Working Paper 25-016

The Rise of Alternatives

Juliane Begenau
Pauline Liang
Emil Siriwardane

Harvard Business School

The Rise of Alternatives

Juliane Begenau
Stanford GSB

Pauline Liang
Stanford GSB

Emil Siriwardane
Harvard Business School

Working Paper 25-016

Funding for this research was provided in part by Harvard Business School.

The Rise of Alternatives*

Juliane Begenau† Pauline Liang‡ Emil Siriwardane§

August 2024

Abstract

Since the 2000s, U.S. public pensions have shifted their risky investments towards alternative assets like private equity and hedge funds, some more aggressively than others. We explore several explanations for these cross-sectional trends, focusing on those implied by the mean-variance models used by most pensions. 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, have weaker empirical support.

*We are grateful to Aleksandar Andonov, Keith Brainard, 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

Juliane Begenau I have nothing to disclose.

Pauline Liang I have nothing to disclose.

Emil Siriwardane I have nothing to disclose.

1 Introduction

Over the last few decades, public pensions in the United States have fundamentally changed how they take risk. At the turn of the century, risky investments—defined as any holdings outside of fixed income and cash—were primarily in 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. These national trends 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. In contrast, the alternative-to-risky share for pensions in South Dakota, Nevada, Georgia, Iowa, and Colorado has hardly changed.

Why have some pensions shifted their risky investments toward alternatives more than others? This paper explores several possible explanations, including disagreement about the risk-adjusted returns of alternatives, risk-seeking motives, and a desire to conceal risk due to agency frictions. The evidence we present suggests that beliefs about alternatives, shaped by investment consultants, peers, and experience, are critical for understanding differences in alternative use across pensions. In the latter part of the paper, we use our cross-sectional results to shed light on the likely forces driving the increase in the aggregate alternative-to-risky share, once again finding a significant role for beliefs.

To structure our empirical analysis, we first look to modern portfolio theory for guidance on the potential factors behind the rise of alternatives in the cross-section. Our hypothesis development is grounded in canonical mean-variance theory (Markowitz, 1952; Merton, 1969), rather than richer models of portfolio choice (e.g., Merton, 1973; Lucas and Zeldes, 2009; van Binsbergen and Brandt, 2016), due to its widespread popularity: 85 out of the top 100 largest U.S. pensions explicitly reference mean-variance optimization in their annual reports, asset-liability studies, or documents available on their websites. According to standard mean-variance models, there are two possible reasons why the level and dispersion of the alternative-to-risky share have increased over time. One reason is that pensions have become more optimistic about the so-called alpha of alternatives relative to public equities, some more so than others.[1] Another reason is that pensions want to take on more risk but are constrained from doing so, leading them to contort the composition of their risky investments toward alternatives instead of increasing their risky share.

Several pieces of evidence affirm the role of beliefs in driving variation in the use and adoption of

[1] 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 expected covariance between the two asset classes divided by the expected variance of public equities.

alternatives across pensions. Most directly, pensions tend to have a higher alternative-to-risky share when their general investment consultant reports a higher alpha of alternatives relative to public equities.[2] This positive correlation could arise if consultants have a causal impact on their clients' portfolios or if consultants are chosen based on their beliefs. To isolate the causal component, we leverage the fact that some consultant selection is based on geographical proximity (Andonov et al., 2023). An instrumental variables design based on this idea indicates that widening disagreement among consultants is responsible for about one-third of the increase in cross-sectional variation in the alternative-to-risky share since the 2000s. A variance decomposition that is agnostic about the direction of causality also shows that a sizable amount of cross-pension variation in the alternative-to-risky share is attributable to beliefs.

Consultants are not the only channel through which pensions form beliefs about alternatives. Drawing on past research linking belief formation to experience (Malmendier and Nagel, 2016; Andonov and Rauh, 2022), we hypothesize that investors who were late to invest in the stock market during the 1990s came to view public equities less favorably than those who invested earlier, as their experience was more heavily influenced by the bursting of the dot-com bubble. Consistent with this idea, states whose pensions invested later in equities during the 1990s subsequently increased their alternative-to-risky shares more during the 2000s. Entry timing alone accounts for nearly one-fifth of the cross-pension variation in changes in the alternative-to-risky share from 2002 to 2021, even after controlling for the total increase in equities during the 1990s, funding levels in 2002, and whether investment restrictions were specified in each state's constitution in the late 1990s.

Motivated by household finance research on belief formation and social networks (Bailey et al., 2018, 2022), we also test whether pensions' beliefs about the alpha of alternatives are influenced by their peers, where peers are based on distance. We show that a 10 pp increase in the alternative-to-risky share of a pension's peers is associated with about a 5 pp increase in its own share. Importantly, this estimate accounts for any herding incentives that may lead nearby peers to invest similarly (Scharfstein and Stein, 1990), making it more likely to reflect belief-driven peer effects. In addition, to account for the possibility that peer effects reflect exposure to common shocks (Angrist, 2014), such as those coming from local labor markets or sharing the same consultant, our estimates are identified from variation in the alternative-to-risky share that exists within consultant and census division. We further address concerns that our peer effect

[2] 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.

estimates are driven by unobserved shocks by showing that they are similarly strong when defining peers based on nearby private-sector pensions, unions, and university endowments—institutions that are governed by different regulations and more plausibly face different shocks.

Next, we evaluate the hypothesis that risk-seeking motives have driven some pensions to increase their alternative-to-risky share more than others. In textbook mean-variance theories, absent a shift in beliefs, the composition of the risky portfolio will only change if pensions want to take more risk but are constrained from doing so. We therefore test whether there is a positive association between long-run changes in the alternative-to-risky share and two measures of risk constraints. The first is based on the difference between each pension's actual and target risky share, with the idea being that the actual risky share of risk-constrained pensions should be consistently above its target. The second is based simply on the cash share of each pension. In both cases, we find almost no correlation with alternative adoption in the cross-section.

We then ask whether pensions with larger incentives to take risk also invest more heavily in alternatives. We consider several proxies for 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 2021 on changes in pension funding yields an R^2 of 1%. It is also not the case that the most underfunded pensions at the turn of the century subsequently increased their alternative-to-risky share. These broad conclusions are robust to controlling for our proxies of risk constraints and to many other measures of risk-seeking incentives, including funding ratios based on market discount rates (Novy-Marx and Rauh, 2011) and whether plan sponsors have failed to make actuarial required contributions.

In the latter part of the paper, we consider the aggregate implications of our cross-sectional results. On its face, the weak cross-sectional link between risk-seeking motives and the alternative-to-risky share casts doubt on their ability to fully explain the aggregate increase in the alternative-to-risky share. Nevertheless, one may wonder whether our results are driven by measurement error in risk-seeking motives or risk constraints. For example, pensions may differ in their ability to take risk due to political factors that are orthogonal to funding and other observable characteristics.[3] We tackle this concern by simulating a mean-

[3]Previous research has shown that public pension investments are often distorted by home bias (Hochberg and Rauh, 2012) and politically-charged boards (Andonov et al., 2018). While these studies focus on fund selection *within* private equity, we analyze portfolio choice *across* asset classes. In Internet Appendix F.2, we find essentially no cross-sectional correlation between board

variance portfolio choice model (Merton, 1969; Campbell and Viceira, 2002), which as mentioned above, is used by a large majority of U.S. pensions for portfolio construction. The model allows us to quantify the strength of the risk-seeking channel in aggregate without needing empirical proxies for pensions' funding constraints or risk-seeking motives. 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 attempt to match the aggregate trend by reducing risk aversion and assuming that pensions are constrained from raising their risky share above its observed level in 2021.

Strikingly, in virtually all simulations (99.6%), it is not possible to lower risk aversion enough to match the observed aggregate change in the alternative-to-risky share. The intuition for this result is simple: by revealed preference, the low alternative-to-risky share in 2001 implies that pensions were bullish on public equities compared to alternatives. Consequently, absent a change in beliefs, they would shift their risky investments toward public equities and not alternatives when their portfolio constraints bind. To be clear, this simulation exercise does not rule out the possibility that U.S. public pensions are reaching for yield or facing binding portfolio constraints. On the contrary, the fact that the *risky* share has increased from 68% to 76% over the last twenty years is a clear indication that effective risk aversion of pensions has indeed declined. However, the mean-variance model highlights that a shift in beliefs is necessary to explain the concurrent increase in the *alternative-to-risky* share.

We then study how much of the aggregate increase in the alternative-to-risky share can be explained by beliefs. The simulated mean-variance model indicates that pension beliefs would need to rise by about 70 basis points (bps) to fully explain the observed aggregated trend. According to our estimates, about 10% of this increase can be causally attributed to the fact that the median consultant's perceived alpha has increased by about 70 bps since 2001. Additionally, using common techniques from network economics (Leontief, 1986; Acemoglu et al., 2016; Herskovic et al., 2020), we show that amplification by the peer network can account for another 20% of the shift in beliefs needed to match the aggregate alternative-to-risky share. Consistent with a broad shift in beliefs about alpha, we further show 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.

In the final part of the paper, we consider the idea that U.S. public pensions are drawn to alternatives

composition and the alternative-to-risky share.

because of their ability to conceal risk and smooth returns (Jurek and Stafford, 2015; Ilmanen et al., 2020; Stafford, 2022), which then allows pension managers and plan sponsors to avoid public scrutiny and enjoy a "quiet life" (Bertrand and Mullainathan, 2003; Dyck et al., 2018). While agency frictions of this kind are certainly possible, we rule in a belief-driven mechanism by showing that consultants' perceptions about alpha explain substitution patterns between private equity and real assets, both of which provide a similar scope for hiding risk (Couts et al., 2020; Stafford, 2022). In addition, our results place restrictions on the nature of any such agency frictions. For example, return-smoothing motives would need to have simultaneously increased over the last twenty years for institutions that vary widely in geography, governance, accounting, and regulation. Moreover, any agency frictions driving alternative adoption must vary in the cross-section of public pensions yet be orthogonal to a host of attributes, including funding, size, and board composition.

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 7.1.2, 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.

Literature Review This paper adds to research studying the portfolio choices of institutional investors. A growing literature emphasizes the importance of beliefs for understanding household investment behavior (Pástor, 2000; Leombroni et al., 2020; Beutel and Weber, 2022; Giglio et al., 2021). Much like with households, our evidence suggests that pension beliefs are shaped by consultants (Foerster et al., 2017), peers (Bailey et al., 2022), and past experience (Malmendier and Nagel, 2016; Bordalo et al., 2022). In an institutional setting, prior work has shown that clients who share a specialized consultant tend to select similar active managers in public (Jenkinson et al., 2016) and private markets (Andonov et al., 2023; Martinez and Qian, 2024), and adopt comparable approaches to benchmarking (Augustin et al., 2024). Though our results are broadly consistent with these studies, we depart from them in two important ways. First, we collect data on the beliefs of general consultants and directly tie them to pensions' broad asset allocation decisions, which generally occur before individual manager selection and are the biggest driver of overall portfolio performance (Ibbotson and Kaplan, 2000; Brown et al., 2010).[4] Second, we estimate the causal impact

[4]There is very little research that uses data of capital market assumptions (CMA) from investment consultants. One notable exception is Couts et al. (2023), who examine the factor structure of return expectations. Dahlquist and Ibert (2024) also use CMAs

of consultant beliefs on pension portfolio structure, showing that consultant beliefs are responsible for a meaningful fraction of the rise of alternatives, both in the cross-section and in aggregate. This causal link is, to our knowledge, new to research on the relationship between institutional investors and their external consultants.

The notion that institutional investors extrapolate from their experience also echoes Andonov and Rauh (2022), who show that pensions' expectations about future portfolio returns are influenced by past performance. We build on their work by tying the rise of alternatives in the cross-section of U.S. pensions to their experience during the 1990s stock market boom and bust cycle. As mentioned above, this event was particularly salient for U.S. public pensions because it was the first time many had meaningfully invested in the stock market. The fact that it still appears to impact the composition of risky investments today suggests a rather long form of institutional memory, consistent with research on risk-taking in the banking sector (Bouwman and Malmendier, 2015).

A related branch of research has focused on why investors appear to take more risk or "reach-for-yield" 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 risen since the turn of the century (Eaton and Nofsinger, 2004; Pennacchi and Rastad, 2011; Mohan and Zhang, 2014; Andonov, Bauer, and Cremers, 2017; Lu, Pritsker, Zlate, Anadu, and Bohn, 2019). In contrast to these previous studies, we ask why pensions have shifted the *composition* of their risky 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. Consistent with intuition from the mean-variance models used by most U.S. pensions (Tobin, 1958), we provide evidence that shifting beliefs are necessary to understand the adoption of alternatives.

The growth of alternatives is not just a U.S. phenomenon. Ivashina and Lerner (2018) document that private- and public-sector institutions across the globe have increased their portfolio share of alternatives since the turn of the century. Our analysis shows that this worldwide trend cannot be explained by a broad expansion into all types of risky assets. Instead, it has largely occurred through a change in the composition of risky investments. Increased optimism about the alpha of alternatives relative to public equities offers a

to study the time-series properties of return expectations and their impact on mutual fund holdings.

compelling explanation for why such diverse institutions have shifted the composition of their risky assets similarly, despite divergent risky shares. The cross-sectional evidence we present, along with canonical portfolio theory, further supports a belief-based view.[5]

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 examines beliefs as an explanation for the rise of alternatives in the cross-section and Section 5 focuses on risk-seeking channels. Section 7 discusses agency-based and supply-side explanations, then concludes. Additional details and results are available in an internet appendix.

2 Data and Motivating Facts

The data for this study come primarily from the Public Plans Database (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. Our definition of alternative assets includes private equity and credit, real assets (real estate, commodities, and infrastructure), hedge funds, and what PPD lists as other alternatives. Unless

[5] Another agency-based interpretation is that all types of institutions have moved to alternatives due to an increased desire to conceal risk. We discuss this possibility in Section 7.1.1.

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.

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.[6]

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 our processing of the PPD data 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 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. This is an imperfect measure of the risky share because high-yield debt securities should be included in risky investments, yet granular data on the credit rating of fixed income investments are limited.[7] Consistent with prior work on

[6] Our mapping takes into account a handful of pension mergers that have occurred during our sample period and is available upon request.

[7] 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.1.3.

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 PPD largely overlap when both are available, reinforcing the quality of the PPD data.[8] Second, and more strikingly, the 8 pp increase 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 and into public equities (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. We explore the cross-section of the risky share in Internet Appendix B.1.2.

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 2021. 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 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,

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

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 sizable amount 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 percentile pension increased its alternative-to-risky share by 15 pp whereas the 75th percentile pension's share 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 Hypothesis Development Using Mean-Variance Theory

Before turning to our main empirical analyses, we now look to theory for guidance on the potential factors driving the sharp rise in the level and dispersion of the alternative-to-risky share. Our hypothesis development is grounded in standard mean-variance theory for a very simple reason: out of the top 100 largest pensions by size, 85 mention mean-variance analysis in their CAFRs, asset-liability studies, or documents available on their websites. Take, for instance, CalPERS, the largest U.S. pension by assets. In materials used to educate pension board members on their approach to portfolio construction, CalPERS writes that they "... utilize mean-variance optimization (MVO) to evaluate the capital market assumptions to maximize desired return (mean) for any given level of undesired risk (variance)."[9]

Historically, public pensions in the U.S. have also not hedged against changes in their investment opportunity set, as prescribed by Merton (1973). According to the 1998 PENDAT survey administered by the Public Pension Coordinating Council (PPCC), 91% of pensions by count and 94% by assets report that their asset allocations do not change frequently with varying economic conditions (i.e., are not tactically set). While richer theories that incorporate intertemporal hedging (Merton, 1973), local labor market risk (Lucas and Zeldes, 2009), or asset-liability matching (van Binsbergen and Brandt, 2016) may be more suitable for a normative analysis of how pensions should optimally invest, we focus on mean-variance theory because our goal is to understand the key forces driving how they actually invest.

In the tradition of Markowitz (1952) and Merton (1969), our specific model features a power-utility investor who lives two periods and begins with wealth W_0. Details and derivations are in Internet Appendix G. 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). 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:

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

[9] The link to the full document for CalPERS can be be found here. To determine the broad use of mean-variance analysis, we searched through CAFRs, asset-liability studies, and any other documents available on the websites of the 100 largest U.S. pensions from 2016 to 2024. We also searched for pension materials on the eVestment database.

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 G.2-G.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:

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

if $\beta\omega_A + \omega_E > 0$. This condition is trivially satisfied in practice. Intuitively, the investor shifts 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.[10] 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 G.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, the classic mean-variance theory used by most public pensions 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. In what follows, we gauge the relative empirical strength of both explanations, though of course these channels are not mutually exclusive.

4 Belief-Based Explanations

In this section, we study whether cross-sectional dispersion in the alternative-to-risky share can be explained by disagreement about their risk-return properties.[11] We first show that the alternative-to-risky share tends to

[10] 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.

[11] Disagreement in our setting is expected, given the challenges associated with measuring the risks of alternatives. These challenges are evident in the lack of consensus in the literature on the risk-adjusted returns of alternatives (e.g., Kaplan and Schoar (2005); Phalippou and Gottschalg (2008); Harris, Jenkinson, and Kaplan (2014); Korteweg and Nagel (2016); Gupta and Van Nieuwerburgh (2019); Couts, Goncalves, and Rossi (2020); Stafford (2022); Korteweg, Panageas, and Systla (2023)).

be higher for pensions whose investment consultants believe alternatives have a high alpha. An instrumental variables design further reveals that this positive correlation is largely due to the causal impact that consultants have on their clients. Our estimates suggest that widening disagreement among consultants about alpha accounts for about a third of the increased cross-sectional dispersion in the alternative-to-risky share since the early 2000s. We then examine two classic factors that influence beliefs: past experiences and peer behavior.

4.1 Evidence from Pension Consultants

4.1.1 Data Collection and Consultant Identification

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 (Andonov et al., 2023). 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.[12]

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 states beliefs about expected returns, volatilities, and correlations of different asset classes (Couts et al., 2023). Depending on the consultant, these beliefs capture risk-return properties for the next five to thirty years. The 14 consultants for whom we have CMAs manage the majority of U.S. pension assets (~75%) but are more limited in total count. Additional details on our CMA data and how we process it can be found in Internet Appendix A.3.

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

Figure 4a provides a sense of how consultant beliefs have evolved in aggregate by plotting the median consultant's reported alpha for each year since 2001.[13] Alternatives include real assets, private equity and credit, 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 indeed risen steadily since the early-2000s, going from 158 basis points in 2001 to 226 bps in 2021. We study the underlying drivers of this increase in alpha (e.g., expected returns vs. betas) later in Section 7.1.1.

4.1.2 Directly Reported Beliefs

We now use CMAs to estimate the causal impact of consultant beliefs on pension portfolios. This causal parameter is useful for at least two reasons. First, it lets us directly measure how much of the increased cross-sectional dispersion in the alternative-to-risky share is attributable to widening disagreement among consultants about alpha. Second, as we discuss later in Section 6.2, it sheds light on how much consultant beliefs have driven the aggregate increase in the alternative-to-risky share.

The key identification challenges can be easily seen through a simple statistical model of how consultant beliefs impact the alternative-to-risky share. Consider a single time period and let $c = c(p)$ be the consultant of pension p. Further suppose that the alternative-to-risky share of pension p is well-approximated as follows:

$$\omega_p^* = \theta \alpha_{c(p)} + \kappa \xi_p + \pi \nu_p, \tag{6}$$

where $\alpha_{c(p)}$ is consultant c's belief about alpha, ξ_p is a pension-specific belief, and ν_p captures any non-belief preferences for alternatives, such as those driven by agency-based or risk-seeking incentives. This equation says that the alternative-to-risky share is determined by three factors: (i) the causal influence of consultant beliefs on portfolio composition, modulated through the parameter θ; (ii) a pension-specific component of beliefs, modulated through κ; and (iii) a non-belief-driven preference for alternatives over public equities.[14]

Suppose we attempt to estimate θ using an OLS regression of ω_p^* on consultant beliefs $\alpha_{c(p)}$. In this

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

[14] We show in Internet Appendix Section C.1 how to derive a version of Equation (6) by linearizing the mean-variance model from Section 3.

case, the population regression coefficient β is given by the familiar formula:

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

where τ_B is the coefficient from a regression of ξ_p on $\alpha_{c(p)}$ and τ_{NB} is the coefficient from a regression of ν_p on $\alpha_{c(p)}$. τ_B will be positive if there is *selection on beliefs,* meaning pensions are more likely to choose consultants who share their beliefs. Similarly, τ_{NB} will be positive if there is *selection on non-beliefs.* This latter form of selection could arise if pensions hire consultants who are bullish on alternatives to justify an agency- or risk-based preference for the asset class. Both forms of selection will cause the OLS coefficient β to overstate the magnitude of the causal parameter θ. Thus, for the purposes of estimating the causal impact of consultant beliefs on the alternative-to-risky share, we need an identification strategy to handle this potential omitted variable bias.

With these identification issues in mind, we estimate the following panel analogue of our simple one-period statistical model:

$$\omega^*_{p,t} = \lambda_t + \beta\alpha_{c(p),t} + \Gamma X_{p,t} + \varepsilon_{p,t}, \tag{7}$$

where $\omega^*_{p,t}$ is the alternative-to-risky share of pension p in year t and $\alpha_{c(p),t}$ is the reported alpha of alternatives of pension p's consultant c in year t. λ_t is a time fixed effect and $X_{p,t}$ is a vector of controls that includes the level of pension funding, (log) size, the expected return on each pension's assets (i.e., its hurdle rate), the ratio of required actuarial contributions to payroll, the ratio of administrative expenses to payroll, the cash share, and the annual return of the fund. $X_{p,t}$ also includes the difference between each pension's target and actual risky share, residualized to its past return. These controls are designed to capture agency-based and risk-seeking motives that may lead some pensions to prefer alternatives over public equities (ν_p). Their logic is discussed further in Sections 5.1 and 5.2. By including them, we seek to minimize the bias on β induced by selection on non-beliefs. Standard errors are double-clustered by consultant and time.

Column (1) of Table 2 presents OLS estimates of Equation (7) when including no controls. The point estimate of 3.43 ($t = 3.01$) is statistically significant and indicates that pensions have higher alternative-to-risky shares when their consultants believe the alpha of alternatives to be high. Figure 4b depicts this relationship using a binscatter plot and confirms our linear specification. One concern with the standard errors in regression (7) is that they are based on a relatively small number of consultant clusters (Cameron et al., 2008). To address this potential issue, we use the wild bootstrap procedure from Cameron et al. (2008) to test the null hypothesis that the estimated β is different from zero. The resulting p-value from this test is reported in the row labeled $p_{wild}(\beta = 0)$ for all specifications in the table and indicates that we can reject the null at conventional confidence levels.

Column (2) shows OLS estimates when adding controls $X_{p,t}$ to the regression. The point estimate of 3.47 is very similar to column (1) and remains statistically significant ($t = 3.21$). Because the controls are designed to capture non-belief preferences for alternatives, the fact that the coefficient remains stable when moving from column (1) to (2) suggests that selection on non-beliefs is likely to be small. We test this idea more formally by applying the bias-adjustment from Oster (2019), under her recommended assumption that the controls $X_{p,t}$ and any unobserved variables related to non-belief preferences contain equal explanatory power for $\alpha_{c(p),t}$. The resulting bias-adjusted coefficient on $\alpha_{c(p),t}$ in this case still equals 3.10. The economic significance of this point estimate is discussed below.

While the controls in column (2) may address selection on non-beliefs, they are unlikely to handle selection on beliefs. Consequently, the OLS estimates of β may still overstate the true causal impact of consultant beliefs on the alternative-to-risky share. To tackle this issue, we build an instrument for $\alpha_{c(p),t}$ based on the findings of Andonov et al. (2023), who show that pensions are more likely to hire geographically proximate consultants. Formally, we define the instrument for each consultant-pension pair as follows:

$$z_{c(p),s,t} = \sum_{j=1}^{C} w_{j,p,s} \alpha_{j,t},$$

where j indexes consultants and $w_{j,p,s}$ is a weight based on the distance between pension p and consultant j in year s. $z_{c(p),s,t}$ should be thought of as a Bartik instrument (Bartik, 1991; Blanchard and Katz, 1992), whose cross-sectional variation is driven entirely by the weights, $w_{j,p,s}$. Goldsmith-Pinkham et al. (2020) show that the validity of this instrument thus depends on whether, conditional on controls $X_{p,t}$, the weights in year s—equivalently, consultant location—are exogenous to pension preferences for alternatives, whether belief-driven ($\xi_{p,t}$) or not ($\nu_{p,t}$).

Given this condition, we construct weights in two different ways. In the first, we set $s = t$, so that $z_{c(p),s,t}$ in year t is defined using weights that are based on the contemporaneous location of each consultant. The logic of this weighting scheme is that consultant location in a given year t is orthogonal to contemporaneous pension preferences for alternatives, *after* conditioning on $X_{p,t}$. One may object that consultants have incentives to be geographically close to large pensions given their fees are typically based on AUM, but we directly control for size in $X_{p,t}$. The second weighting scheme sets $s = 2005$, the year before the shift to alternatives accelerated (Figure 2c). For both reference dates s, weights are defined using indicator variables for whether consultant j was a registered investment advisor in p's state in any year prior to s. Data on registered investment advisors come from the SEC's IAPD and Finra's Broker-Check platform.[15] The weights are then scaled to sum to one based on the underlying indicator variables. The instrument will be relevant if pensions are more likely to choose a consultant who is located in their state.

Column (3) of Table 2 presents instrumental variable (IV) regression estimates of Equation (7) when instrumenting $\alpha_{c(p),t}$ with the $z_{c(p),s,t}$ that is based on contemporaneous distances ($s = t$). The large first-stage F-statistic validates the relevance of our instrument and confirms that pensions are more likely to hire nearby consultants, consistent with Andonov et al. (2023). The estimated β_{IV} equals 3.24 ($t = 3.40$) and is statistically significant when using standard errors that are clustered by consultant and year or when using the wild bootstrap procedure from Cameron et al. (2008). The third row from the bottom of the table shows that the IV estimate is slightly lower than the OLS estimate, indicating that selection is positive yet small on average. This suggests that the cross-sectional correlation between the alternative-to-risky share and consultant beliefs arises mainly due to a causal channel. This causal relationship almost surely operates through beliefs, as it is implausible that consultants who are bullish on alternatives could somehow cause an increase in risk-seeking or agency-based incentives that might lead pensions to invest more in alternatives (i.e., ν_p).

For robustness, column (4) shows IV estimates when the instrument is based on where consultants were located in 2005. For reasons discussed above, this instrument is arguably more valid given consultants are less likely to have chosen where to operate in 2005 based on *future* pension preferences for alternatives, especially after conditioning on observable pension characteristics $X_{p,t}$. The point estimate here equals 3.93 and remains statistically significant ($t = 3.94$). To contextualize this magnitude, note that the spread between the 90th and 10th percentile consultant-reported α expanded by roughly 100 basis points from 2001 to 2021. Accordingly, the point estimate in column (4) suggests this widening contributed to a 4 pp increase in the cross-sectional dispersion of the alternative-to-risky share. During the same period, the gap in the alternative-to-risky share between the 10th and 90th percentile pension grew by about 13 pp. Therefore,

[15] We thank Mark Egan for sharing these data with us. We further assume that if the consultant has a pension client in the state based on annual reports, they were also a registered investment advisor in the state.

assuming our IV is valid, about one-third of the variation in adoption rates of alternatives across pensions can be attributed to widening disagreement among consultants regarding the alpha of alternatives.

4.1.3 Consultant Fixed Effects

Next, to complement our IV design, we use a series of fixed-effects regressions to determine how much of the cross-sectional dispersion in the alternative-to-risky share is attributable to disagreement about alpha, regardless of whether the disagreement is driven by consultants or pension-specific beliefs. Our workhorse panel specification mirrors Equation (7) but replaces consultant-reported beliefs with consultant fixed effects, λ_c:

$$\omega^*_{p,t} = \lambda_t + \lambda_c + \Gamma X_{p,t} + \varepsilon_{p,t}. \tag{8}$$

Compared to Equation (7), the advantage of this specification is that it can be estimated for our full sample of pensions matched with consultants, as CMA data are available for 62% of the pension system-year pairs in our sample.

The objects of interest in the regression are the consultant fixed effects λ_c. In Internet Appendix C.2, we formally show that the incremental R^2 from adding these fixed effects to the regression provides a lower bound on the amount of belief-driven dispersion in the cross-section, assuming there is no selection into consultants based on non-beliefs after conditioning on observables $X_{p,t}$. This condition would be violated if, for example, there is an agency-based preference for alternatives that is both orthogonal to $X_{p,t}$ and leads pensions to hire consultants who believe alpha is high. Intuitively, when there is no selection on non-beliefs, variation in consultant fixed effects can only be driven by consultant beliefs, pension beliefs, or selection based on beliefs. The upshot is that we do not need to perfectly identify the causal impact of consultants in order to put a lower bound on the size of belief-driven variation in the cross-section, so long as there is little to no selection on non-beliefs. We maintain this weaker assumption when interpreting the estimates of λ_c and provide additional support for it in Internet Appendix C.2.3, where we study the portfolio composition of public- and private-sector clients who share a consultant.[16]

The first row of Table 3 presents the adjusted R^2 obtained from estimating the model with time fixed effects λ_t and pension attributes $X_{p,t}$. The main result is in the second row, which shows that consultant

[16]The basic idea is that public- and private-sector pensions face different agency frictions because they have different governance structures, accounting rules, and regulations. Thus, if public- and private-sector clients who share a consultant invest similarly, it is more likely driven by beliefs rather than selection based on agency frictions.

fixed effects λ_c add a considerable amount of explanatory power to the model. The increase in adjusted R^2 from their inclusion suggests that beliefs are responsible for at least 16 pp of the cross-sectional dispersion in the alternative-to-risky share. The actual amount attributable to beliefs will be much higher if there is meaningful variation in pension-specific beliefs within consultants (see Internet Appendix C.2). We also formally test the null hypothesis that the consultant effects are equal. The F-statistic and the p-value from this test are reported in the table, and indicate that we can easily reject the null of equal consultant effects at conventional confidence levels.

To further visualize the variance of the consultant fixed effects, Figure 5 plots their distribution 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 50%.

One concern with the controls $X_{p,t}$ is that they may not adequately control for non-belief motives that could drive pension preferences for alternatives. An obvious example is that some state legislatures allowed for local pensions to invest in alternatives later than others (e.g., South Carolina). In addition, agency frictions related to home bias (Hochberg and Rauh, 2012) or pension governance (Andonov et al., 2018) also likely vary by state. To the extent that these local factors also drive consultant selection, they may not be fully captured by $X_{p,t}$, potentially invalidating our bounding argument. A natural way to handle state-level confounders of this kind is to replace time fixed effects in Equation (8) with state-by-time fixed effects. The third row of Table 3 shows benchmark regression results when including only state-level fixed effects and controls, then row (4) shows results when also including consultant effects. Even in this more stringent specification, the incremental R^2 from adding consultant fixed effects in row (4) suggests that beliefs are responsible for at least 12 pp of the cross-sectional dispersion in the alternative-to-risky share. This result offers strong evidence of the importance of beliefs (via consultants) because the state-by-time fixed effects absorb any state-level preference for alternatives, whether they are belief-driven or not.

Rows (5) to (16) of Table 3 repeat this exercise by running regression (8) for different types of alternatives, not just the alternative-to-risky share. Rows (6), (10), and (14) show there are also large consultant effects for private equity, hedge fund, and real asset investments. This conclusion is robust to adding state-by-time fixed effects. The incremental R^2 from adding consultant effects indicates that beliefs account for at least 14-19% of the cross-pension variation in the use of specific types of alternatives. The private equity effects have a correlation of 8% and 35% 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. To the extent that selection on beliefs and non-beliefs is small, this result suggests consultants disagree on the alpha of different alternatives and advise their clients accordingly.

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, in the context of public pensions, Andonov and Rauh (2022) find that beliefs about portfolio-level returns are influenced by past performance. We build on their work by exploring how past experience influences the *composition* of risky assets, focusing primarily on experience during the 1990s.

The 1990s were a unique period for U.S. pensions, as many rapidly increased their overall risky share and invested significant amounts in the stock market for the first time (see Figure 1b).[17] At the same time, public markets during the 1990s famously went through a large boom-and-bust cycle that culminated with the bursting of the dot-com bubble between 2000 and 2002. Thus, for pensions that shifted to equities later in the decade, their early experience investing in the stock market was more heavily influenced by the bursting of the bubble compared to those that shifted earlier. We hypothesize that this negative and salient experience caused these pensions to become bearish on public equities relative to alternatives. Consequently, after the turn of the century, they increased their alternative-to-risky share more than pensions that had a relatively positive experience with the stock market in the 1990s.

One way to test this hypothesis is to compare two pensions, *A* and *B*, that increased their equity share by the same amount between 1990 and the bursting of the dot-com bubble in 2002. The key difference is that pension *A* increased its share later in the bull market and was therefore more influenced by the bubble's burst. We implement this idea using data on 1990s portfolio composition from PENDAT surveys conducted

[17] According to the Annual Survey of Public Pensions (ASPP), virtually every dollar that left fixed income during the 1990s went into public equities (see Internet Appendix B.1.1.

by the Public Pension Coordinating Council (see Mitchell et al., 2001). As detailed in Internet Appendix A.4, the PENDAT data cover virtually all U.S. pension assets and allow us to compute the state-level share of pension assets invested in equities at the beginning of the bull market in 1990 and midway through it in 1996. After merging these data with the PPD, we then run the following state-level regression:

$$\Delta\omega^*_{s,2002\to2021} = c + \beta_1 \Delta\omega^E_{s,1996\to2002} + \beta_2 \Delta\omega^E_{s,1990\to2002} + \Gamma X_s + \varepsilon_s \tag{9}$$

where $\Delta\omega^*_{s,2002\to2021}$ is the change in state s's alternative-to-risky share from 2002 to 2021, $\Delta\omega^E_{s,1996\to2002}$ is the change in s's equity share between 1996 and 2002, and $\Delta\omega^E_{s,1990\to2002}$ is the change in s's equity share between 1990 and 2002. X_s is a vector of state-level controls that we describe below. The object of interest in the regression is β_1. Because we control for the total change in equities between 1990 and 2002 ($\Delta\omega^E_{s,1990\to2002}$), β_1 isolates the timing of the change. Under our hypothesized mechanism, it should therefore be positive. Robust standard errors are used in all regressions.

As a baseline, column (1) of Table 4 runs the regression only with the change in the equity share between 1996 and 2002, $\Delta\omega^E_{s,1996\to2002}$.[18] As expected, β_1 is positive and statistically significant. Its magnitude means that a one-standard deviation increase in $\Delta\omega^E_{s,1996\to2002}$ is associated with a 6.84 pp increase in $\Delta\omega^*_{s,2002\to2020}$. The standard deviation of the latter equals 16.65 pp. Moreover, the change in equity share between 1996 and 2002 alone explains 19% of the variation in state-level changes in the alternative-to-risky share between 2002 and 2021.

Ideally, regression (9) would use target equity shares, not actual equity shares, as the former reflect only active portfolio decisions. Unfortunately, target shares are not available in the 1990 wave of the PENDAT survey and are less populated at the pension level in the 1996 wave. For this reason, when constructing the regression variables, we use actual equity shares for 1990 and 1996 from PENDAT and target shares for 2002 from PPD. In column (2), for robustness, we instead use 1996 target shares from PENDAT. Reassuringly, the point estimate is comparable to that in column (1) and remains statistically significant.

Column (3) adds the key control for our test: the overall change in equities between 1990 and 2002. As mentioned above, the inclusion of this control is important because it holds fixed the total increase in equities during the boom-and-bust cycle and therefore isolates the timing of entry into equities. The point

[18] Portfolio shares in PENDAT are not available for DC, Maine, and Nebraska. Additionally, PPD data are not available for New Jersey, Rhode Island, and Vermont in 2002, which is why the sample size is 45. We find similar results when predicting changes in the alternative-to-risky share between 2006 and 2021, where New Jersey, Rhode Island, and Vermont are also included.

estimate in this case is larger compared to column (1) and is statistically significant. The point estimate on $\Delta\omega^{E}_{s,1990\to 2002}$ is statistically insignificant and adds very little explanatory power to the regression, further highlighting the importance of entry timing into the stock market during the 1990s.

In column (4), we add several controls to the regression, including changes in hurdle rates and in the fraction of retired members (Andonov et al., 2017), the level of BEA-adjusted funding in 2002, and the log of AUM in 2002. All changes are measured between 2002 and 2021. These variables are designed to control for variation in risk-seeking motives, especially those potentially induced by poor stock-market performance during the 1990s. Additionally, we include the fraction f_s of each state's assets that were affected by constitutional investment restrictions in 1998, constructed using PENDAT data. This control is important because it accounts for the possibility that states increasing their alternative-to-risky share the most during the 2000s were simply those that lifted constitutional investment restrictions later (e.g., South Carolina). While the point estimate on f_s is positive, it is not statistically different from zero. More importantly, its inclusion does little to change the estimated β_1.

Finally, as a placebo test, we replace the change in the alternative-to-risky share from 2002 to 2021 with the change in the risky share as the outcome variable in the regression. The estimate β_1 is nearly a quarter of the magnitude compared to column (4) and is not statistically significant. This finding reinforces the idea that 1990s experience in the stock market influenced pension beliefs about alternatives versus public equities, rather than their overall willingness to take risk.

4.3 Peer Effects

Finally, we investigate the extent to which cross-sectional variation in alternative use 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.[19] 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:

$$\omega^{*}_{A,p,t} = \lambda_{cdt} + \beta_z n_{pt} + \theta^{'}\mathbf{X}_{pt} + \varepsilon_{pt}, \tag{10}$$

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 4.1.3). 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 5 reports our baseline estimate of the regression specification in (10). 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 5.2, 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

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

for the first channel by testing whether peer effects are still present in subsets of the data where herding incentives are likely weaker.

We analyze three different proxies for herding incentives. The first, presented in column (2), is based on the idea that CIOs with more job security have less of an incentive to herd. Job security is measured 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. To mitigate this measurement issue, we create an indicator variable based on whether CIO tenure is at least five years and only estimate Regression (10) 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 (10), 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 proxy for herding incentives with 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.

The third measure of herding incentives that we consider is an indicator variable based on whether a pension is in the top quintile of annual overall performance. Similar to column (3), the idea is that high-performing pensions have less negative public scrutiny, thereby reducing their incentive to herd as in Scharfstein and Stein (1990). The negative and statistically significant interaction term in column (4) 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 of about 0.5 for this subgroup of pensions.

For robustness, column (5) replaces 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.

Our preferred interpretation of the results in Table 5 is that peers influence how pensions form beliefs about the risk-return properties of alternatives. As mentioned above, it is hard to distinguish this interpretation from one in which peer effects are driven by exposure to common non-belief-driven shocks, such as those impacting member life-expectancies. To further reinforce the belief-driven channel, column (6) measures peer behavior based on the average alternative-to-risky share of endowments, corporate pensions, and unions in each state. These institutions differ widely from public pensions in their funding structure, governance, and regulation, thereby reducing the likelihood they are exposed to common shocks. In the regression, we replace the consultant-by-census-division-by-year fixed effect with a consultant-by-year fixed effect to ensure there is sufficient residual variation in this state-level measure of peer behavior. Consistent with beliefs driving peer effects, the estimated β is positive and statistically significant, and is comparable to the elasticity for high-performing pensions found in column (5). The idea that public pensions form beliefs about alternatives based on the behavior of local private-sector institutions seems plausible given the relatively high historical performance of these institutions, most notably endowments (Lerner et al., 2007, 2008). Indeed, university endowments adopted alternatives starting in the 1990s, following the so-called Yale Endowment model pioneered by David Swensen, who was famously bullish on alternatives (Goetzmann and Oster, 2012).

5 Risk-Seeking Explanations

In the mean-variance model from Section 3, pensions may shift their risky investments toward alternatives if they want to take on more risk but are constrained from doing so. In this section, we ask whether this mechanism can explain cross-pension variation in the adoption of alternatives. Our analysis consists of two parts. First, we develop several measures of risk constraints and link them to long-run changes in the alternative-to-risky share. Second, we propose a few related mechanisms through which effective pension risk aversion might decline or vary in the cross-section, including factors such as pension underfunding. We then connect various proxies for risk-seeking motives to investment in alternatives, both before and after controlling for risk constraints.

5.1 Binding Risk Constraints

A necessary condition for risk aversion to impact the alternative-to-risky share in textbook portfolio choice models is the presence of binding constraints on the amount of risk that investors can take. Studying this mechanism empirically is, however, not straightforward. The main challenge we face is that binding risk 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.[20] The basic idea is as follows. Market fluctuations will naturally move each pension's actual risky share from its target, after which pensions must rebalance to bring the two in line. For example, in 2022, CalPERS 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.

Next, we test whether pensions facing binding risk constraints are also those that have increased their alternative-to-risky share. We do so via the following cross-sectional regression:

$$\Delta\omega^*_{A,target,p} = a + \theta \bar{l}_p + \varepsilon_p \tag{11}$$

where $\Delta\omega^*_{A,target,p}$ is the change in the alternative-to-risky share for pension p between 2002 and 2021 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 6a 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 0.90 ($t = 1.19$) and the linear specification yields a fairly low R^2 of 2%. The regression coefficient indicates that a one-standard deviation increase in the average constraint is associated with a 2 pp increase in the alternative-to-risky share. The standard deviation of the change in the alternative-to-risky share for this sample equals 16 pp.

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

One concern with this analysis is that deviations of actual from target 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 0.89 ($t = 1.15$).

Cash holdings provide another way to proxy for binding risk constraints, with the intuition being that a constrained pension will want to hold as little of its fixed income assets in cash as possible. Figure 6b tests this idea by replacing $\bar{l}_p$ in regression (11) with the contemporaneous change in each pension's target cash share. The figure offers some cross-sectional evidence that declines in cash holdings are associated with an increases in the alternative-to-risky share, but the correlation between the two is again rather weak. The regression yields an estimated coefficient equal to -0.68 ($t = -1.44$) and a low R^2 of 2%. In terms of magnitude, the regression coefficient implies that a one standard deviation decrease in the change in cash holdings is associated with a 2pp increase in the alternative-to-risky share.

As a complementary test, we instead regress the 2021 level of the alternative-to-risky share on the contemporaneous level of target cash share. Assuming cash holdings are a reasonable proxy for the tightness of risk constraints, then under the view that risk-constrained pensions are more likely to invest their risky assets in alternatives, we should find a negative correlation between the two. Consistent with this hypothesis, the coefficient from this regression equals -0.84 ($t = -2.18$). Still, the R^2 in the regression is 3% and the regression coefficient implies that a one-standard deviation decrease in the cash share is only associated with a 3 pp increase in the alternative-to-risky share. Thus, while we find some statistical evidence that risk-constrained pensions have shifted toward alternatives and away from public equities, the strength of this relationship is economically weak.

5.2 Risk-Seeking Motives

Next, we examine whether pensions with higher risk-seeking motives are also those that have shifted more aggressively out of public equities and into alternatives. Our baseline regression for this analysis is as follows:

$$\Delta\omega^*_{A,target,p} = a + \beta \Delta X_p + \varepsilon_p, \tag{12}$$

where $\Delta\omega^*_{A,target,p}$ is the change in the alternative-to-risky share for pension p between 2002 and 2021. The explanatory variable ΔX_p is the contemporaneous change in one of several proxies for pension risk aversion. Standard errors are clustered by state for regressions run at the system level and are otherwise robust.

Our first proxy for risk-seeking motives is based on the idea that pension underfunding creates incentives for U.S. pensions to "reach-for-yield" via alternatives (Lu et al., 2019; Gillers, 2021), similar to those facing corporate equity holders when approaching default (Jensen and Meckling, 1976).[21] Figure A9a shows that these incentives have risen for U.S. pensions as interest rates have fallen over the last twenty years. At the beginning of the 2000s, U.S. pensions were more than fully funded according to standards set forth by Statement 25 of the Governmental Accounting Standards Board (GASB), but by 2021 their assets were only 75% of liabilities.[22] Figure A9b further shows that pension funding varies strongly in the cross-section, with some pension funding ratios sitting below 50% as of 2021.

Column (1) of Table 6 shows that deterioration in pension funding is positively 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 statistically significant ($p = 0.26$) and implies that moving from the 90th to 10th percentile of the change in funding is associated with a 4% 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 40%. 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 D.1.1 that the poor model fit is not driven by a non-linear relationship between funding and alternative usage.

An 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'

[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.

[22]Under 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.

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 6 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.

One explanation for the weak cross-sectional relationship between the alternative-to-risky share and the proxies for risk-seeking motives considered in columns (1)-(4) of Table 6 is that we do not account for the extent to which each pension is risk-constrained. To see why this could potentially be an issue, suppose that changes in pension risk aversion are well-proxied by funding status. If pensions are mean-variance investors as in the model from Section 3, then underfunding will only impact the alternative-to-risky share for those that are also risk-constrained. Consequently, if only a subset of pensions are constrained in the cross-section, then the relationship between funding and changes in the alternative-to-risky share will be attenuated.

A natural way to handle this potential issue is to control directly for the tightness of risk constraints in regression (12). The logic of the controls is to compare two pensions that are equally constrained. To the extent that risk-seeking motives drive alternative use, the pension with larger risk-seeking motives should then shift more aggressively to alternatives. Columns (5)-(7) implement this idea by controlling for the two measures of constraint slack developed in Section 5.1: (i) the average difference between actual and target risky shares; and (ii) the contemporaneous change in the target cash share. The regressions are all run at the system level because this is the unit of observation for which our measures of constraint slack are constructed. In all cases, controlling for constraint slack does little to change the point estimates observed in columns (1), (3), and (4).

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.[23]

Robustness. In Internet Appendix D.1, we probe the robustness of this conclusion in a few ways. First, we document that the initial levels of the covariates in Table 6 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 between 2002 to 2021. There is also essentially no relationship between the change in the alternative-to-risky share and pension size or failure to make actuarial required contributions. Second, we show that the findings in Table 6 are 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.

6 Aggregate Implications

Our analysis thus far points to beliefs as an important factor driving cross-sectional variation in the use and adoption of alternatives, with risk-seeking motives playing a lesser role. In this section, we investigate the ability of these two mechanisms to explain the aggregate increase in the alternative-to-risky share. We

[23] Although not our primary focus, for completeness we explore the role of funding and accounting-based mechanisms in explaining the increase of the overall risky share in Internet Appendix D.2.

present four complementary pieces of evidence. First, we simulate the mean-variance model from Section 3 to study the conditions under which reaching-for-yield behavior—equivalently, a decline in risk aversion—can match the evolution of the aggregate portfolio from 2001 to 2021. Second, we use the simulated model to assess how much beliefs about alpha would need to shift to explain the aggregate. As a part of this exercise, we also incorporate the impact of consultants and peers that were estimated in Sections 4.1 and 4.3. Finally, we examine the behavior of U.S. endowments, U.S. corporate pensions, and U.K. corporate pensions to shed light on the likely drivers of the alternative-to-risky share for U.S. public pensions.

6.1 Simulating Reaching-for-Yield in the Model

Our results from Section 5 suggest that risk-seeking motives and binding portfolio constraints can explain only a small amount of cross-pension variation in alternative use. On its face, this result casts doubt on the ability of these factors to fully explain the increase in the aggregate alternative-to-risky share as well. However, one potential caveat to this conclusion revolves around our 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 underfunding may vary at the state or local level. Pension regulation is another potential confounding factor. In the late 1980s and 1990s, many state and local legislatures lifted restrictions on the ability of pensions to take risk, thereby decreasing the effective risk aversion of pensions in aggregate and potentially driving the rise in the alternative-to-risky share during the 2000s. Such complexities could introduce measurement error into our proxies for risk aversion or binding portfolio constraints, potentially attenuating their relationship with pension investment behavior in the cross-section.

To address these concerns, we use the mean-variance model from Section 3 to simulate how a decline in risk aversion coupled with a binding portfolio constraint would impact the composition of the aggregate pension portfolio. We do so for a wide range of beliefs that are held fixed throughout each simulation. This strikes us as a reasonable exercise given that, as previously mentioned, the large majority of U.S. pensions use some form of mean-variance analysis for portfolio construction.

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_{2001} 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^* so that $\sigma_\eta^2 > 0$, $\gamma_{2001} \geq 1$, 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_{2021} needed to generate the observed $\Delta\omega_A^*$ in the data. Panels A and B in Table 7 summarize the simulation approach. In all cases, we impose a constraint that alternatives cannot be shorted, meaning $\omega_A \geq 0$.

Panel C of Table 7 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), pensions would want to shift towards public equities and *not* alternatives when they become 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 A21 in the Internet Appendix). Panel C of Table 7 also shows that 0.4% of simulations (419 out of 100,000) succeed in matching the observed

[24] See Figure A19 in the Internet Appendix for the histograms of admissible beliefs.

$\Delta\omega_A$ with a decline in risk aversion. However, in Internet Appendix E.1, we show that the implied shadow cost of the risk constraint is implausibly high, as it implies that pensions would forgo over 700 basis points of returns per year to relax the constraint.

It is important to stress that the simulation exercise above does not rule out the possibility that U.S. public pensions are reaching for yield or facing binding portfolio constraints. On the contrary, the fact that the *risky* share has increased over the last twenty years is a clear indication that pension risk aversion has indeed declined. Rather, the model highlights that neither of these channels can sufficiently explain the increase in the *alternative-to-risky* share without a concurrent shift in beliefs. The cross-sectional evidence in Sections 5.1 and 5.2 further supports this interpretation, as does the positive evidence we provided in Section 4.

6.2 The Role of Beliefs

Next, we modify the simulation framework from above to study how much beliefs would need to shift to explain the aggregate increase in the alternative-to-risky share. To do so, we no longer assume that pensions face portfolio constraints and instead allow beliefs about alternatives to change. We start from the same set $S^* = 100,000$ of random initial beliefs and risk aversion parameters that match the aggregate pension portfolio in 2001. Within each simulation, we then solve for the increase in perceived alpha α_p of alternatives relative to public equities that is needed to match the alternative-to-risky share ω_A^* in 2021. We then solve for the risk aversion γ_{2021} in 2021 needed to match the aggregate risky share.[25] Unsurprisingly, an increase in α_p can rationalize $\Delta\omega_A^*$ for all $100,000$ initial beliefs. Figure A22a in the Internet Appendix shows how the implied increase in α_p varies across simulations. We winsorize the distribution at the 97.5th percentile to make the plot more readable. The average required increase in perceived alpha across simulations is 70 bps.

The Direct Impact of Consultants In Internet Appendix E.2.2, we show how to further enrich this simulation by incorporating the causal impact of consultants estimated from Section 4.1. The basic idea is to express the total required change in alpha ($\Delta\alpha_p$) within each simulation as the sum of two components: $\Delta\alpha_p = \zeta\Delta\alpha_c + \Delta\xi_p$, where $\Delta\alpha_c$ is the observed change in consultant-perceived alpha in Figure 4a and $\Delta\xi_p$ is the change in the pension sector's private beliefs about the asset class. ζ is the causal influence of consultant

[25] See Internet Appendix E.2 for the complete set of simulation results.

beliefs on pension beliefs. In each simulation, it can be computed by feeding in the estimated elasticity ϕ of the alternative-to-risky share with respect to consultant beliefs from Section 4.1 into the mean-variance model from Section 3. ζ varies across simulations because it depends on ϕ, as well as the initial beliefs and risk aversion parameters used in each simulation.

Figure 4a shows that the median consultant's perceived alpha rose by roughly 70 bps between 2001 and 2021. Given a value of ζ and setting $\Delta\alpha_c \approx 70$ bps, we can then back out the change in private pension beliefs $\Delta\xi_p$ needed to match the observed increase in the alternative-to-risky share. By construction, the implied $\Delta\xi_p$ in each simulation is therefore consistent with the composition of the aggregate pension portfolio in 2001 and 2021, as well as the causal impact of consultants on the alternative-to-risky share estimated in Section 4.1. The average required increase in $\Delta\xi_p$ across simulations appears plausible and is about 60 basis points (see Figure A24). In other words, the causal estimates from Section 4.1 suggest that about 10-15% of the aggregate trend in the alternative-to-risky share can be attributed to consultants becoming more bullish about alternatives.

Peer Effects Our results from Section 4.3 further imply that the required increase of $\Delta\xi_p \approx 60$ bps in pension-specific beliefs could be partly driven by peer effects. For example, if one pension has an idiosyncratic belief shock about alpha, the existence of peer effects means that the shock will impact the behavior of other pensions and thus the aggregate portfolio composition. In Internet Appendix E.2.3, we use standard techniques from network economics to translate our estimates of peer effects into multipliers that measure how the peer network amplifies belief shocks in aggregate (e.g., Leontief, 1986; Acemoglu et al., 2016; Herskovic et al., 2020). Our analysis suggests that amplification by the peer network can account for about 20% of the increase in pension-specific beliefs needed to match the aggregate alternative-to-risky share.

6.3 Evidence from Other Institutions

We conclude our aggregate analysis by exploring the idea that the perceived alpha of alternatives 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, and U.S. endowments. Data for U.S. corporate pensions are 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.[26] 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 more than quadrupled 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 aggregate data in Section 7.1.1, particularly those related to agency frictions.

7 Other Explanations 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

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

on return-smoothing motives and growth in the supply of alternatives.

7.1 Other Explanations

7.1.1 Return Smoothing

It is well known that risk-adjusting the returns to alternatives is challenging (Lo, 2001; Jurek and Stafford, 2015). This is especially true in illiquid asset classes like private equity and real estate, where the lack of mark-to-market pricing means realized returns are mechanically less volatile and less correlated with public markets (Stafford, 2022; Couts et al., 2020). Consequently, the realized alpha of alternatives may appear artificially large and cause investors to form overly bullish beliefs about future alpha.

In recent years, many observers have posited that these measurement issues are precisely why institutional investors, like pensions, are attracted to alternatives (Ilmanen et al., 2020). For example, agency frictions may lead some institutions to prefer the "quiet life" that unmarked assets like private equity provide over public markets (Bertrand and Mullainathan, 2003; Stafford, 2022). Under this agency-based view, pensions and consultants may publicly state that they believe alternatives offer diversification benefits and alpha over public markets, even though they privately recognize that these purported advantages are not real.

One way to support the idea that true beliefs about alpha influence portfolio choice is to study the choice between real assets and private equity. Because both asset classes offer a similar ability to conceal risk, it is less likely that substitution patterns between the two are driven by return-smoothing or other agency motives. Indeed, in Internet Appendix F.1, we document that pensions are more likely to invest in real assets over private equity when their consultant reports real assets to have a higher alpha, even after accounting for any time-varying pension characteristics that may drive a desire to take on risk or invest in unmarked assets.

While this evidence is certainly helpful, we recognize that it is typically challenging to discern from observational data whether beliefs genuinely reflect one's true convictions or are influenced by underlying agency frictions. Our findings do, however, place restrictions on the specific nature of any such friction. For one, in aggregate, the return-smoothing motives must have increased over the last twenty years for a variety of institutions operating in different countries and thus facing different regulatory and public pressures (Section 6.3). Second, the agency friction driving the desire to smooth returns must vary in the cross-section of pensions, but in a way that is orthogonal to funding, size, member demographics, board composition (Andonov et al., 2018), and several other pension attributes (Section 5.2 and Internet Appendix F.2). Third, it must lead a subset of private- and public-sector pensions to choose consultants who report a high alpha of alternatives (Section 4.1.2). The expressed beliefs of consultants may themselves be the result of agency frictions, though these frictions would also need to vary across consultants.[27]

7.1.2 Supply-Side Explanations

Up to this point, 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 8 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 in aggregate, 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). In addition, Figure 8 indicates that the rotation by U.S. pensions towards alternatives has far outpaced global supply, resulting in a portfolio that heavily overweights alternatives relative to public markets in aggregate.

7.2 Conclusion

Since the early 2000s, public pensions in the United States have substantially altered the composition of their risky investments, shifting out of public equities and into alternative investments like private equity, real estate, and hedge funds. Explanations based on a desire to take risk, such as those caused by underfunding, 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

[27] In Internet Appendix C.3, we show that the rise in consultant-perceived alpha has occurred through an increase in the expected return of alternatives. The beta and expected excess return to public equities have stayed relatively flat. In addition, consultants have consistently viewed alternatives as riskier than public markets in terms of total volatility. These latter two facts are challenging for the view that increases in consultants' perceived alpha reflect a desire to, say, cater to a pension sector that increasingly wants to hide risk.

equities, some more so than others. This belief-based perspective follows from textbook portfolio theory and helps rationalize the long-run increase in the alternative-to-risky share, both in aggregate and in the cross-section. While our study provides suggests beliefs about alternatives are shaped by consultants, peers, and experience during the 1990s, 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)).

References

Acemoglu, D., U. Akcigit, and W. Kerr (2016). Networks and the macroeconomy: An empirical exploration. *Nber macroeconomics annual 30*(1), 273–335.

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

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

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

Augustin, N., M. Binfarè, and E. Fermand (2024). Benchmarking private equity portfolios: Evidence from pension funds. *Available at SSRN 4590271*.

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.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

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.

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).

Blanchard, O. and L. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 23*(1), 1–76.

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.

Bouwman, C. H. S. and U. Malmendier (2015, May). Does a bank's history affect its risk-taking? *American Economic Review 105*(5), 321–25.

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

Brown, K. C., L. Garlappi, and C. Tiu (2010). Asset allocation and portfolio performance: Evidence from university endowment funds. *Journal of Financial Markets 13*(2), 268–294.

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.

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*.

Dahlquist, M. and M. Ibert (2024). Equity return expectations and portfolios: Evidence from large asset managers. *The Review of Financial Studies 37*(6), 1887–1928.

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

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

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*.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

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

Goetzmann, W. N. and S. Oster (2012, June). Competition among university endowments. Working Paper 18173, National Bureau of Economic Research.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

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.

Herskovic, B., B. Kelly, H. Lustig, and S. Van Nieuwerburgh (2020). Firm volatility in granular networks. *Journal of Political Economy 128*(11), 4097–4162.

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

Ibbotson, R. G. and P. D. Kaplan (2000). Does asset allocation policy explain 40, 90, or 100 percent of performance? *Financial Analysts Journal 56*(1), 26–33.

Ilmanen, A., S. Chandra, and N. McQuinn (2020). Demystifying illiquid assets: Expected returns for private equity.

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

Jenkinson, T., H. Jones, and J. V. Martinez (2016). Picking winners? investment consultants' recommendations of fund managers. *The Journal of Finance 71*(5), 2333–2370.

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

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

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

Korteweg, A. 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.

Leontief, W. (1986). *Input-Output Economics*. Oxford University Press.

Lerner, J., A. Schoar, and J. Wang (2008, September). Secrets of the academy: The drivers of university endowment success. *Journal of Economic Perspectives 22*(3), 207–22.

Lerner, J., A. Schoar, and W. Wongsunwai (2007). Smart institutions, foolish choices: The limited partner performance puzzle. *The Journal of Finance 62*(2), 731–764.

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.

Lo, A. W. (2001). Risk management for hedge funds: Introduction and overview. *Financial Analysts Journal 57*(6), 16–33.

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.

Lucas, D. J. and S. P. Zeldes (2009, May). How should public pension plans invest? *American Economic Review 99*(2), 527–32.

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.

Martinez, J. V. and Y. Qian (2024). Investment consultants in private equity. *Available at SSRN 4322963*.

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.

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

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.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics 37*(2), 187–204.

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. *American 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.

van Binsbergen, J. H. and M. W. Brandt (2016). *Optimal Asset Allocation in Asset Liability Management*, Chapter 8, pp. 147–168. John Wiley & Sons, Ltd.

Figure 1: The Risky Share

(a) Aggregate, Post-2000

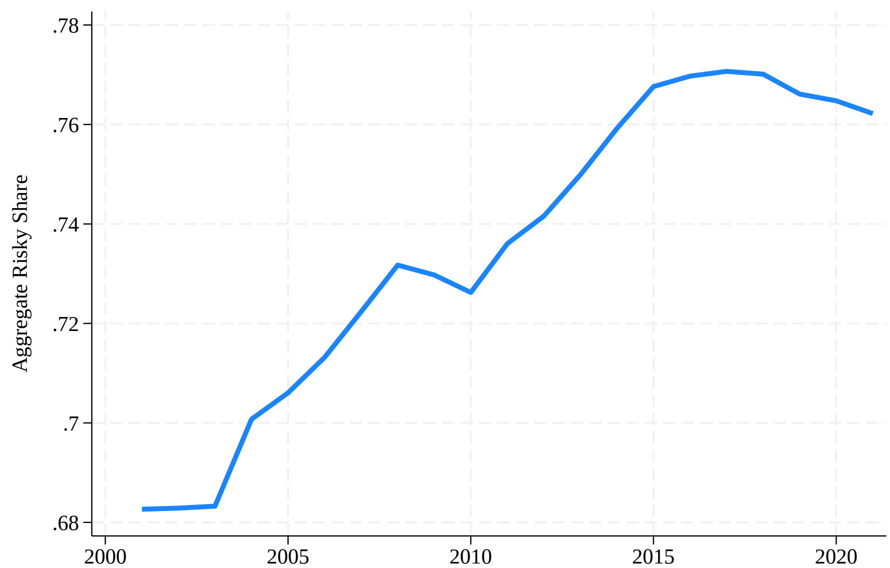

(b) Aggregate, Post-1968

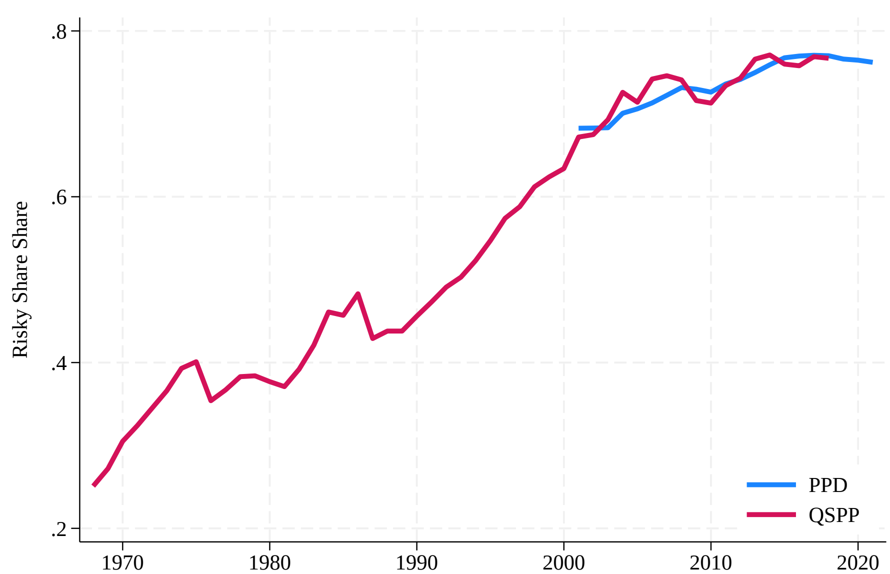

Notes: Panel (a) plots the aggregate target risky share for U.S. public pensions based on data from the PPD, where the risky share is defined as any holding outside of fixed income and cash. Panel (b) adds a longer-history of the risky share using data from the U.S. Census Bureau's Quarterly Survey of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

Figure 2: The Composition of the Aggregate Risky Portfolio

(a) Raw Portfolio Weights

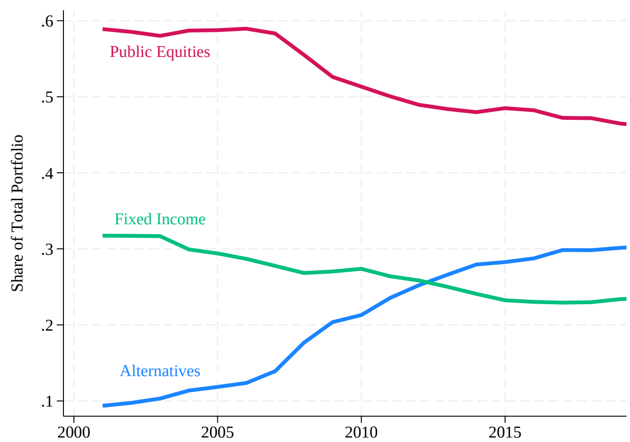

(b) Subcategories of Alternatives

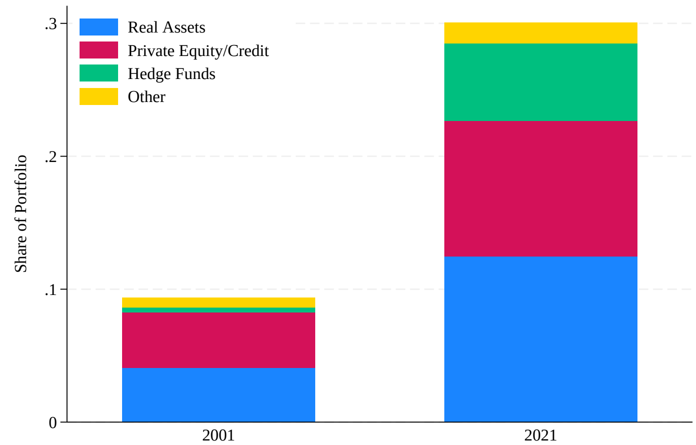

(c) Alternative-to-Risky Share

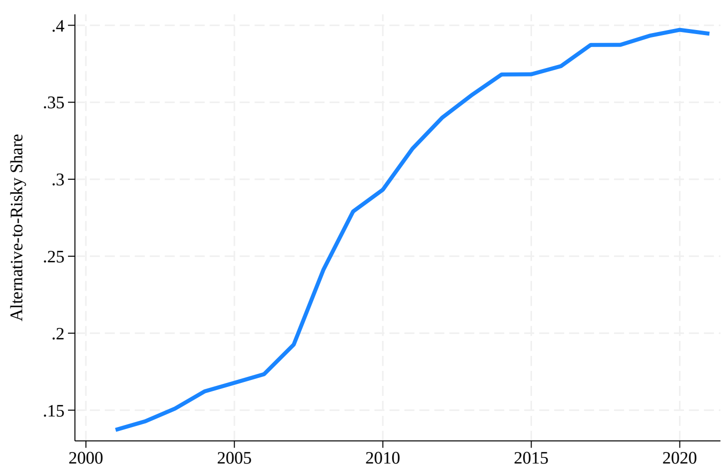

Notes: Panel (a) of the figure plots the target share of public equities, alternatives, and fixed income (including cash) for the aggregate U.S. public pension portfolio. Panel (b) shows the aggregate portfolio shares of different categories of alternatives for 2001 and 2021. Panel (c) plots the share 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

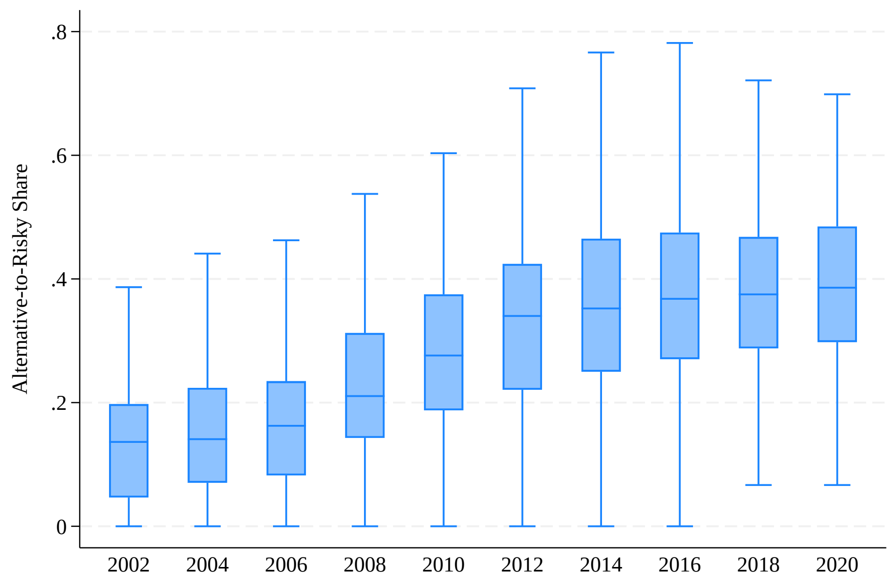

(b) Change from 2002 to 2021

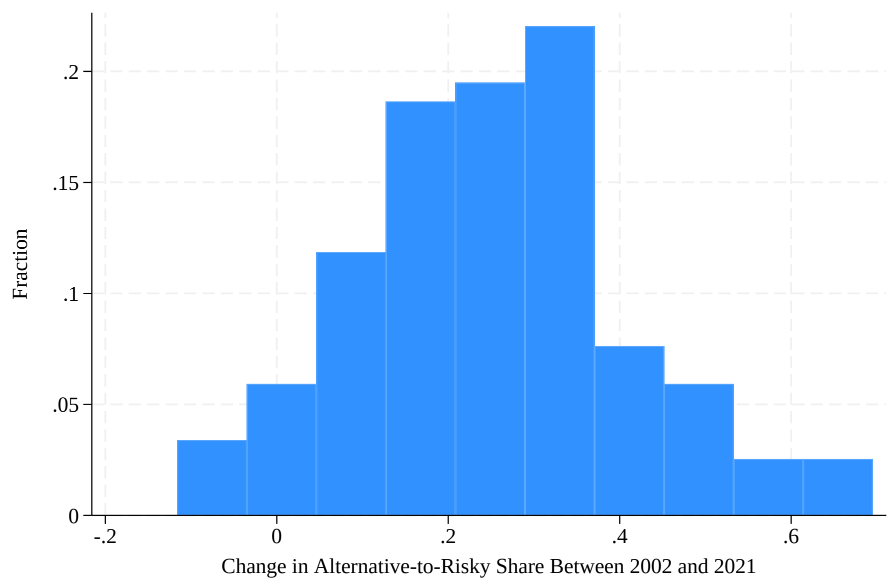

Notes: Panel (a) of the figure depicts the distribution of the alternative-to-risky share across pension systems through time. Each box plot summarizes the distribution for the corresponding year on the *x*-axis. Only even years are plotted to make the graph more readable. Panel (b) plots the distribution 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

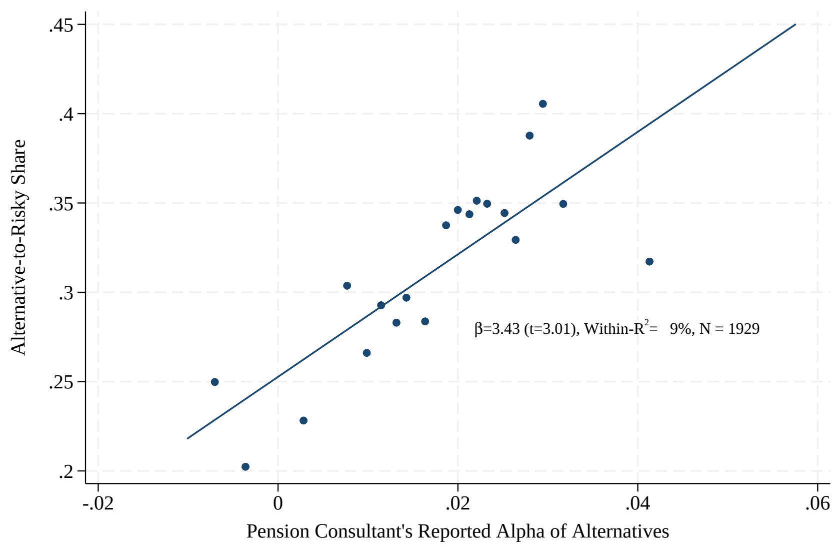

Notes: Each year in Panel (a) shows the median perceived alpha of alternatives relative to large-cap US equities across consultants. The data for the plot are based on hand-collected capital market assumption (CMA) reports that were obtained under confidentiality agreements. Panel (b) is a binscatter 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. The reported t-statistic in panel (b) based on standard errors that are clustered by consultant and pension. See Section 4.1.2 for details.

Figure 5: Consultant Fixed Effects

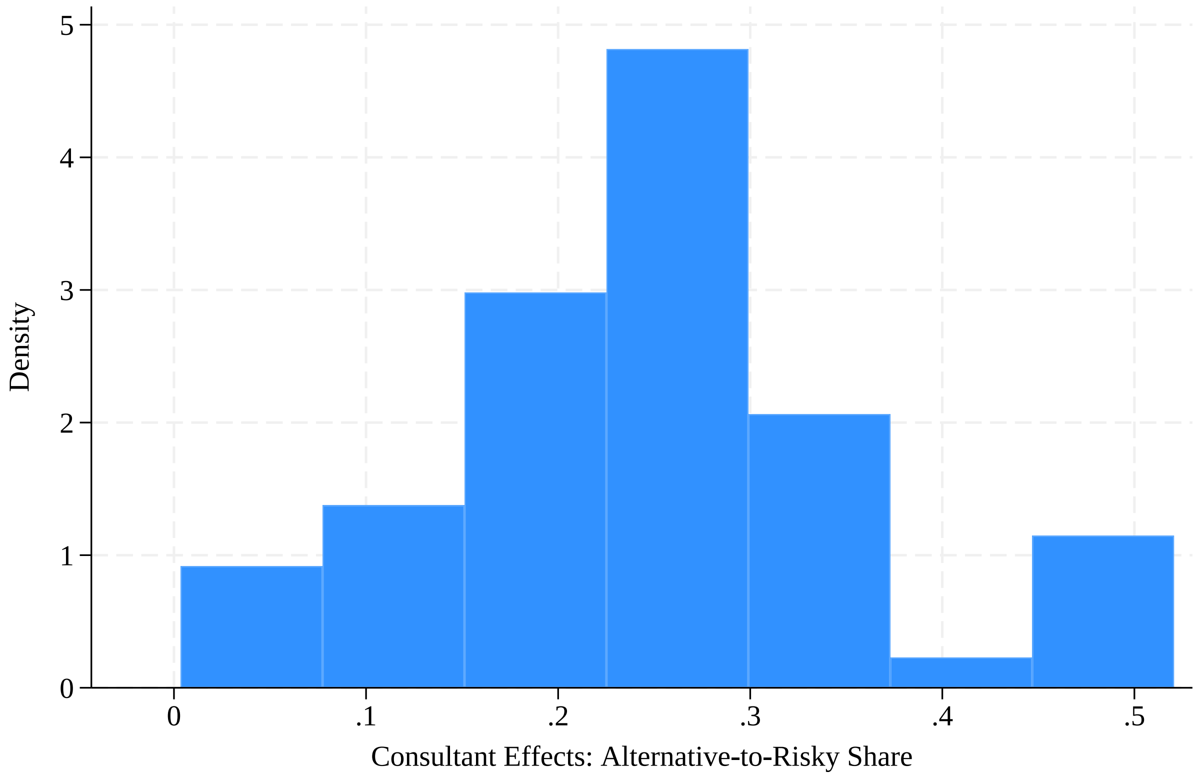

Notes: This figure shows the distribution of consultant effects λ_c estimated from the following regression:

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

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: Alternative Adoption and Risk Constraints

(a) Actual-Target

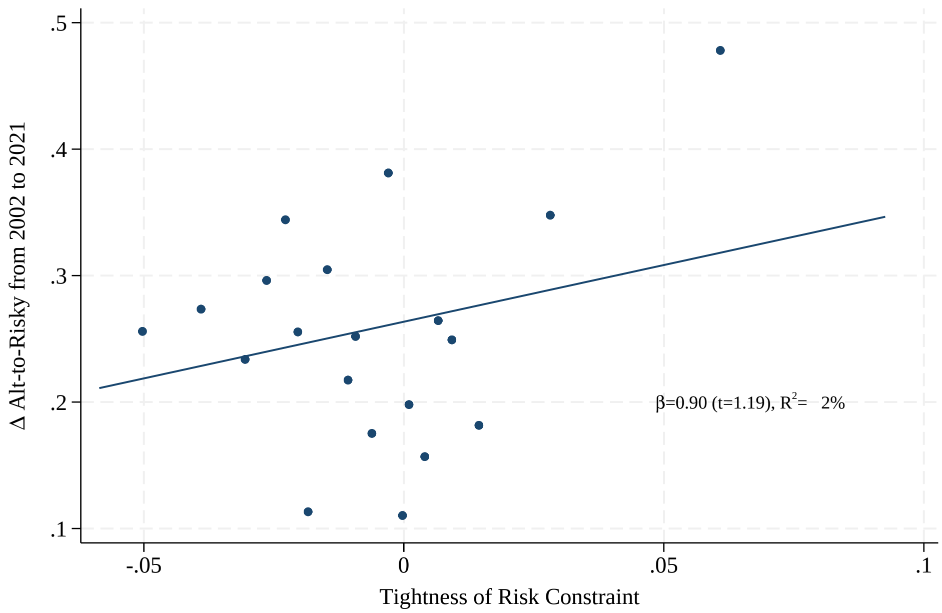

(b) Cash Share

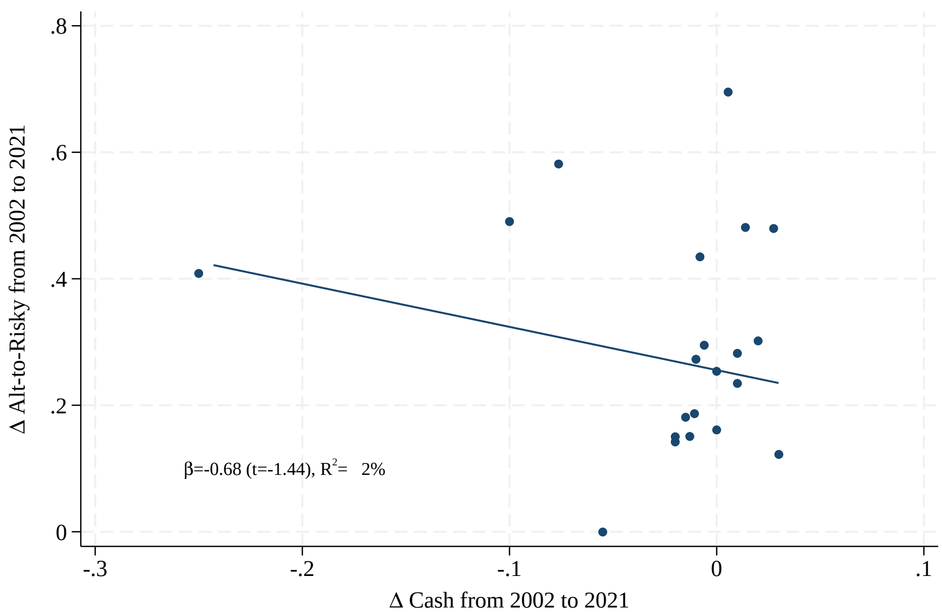

Notes: Panels (a) and (b) of this figure respectively show a binscatter of the change in the alternative-to-risky share from 2002 to 2021 against the average tightness of portfolio constraints over the same period and the change in the target cash share. In panel (a), constraint slack in a given year 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. Both plots reflects only the sample of pensions for which both actual and target shares are available. The OLS regression line is also shown in both plots, with standard errors clustered by state. Data are at the pension-system level. See Section 5.1 for details.

Figure 7: Portfolio Behavior Across Institutions

(a) Alternative-to-Risky Share

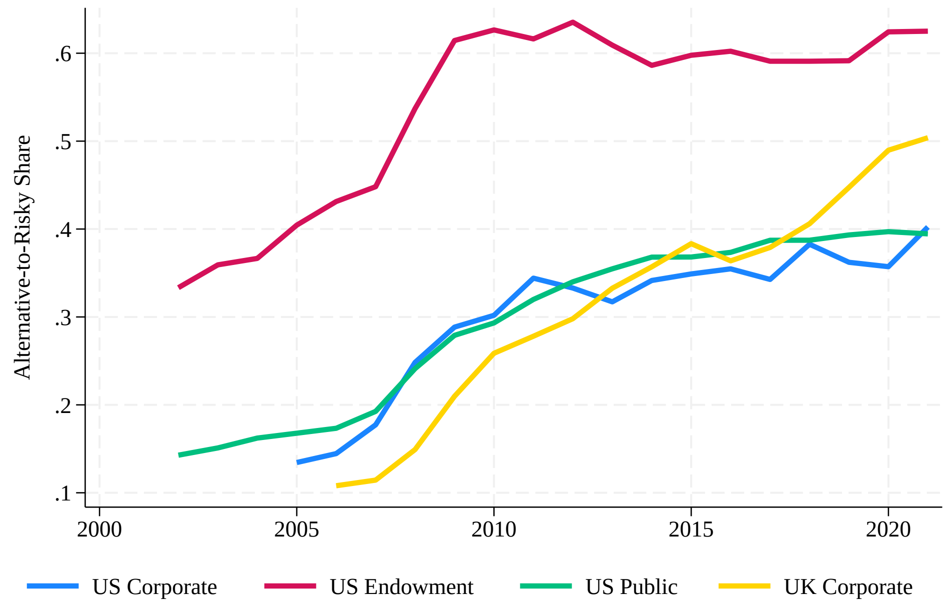

(b) Risky Share

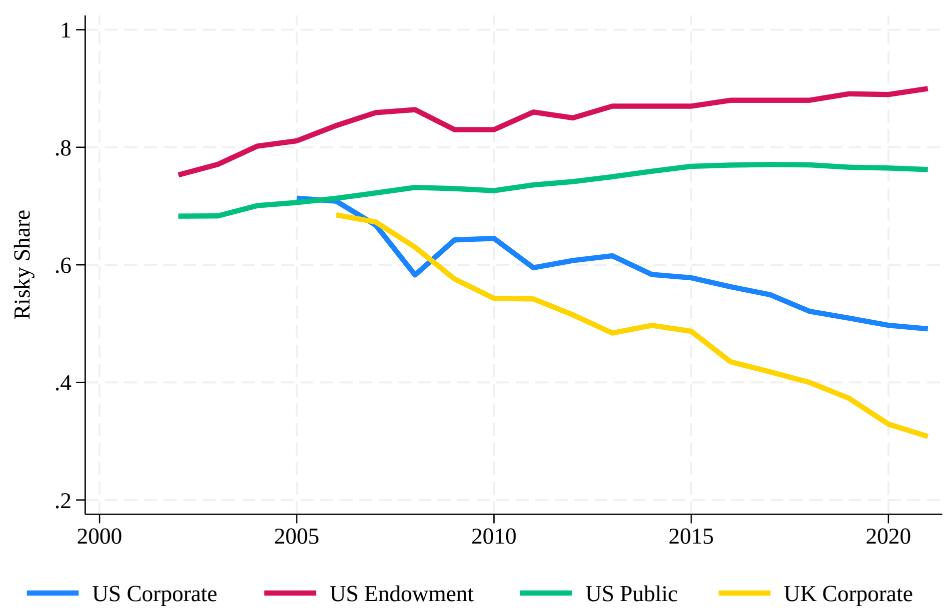

Notes: Panels (a) and (b) of the panel plot the alternative-to-risky and risky shares, respectively, of different institutions through time. U.S. Endowment data is from NACUBO historical endowment studies, U.K. corporate defined-benefit pension data is from the U.K. Pension Protection Authority, and U.S. corporate pension data is from the Milliman Corporate 100 report. See Section 6.3 and Section A of the Internet Appendix for more details on the data.

Figure 8: The Supply of Alternatives

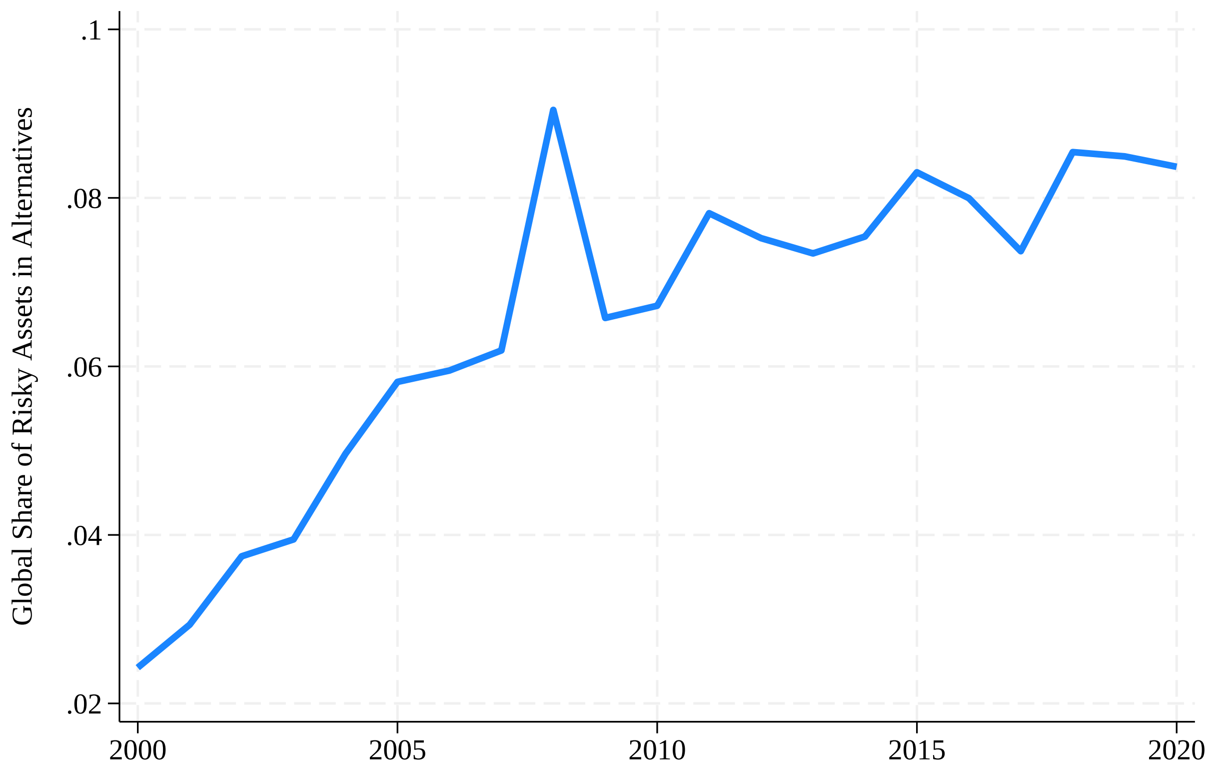

Notes: This figure shows the global supply of alternatives relative to all risky assets. The global stock of alternatives equals the total net asset value of all private-capital funds from Preqin plus the AUM of global hedge funds from Hedge Fund Research. The total stock of risky assets equals the stock 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: Consultant Beliefs and Pension Allocations

	Alternative-to-Risky Share$_{p,t}$			
	(1)	(2)	(3)	(4)
$\alpha_{c(p),t}$	3.43**	3.47**	3.24**	3.93**
	(3.01)	(3.21)	(3.40)	(3.94)
Estimation Type	OLS	OLS	IV	IV
Controls	No	Yes	Yes	Yes
Fixed-Effects	Time	Time	Time	Time
$p_{wild}(\beta = 0)$	0.01	0.01	0.02	0.01
IV Location Date			t	2005
First-Stage F			18.77	10.73
β_{IV}/β_{OLS}			0.93	1.12
Within-R^2	0.09	0.13		
N	1,929	1,915	1,915	1,616

Notes: This table summarizes linear regressions of pension p's target alternative-to-risky share in year t on the reported alpha of alternatives of pension p's consultant in year t. Pension controls include the level of pension funding, (log) size, hurdle rate, the ratio of required actuarial contributions too payroll, the ratio of administrative expenses to payroll, the cash share, annual return of the fund, and the residualized difference between each pension's target and actual risky share. Standard errors are clustered at the consultant and year levels, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10.

Table 3: Consultants and the Composition of Risky Investments

			Fixed Effects							
	Dep. Variable:	Controls	Time	State-Time	Consultant	F	p	Adj. R^2	C	N
(1)	Alternatives	x	x					0.33		3,087
(2)	Alternatives	x	x		x	18.32	0.00	0.50	57	3,082
(3)	Alternatives	x		x				0.50		2,689
(4)	Alternatives	x		x	x	12.85	0.00	0.62	54	2,683
(5)	Private Equity/Credit	x	x					0.20		3,087
(6)	Private Equity/Credit	x	x		x	16.71	0.00	0.38	57	3,082
(7)	Private Equity/Credit	x		x				0.32		2,689
(8)	Private Equity/Credit	x		x	x	12.80	0.00	0.49	54	2,683
(9)	Hedge Funds	x	x					0.15		3,087
(10)	Hedge Funds	x	x		x	11.46	0.00	0.29	57	3,082
(11)	Hedge Funds	x		x				0.30		2,689
(12)	Hedge Funds	x		x	x	9.22	0.00	0.43	54	2,683
(13)	Real Assets	x	x					0.16		3,087
(14)	Real Assets	x	x		x	16.36	0.00	0.35	57	3,082
(15)	Real Assets	x		x				0.30		2,689
(16)	Real Assets	x		x	x	13.26	0.00	0.47	54	2,683

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

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

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

Table 4: Experience in the 1990s

	Δ Alt-to-Risky$_{2002\to 2021}$				Δ Risky$_{2002\to 2021}$
	(1)	(2)	(3)	(4)	(5)
Δ Equity$_{1996\to 2002}$	0.46**	0.37**	0.60**	0.58**	0.16
	(2.93)	(2.69)	(2.53)	(2.38)	(1.16)
Δ Equity$_{1990\to 2002}$			-0.21	-0.13	-0.11
			(-0.85)	(-0.43)	(-0.57)
Fraction of Restricted Assets$_{1998}$				0.07	0.05
				(0.91)	(0.76)
1996 Share Type	Actual	Target	Actual	Actual	Actual
Controls	No	No	No	Yes	Yes
R^2	0.19	0.15	0.20	0.25	0.39
N	45	45	45	44	44

Notes: This table summarizes linear regressions of state-level portfolio composition changes between 2002 and 2021 on prior changes in the public-equity portfolio share:

$$\Delta\omega^*_{s,2002\to 2021} = c + \beta_1 \Delta\omega^E_{s,1996\to 2002} + \beta_2 \Delta\omega^E_{s,1990\to 2002} + \Gamma X_s + \varepsilon_s.$$

In columns (1)-(4) the dependent variable is the state-level change in the alternative-to-risky share and in column (5) it is the change in the risky share, both computed using target shares in PPD between 2002 and 2021. All regressions include the change in the share of public equities between 1996 and 2002 ($\Delta\omega^E_{s,1996\to 2002}$), computed using PENDAT data from 1996 and target shares from PPD in 2002. In column (2), the target share of public equities in 1996 is used to construct $\Delta\omega^E_{s,1996\to 2002}$ and the actual share is used in all other columns. Columns (3)-(5) also include the change in the equity share between 1990 and 2002 ($\Delta\omega^E_{s,1990\to 2002}$), computed using 1990 actual shares in PENDAT and 2002 target shares in the PPD. Target shares in 1990 are not available in PENDAT and are less populated than actual shares at the pension level in 1996. Columns (4)-(5) add the fraction of assets that are restricted by each state's constitution as of 1998. The vector of controls X_s includes the change in hurdle rates, the change in the fraction of retired members, the level of BEA-adjusted funding in 2002, and the log of assets under management in 2002, with all changes measured between 2002 and 2021. Robust standard errors are used in all regressions, and t-statistics are listed below point estimates in parentheses. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. See Section 4.2 for more details on the data.

Table 5: Peer Effects

	Alternative-to-Risky Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Peers' Alt-to-Risky Share	0.69**	0.61**	0.70**	0.73**		
	(3.11)	(3.46)	(3.15)	(3.34)		
× Established-CIO		0.01				
		(0.04)				
× Well-Funded			-0.14			
			(-1.18)			
× High-Performing				-0.26**		
				(-2.44)		
Lagged Peers' Alt-to-Risky Share					0.71**	
					(3.03)	
Non-Pension Peers' Alt-to-Risky Share						0.41**
						(2.20)
FE	(c,y,d)	(c,y,d)	(c,y,d)	(c,y,d)	(c,y,d)	(c,y)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.14	0.14	0.15	0.16	0.15	0.02
Total R^2	0.68	0.59	0.68	0.69	0.67	0.45
N	2,039	900	2,039	2,039	1,915	2,467

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. In columns (1)-(5), 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. In column (6), the alternative-to-risky share of each pension's peers is defined as the average alternative-to-risky share of unions, private-sector pensions, and endowments in its state. Columns (1)-(5) include consultant-by-time-by-census-division fixed effects and column (6) includes consultant-by-year fixed effects. All regressions include controls for funding ratio, asset hurdle rate, (log) size, required actuarial contributions relative to payroll, and administrative expenses relative to payroll. Columns (1) to (4) show the baseline relationship between the alternative-to-risky share for pension p in year t and its peers' share, as well as the interaction with several characteristics. In all cases, the regression includes both the characteristic and its interaction with peer target share of alternatives. The indicator variable "Established-CIO" in column (2) equals one if the pension has had the same CIO for at least five years. The regression in column (2) is estimated only using data after 2005. The indicator for "Well-Funded" in column (3) equals one if the pension is in the top quartile of funding for the entire sample. The indicator for "High-Performing" in column (4) equals one if the pension is in the top 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 6: Risk-Seeking Explanations for the Rise in Alternatives

	ΔAlternative-to-Risky Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ GASB 25 Funding Ratio	-0.10				-0.09		
	(-1.13)				(-1.02)		
Δ BEA-Adjusted Funding Ratio		-0.07					
		(-0.34)					
Δ Liability Discount Rate			-0.38			-1.14	
			(-0.12)			(-0.32)	
Δ Fraction of Retired Members				0.13			0.12
				(0.75)			(0.66)
Aggregation	System	State	System	System	System	System	System
Constraint Control	No	No	No	No	Yes	Yes	Yes
Total R^2	0.01	0.00	0.00	0.00	0.02	0.01	0.00
N	118	47	117	115	116	115	113

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 2021. Risky investments are any holdings outside of fixed income and cash. Alternatives encompass investments in hedge funds, private equity, private debt, and real assets. GASB 25 funding ratios are based on liabilities that equal future promised benefits discounted at the assumed long-term rate of return for each plan. BEA-adjusted funding ratios instead discount future benefits using AAA-rated corporate borrowing rates. The liability discount rate is the one used for computing GASB 25 funding ratios and is also the system's return hurdle rate. The row labeled "Aggregation" specifies whether the regression is run at the system or state level. The row "Constraint Control" indicates whether controls are included for the average difference between each pension's actual and its target risky share between 2002 and 2020, as well as the change in their target cash share over the same period. Standard errors are clustered by state for regressions run at the system level and are robust for regressions run at the state level. t-statistics are listed below point estimates. ** indicates a p-value of 0.05 and * indicates a p-value of 0.10. See Section 2.1 for how we filter the data and Section 5.1 for more details on how we define the tightness of risk constraints.

Table 7: 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 6.1, 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 6.1 for more details.