 6.1, particularly those related to agency frictions.

5 Risk-Seeking Explanations

We now turn to evaluating the second mechanism suggested by the model in Section 3, namely that U.S. public pensions want to take more risk but are constrained from doing so. Our analysis proceeds in three steps. First, we posit a few related ways through which effective pension risk aversion might decline or vary in the cross-section, including those created by pension underfunding. We then connect various proxies for risk-seeking motives to investment in alternatives. Second, we propose a method to measure binding portfolio constraints and test whether it explains cross-pension variation in the alternative-to-risky share. Third, we use the model from Section 3 to simulate a decline in risk aversion plus a binding portfolio constraint and gauge whether it can credibly generate the portfolio shifts that are observed empirically.

5.1 Funding, Hurdle Rates, and Accounting-Based Explanations

The first hypothesis we test stems from the observation that U.S pensions have become increasingly underfunded as interest rates have declined. This fact is apparent in Figure A7 of the internet appendix, which plots the evolution of the national funding ratio when measured according to Statement 25 of the Governmental Accounting Standards Board (GASB). At the beginning of the 2000s, U.S. pensions were more than fully funded by GASB 25 standards, but by 2020 their assets were only 70% of liabilities.[20]

The deterioration in funding may create incentives for U.S. pensions to "reach-for-yield" (Lu et al. (2019)), similar to those facing corporate equity holders when approaching default (Jensen and Meckling (1976)).[21] Motivated by this logic, we proxy for pension risk aversion using the degree of underfunding. We then test whether pensions that have become more underfunded have concurrently shifted into alternatives using the following regression:

$$\Delta\omega^*_{A,target,p} = a + \beta\Delta X_p + \varepsilon_p,$$

where $\Delta\omega^*_{A,target,p}$ is the change in the alternative-to-risky share for pension p between 2002 and 2020. The explanatory variable ΔX_p is the contemporaneous change in each pension's GASB 25 funding ratio, though we also explore other proxies for risk aversion below. One of these proxies (market-based funding ratios) is only available through 2020, hence why we compute changes using that end point. Standard errors are clustered by state for regressions run at the system level and are otherwise robust.

Consistent with popular accounts (Gillers (2021)), Column (1) of Table 4 shows that deterioration in funding is correlated with an increase in the use of alternatives. However, the link between the two is weak in both statistical and economic terms. The estimated point estimate is not statisti-

[20] Under so-called GASB 25 standards, funding ratios are computed from: (i) smoothed asset values that minimize the impact of short-term market fluctuations and (ii) liabilities that equal future benefits discounted using each plan's assumed long-run investment return.

[21] By reach-for-yield, we mean the propensity to take more risk as interest rates fall, though this is one of many ways that the term has been used in the literature. For example, Lian et al. (2018) show that investors may exhibit reach-for-yield behavior for behavioral reasons, as opposed to institutional or agency frictions.

cally significant ($p = 0.29$) and implies that moving from the 90th to 10th percentile of the change in funding is associated with a 3% increase in the alternative-to-risky share. Over this period, the actual spread between the 90th and 10th percentile of the change in the alternative-to-risky share was 39%. Moreover, the R^2 of 1% in the regression indicates that most of the cross-pension variation in the adoption of alternatives cannot be explained by deterioration in funding. We show further in Internet Appendix C.1.1 that the poor model fit is not driven by a non-linear relationship between funding and alternative usage.

One issue with the analysis in column (1) is that funding is measured according to GASB 25 standards. As Brown and Wilcox (2009), Novy-Marx and Rauh (2011), and others have pointed out, GASB 25 funding ratios overstate the true economic funding gap because liabilities are discounted using the expected rate of return on pension assets (asset hurdle rates), rather than more appropriate discount rates based on federal or local government debt. As one way to handle this measurement issue, column (2) replaces GASB 25 funding ratios with those provided by the U.S. Bureau of Economic Analysis (BEA), who uses the yield-curve for AAA-rated corporate bonds to discount liabilities. The regression is run at the state level because pensions' BEA-funding ratios are only available at the state level. The point estimate in this case is effectively zero and the R^2 is once again low, suggesting that the measurement of funding is unlikely to drive its weak relationship with the adoption of alternatives.

A related hypothesis concerning the increased adoption of alternative investments centers on nominal return targets, or hurdle rates. Specifically, if U.S. pensions have fixed or relatively unchanging nominal hurdle rates, then a decrease in interest rates poses a challenge to meeting these targets. Consequently, pensions may seek to take on extra risk in search of higher expected returns. We test this proposition directly in column (3) of Table 4 by regressing changes in the alternative-to-risky share on contemporaneous changes in each pension's hurdle rate. If anything, the regression coefficient suggests a negative, not positive, correlation between alternative adoption and changes in liability discount rates. However, much like funding, the relationship is small in both statistical and economic terms and explains almost none of the variation in changes in the alternative-to-risky share.

Andonov et al. (2017) point out that hurdle rates may be sticky precisely because lowering them causes U.S. pensions to appear more underfunded by GASB 25 accounting. They further argue that this accounting rule creates incentives for pensions with a high fraction of retired members to take more risk. The reason why is that underfunded pensions are typically required to amortize any unfunded accrued liabilities by making additional "catch-up" contributions. The exact size of the required contributions depends on the dollar amount of the funding deficit. Because more mature pensions have larger accrued liabilities, their incentive to invest in risky assets is high because doing so would maintain their liability discount rate, lower their GASB 25 funding gap, and ultimately reduce the size of required contributions. Column (4) builds on Andonov et al. (2017) by testing whether pensions that have had more members retire are also those that have shifted more to alternatives. While there is a positive relationship between the two, the regression coefficient is not statistically different from zero.

To summarize, we find weak empirical support for funding status, hurdle rates, or regulatory incentives as explanations for why some U.S. pensions have increased their alternative-to-risky share more than others since the 2000s.

Robustness. The results in Table 4 cut against the idea that pensions have shifted their risky investments toward alternatives in response to underfunding or for accounting reasons.[22] In Internet Appendix C.1, we reinforce this conclusion in a few complimentary ways. First, we document that the initial levels of the covariates in Table 4 also do not predict subsequent changes in the alternative-to-risky share. For example, it is not the case that pensions who were more underfunded in 2002 shifted more aggressively to alternatives from 2002 to 2020. There is also essentially no relationship between the change in the alternative-to-risky share and pension size or the failure to make actuarial required contributions. Second, we show that the findings in Table 4 are

[22]Campbell and Sigalov (2022) show reach-for-yield behavior can arise if investors are constrained to maintain current wealth in expectation ("sustainable spending"). We view the model of Campbell and Sigalov (2022) as a better description of university endowments rather than U.S. public pensions, mainly because pensions are increasingly facing calls to run down wealth for payments to existing beneficiaries and simultaneously phase out defined benefit plans for active employees (Giesecke and Rauh, 2022).

not driven by our choice of a linear regression specification. Third, the relationship between the alternative-to-risky share and funding status remains weak when employing panel regressions in levels or focusing solely on levels in more recent data. Fourth, after assigning pensions into bins based on their funding and hurdle rates (proxies for risk aversion), we find no distinct trends in the alternative-to-risky share across bins.

The Risky Share. Although not our primary focus, we explore the role of funding and accounting-based mechanisms in explaining the increase of the overall risky share in Internet Appendix C.3. We find only a weak cross-sectional correlation between underfunding and the rise in the risky share. For instance, a cross-sectional regression of changes in the risky share on changes in BEA-adjusted funding ratios yields an R^2 of less than 5%. In line with the findings of Andonov et al. (2017), we do observe a positive relationship between the risky share and changes in the fraction of retirees and liability discount rates, albeit with relatively low R^2 values (less than 10%). These results suggest that funding and accounting-based explanations are unlikely to fully explain why U.S. pensions have taken on more risk in recent decades. 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).

5.2 Binding Portfolio Constraints

Next, we investigate the second necessary condition for risk aversion to impact the alternative-to-risky share in textbook portfolio choice models (Section 3): binding portfolio constraints. The main challenge we face is that the binding constraints are difficult to distinguish from pure risk preferences, as a pension with a low risky share may not be constrained from taking additional risk and could simply have high risk aversion. We overcome this measurement issue by studying deviations $l_{pt} = \text{actual}_{pt}^{risky} - \text{target}_{pt}^{risky}$ of actual from target risky shares.[23] The basic idea is as follows. Market fluctuations will naturally move each pension's actual risky share from its target,

[23] We plot and discuss the time-series properties of l_{pt} in Internet Appendix C.1.4.

after which pensions must rebalance to bring the two in line. For example, in 2022, Calpers – the largest pension in the U.S. – targeted a public equity allocation of 42% and allowed for a range of 7 pp around the target. If a pension wants to take risk but is constrained from doing so (e.g., by statute), it will want to rebalance quickly when its actual risky share falls below target. Conversely, it should be slower to rebalance when its actual risky share is above target because it prefers to take on the extra risk. Under this logic, positive values of l_{pt} are indicative of binding risk constraints. We operationalize this idea using the following cross-sectional regression:

$$\Delta\omega^*_{A,target,p} = a + \theta \bar{l}_p + \varepsilon_p$$

where $\Delta\omega^*_{A,target,p}$ is the change in the alternative-to-risky share for pension p between 2002 and 2020 and $\bar{l}_p$ is the average l_{pt} over the same period. The coefficient of interest in the regression is θ. A large $\bar{l}_p$ is an indication that pension p has consistently faced binding risk constraints, in which case θ should be positive if this constraint induces them to shift more to alternatives.

Figure 9 visualizes the regression using a binned scatter plot of $\Delta\omega^*_{A,target,p}$ and $\bar{l}_p$. The plot provides some indication that risk-constrained pensions have increased their alternative-to-risky share, though the correlation between the two is rather weak. The estimated coefficient in the plot equals 1.33 ($t = 1.64$) and the linear specification yields a fairly low R^2 of 4%. The regression coefficient indicates that a one-standard deviation increase in the average constraint is associated with a 3 pp increase in the alternative-to-risky share. The standard deviation of the change in the alternative-to-risky share for this sample equals 15 pp.

One concern with this analysis is that deviations of target from actual risky shares reflect both market fluctuations and any potential portfolio constraints. To isolate the latter, we first regress l_{pt} on each year's portfolio return then compute $\bar{l}_p$ based on the regression residuals. The regression coefficient in this case is basically unchanged at 1.35 ($t = 1.63$). Overall, these results indicate that portfolio constraints—to the extent that we can measure them—are unlikely to explain why some pensions have shifted more aggressively to alternatives than others.

5.3 Simulation Evidence

One potential issue with our analysis thus far is the measurement of risk aversion and portfolio constraints. For example, funding may not perfectly capture a pension's effective risk aversion because the link between legally mandated contributions and the funding gap may vary at the state or even local level. Such complexities would introduce measurement error into our proxies for risk aversion or binding portfolio constraints, thereby attenuating the relationship with pension investment behavior.

We address this concern via a simulation exercise. Using the portfolio choice model from Section 3, we simulate how a decline in risk aversion coupled with a binding portfolio constraint would impact a pension's portfolio composition. We do so for a wide range of beliefs that are held fixed throughout each simulation. The goal of the exercise is therefore to understand how much risk-seeking motives and binding portfolio constraints can explain the change in pension behavior between 2001 and 2021, holding beliefs fixed.

We first assume that the aggregate pension portfolio was unconstrained in 2001 but that a minimum constraint on fixed income becomes binding by 2021. This means that the observed fixed income share in 2021 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 on alternatives relative to equities, $\alpha \sim U(0, 0.05)$; (iv) the CAPM-beta of alternatives relative to public equities, $\beta \sim U(0, 1.5)$. Realized excess returns on alternatives are therefore given by:

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

where σ_η^2 is the idiosyncratic variance of alternatives. Together, these belief parameters fully define the expected excess return and variance-covariance matrix that determine optimal asset allocation. For each parameter draw, we set σ_η^2 to match the aggregate alternative-to-risky share ω_A^* in 2001 according to Equations (2) and (3). Finally, we select γ_1 as the risk aversion that would also match the overall risky share in 2001, again assuming that pensions are not constrained. These last two steps ensure that each random belief set is consistent with the 2001 aggregate pension portfolio. We restrict the set of beliefs S^* to imply positive idiosyncratic variance $\sigma_\eta^2 > 0$, $\gamma_1 \geq 1$, and $\sigma_A^2 \leq 0.25$. We then draw random belief sets until we reach $S^* = 100,000$ that satisfy these conditions.[24] Given this set of admissible initial beliefs, we then assume the portfolio becomes constrained in 2021 and solve for the new risk aversion γ_2 needed to generate the observed $\Delta\omega_A^*$ in the data. Panels A and B in Table 5 summarize the simulation approach. In all cases, we impose a constraint that alternatives cannot be shorted, meaning $\omega_A \geq 0$.

Panel C of Table 5 contains a breakdown of our simulation outcomes. The headline result is that in roughly 99.6% of simulations ($99,581$ out of 100,000), there is *no* change in risk aversion that is able to match the observed increase in the alternative-to-risky share. The intuition for why is simple and essentially follows from revealed preference: in 2001, public equities were a large portion of the risky portfolio. Consequently, for the vast majority of beliefs that match this initial composition ($99,080$ out of 100,000), the investor would want to shift towards public equities and *not* alternatives when she becomes constrained. In other words, $\Delta\omega_A^*$ is counterfactually negative for these simulations. There is a much smaller subset of simulations (501 out of 100,000) for which a decline in risk aversion would 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. For this small subset of cases, the median alternative-to-risky share increases to a level of 28%, compared to 39% in the data (see Figure A18 in the Internet Appendix).

Panel C of Table 5 shows that only 0.4% of simulations (419 out of 100,000) can generate the observed $\Delta\omega_A$ with a decline in risk aversion. Figure 10a indicates that the implied average reduction in γ is about 4.51 for these cases. For this subset of simulations, we calculate the shadow cost of the binding constraint on fixed income. To do so, define $M = \mathbb{E}[r_p] + \frac{1}{2}(1-\gamma)\mathbb{V}[r_p]$ as the

[24] See Figure A17 in the Internet Appendix for the histograms of admissible beliefs.

function that investors maximize. The shadow cost of the portfolio constraint is then given by:

$$ShadowCost = M_{counterfactual} - M_{portfolio}, \tag{9}$$

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 portfolio constraint. The shadow cost can therefore be interpreted as the fee that pensions would be willing to pay to relax the constraint and invest in their perceived tangency portfolio. The unit of this shadow cost is the same as the log return. Figure 10b 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.

We want to stress that the simulation exercise above does not rule out explanations based on risk-seeking (e.g., reach-for-yield) and binding portfolio constraints. Rather, the model highlights that neither of these channels can explain the rise of alternatives without a concurrent shift in beliefs. The reduced-form evidence in Sections 5.1 and 5.2 further supports this interpretation, as does the positive evidence we provided in Section 4.

6 Discussion and Conclusion

Our empirical findings collectively suggest that the rise of alternatives has been fueled by a shift in beliefs about their alpha relative to public equities. In contrast, there is weak and inconsistent support for explanations based on risk-seeking motives. Before concluding, we now discuss other potential explanations based on agency frictions and a shift in the supply of alternatives.

6.1 Agency-Based Explanations

While our findings cannot rule out the possibility that agency frictions have driven the rise of alternatives, they do place restrictions on the specific nature of any such friction. First, the friction must operate across countries and institution types, and vary in the cross-section of pensions, but in a way that is orthogonal to funding, size, membership age, and other pension attributes (Sections 4.4 and 5.1). Second, it must cause many institutions to substitute between alternatives and public equities without necessarily increasing their overall risky shares (Section 4.4). Third, it must lead a subset of private and public sector pensions to choose consultants who report a high alpha of alternatives (Section 4.1.3). The expressed beliefs of consultants may themselves be the result of agency frictions, though these frictions would also need to vary across consultants.

What is an example of such a friction? One possibility is that some investors may prefer the quiet life that unmarked assets like private equity provide over public markets (Bertrand and Mullainathan (2003); Stafford (2022)). For this reason, our analyses in Section 4.1.1 also includes a test of whether consultant beliefs impact the choice between private equity and real assets. Because both asset classes offer a similar ability to hide risk, this test holds fixed any desire to invest in unmarked assets and more cleanly isolates the link between beliefs and portfolio choice. Consistent with a belief-driven explanation, we find that pensions invest more in private equity over real assets if their consultant reports it to have a higher alpha, after absorbing any unobserved factors that may lead pensions to prefer alternatives overall. With that said, agency frictions of this kind are certainly plausible and further research is needed to investigate how large a role they have played in the rise of alternatives.

6.2 Supply-Side Explanations

Thus far, we have exclusively considered demand-based explanations for the rise of alternatives, though supply-side factors are also potentially important. For example, the development of the private equity industry over the last thirty years has arguably made it easier for institutional investors to take equity stakes in firms that are not publicly listed. This implies that the portfolio trends observed in U.S. public pensions might be a passive reflection of this technological change. Figure 11 provides one way to assess this mechanism by showing how the global supply of alternative assets has evolved relative to all risky assets since 2000. The supply of alternatives is defined as the net asset value of all private-market funds based on data from Preqin and the global AUM of hedge funds from the Hedge Fund Research database. Risky assets equal the supply of alternatives plus the worldwide market capitalization of all publicly traded firms according to the World Bank. The plot shows that the global alternative-to-risky share rose from 2% in 2000 to just over 8% in 2020. Even though this increase is consistent with supply playing a role in the rise of alternatives, it is important to note that supply-side explanations cannot account for the wide cross-sectional heterogeneity in the adoption of alternatives across U.S. public pensions (Section 2.2.3).[25]

6.3 Conclusion

Since the early 2000s, there has been a notable shift in the composition of risky investments for public pensions in the United States, with a significant emphasis on alternative investments like private equity, real estate, and hedge funds. The conventional explanations for this trend, including underfunding, portfolio constraints, and the need to meet nominal return targets, have limited empirical support. Instead, we propose a new perspective rooted in beliefs: U.S. pensions increasingly perceive alternative investments to provide a more favorable risk-return profile than public equities. This belief-based perspective helps rationalize the long-run increase in the alternative-to-risky share among U.S. public pensions and the variation observed across different pensions. While our study provides some insights into the drivers of beliefs, more research is needed to fully understand the process by which pensions form beliefs. This question is critical for assessing the welfare implications of alternatives for pension beneficiaries, especially given the costs and complexity of investing in this asset class (Metrick and Yasuda (2010); Phalippou, Rauch, and Umber (2018); Begenau and Siriwardane (2022)).

[25] In addition, Figure 11 indicates that the rotation by U.S. pension towards alternatives has far outpaced global supply, resulting in a portfolio that heavily overweights alternatives relative to public markets.

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., M. Bonetti, and I. Stefanescu (2023). Choosing pension fund investment consultants. *SSRN Working paper*.

Andonov, A. and J. D. Rauh (2021, 11). The Return Expectations of Public Pension Funds. *The Review of Financial Studies 35*(8), 3777–3822.

Angrist, J. D. (2014). The perils of peer effects. *Labour Economics 30*, 98–108.

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

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

Becker, B. and V. Ivashina (2015). Reaching for yield in the bond market. *The Journal of Finance 70*(5), 1863–1902.

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

Bertrand, M. and S. Mullainathan (2003). Enjoying the quiet life? corporate governance and managerial preferences. *Journal of Political Economy 111*(5), 1043–1075.

Bertrand, M. and A. Schoar (2003, 11). Managing with Style: The Effect of Managers on Firm Policies*. *The Quarterly Journal of Economics 118*(4), 1169–1208.

Betermier, S., N. Byrne, J.-S. Fontaine, H. Ford, J. Ho, and C. Mitchell (2021, August). Reaching for yield or resiliency? Explaining the shift in Canadian pension plan portfolios. Staff Analytical Notes 2021-20, Bank of Canada.

Beutel, J. and M. Weber (2022). Beliefs and portfolios: Causal evidence. *Chicago Booth Research Paper* (22-08).

Bordalo, P., N. Gennaioli, and A. Shleifer (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives 36*(3), 223–244.

Borio, C. and H. Zhu (2012). Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism? *Journal of Financial stability 8*(4), 236–251.

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.

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

Campbell, J. Y. and L. M. Viceira (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic.

Casella, G. (1992). Illustrating empirical bayes methods. *Chemometrics and intelligent laboratory systems 16*(2), 107–125.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2023). On binscatter.

Couts, S. J., A. S. Goncalves, and J. Loudis (2023). The subjective risk and return expectations of institutional investors. *SSRN Working paper*.

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

Gabaix, X., R. S. Koijen, F. Mainardi, S. Oh, and M. Yogo (2022). Asset demand of us households. *Available at SSRN 4251972*.

Giesecke, O. and J. Rauh (2022). Trends in state and local pension funds. *Annual Review of Financial Economics 15*.

Giglio, S., M. Maggiori, J. Stroebel, Z. Tan, S. Utkus, and X. Xu (2023). Four facts about esg beliefs and investor portfolios. Technical report, National Bureau of Economic Research.

Gillers, H. (2021). An ohio pension manager risks running out of retirement money. his answer: Take more risks. https://tinyurl.com/mws6wkpj.

Gupta, A. and S. Van Nieuwerburgh (2019). Valuing private equity investments strip by strip. *Working Paper*.

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

Ivashina, V. and J. Lerner (2018). Looking for alternatives: Pension investments around the world, 2008 to 2017. Working Paper.

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.

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

Korteweg, A. (2019). Risk adjustment in private equity returns. *Annual Review of Financial Economics 11*(1), 131–152.

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

Korteweg, A., S. Panageas, and A. Systla (2023). Private equity for pension plans? evaluating private equity performance from an investor's perspective. Working Paper, University of Southern California.

Leombroni, M., M. Piazzesi, C. Rogers, and M. Schneider (2020). Inflation and the price of real assets.

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

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

Lu, Y., K. Mullally, and S. Ray (2023). Paying for performance in public pension plans. *Management Science 69*(8), 4888–4907.

Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *The Quarterly Journal of Economics 131*(1), 53–87.

Markowitz, H. (1952). Portfolio analysis. *Journal of Finance 8*, 77–91.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics 51*(3), 247–257.

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

Metrick, A. and A. Yasuda (2010). The economics of private equity funds. *The Review of Financial Studies 23*(6), 2303–2341.

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

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.

Pástor, L. (2000). Portfolio selection and asset pricing models. *The Journal of Finance 55*(1), 179–223.

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

Phalippou, L. and O. Gottschalg (2008). The performance of private equity funds. *The Review of Financial Studies 22*(4), 1747–1776.

Phalippou, L., C. Rauch, and M. Umber (2018). Private equity portfolio company fees. *Journal of Financial Economics 129*(3), 559–585.

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

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

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

Figure 1: The Risky Share

(a) Aggregate, Post-2000

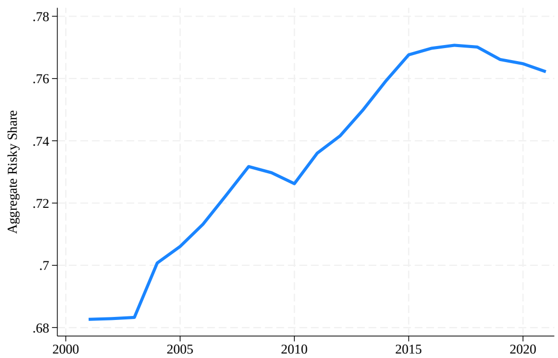

(b) Aggregate, Post-1968

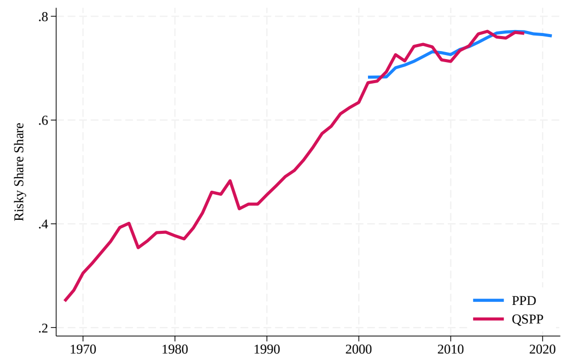

(c) Cross-Sectional Distribution

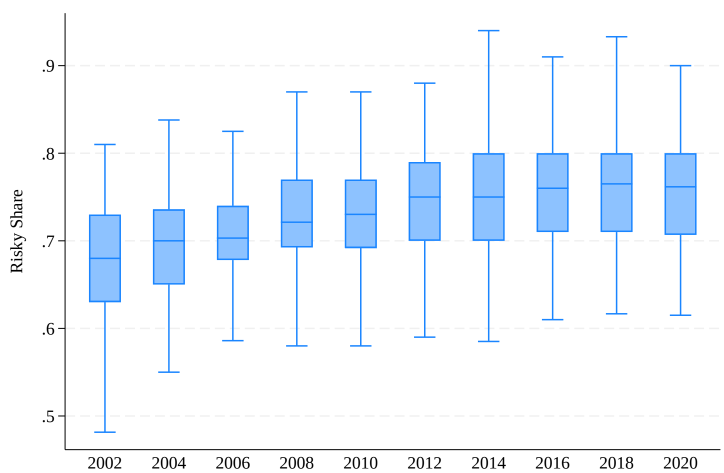

Notes: Panel (a) plots the aggregate target risky share for U.S. public pensions based on data from the PPD, where the risky share is defined as any holding outside of fixed income and cash. Panel (b) adds a longer-history of the risky share using data from the U.S. Census Bureau's Quarterly Survey of Public Pensions (QSPP). The risky share in the QSPP similarly excludes fixed income and cash but is instead based on actual, not target weights. Panel (c) visualizes the cross-section of the risky share using the PPD data. See Section 2 for complete details.

Figure 2: The Composition of the Aggregate Risky Portfolio

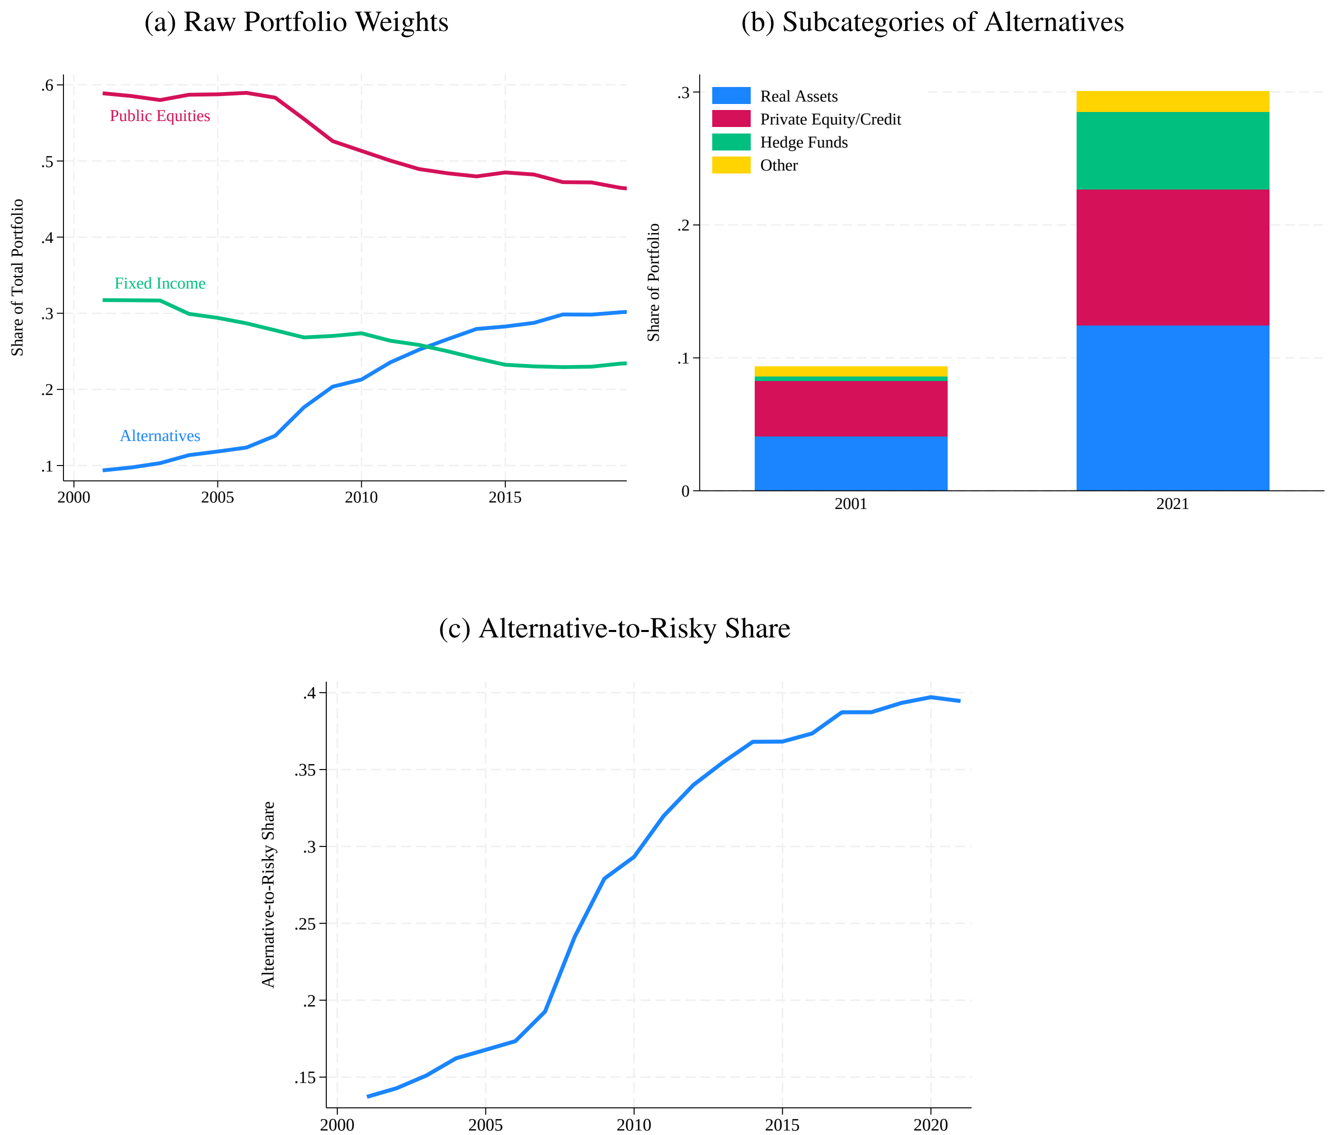

Notes: Panel (a) of the figure plots the target share of public equities, alternatives, and fixed income (including cash) for the aggregate U.S. public pension portfolio. Panel (b) shows the aggregate portfolio shares of different categories of alternatives for 2001 and 2021. Panel (c) plots the share of alternatives in the risky portfolio. All data are based on the PPD. Risky investments are defined as any holding outside of fixed income and cash.

Figure 3: Alternative-to-Risky Share in the Cross-Section of Pensions

(a) Distribution of the level through time

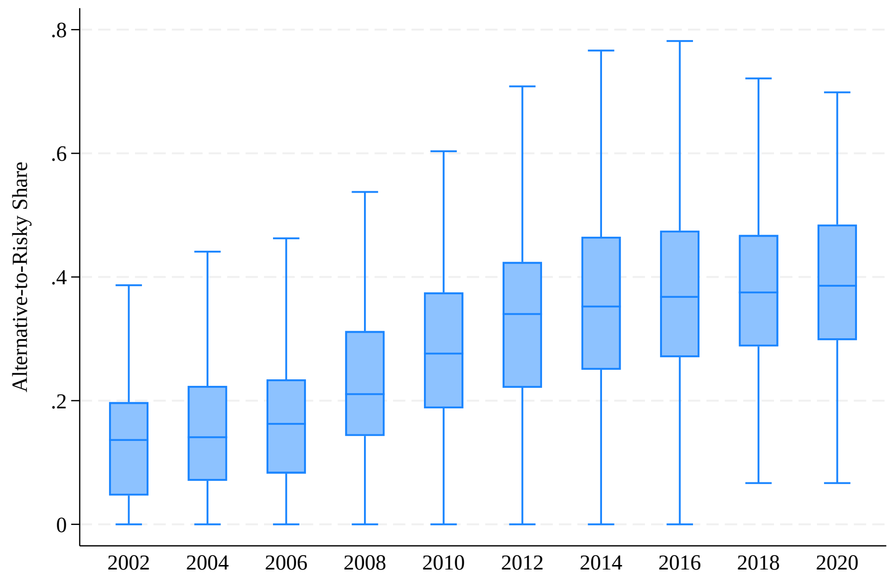

(b) Change from 2002 to 2021

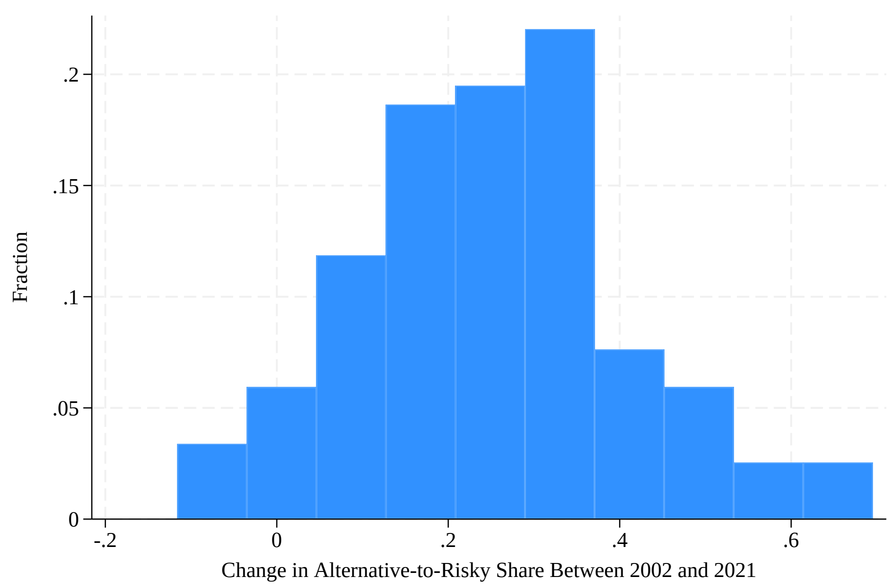

Notes: Panel (a) of the figure depicts the distribution of the alternative-to-risky share across pension systems through time. Each box plot summarizes the distribution for the corresponding year on the *x*-axis. Only even years are plotted to make the graph more readable. Panel (b) plots the distribution of the change in alternative-to-risky share across U.S. pension systems between 2002 and 2021. All data are from the PPD.

Figure 4: Consultant-Reported Beliefs about the Alpha of Alternatives

(a) Median Alpha Over Time

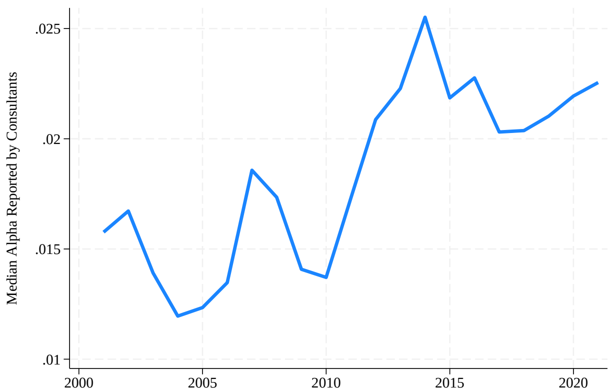

(b) Clients' Alternative-to-Risky Share

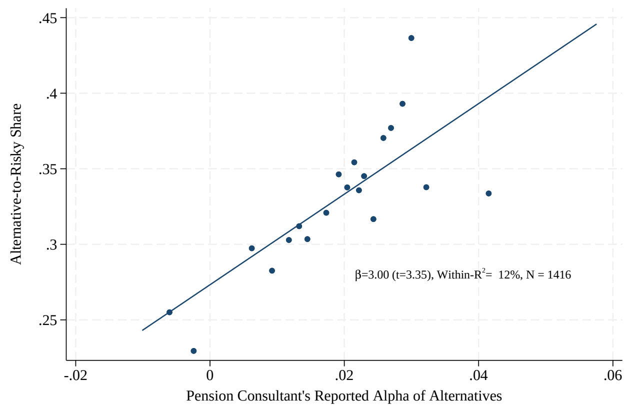

(c) Clients' Use of Different Alternatives

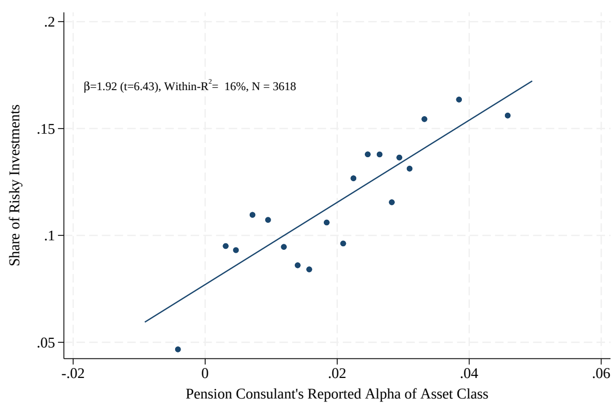

Notes: Each year in Panel (a) shows the median perceived alpha of alternatives relative to large-cap US equities across consultants. The data for the plot are based on hand-collected capital market assumption (CMA) reports that were obtained under confidentiality agreements. Panel (b) is a binscatter of each pension's alternative-to-risky share against its consultant's reported alpha, after controlling for state-by-year fixed effects, funding status, (log) size, and hurdle rates. Panel (c) is 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 in panels (b) and (c) are based on standard errors that are clustered by consultant and pension. See Section 4.1.1 for details.

Figure 5: Consultant Fixed Effects

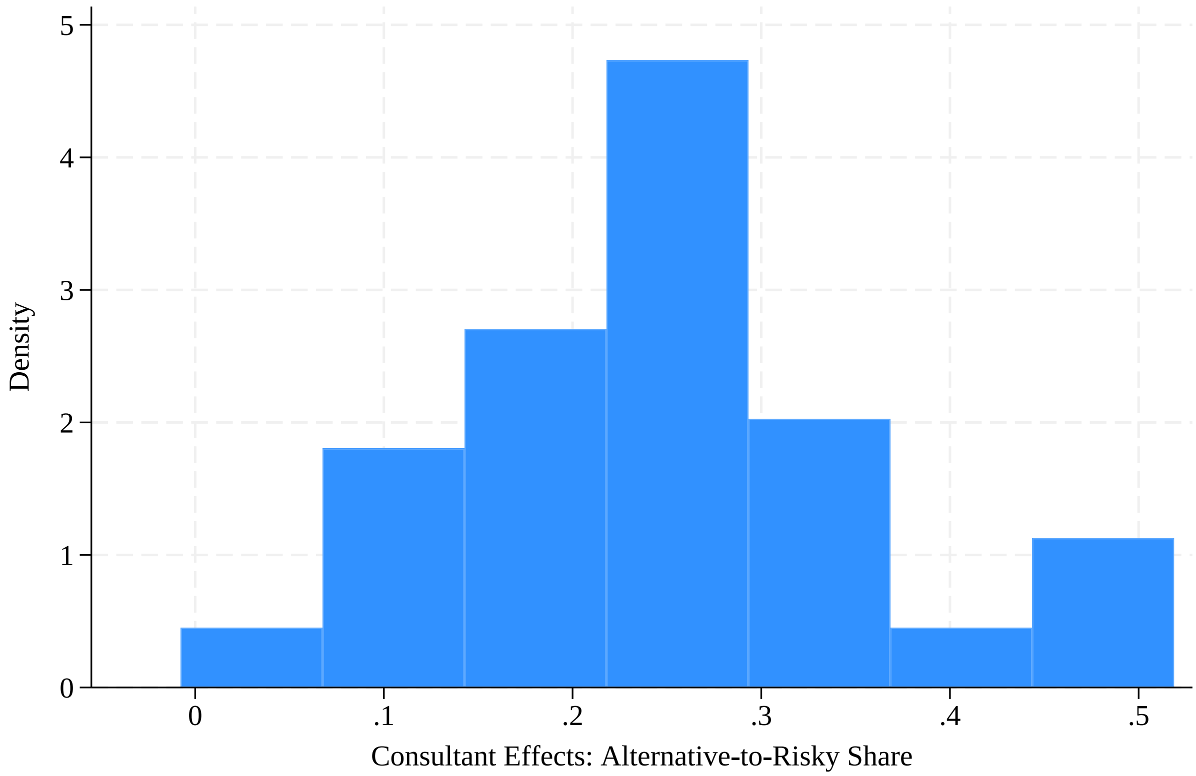

Notes: This figure shows the distribution of consultant effects λ_c estimated from the following regression:

$$\omega^*_{A,p(c),t} = \lambda_t + \sum_k \beta^k_t \times X^k_{pt} + \lambda_c + \varepsilon_{pct}$$

where $\omega^*_{A,p(c),t}$ is the alternative-to-risky share for pension p matched with consultant c in fiscal year t. λ_t is a time fixed effect and X^k_{pt} is the set of covariates including (log) size, GASB 25 funding, asset hurdle rates, the ratio of required actuarial contributions to payroll, and the ratio of administrative expenses to payroll. The estimated consultant fixed effects λ_c are then shrunk towards their mean using an empirical Bayes estimate (Casella, 1992) to account for sampling error. The average alternative-to-risky share is added back to the fixed effects to facilitate interpretation.

Figure 6: U.S. Public vs Private-Sector Clients by Consultant

(a) Alternative-to-Risky Share

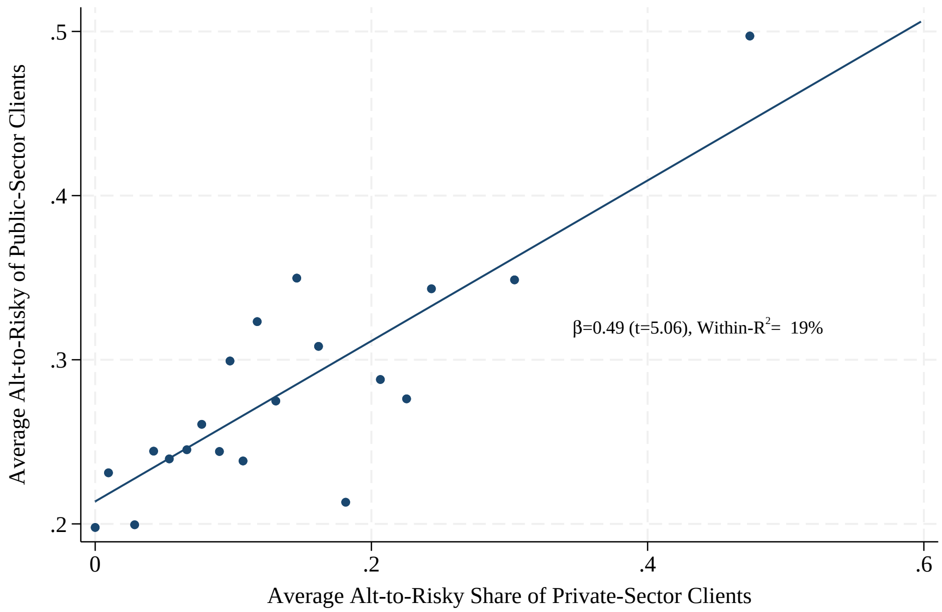

(b) Risky Share

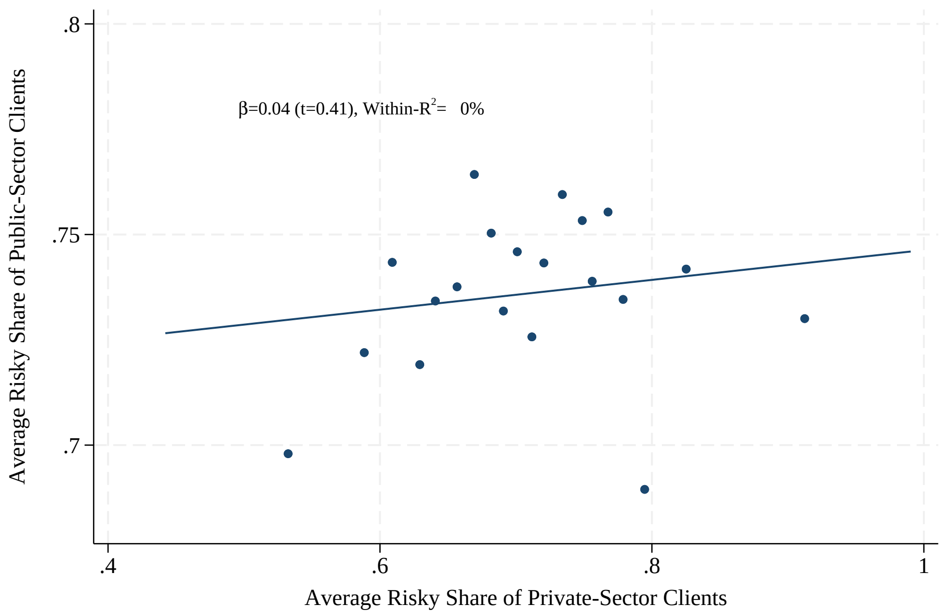

Notes: For each consultant c in year t, we compute the average alternative-to-risky and risky share of U.S. public pension clients, denoted by $\bar{u}^a_{c,t}$ and $\bar{u}^r_{c,t}$, respectively. In addition, we compute the average alternative-to-risky and risky share of U.S. private-sector clients, 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 7: Portfolio Behavior Across Institutions

(a) Alternative-to-Risky Share

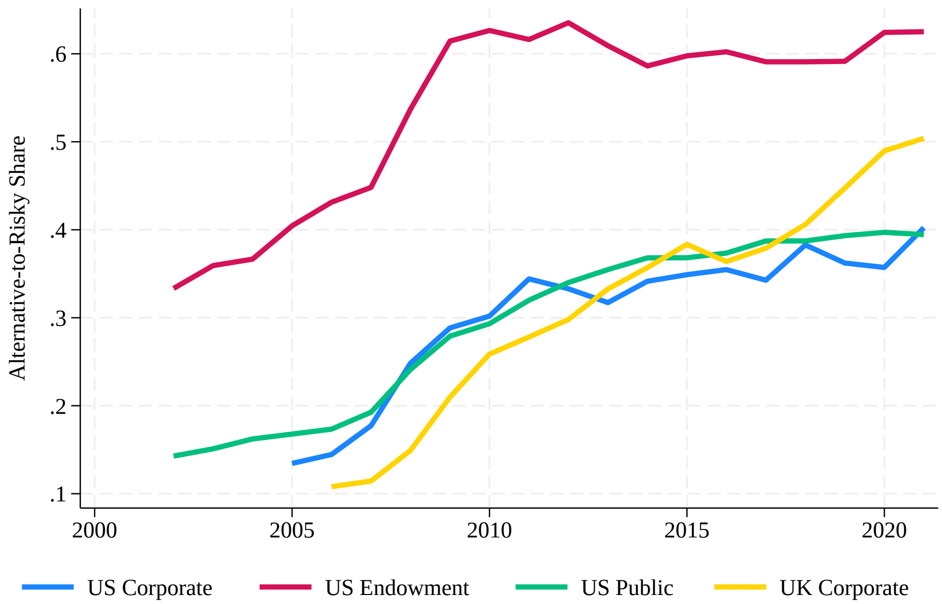

(b) Risky Share

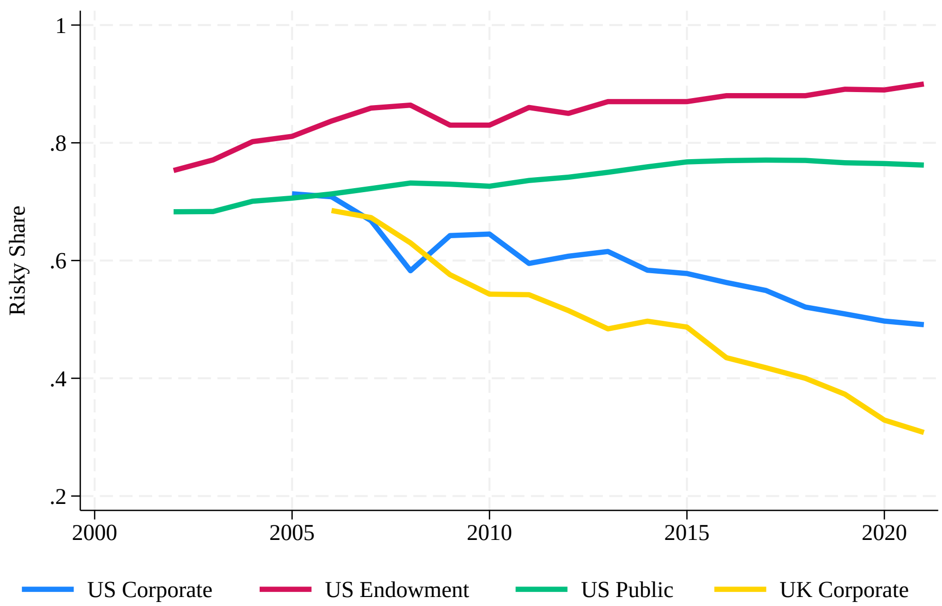

Notes: Panels (a) and (b) of the panel plot the alternative-to-risky and risky shares, respectively, of different institutions through time. U.S. Endowment data is from NACUBO historical endowment studies, U.K. corporate defined-benefit pension data is from the U.K. Pension Protection Authority, and U.S. corporate pension data is from the Milliman Corporate 100 report. See Section 4.4 and Section A of the Internet Appendix for more details on the data.

Figure 8: Experience in the 1990s and the Rise of Alternatives

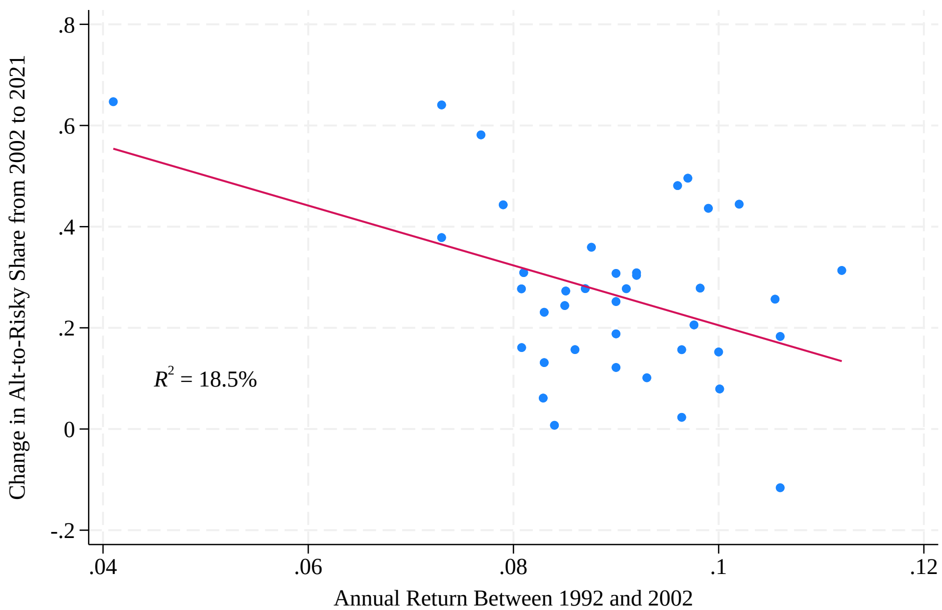

Notes: This figure plots each pension's change in the alternative-to-risky share between 2002 and 2021 against its geometric average return between 1992 and 2002. The sample size is lower for this plot because not all pensions report 10-year returns on their annual reports and we do not have a complete panel of annual reports from 1992 to 2002. Data are from the PPD.

Figure 9: Alternative Adoption and Portfolio Constraints

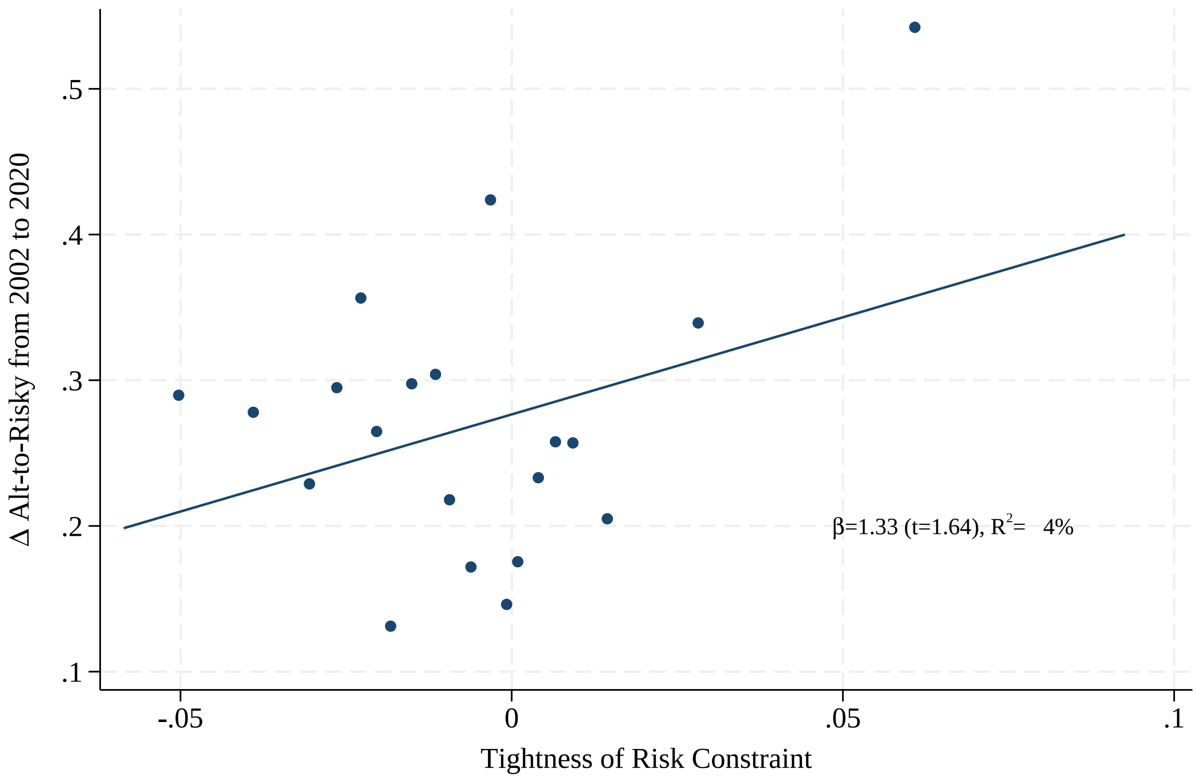

Notes: This figure shows a binscatter of the change in the alternative-to-risky share from 2002 to 2020 against the average tightness of portfolio constraints over the same period. Constraint slack in a given year t for a pension p is defined as the difference l_{pt} between its actual and its target risky share. Higher values of $\bar{l}_p$ indicate a tighter constraint. The plot reflects only the sample of pensions for which both actual and target shares are available. The OLS regression line is also shown in the plot, with standard errors clustered by state. Data are at the pension-system level. See Section 5.2 for details.

Figure 10: Simulating a Decline in Risk Aversion Plus Binding Portfolio Constraints

(a) Implied Change in Risk Aversion

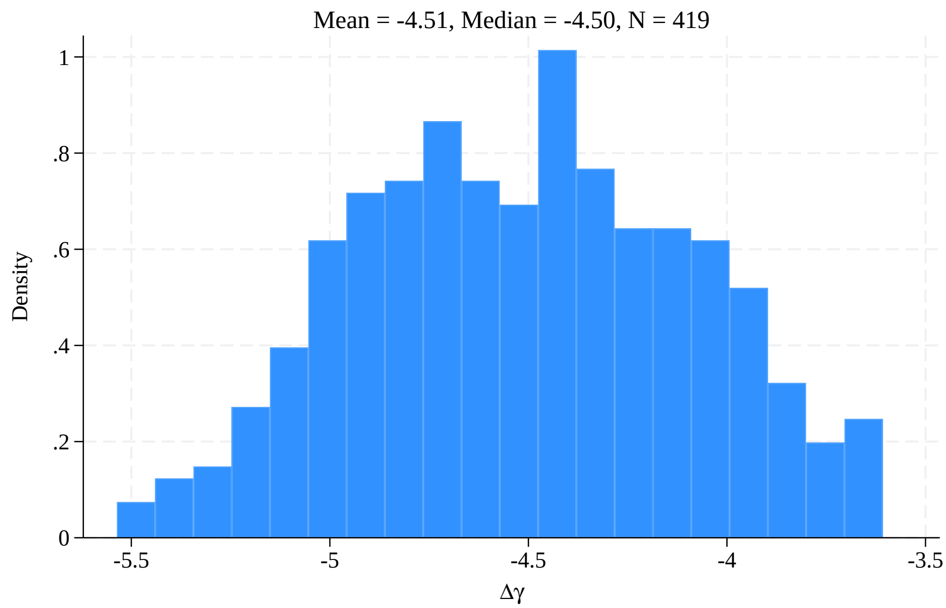

(b) Implied Shadow Cost

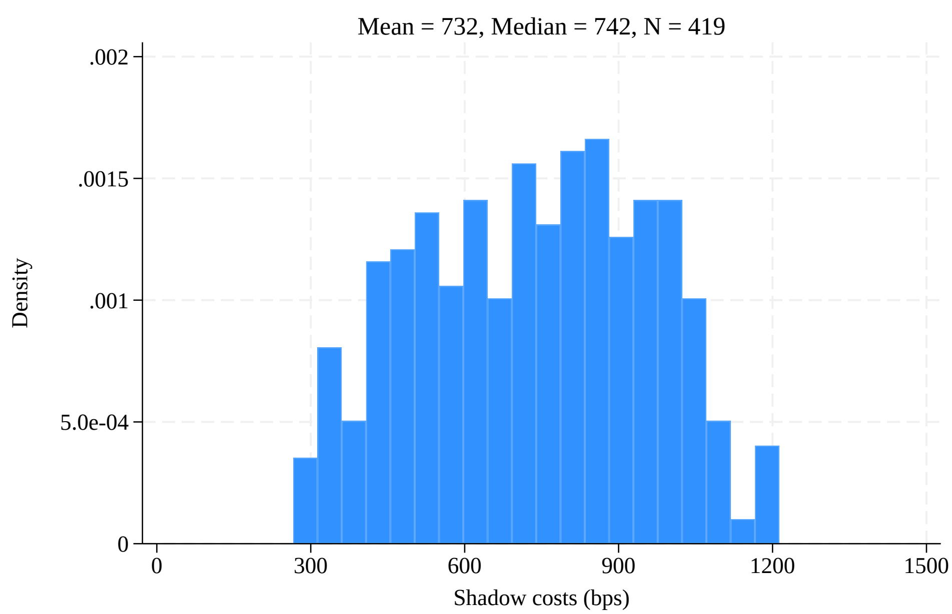

Notes: This figure shows simulation results based on the model in Section 3, in which pensions face a minimum constraint on fixed income investment. Panel (a) of the figure shows how the change in risk aversion required to match portfolio weights in 2001 and 2021 varies across simulations of initial 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 5.3 for details.

Figure 11: The Supply of Alternatives

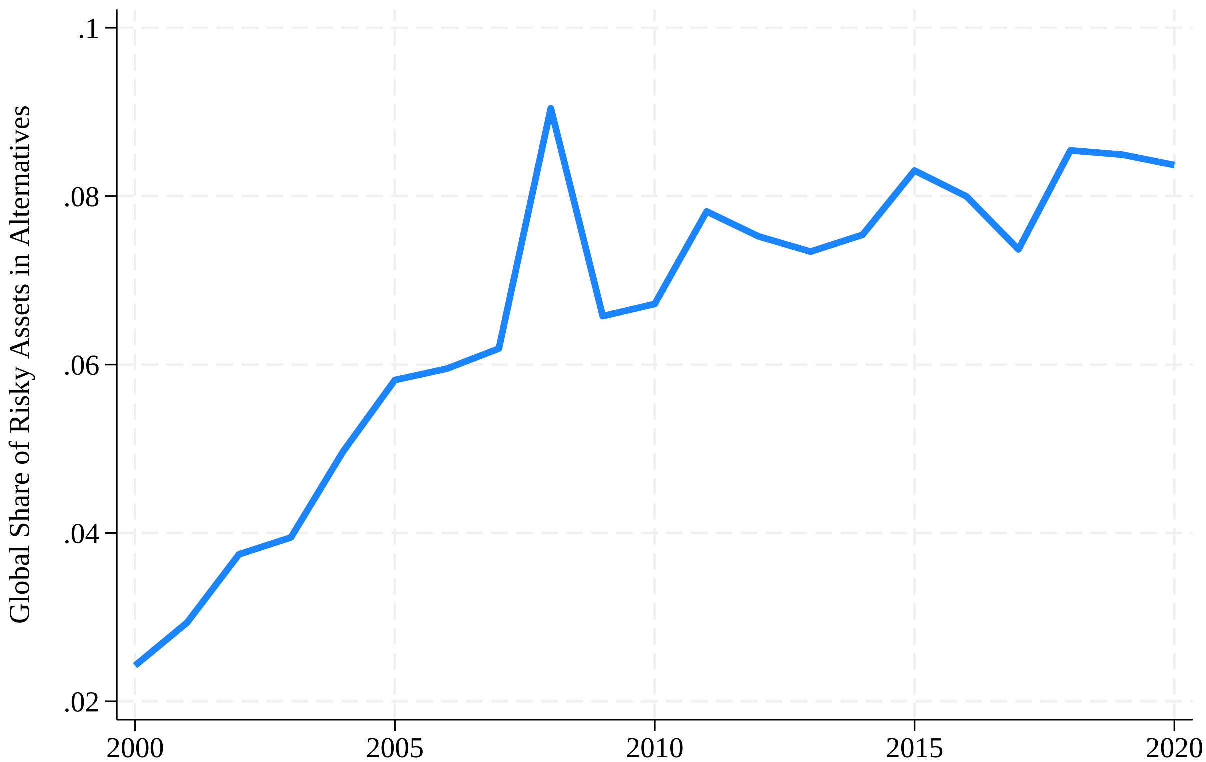

Notes: This figure shows the global supply of alternatives relative to all risky assets. The global stock of alternatives equals the total net asset value of all private-capital funds from Preqin plus the AUM of global hedge funds from Hedge Fund Research. The total stock of risky assets equals the stock of alternatives plus the market capitalization of all global public stock markets from the World Bank.

Table 1: National Summary Statistics

	Subsample			
	2001-2005	2006-2010	2011-2015	2016-2021
Number of Systems	165	190	202	208
Members (mm)	21	24	25	27
Percent Retired	28	31	35	37
AUM ($ bn)	2,120	2,633	3,154	4,065
GASB 25 Funding (%)	91	81	73	73
Assumed Asset Return (%)	8.0	7.9	7.6	7.2
Annual Investment Return (%)	5.3	6.2	9.1	10.0
National Coverage (%)				
Public DB Pensions	87	90	92	91
All Private and Public Pensions	25	25	24	22
Portfolio Composition (%)				
Fixed Income	31	27	25	23
Public Equities	59	54	49	47
Alternatives	11	18	27	30

Notes: This table provides national-level summary statistics for defined benefit public pensions in the United States from 2001 to 2021. Individual pension plans are aggregated to pension systems when the assets of multiple plans are legally pooled and managed together. GASB 25 funding is defined as the ratio of actuarial assets to liabilities, where liabilities are computed by discounting future promised benefits using each plan's assumed long-run rate of return on assets. The rows listed below *National Coverage* report the average annual ratio of assets in PPD to defined benefit pension assets listed in the U.S. Census Bureau's Annual Survey of Public Pensions and total public and private-sector pension assets listed in the Flow of Funds. The rows listed below *Portfolio Composition* show the percent of the aggregate public pension portfolio that is invested in fixed income (including cash), public equities, and alternatives, respectively. Alternatives encompass investments in hedge funds, private equity, private debt, and real assets (e.g., real estate private equity). Data are based primarily on the Public Plans Database (PPD) that is maintained by the Center for Retirement Research at Boston College.

Table 2: Consultants and the Composition of Risky Investments

	Dep. Variable:	Controls	Fixed Effects: Time	State-Time	Consultant	Pension	F	p	Adj. R^2	C	N
(1)	Alternatives		x						0.31		3,121
(2)	Alternatives	x	x						0.32		3,087
(3)	Alternatives	x	x		x		17.57	0.00	0.49	57	3,082
(4)	Alternatives	x		x	x	x	7.94	0.00	0.82	48	2,682
(5)	Private Equity/Credit		x						0.10		3,121
(6)	Private Equity/Credit	x	x						0.19		3,087
(7)	Private Equity/Credit	x	x		x		15.94	0.00	0.37	57	3,082
(8)	Private Equity/Credit	x		x	x	x	5.69	0.00	0.75	48	2,682
(9)	Hedge Funds		x						0.14		3,121
(10)	Hedge Funds	x	x						0.13		3,087
(11)	Hedge Funds	x	x		x		11.30	0.00	0.28	57	3,082
(12)	Hedge Funds	x		x	x	x	4.72	0.00	0.65	48	2,682
(13)	Real Assets		x						0.13		3,121
(14)	Real Assets	x	x						0.13		3,087
(15)	Real Assets	x	x		x		16.45	0.00	0.33	57	3,082
(16)	Real Assets	x		x	x	x	5.07	0.00	0.70	48	2,682

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

$$\omega^*_{A,p(c),t} = \lambda_t + \sum_k \beta^k_t \times F^k_{pt} + \lambda_c + \varepsilon_{pct}$$

where $\omega^*_{A,p(c),t}$ is one of several target asset allocations for pension system p matched with consultant c in fiscal year t. λ_t is a year fixed effect, F^k_{pt} is characteristic k of pension p in year t, and λ_c is a consultant fixed effect. Control variables include (log) size, GASB 25 funding, asset hurdle rates, the ratio of required actuarial contributions to payroll, and the ratio of administrative expenses to payroll. Some regressions also include a pension fixed effect. The listed F-statistic is the result of testing the null hypothesis that the consultant effects λ_c are equal to each other. C is the number of consultant effects that are included in the test. In rows (4), (8), (12), and (16), the time fixed effect λ_t is replaced with a state-by-time fixed effect. Sample sizes differ across models because we drop singleton groups for any included fixed effects. All allocations are relative to the overall share of risky investments.

Table 3: Peer Effects

	Alternative-to-Risky Share				
	(1)	(2)	(3)	(4)	(5)
Peers' Alt-to-Risky Share	0.69**	0.61**	0.70**	0.73**	
	(4.68)	(3.96)	(4.77)	(5.02)	
× Established-CIO		0.01			
		(0.04)			
× Well-Funded			-0.14		
			(-1.14)		
× High-Performing				-0.26**	
				(-2.34)	
Lagged Peers' Alt-to-Risky Share					0.71**
					(4.60)
Consultant× Year × Division FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.14	0.14	0.15	0.16	0.15
Total R^2	0.68	0.59	0.68	0.69	0.67
N	2,039	900	2,039	2,039	1,915

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

Table 4: Funding-Based Explanations for the Rise in Alternatives

	ΔAlternative-to-Risky Share			
	(1)	(2)	(3)	(4)
Δ GASB 25 Funding Ratio	-0.09			
	(-1.08)			
Δ BEA-Adjusted Funding Ratio		-0.00		
		(-0.01)		
Δ Liability Discount Rate			-1.55	
			(-0.49)	
Δ Fraction of Retired Members				0.15
				(0.80)
Aggregation	System	State	System	System
Total R^2	0.01	0.00	0.00	0.00
N	119	47	118	119

Notes: This table shows regressions of changes in the alternative-to-risky share on contemporaneous changes in several covariates. All changes are computed between 2002 and 2020. Risky investments are any holdings outside of fixed income and cash. Alternatives encompass investments in hedge funds, private equity, private debt, and real assets. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate borrowing rates. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled "Aggregation" specifies whether the regression is run at the system or state level. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. See Section 2.1 for how we filter the data.

Table 5: Simulation Exercise: Parametrization

Panel A: Randomly drawn parameters

Parameter	Definition	Sampling Distribution
μ_E	Expected excess return of public equities	$U[0.02, 0.07]$
σ_E^2	Variance of public equities	$U[0.02, 0.09]$
α	Alpha of alternatives	$U[0, 0.05]$
β	Beta of alternatives	$U[0, 1.5]$

Panel B: Inferred Parameters

Parameter	Definition	Method
μ_A	Expected excess return of alternatives	$\mu_A = \alpha + \beta\mu_E$
σ_A^2	Variance of alternatives	$\sigma_A^2 = \beta^2\sigma_E^2 + \sigma_\eta^2$
σ_{AE}	Covariance of public equities and alternatives	$\sigma_{AE} = \beta\sigma_E^2$
σ_η^2	Idiosyncratic variance of alternatives	Chosen to match initial risky portfolio composition (Eq. 2 to 3)
ω_f^{min}	Minimum constraint on riskless asset	Based on observed fixed income share in 2021

Panel C: Sample Size	Sample size
Admissible beliefs S^*:	$100,000$
Risk aversion able to explain $\Delta\omega_A^*$:	419
Risk aversion unable to explain $\Delta\omega_A^*$:	$99,581$
Pensions shift to public equities when constrained	$99,080$
Pensions shift to alternatives when constrained, but γ_2 binds at 1	501

Panel A of this table summarizes the distributions from which beliefs are drawn in Section 5.3, where we simulate a decline in risk aversion paired with a binding portfolio constraint. Panel B summarizes how the remaining parameters are backed out from the model in Section 3. Panel C provides a breakdown of when simulations are able to generate the observed aggregate portfolio shifts in the alternative-to-risky share, $\Delta\omega_A^*$, from 2001 to 2021 with reasonable declines in risk aversion. See Section 5.3 for more details.