Subjective Beliefs and the Portfolio Allocations of Institutional Investors

Aleksandar Andonov[*] Spencer J. Couts[†] Andrei S. Gonçalves[‡]

Johnathan A. Loudis[§] Andrea Rossi[¶]

This Version: December 2025[‖]

Abstract

We study how subjective beliefs shape the portfolio allocations of institutional investors. Linking the multi-asset allocations of U.S. public pension funds to the long-term capital market assumptions of their consultants, we examine the extent to which differences in subjective expected returns, volatilities, and correlations map into differences in portfolio weights. We embed these belief inputs in a mean-variance framework that incorporates fund-consultant belief wedges, heterogeneous risk aversion, non-negative weight constraints, and a benchmarking incentive due to frictions. We find that pension fund allocations are significantly linked to belief-implied mean-variance efficient allocations across pension funds, across asset classes, and over time. Accounting for frictions is essential: it dramatically increases the pass through and explanatory power of beliefs to portfolio allocations. Overall, our results show that beliefs play a central role in institutional portfolio decisions and that frictions critically shape their transmission into observed allocations.

JEL Classification: G10, G11, G12, G23

Keywords: Subjective Beliefs, Portfolio Allocations, Institutional investors, Pension Funds, Consultants

[*]Amsterdam Business School, University of Amsterdam. E-mail: a.andonov@uva.nl.

[†]Price School of Public Policy, University of Southern California. E-mail: couts@usc.edu.

[‡]Fisher College of Business, The Ohio State University. E-mail: goncalves.11@osu.edu.

[§]Mendoza College of Business, University of Notre Dame. E-mail: jloudis@nd.edu.

[¶]Eller College of Management, University of Arizona. E-mail: rossi2@arizona.edu.

[‖]We are very grateful to the institutions that shared data on Capital Market Assumptions with us. This paper benefited a lot from conversations with members of these institutions and from helpful comments from Francesca Bastianello, Thummim Cho, Ricardo Delao, Lukas Kremens, Eben Lazarus, Sean Myers, Felix Wilke as well as seminar participants at the 2025 Annual Valuation Workshop, University of Arizona, University of Amsterdam, and Nova School of Business and Economics.

Introduction

Investors' risk and return expectations are fundamental determinants of portfolio allocations and, ultimately, of asset prices (see Adam and Nagel (2023) and Giglio et al. (2025)). Specifically, in canonical portfolio choice models, investors allocate wealth to maximize expected utility given their subjective beliefs about the distribution of future returns. Consequently, linking return beliefs to portfolio allocations provides a natural way to study how beliefs are translated into investment decisions and, through them, into asset prices.

A growing empirical literature examines this link by regressing investors' equity allocations onto their subjective expected equity returns (e.g., Giglio et al. (2021) and Dahlquist and Ibert (2024)). While intuitive, this approach faces three important limitations. First, by focusing on equities, it overlooks the multi-asset nature of typical portfolios, where optimal allocations depend on the joint distribution of returns across asset classes. Second, it abstracts from key determinants of optimal portfolio choice, including subjective volatilities, perceived diversification benefits, and frictions in the institutional and informational environment. And third, even though an increasingly large share of global assets is managed by institutional investors, most existing evidence is based on retail investors (we detail exceptions later).

We address these limitations by studying how subjective beliefs shape the portfolio allocations of institutional investors across major asset classes. Our analysis combines the allocations of U.S. public pension funds with the long-term capital market assumptions (CMAs) of their general investment consultants, which are firms hired to advise these funds on strategic asset allocation (Begenau, Liang, and Siriwardane (2025)). This unique linkage allows us to directly observe how beliefs about expected returns, volatilities, and cross-asset correlations are reflected in portfolio choices. Moreover, using a mean-variance framework that accounts for fund-consultant belief wedges, we quantify the extent to which variation in institutional portfolios can be explained by variation in beliefs in the presence of frictions.

Three main findings stand out. First, even ignoring frictions, there is typically a positive and statistically significant link between pension fund allocations and the mean-variance

efficient (MVE) allocations implied by the beliefs of their consultants once we account for the perceived diversification benefits of different asset classes. The explanatory power of this relationship is non-trivial, but the estimated pass-through from MVE allocations to pension fund allocations is almost two orders of magnitude smaller than the theoretical ideal value of one. This low pass-through arises in part because unconstrained MVE allocations tend to imply extreme portfolio positions that differ substantially from those taken by pension funds. Second, the link between pension fund allocations and MVE allocations strengthens considerably once we incorporate non-negative weight constraints and a benchmarking incentive (which we show can emerge from institutional or informational frictions). In particular, MVE allocations better align with actual allocations in levels and their explanatory power and slope coefficients increase sharply, especially when we allow for risk-aversion heterogeneity. Third, under our full specification, which accounts for constraints, benchmarking, and risk aversion heterogeneity, between 48% and 88% of the variation in pension fund allocations across pension funds, across asset classes, and over time can be explained by variation in MVE allocations. Overall, our results demonstrate that risk and return expectations play a central role in institutional portfolio decisions, but isolating their effect requires recognizing the constraints and frictions that attenuate their transmission to observed allocations.

Unpacking our Analysis

We observe the annual portfolio allocations of U.S. public pension funds from 2001 to 2023 for six asset classes (cash, fixed income, public equity, private equity, real estate, and hedge funds). Because pension fund beliefs are unobserved, we proxy for them using the CMAs of their general consultants. Taking cash as the baseline asset, we use each consultant-year CMA to construct subjective expected excess returns (μ) and covariance matrices of excess returns (Σ). Combined with a mean-variance framework that allows for fund-consultant belief wedges, these belief inputs generate MVE allocations that vary by consultant, asset class, and year. We construct pension fund allocations at the same level of variation by aggregating, for each consultant-year, the allocations of all pension funds advised by that consultant.

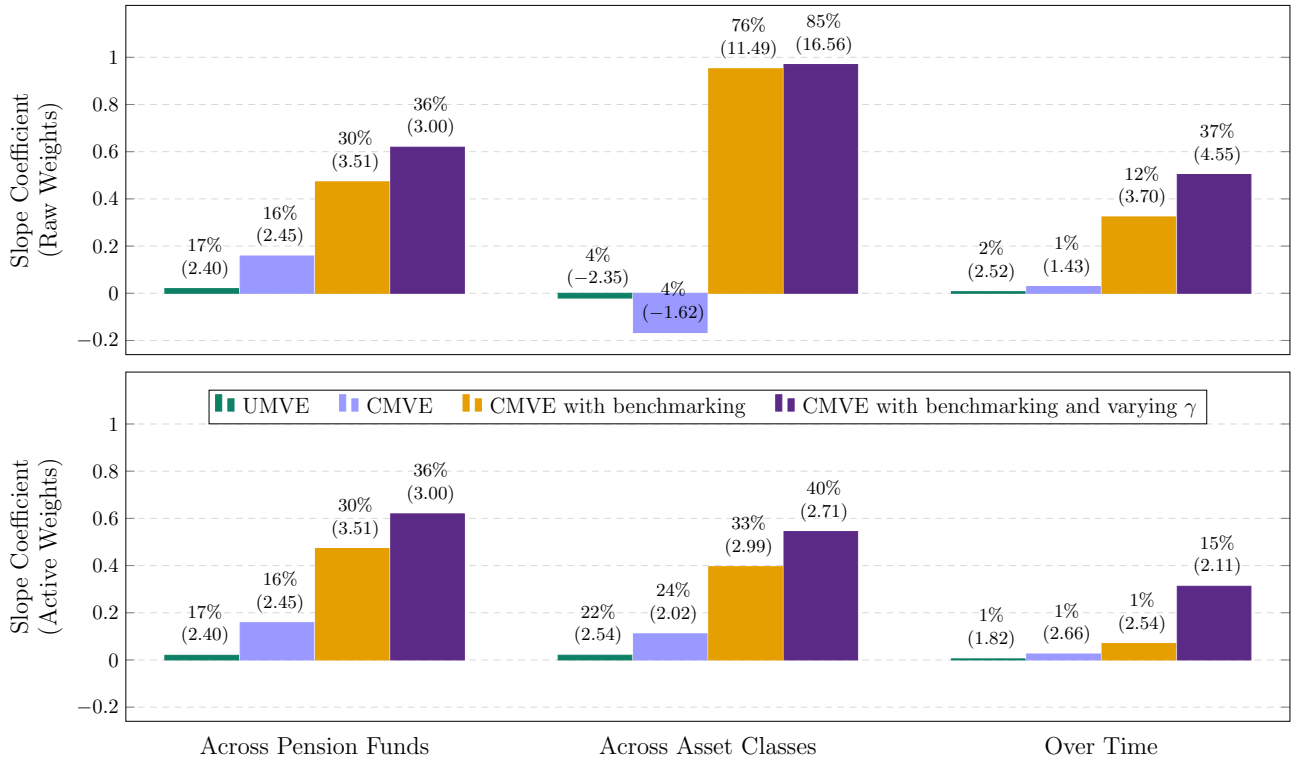


Figure 1
Regressing Pension Fund Allocations on MVE Allocations

The figure reports slope coefficients (with t_{stat} and R^2_{within}) from regressions of pension fund allocations (ω) on MVE allocations (ω^*), with active weights reflecting deviations from the benchmark. UMVE is unconstrained, CMVE imposes non-negative weights, and "with benchmarking" adds a benchmark incentive to the mean-variance problem. Risk aversion is fixed ($\gamma = 3$) or estimated to vary across funds and over time. Regressions use fixed effects to focus on variation across pension funds (at the consultant level), across asset classes, or over time. All numbers come from Table 4, with the table header providing more details.

Our main analysis regresses pension fund allocations (ω) onto the corresponding MVE allocations (ω^*). Figure 1 summarizes the results. The upper panel reports regressions using raw portfolio weights, while the lower panel uses active weights, defined as deviations from benchmark weights (constructed from aggregate pension fund allocations from the previous year). For variation across pension funds, the results using raw and active weights coincide because benchmark weights do not vary across pension funds.

The UMVE specification in Figure 1 starts from an unconstrained mean-variance solution without frictions. In this case, there is a positive and statistically significant link between ω and ω^* across all dimensions (except for raw weights across asset classes). Moreover, some R^2

values are relatively high, a result that we later show is mainly due to perceived diversification benefits, as R^2 values become much weaker if we set Σ to a diagonal matrix. However, the slope coefficients are small (almost two orders of magnitude below the theoretical ideal of one). This result is partially due to unconstrained MVE allocations varying far more than pension fund allocations, which compresses the estimated pass-through from ω^* to ω.

The CMVE specification in Figure 1 imposes non-negative weight constraints, leading to MVE allocations with more realistic variation. As a result, the slope coefficients rise across all three dimensions of variation, although the improvements are relatively modest for variation over time. Nevertheless, the constrained model still results in relatively small slope coefficients. Moreover, constrained MVE weights completely fail to explain differences in pension fund allocations across asset classes unless we control for benchmark allocations.

The next specification in Figure 1 incorporates a benchmarking incentive, which we show can arise from institutional frictions (a tracking error constraint) or informational frictions (a Black and Litterman (1991, 1992) approach to deal with μ uncertainty). We find that it substantially strengthens the empirical link between pension fund allocations and MVE allocations. In particular, benchmarking sharply increases slope coefficients and R^2 values for variation across pension funds, across asset classes, and over time. Moreover, even when we control for benchmark weights through the use active weights, a strong positive link remains. This indicates the improvement in fit is not solely driven by the inclusion of benchmark weights into ω^* but instead reflects a better mapping from beliefs to portfolio choices.

The last specification in Figure 1 also incorporates variation in risk aversion (across consultants and over time), capturing differences in how pension funds translate beliefs into portfolio allocations. Analogous to Couts et al. (2025), we estimate risk aversion for each consultant-year from regressions of expected asset class returns onto return covariances between asset classes and the portfolio of pension funds matched to the respective consultant-year (with some adjustments for our non-negative weight constraints and benchmarking incentive). This final specification yields the closest alignment between ω and ω^*, with the highest slope coefficients and explanatory power across all three sources of variation.

Overall, Figure 1 shows that pension fund allocations are systematically related to the belief-implied MVE allocations of their consultants. Moreover, accounting for perceived diversification benefits, non-negative weight constraints, benchmarking incentives, and heterogeneity in risk aversion is essential for uncovering the strength of this link. These results are highly robust to changes in our baseline empirical procedures. For instance, they continue to hold if we estimate the regressions using fund-level allocations or if we directly control for benchmark weights (instead of using active weights) or for fund characteristics studied in the prior literature, which partially capture variation in fund incentives that could affect allocations other than through a mean-variance channel (e.g., a plan's funding status).

Since we use consultant CMAs to construct ω^*, the economic interpretation of our results requires pension fund beliefs to correlate with the CMAs of their consultants.[1] To verify this assumption, we compare consultant μs with the actuarially assumed rates of return (μ_{AARR}) public pension funds are required to report annually since 2014. As discussed in Andonov and Rauh (2022), these μ_{AARR} values capture beliefs as well as incentives since they are used to discount pension liabilities. We find a strong positive correlation between μ_{AARR} values and consultant μs, and two facts suggest this correlation arises from shared beliefs rather than reporting incentives. First, consultants distribute identical CMAs to a heterogeneous client base, with public pension funds representing only 17.3% of the clients and 7.2% of the regulatory assets under advisory for the average consultant. So, CMAs are unlikely to reflect the incentives of public pension funds. In fact, consultant μs are lower than the μ_{AARR} values reported by their pension fund clients, showing consultants do not cater to pensions' incentive to inflate liability discount rates. Second, CMAs are based on transparent economic and statistical analyses rather than opaque processes that would facilitate the incorporation of client incentives (Andolfatto and Bastianello (2025)). Indeed, μs aggregated across consultants provide unbiased predictors of future asset class returns (Couts, Gonçalves, and Loudis (2024)), inconsistent with CMAs reflecting client incentives.

[1] We interpret the link between ω and ω^* as pension fund allocations responding to their own beliefs. This interpretation holds whether consultant CMAs and pension beliefs correlate causally (e.g., CMAs affect pension fund beliefs) or not (e.g., pension funds select consultants that share their belief formation process).

Importantly, constructing ω^* using consultant CMAs helps mitigate the reverse causality concern that beliefs are driven by pension funds' desired allocations. As discussed above, consultant μs are unbiased predictors of future returns, and pension funds represent only a small fraction of the clients and regulatory assets under advisory of their consultants. As such, consultant CMAs are not formed to justify pension funds' desired allocations. Moreover, on average, fund-consultant relationships last 16 years, with the link between ω and ω^* becoming stronger (not weaker) as relationships mature. So, our results are also not driven by pension funds ex post selecting consultants whose beliefs justify their desired allocations.

However, our use of consultant CMAs to proxy for pension fund beliefs induces belief wedges that lead to a downward bias in the slope estimates from regressions of ω onto ω^*, implying the reported R^2 values understate the fraction of ω variability explained by ω^*. So, we also perform a variance decomposition for $\omega = \omega^* + \omega^\alpha$, where ω^α are MVE allocation wedges constructed from the difference between ω and ω^*. We find that a large share of the variation in pension fund allocations is explained by MVE allocations under our full specification that incorporates non-negative weight constraints, benchmarking incentives, and varying risk aversion. For raw weights, the shares vary from 58% (across pension funds) to 88% (across asset classes) whereas for active weights they vary from 48% (over time) to 74% (across asset classes). So, MVE allocations go a long way in explaining pension fund allocations once we account for institutional and informational frictions.

Finally, we also directly study the MVE allocation wedges, ω^α. We quantify their magnitudes through the annual fee a pension fund would be willing to pay to move from the allocation ω to ω^* (if the pension fund had our mean-variance utility and the same beliefs as its consultant). The median annual fee is around 0.15% for our full specification, which is a little less than one third of the average aggregate investment management cost borne by pension funds annually over our sample period. So, the MVE allocation wedges we identify are economically meaningful, though perhaps smaller than some might have expected. We also show that these allocation wedges can be at least partially accounted for by variables that the prior literature links to pension-fund incentives and other considerations.

Contribution to the Literature

Our primary contribution is to the subjective beliefs literature in asset pricing (surveyed by Adam and Nagel (2023) and Giglio et al. (2025)).[2] We advance this literature by jointly addressing three major challenges highlighted in Adam and Nagel (2023): (i) the need for more evidence linking investors' expectations to their portfolio decisions; (ii) the substantial heterogeneity in subjective beliefs across different groups of market participants; and (iii) the importance of subjective risk perceptions alongside perceived expected returns. We tackle these issues by studying the link between the subjective beliefs and portfolio allocations of institutional investors (public pension funds) whereas most existing work focuses on individual investors.[3] Moreover, we explicitly incorporate subjective risk perceptions within a mean-variance framework, which is new in this literature. This framework then allows us to quantify the importance of two central frictions (non-negative weight constraints and benchmarking incentives) for the pass-through of beliefs to portfolio allocations.

A small but growing set of papers also examines the link between institutional investors' subjective beliefs and their portfolio choices. Dahlquist and Ibert (2024), Andolfatto and Bastianello (2025), and Ibert (2025) study mutual funds investing in equities and fixed income, with the first two emphasizing the role of long-term expected returns and their components, and the latter highlighting short-term tactical views. In addition, Bastianello and Peng (2025) use Bank of America surveys of global fund managers to document a time-series relation between perceived mispricing and stated portfolio allocations. For pension funds, Andonov and Rauh (2022) examine how expected returns relate to cross-sectional differences in allocations, while Begenau, Liang, and Siriwardane (2025) show that pension funds' allocations to alternative asset classes reflect the CAPM alphas implied by the CMAs of their consultants.

[2]While most of this literature focuses on individual investors, several recent papers study the properties of the subjective beliefs of financial analysts and institutional investors (e.g., Delao and Myers (2021), Bastianello (2024), Couts, Gonçalves, and Loudis (2024), Couts et al. (2025), Dahlquist and Ibert (2024, 2025), Décaire and Graham (2024), Delao, Han, and Myers (2024), Kremens, Martin, and Varela (2024), Loudis (2024), Andolfatto and Bastianello (2025), and Bastianello, Décaire, and Guenzel (2025)).

[3]Some recent papers studying the pass-through of subjective beliefs to the portfolio allocations of individual investors are Giglio et al. (2021), Beutel and Weber (2023), and Egan, MacKay, and Yang (2024).

Our work fundamentally differs from these studies. We integrate CMAs into a full mean-variance framework to directly quantify the link between pension fund allocations and subjective MVE allocations, rather than isolating the influence of individual components of beliefs. Crucially, our framework incorporates not only subjective risk perceptions (whose importance we document) but also institutional and informational frictions in the form of non-negative weight constraints and benchmarking incentives, both of which must be modeled to fully understand the pass-through of beliefs to portfolio allocations. We further exploit the multi-asset-class nature of pension fund portfolios to examine the link between beliefs and portfolio allocations across funds, across asset classes, and over time. This structured analysis reveals that while beliefs play a central role in institutional portfolio decisions, recovering their effect requires explicitly accounting for the constraints and frictions that shape how beliefs translate into observable portfolio positions.

Our work also has important implications for several literatures beyond subjective beliefs. Most notably, it offers direct evidence on the quantitative importance of the type of institutional frictions emphasized in the limits to arbitrage literature (see Gromb and Vayanos (2010) for a review). Our results show that the presence of such frictions means that asset demand appears considerably more inelastic than what would be implied by models in which beliefs translate frictionlessly into portfolio allocations. This insight naturally connects our findings to the mechanisms emphasized in the demand-based asset pricing literature, where inelastic demand plays a central role in explaining price sensitivity to flows and the volatility of asset prices (see, e.g., Koijen and Yogo (2019) and Gabaix and Koijen (2022)).

The rest of this paper is organized as follows. Section 1 introduces the data on pension fund allocations and consultant CMAs. Section 2 details our mean-variance framework and the resulting MVE allocations. Section 3 presents the results from regressions of observed pension fund allocations onto MVE allocations. Section 4 decomposes the variability of pension fund allocations into the contributions of MVE allocations and allocation wedges. Section 5 examines these allocation wedges in greater detail. Section 6 concludes. The Internet Appendix contains additional data details as well as supplementary results and derivations.

1 Data and Descriptive Statistics

Subsections 1.1 and 1.2 cover our data on pension funds and the capital market assumptions (CMAs) of their consultants, respectively. In turn, Subsection 1.3 provides summary statistics for portfolio allocations and belief elements and Subsection 1.4 studies the link between the expected returns of pension funds and those of their consultants.

1.1 Dataset of Pension Funds

Our dataset of U.S. public pension funds covers the period from 2001 to 2023 and combines information on funds' assets under management (AUM), portfolio allocations, and general consultants.[4] Following Andonov, Bonetti, and Stefanescu (2025), we start from the Public Plans Database (PPD), published by the Center for Retirement Research at Boston College, which provides data on 140 pension funds (2,949 fund-years). We then extend the PPD sample by hand-collecting data on 41 unique local county and city pension funds (704 fund-years) that are not part of the PPD dataset.[5] Our combined pension fund dataset contains 181 unique funds for a total of 3,653 fund-years. General consultants are present in 94.9% of these fund-years, and each fund-year is associated with at most one general consultant.

Following Begenau, Liang, and Siriwardane (2025), we measure pension fund allocations using their target allocations when available and their actual allocations when target allocations are not available (which is only the case for 2.2% of the fund-years used in our baseline analysis). The motivation for giving priority to target allocations is that they better reflect active decisions by pension funds since targets are unaffected by market movements whereas

[4]Following the prior literature (e.g., Andonov, Bonetti, and Stefanescu (2025) and Begenau, Liang, and Siriwardane (2025)), we define pension funds by aggregating multiple defined benefit plans to their retirement system level. For instance, the New Jersey Division of Investment is a single pension fund in our dataset, but it is responsible for the asset management of multiple define benefit plans, including the New Jersey PERS, the New Jersey Police & Fire, and the New Jersey Teachers.

[5]The hand-collected data on pension funds is based on their Comprehensive Annual Financial Reports (CAFRs), websites, and Pensions & Investments newswire articles. We use these data not only to supplement the PPD dataset, but also to verify and adjust the PPD data on asset allocations and general consultants. In particular, we apply the adjustments outlined in Andonov, Bonetti, and Stefanescu (2025) and Begenau, Liang, and Siriwardane (2025), which cover a range of issues, including misclassified asset categories.

actual weights vary with market movements given typical rebalancing policies. However, Internet Appendix C.5 provides consistent results using only actual pension fund allocations. Hereafter, we omit the terms "target" and "actual" except where they are essential to the discussion.

We classify the portfolio allocations to the six asset classes we study: cash (Cash), fixed income (FI), public equity (EQ), private equity (PE), real estate (RE), and hedge funds (HF). In all cases, the sum of the allocations to these six asset classes covers the vast majority of the fund AUM for the typical fund-year observation (e.g., in our baseline analysis, portfolio weights sum to 97.8% on average). The remaining assets that are not classified into one of these six asset classes account for, on average, 2.2% of pension assets. To ensure allocations add to exactly 100%, we divide each allocation by the sum of the allocations to the asset classes present in our analysis within each fund-year.

We also collect a few other fund-level variables beyond general consultant identity, AUM, and portfolio allocations. However, since these variables are only used in some parts of our analysis, we do not require them to be available for a fund-year observation to enter our sample. Accordingly, we introduce these variables in the subsections where they are used.

1.2 Dataset of Consultant Capital Market Assumptions (CMAs)

Our analysis requires linking pension fund portfolio allocations to beliefs. Since we do not directly observe the beliefs of pension funds, we instead use the beliefs of their general consultants obtained through CMAs as described below. Subsection 1.4 provides results that justify this approach by comparing consultant expected returns with those reported by pension funds since 2014 (but, in contrast to consultants, pension funds do not report correlations and volatilities, which are essential for our mean-variance analysis).

Our raw CMAs dataset is the same one used in Couts, Gonçalves, and Loudis (2024). It contains the expectations of 45 investment institutions about future returns, volatilities, and correlations across multiple asset classes in an unbalanced panel covering from 1987 to 2022. These CMAs are collected manually using a combination of direct CMA requests to the

underlying institutions, online searches for publicly available CMAs, and requests made to pension funds for the CMAs of their consultants. Importantly, these CMAs reflect long-term expectations so that their forecasting horizons vary from 4 years to 30 years, with the modal and median horizons being 10 years. See Couts, Gonçalves, and Loudis (2024) for further details about CMAs and the data collection process.

There are two key differences between our final CMAs dataset relative to the one used in Couts, Gonçalves, and Loudis (2024). First, we map the asset classes in the original CMAs to the asset classes available in our pension fund allocation dataset (Cash, FI, EQ, PE, RE, and HF) instead of the twenty asset classes they study (see Internet Appendix B for the mapping procedure). Second, we retain only CMAs from investment institutions that are general pension fund consultants for at least one fund-year in our pension fund dataset. Consequently, our final CMAs dataset contains 15 consulting firms rather than their 45 institutions. Hereafter, we refer to these 15 consulting firms simply as "consultants" and to the final CMAs dataset as the "consultant CMAs dataset."

We match each fund-year observation to a consultant CMA as follows. For each pension fund allocation, we identify its general consultant and the corresponding allocation year. We then link that allocation to the consultant's CMA from the preceding year, when available. As we detail in Internet Appendix A, this procedure mitigates look-ahead bias (i.e., using beliefs formed after the allocation date) while ensuring that beliefs are not generated too far in advance of the fund's portfolio allocation decisions.

1.3 Sample Coverage and Summary Statistics

Table 1 provides annual information on the pension fund dataset and the matching to CMAs. The first three columns show, for each year, the number of pension funds and their general consultants as well as the total AUM in our pension fund dataset. The dataset covers 3,653 fund-years, representing 670 consultant-years, and a total AUM of $73.1 trillion. The next three columns show the fraction of funds, consultants, and AUM matched to consultant CMAs each year. Overall, 57% of the fund-years, 30% of the consultant-years, and 65% of

the AUM in the pension fund dataset is matched to CMAs.[6] Moreover, these fractions tend to grow over time, being substantially lower in the early years and much higher in the more recent years. The next five columns show the number of matched CMAs that cover each asset class in our analysis (except for Cash since it is covered by all CMAs by design). All asset classes are covered by almost all CMAs, except for hedge funds, which are covered by only 123 (out of 199) CMAs.

Our main analysis requires beliefs on all asset classes to be available so that we can construct mean-variance efficient weights (discussed in Section 2). So, the last column in Table 1 provides the number of matched CMAs that cover all asset classes. The requirement to cover hedge funds substantially reduces the sample for our main analysis. The last row in Table 1 quantifies this effect by showing the fraction of funds, consultants, and AUM matched to CMAs that have all asset classes available. In this case, 38% of the fund-years, 18% of the consultant-years, and 37% of the total AUM in the pension fund dataset is matched to CMAs. Clearly, requiring CMAs to cover all asset classes lowers the fraction of the pension fund dataset covered in our main sample. However, it remains the case that our main sample covers a substantial portion of the pension fund industry, especially in the more recent years. Moreover, Internet Appendix C.7 shows that our main results are similar overall even if we keep pension fund observations matched to CMAs that do not cover Hedge Funds (in these cases, we remove Hedge Funds from pension fund allocations and the mean-variance optimization we use to construct belief-implied weights).

Table 2 provides summary statistics for the key variables used in our analysis (focusing on our main sample of pension funds matched to CMAs with full coverage). The first five rows (under the ω_i label) show the portfolio allocations of pension funds to the five non-

[6]Note that the fraction of fund-years matched to CMAs is almost twice as large as the fraction of consultant-years matched to CMAs. This difference arises because our consultant CMAs dataset includes most of the largest consultants in the pension fund dataset. To illustrate this point, Internet Appendix Table IA.1 lists all general consultants that ever appear in the second column of Table 1, along with the number of fund-years each consultant covers. The table also reports, for each consultant, the fraction of fund-years and AUM matched to CMAs, which (in total) correspond exactly to the respective fractions reported in Table 1. The fifteen consultants included in our CMAs dataset are among the largest in the pension fund dataset, in the sense that they account for many of the fund-years in which a general consultant is present.

cash asset classes in our analysis (with a total of 1,422 fund-years).[7] The allocation variable, $\omega_{i,n,t}$, varies by pension fund (i), asset class (n), and year (t). Allocations vary strongly across asset classes and within each asset class. However, since our beliefs data are at the consultant level, we explain portfolio allocations at the consultant level (with a total of 122 consultant-years). So, the next five rows in the table (under the ω_c label) focus on the consultant allocation variable $\omega_{c,n,t} = \Sigma_{i \in I_{c,t}} w_{i,t} \cdot \omega_{i,n,t}$, where $I_{c,t}$ captures the set of pension funds matched to consultant c in year t and $w_{i,t}$ reflects the AUM weight of pension fund i at year t. Strikingly, the level of variability for consultant allocations (within and across asset classes) is comparable to the level of variability for pension fund allocations, indicating that aggregating allocations to the consultant-level does not remove variability from the data in a major way (consistent with Andonov, Bonetti, and Stefanescu (2025)). In line with this result, Internet Appendix C.2 replicates some of our main findings using allocations at the pension fund level, showing that consultant beliefs explain a large fraction of the total variation in pension fund allocations even though these beliefs do not vary across pension funds matched to the same consultant-year.

The remaining rows of Table 2 provide summary statistics for belief elements from consultant CMAs. As it becomes clear in Section 2, the relevant belief elements for portfolio allocation decisions are those for returns in excess of cash returns. So, for each consultant-year, we follow Couts, Gonçalves, and Loudis (2024) in constructing the subjective expected excess return vector, $\mu_{c,t}$, and the subjective covariance matrix of excess returns, $\Sigma_{c,t}$. The five rows under the μ_c label summarize the expected excess return values. As with portfolio allocations, expected excess returns display large variation across asset classes and within each asset class. The remaining rows summarize the elements of the covariance matrix of excess returns. Again, for each belief element, we observe large variation both across asset classes and within each asset class. This large variation in CMA-implied belief elements beyond μ is consistent with Couts, Gonçalves, and Loudis (2024) and Couts et al. (2025).

[7] The cash allocation is omitted from Table 2 because our analysis linking pension fund allocations to mean-variance efficient allocations focuses on the non-cash asset classes since the sum of allocations to all asset classes is matched by construction as it is 100%.

1.4 Link Between the Expected Returns of Pension Funds and their Consultants

As discussed above, we match pension fund allocations to the beliefs of their consultants because we do not directly observe the beliefs of pension funds (similar to Begenau, Liang, and Siriwardane (2025)). However, starting in 2014, pension funds' Comprehensive Annual Financial Reports (CAFRs) must contain their actuarially assumed rates of return (μ_{AARR}) at the asset class level (albeit they do not report Σ elements).[8] As discussed in Andonov and Rauh (2022), these μ_{AARR} values capture beliefs as well as incentives since they are used to discount pension liabilities. So, we validate our approach by verifying that consultant μs positively correlate with μ_{AARR}, and then explain why this correlation reflects shared beliefs rather than shared reporting incentives.

Figure 2 plots μ_{AARR} values from pension funds aggregated to the level of their consultants, $\mu^{\text{AARR}}_{c,n,t} = \sum_{i \in I_{c,t}} w_{i,t} \cdot \mu^{\text{AARR}}_{i,n,t}$, against expected returns directly from their consultants, $\mu_{c,n,t}$. There is a strong positive link between μ_{AARR} and μ for the subsample over which we observe both (42 consultant-year observations). Interestingly, we also see that μ_{AARR} tends to be higher than μ, which is consistent with the view that μ_{AARR} reflects true pension fund expected returns plus a distortion due to the incentive to report inflated expected returns to discount liabilities more aggressively (Novy-Marx and Rauh (2011)).

The figure also reports the correlations between μ_{AARR} and μ as well as the slope coefficients from regressions of μ_{AARR} onto μ. We consider variation across funds (using $n \times t$ fixed effects), across asset classes (using $c \times t$ fixed effects), and over time (using $c \times n$ fixed effects). In all three dimensions, there is a strong and statistically significant link between μ_{AARR} and μ. The strongest link is across asset classes, with a slope coefficient of 0.96 ($t_{stat} = 31.40$) and a correlation of 0.97. However, even the weakest link (across funds) is still strong, with a slope coefficient of 0.36 ($t_{stat} = 2.86$) and a correlation of 0.33. The weaker relation between

[8] Under the U.S. Governmental Accounting Standards Board Statement 67 (GASB 67), pension funds are required to disclose "...the long-term expected real rate of return for each major asset class," which CAFRs typically refer to as actuarially assumed rates of return. We follow Andonov and Rauh (2022) in collecting these data. However, we do not have μ_{AARR} values before 2014, the year the GASB 67 requirement started. Moreover, pension funds do not report Σ elements, which are also required for our mean-variance analysis.

μ_{AARR} and μ across funds (relative to their connection across asset classes and over time) is likely due to pension fund incentives varying more across pension funds than along other dimensions. If the link between μ^{AARR} and μ is primarily driven by beliefs rather than incentives (as we argue below), it is natural that this link becomes weaker in dimensions where incentives differ most.

It is important to highlight that we do not claim (and our analysis does not require) that CMAs have a causal effect on the beliefs of pension funds. Our goal is to shed light on the link between pension fund allocations and their beliefs, not the link between pension fund allocations and the beliefs of their consultants. We use consultant beliefs simply as a proxy for pension fund beliefs, requiring only that the two are correlated in equilibrium. Since Figure 2 shows that they are correlated, we (as econometricians) can treat consultant beliefs as noisy signals for pension fund beliefs (as we formalize in the next section). For example, we interpret our subsequent analyses as showing that pension fund allocations are linked to their beliefs, and this interpretation remains valid whether the correlation between μ and μ_{AARR} arises because consultants influence the views of pension funds (as argued by Begenau, Liang, and Siriwardane (2025)) or because of other reasons, such as pension funds selecting consultants that share their belief formation process.

Nevertheless, the economic interpretation of our mean-variance analysis requires pension fund beliefs to correlate with the CMAs of their consultants through shared beliefs rather than shared reporting incentives extraneous to risk and return expectations. This requirement would be violated if consultant CMAs were designed to cater to the incentives of their pension fund clients. However, we argue that this is not the case for two reasons. First, consultants distribute identical CMAs to a broad and heterogeneous client base. Specifically, the average consultant in our sample advises 275 clients, with public pension funds representing only 17.3% of their clients and 7.2% of their regulatory assets under advisory.[9] So, consultant

[9]These values are computed using the Type of Clients table in Form ADV Part 1A, Item 5, line (i), available starting in 2015. The consulting firms in our sample serve a wide range of investors, including public pension funds, corporate pension plans (both Defined Benefit and 401K), endowment funds, foundations, healthcare systems, and high-net-worth individuals.

CMAs are unlikely to reflect the incentives of their public pension fund clients at the expense of all other clients. In fact, Figure 2 shows that consultant μs tend to be lower than the μ_{AARR} values reported by their pension fund clients, suggesting consultants do not cater to the incentive of their pension fund clients to inflate liability discount rates.[10] Second, CMAs are based on transparent economic and statistical analyses rather than opaque processes that would facilitate the incorporation of client incentives (Andolfatto and Bastianello (2025)). Indeed, μs aggregated across consultants provide unbiased predictors of future asset class returns (Couts, Gonçalves, and Loudis (2024)), which would not be the case if consultant CMAs merely reflected the distortionary incentives of their clients.

An alternative way for consultant μs to correlate with μ_{AARR} values through incentives rather than beliefs is selection. Specifically, pension funds may choose their μ_{AARR} values and desired allocations based on distortionary incentives and then select consultants whose CMAs can ex post justify their μ_{AARR} values and desired allocations. As we detail in Section 3.3, this channel also cannot explain our results. Therefore, if the link between μ and μ_{AARR} reflects selection, it must operate through pension funds selecting consultants who share their belief formation process, rather than pension funds selecting consultants whose CMAs line up with their distortionary incentives.[11]

Overall, we see a link (causal or not) between the expected returns of consultants and those reported by their pension fund clients. However, the link is not perfect since the correlations between μ_{AARR} and μ as well as the slope coefficients from regressions of μ_{AARR} onto μ are below one. Our mean-variance framework detailed in the next section accounts for this feature of the data.

[10] In untabulated results, we also find that, for each asset class we study, the average μ of the consultants in our analysis is similar to the average μ from other consultants that are not general consultants for any US public pension fund over our sample. This result reinforces the finding that consultants do not inflate expected returns to cater to the incentives of their pension fund clients at the expense of all other clients.

[11] If a pension fund wanted to manipulate its μ_{AARR} due to reporting incentives, changing its general consultant would be a relatively difficult and even unnecessary way to do so. First, general consultants are chosen via public procurement solicitation with a semi-transparent bidding process and evaluation procedure, and the contracts typically have a pre-determined minimum length (e.g., 5 years). Second, GASB 67 only lists general principles for the determination of μ_{AARR}, leaving pension funds ample discretion on how to determine them (including whether and to what extent they rely on their consultant's CMAs).

2 The Mean-Variance Efficient (MVE) Allocations from CMAs

Our empirical analyses combine the CMAs-based beliefs from the previous section with a mean-variance framework that incorporates frictions and belief wedges. Subsection 2.1 presents the mean-variance framework and its portfolio allocation solution. Then, Subsections 2.2 and 2.3 summarize the mean-variance efficient (MVE) allocations we obtain under a fixed risk aversion parameter and heterogeneous risk aversion estimates, respectively.

We use the same subscripts as in the prior section. In particular, i indexes investors (pension funds), c indexes institutions providing beliefs (consultants), n indexes asset classes (FI, EQ, PE, RE, and HF), and t indexes time (years). In addition, hereafter $R_{n,t}$ reflects the gross return on asset class n at time t while $r_{n,t} = R_{n,t} - R_{Cash,t}$ reflects the respective excess return. We also suppress time indices inside moments when convenient. For instance, $\mu_{c,n,t} \equiv \mathbb{E}_{c,t}[r_n]$ and $\sigma^2_{c,n,t} \equiv \mathbb{V}ar_{c,t}[r_n]$. Moreover, we often refer to r as a return and μ as an expected return (i.e., we omit the "excess" terminology for simplicity).

2.1 The Mean-Variance Framework and Solution

Consider an investor who chooses a vector of portfolio weights $(\omega_{i,t})$ each period to maximize the long-term mean-variance utility

$$\underset{\{\omega_{i,t}\}}{Max} \; \mathbb{E}_{i,t}[r_i] \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i] \; - \; \frac{\gamma^b_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i - r^b] \tag{1}$$

$$s.t. \quad (1 - \mathbf{1}'\omega_{i,t}) \geq 0$$

$$\omega_{i,t} \geq \mathbf{0}$$

where $r_{i,t+h_i} = \omega'_{i,t} r_{t+h_i}$ is the long-term excess return on the wealth portfolio of investor i and $r^b_{t+h_i} = \omega^{b\prime}_t r_{t+h_i}$ is the long-term excess return on a benchmark portfolio.

There are two important initial observations about the mean-variance allocation problem in Equation 1. First, it allows for arbitrary $h_i \geq 1$, and thus it accommodates beliefs with heterogeneous horizons (a feature of our beliefs data). Second, it does not impose the

constraint $\mathbf{1}'\omega_{i,t} = 1$ because $\omega_{i,t}$ does not include the weight on cash, which is given by $\omega_{Cash,i,t} = (1 - \mathbf{1}'\omega_{i,t})$. For simplicity, hereafter we refer to the assets covered in $\omega_{i,t}$ as risky assets even though, technically, cash can also be risky.[12]

The investor's mean-variance problem has two constraints, which are crucial to explain pension fund portfolio allocations empirically (as Section 3 demonstrates). The first, $\omega_{Cash,i,t} \geq 0$, ensures the weight on cash is non-negative (as in Black (1972)). The second, $\omega_{i,t} \geq \mathbf{0}$, ensures the weights on the risky assets are also non-negative (as in Miller (1977)). So, we refer to our two constraints as imposing "non-negative weights."

Moreover, the investor's mean-variance problem has a benchmarking incentive through the $\mathbb{V}ar_{i,t}[r_i - r^b]$ term in the objective function, which is also empirically important (as we show in Section 3). So, while $\gamma_{i,t} \geq 0$ reflects the standard risk aversion parameter, $\gamma^b_{i,t} \geq 0$ reflects aversion to deviations from the benchmark. The prior literature has proposed multiple channels for this type of benchmarking incentive (e.g., Evans et al. (2024) and Buffa, Vayanos, and Woolley (2022)). So, we intentionally set up the portfolio allocation problem in Equation 1 to be flexible on the economic mechanism inducing the benchmarking incentive (in the spirit of Brennan (1993) and Basak and Pavlova (2013)).

However, Internet Appendix D.1 formalizes two economic foundations for this benchmarking incentive through frictions. In the first, we consider a standard constrained mean-variance problem with an extra tracking error constraint. In the second, we consider a standard constrained mean-variance problem in which a Bayesian investor is uncertain about expected returns, forming them as in Black and Litterman (1991, 1992). As discussed in Internet Appendix D.1, the economic interpretation for γ^b differs across these two portfolio allocation problems, with γ^b arising from institutional frictions in the former case and from informational frictions in the latter case. Nevertheless, the resulting portfolio allocation solutions are

[12]Note that the use of cash as the reference asset in Equation 1 has two potential interpretations and we do not attempt to differentiate between them in our analysis. The first is that the investor perceives cash to be the risk-free asset, in which case the maximization in Equation 1 is equivalent to maximizing $\mathbb{E}_{i,t}[R_i] - 0.5 \cdot \gamma_{i,t} \cdot \mathbb{V}ar_{i,t}[R_i] - 0.5 \cdot \gamma^b_{i,t} \cdot \mathbb{V}ar_t[R_i - R^b]$ subject to non-negative weights that add to one. The second is that a risk-free asset does not exist, but that the investor cares about returns in excess of cash instead of caring about raw returns. Both are plausible interpretations for Equation 1

identical to those from Equation 1. Consequently, the "benchmarking incentive" we consider can viewed as arising from frictions (institutional or informational).

Let the vector $\mu_{i,t} = \mathbb{E}_{i,t}[r]$ and matrix $\Sigma_{i,t} = \mathbb{V}ar_{i,t}[r]$ capture the investor's beliefs. The solution to the mean-variance problem in Equation 1 yields an expression for the optimal $\omega_{i,t}$ as a function of $\mu_{i,t}$ and $\Sigma_{i,t}$ (derived in Internet Appendix D.2). If we (the econometricians) observed $\mu_{i,t}$ and $\Sigma_{i,t}$, then we could use that expression to study the explanatory power of beliefs for investor portfolio allocations ($\omega_{i,t}$). However, we only observe $\mu_{c,t}$ and $\Sigma_{c,t}$, which are the beliefs of the consultant (c) matched to the given investor (i), with such beliefs being directly related to (but not perfectly revealing of) the investor's beliefs (as discussed in Subsection 1.4). So, we instead study the explanatory power of $\mu_{c,t}$ and $\Sigma_{c,t}$ for the consultant-level portfolio allocation $\omega_{c,t} = \sum_{i \in I_{c,t}} w_{i,t} \cdot \omega_{i,t}$, where $I_{c,t}$ is the set of pension funds matched to consultant c at time t and $w_{i,t}$ captures aggregation weights (proportional to AUM in our empirical analysis).

To derive an expression linking $\omega_{c,t}$ to $\mu_{c,t}$ and $\Sigma_{c,t}$, we define the belief wedges

$$\epsilon_{\mu,i,t} \ = \mu_{i,t} \ - \ \mu_{c,t} \tag{2}$$

$$\epsilon_{\Sigma,i,t} \ = \Sigma_{i,t} \ - \ \Sigma_{c,t} \tag{3}$$

and take a weighted average (across $i \in I_{c,t}$) of the solution to the mean-variance problem in Equation 1 to obtain (see Internet Appendix D.2)

$$\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t} \ = \alpha_{c,t} \ + \ \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r, r_c] \ + \ \gamma_{c,t}^{b} \cdot \mathbb{C}ov_{c,t}[r, r_c - r^b] \tag{4}$$

where $r_c = \omega_{c,t}^{'} r$ is the wealth portfolio of pension funds matched to consultant c and the $\gamma_{c,t}$ and $\gamma_{c,t}^{b}$ parameters are aggregates of the risk aversion parameters $\gamma_{i,t}$ and $\gamma_{i,t}^{b}$ (across $i \in I_{c,t}$).

Equation 4 is the standard risk premium representation of the mean-variance portfolio allocation solution, with three modifications. First, it includes an adjustment for the constraints through the scalar $\ell_{c,t}^{cash}$ and the vector $\ell_{c,t}$, which are weighted averages across $i \in I_{c,t}$ of the lagrange multipliers on the non-negative weight constraints for cash and the risky as-

sets ($\ell_{i,t}^{cash}$ and $\ell_{i,t}$). Second, it includes a premium for exposure to wealth returns in excess of benchmark returns, $\mathbb{C}ov_{c,t}[r, r_c - r^b]$, which arises from the benchmarking incentive ($\gamma^b \geq 0$). And third, it includes an alpha term, $\alpha_{c,t}$, which reflects belief wedges as it is a weighted average of $\alpha_{i,t} = \gamma_{i,t} \cdot \epsilon_{\Sigma,i,t}\omega_{i,t} + \gamma_{i,t}^b \cdot \epsilon_{\Sigma,i,t}(\omega_{i,t} - \omega_{i,t}^b) - \epsilon_{\mu,i,t}$ across $i \in I_{c,t}$ (with $\alpha_{c,t} = \mathbf{0}$ if $\mu_{i,t} = \mu_{c,t}$ and $\Sigma_{i,t} = \Sigma_{c,t}$ $\forall$ $i \in I_{c,t}$).

Defining $\omega_{c,t}^{\alpha} = -\Sigma_{c,t}^{-1}\alpha_{c,t}/(\gamma_{c,t} + \gamma_{c,t}^b)$, Equation 4 yields a simple closed-form expression for the optimal portfolio allocation:

$$\omega_{c,t} = \omega_{c,t}^{\alpha} + \omega_{c,t}^{*} = \omega_{c,t}^{\alpha} + (1 - b_{c,t}) \cdot \omega_{c,t}^{mv} + b_{c,t} \cdot \omega_t^b \tag{5}$$

where

$$\omega_{c,t}^{mv} = \frac{1}{\gamma_{c,t}} \cdot \Sigma_{c,t}^{-1}(\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t}) \qquad \text{and} \qquad b_{c,t} = \frac{\gamma_{c,t}^b}{\gamma_{c,t} + \gamma_{c,t}^b} \tag{6}$$

Intuitively, the optimal portfolio allocation has two components. The first (ω^{α}) arises from belief wedges between investors and consultants whereas the second (ω^*) reflects the constrained mean-variance efficient (MVE) allocation with a benchmarking incentive from the perspective of the consultant. Moreover, ω^* is a weighted average of the constrained MVE allocation without benchmark weights (ω^{mv}) and the benchmark allocation (ω^b). Finally, b captures the strength of the benchmarking incentive since $\omega^* = \omega^{mv}$ if $b = 0$ and $\omega^* = \omega^b$ if $b = 1$.

To simplify terminology, hereafter we refer to ω^* as the "MVE allocation," which we treat as shorthand for the "constrained MVE allocation with a benchmarking incentive from the perspective of the consultant." Similarly, we refer to ω^{α} as the "MVE allocation wedge" since it reflects the wedge between observed allocations and MVE allocations.

In our empirical analysis, we measure the benchmark weights using the AUM-weighted average allocation of all pension funds (present in the first column of Table 1) in the prior year: $\omega_t^b = \sum_{i \in I_{t-1}} w_{i,t-1} \cdot \omega_{i,t-1}$. There are three important observations about these benchmark weights. First, they are designed to not vary across pension funds as I_t reflects all pension funds in our full dataset in year t (so, benchmark weights cannot explain portfolio weight variation across pension funds). Second, they rely on an AUM-weighted average of the allo-

cations of all pension funds, which puts more weight in large pension funds, consistent with how we define consultant-year allocations. And third, they use pension fund allocations in the prior year since it is hard for a fund to benchmark its allocation to the contemporaneous allocation of its peers.

Given that $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$, with ω^b measured from lagged pension fund allocations, benchmark weights can help MVE allocations to explain pension fund allocations simply through allocation persistence. Consequently, to understand the effect of beliefs that is incremental to benchmark weights, our empirical analysis in Section 3 considers different approaches to control for benchmark weights. The simplest one is to also study active weights, which are deviations from the benchmark weights ($\widetilde{\omega}_{c,t} = \omega_{c,t} - \omega^b_t$). Based on Equation 5, these active weights can be written as

$$\widetilde{\omega}_{c,t} = \widetilde{\omega}^{\alpha}_{c,t} + \widetilde{\omega}^{*}_{c,t} = \widetilde{\omega}^{\alpha}_{c,t} + (1-b_{c,t})\cdot\widetilde{\omega}^{mv}_{c,t} \tag{7}$$

with $\widetilde{\omega}^{\alpha}_{c,t} = \omega^{\alpha}_{c,t}$ capturing MVE allocation wedges and $\widetilde{\omega}^{*}_{c,t} = (1-b_{c,t})\cdot(\omega^{mv}_{c,t} - \omega^b_t)$ reflecting the MVE active weights. So, the MVE active weights of a model with a benchmarking incentive are proportional to the pure MVE active weights ($\widetilde{\omega}^{mv}$). This implies that there is no extra component from benchmarking incentive in the explanatory power of MVE active weights for pension fund active weights.

2.2 The MVE Allocations with $\gamma = 3$

As discussed in Section 1, our CMAs dataset allows us to measure $\mu_{c,t}$ and $\Sigma_{c,t}$. Moreover, as shown in Internet Appendix D.3, the lagrange multipliers ($\ell^{cash}_{c,t}$ and $\ell_{c,t}$) can be obtained by numerically solving the optimization problem in Equation 1 given $\mu_{c,t}$, $\Sigma_{c,t}$, ω^b_t, $\gamma_{c,t}$, and $\gamma^b_{c,t}$. So, to obtain $\omega^*_{c,t}$ from Equation 5 we only need the risk aversion parameter ($\gamma_{c,t}$) and the benchmarking incentive parameter ($b_{c,t} = \gamma^b_{c,t}/(\gamma_{c,t} + \gamma^b_{c,t})$).

To isolate the pure effect of variation in beliefs, three of the main specifications we study impose $\gamma_{c,t} = \gamma$ and $b_{c,t} = b$, thereby eliminating variation in risk aversion and benchmarking incentives. In these specifications, we set $\gamma = 3$, a moderate risk aversion that lies between

the median and mean of the distribution of risk aversion estimates we report in the next subsection. We also explore $\gamma = 5$ and $\gamma = 10$ since risk aversion estimates in the experimental literature tend to fall between 3 and 10 (see the discussion in Giglio et al. (2021)). The resulting (untabulated) findings are similar to the ones we report.

In contrast to risk aversion, there is no clear reference value for b in the prior literature. We consider specifications with $b = 0$ (so that there is no benchmarking incentive) as well as a specification where we calibrate b to match the average subjective tracking error in the data. Specifically, the pension fund portfolio allocation $\omega_{c,t}$ implies the subjective tracking error $TE_{c,t} = \sqrt{\mathbb{V}ar_{c,t}[r_c - r^b]} = \sqrt{\widetilde{\omega}'_{c,t}\Sigma_{c,t}\widetilde{\omega}_{c,t}}$ whereas the MVE allocation $\omega^*_{c,t}$ implies the subjective tracking error $TE^*_{c,t} = \sqrt{\widetilde{\omega}^{*\prime}_{c,t}\Sigma_{c,t}\widetilde{\omega}^*_{c,t}} = (1-b)\cdot\sqrt{\widetilde{\omega}^{mv\prime}_{c,t}\Sigma_{c,t}\widetilde{\omega}^{mv}_{c,t}}$. We set b such that the average $TE^*_{c,t}$ matches the average $TE_{c,t}$. This b calibration process is conservative in that it targets the overall level of pension fund deviations from benchmark weights but not the directions of these deviations.[13] Estimating b as a free parameter to match pension fund allocations would lead to MVE allocations that explain pension fund allocations even better.

The first three rows of each asset class block in Table 3 provide summary statistics for the MVE weights (ω^*) under the specifications with fixed γ and b. The first specification, "UMVE ($b = 0$ & $\gamma = 3$)", imposes no constraints ($\ell^{cash}_{c,t} = 0$ and $\ell_{c,t} = \mathbf{0}$) and no benchmarking incentive ($b = 0$). In this case, the MVE weights vary excessively across and within asset classes. For example, the average ω^* varies from -37.3% for public equities to 216.4% for hedge funds, with the lowest within asset class standard deviation being 73.2% (for private equity). This level of variation is highly counterfactual when compared to the distribution of actual pension fund weights, summarized in the last row of each asset class block.

The second specification, "CMVE ($b = 0$ & $\gamma = 3$)," has non-negative weight constraints but still no benchmarking incentive.[14] In this case, the level of portfolio weight variation across and within asset classes declines substantially, becoming more comparable to the

[13] To see this fact, consider the stylized case in which $\widetilde{\omega}^*_{c,t} = -\widetilde{\omega}_{c,t}$. Even though in this case MVE active allocations are perfectly negatively correlated with pension fund active allocations, we have $TE^*_{c,t} = TE_{c,t}$.

[14] Internet Appendix Table IA.4 provides descriptive statistics for the lagrange multipliers for the constrained MVE specifications we consider.

level of variation observed in actual pension fund allocations. Yet, the distribution of MVE allocations is far from perfectly replicating the distribution of actual pension fund allocations. For example, the average MVE allocations are substantially lower than the average pension fund allocations for public equities and fixed income (the more traditional asset classes) and substantially higher than the average pension fund allocations for private equity, real estate, and hedge funds (the alternative asset classes). This result suggests that, on average, institutions perceive traditional (alternative) asset classes as having a higher (lower) benefit than implied by the beliefs of their consultants under a standard mean-variance objective.

The third specification, "CMVE ($b \neq 0$ & $\gamma = 3$)," has non-negative weight constraints and a benchmarking incentive, with $b = 0.89$ calibrated as discussed above. In this case, the variation in portfolio weights (both across and within asset classes) is comparable to that observed in actual pension fund allocations. Moreover, the mismatch in average allocations is substantially smaller than in the model without a benchmarking incentive. For example, the average MVE allocations to public equity, private equity, and real estate are 44.8%, 9.7%, and 12.6%, respectively, which closely align with the corresponding pension fund averages of 47.2%, 9.7%, and 10.5%. Nonetheless, a non-trivial discrepancy remains for fixed income and hedge funds: the average MVE allocation to fixed income falls short of the pension fund average by more than six percentage points, while the reverse holds for hedge funds. On balance, this pattern suggests that institutions perceive fixed income (hedge funds) to offer higher (lower) benefits than implied by the beliefs of their consultants, even when optimizing under a mean-variance framework with a benchmarking incentive.

2.3 The MVE Allocations with Estimated $\gamma_{c,t}$

Our full specification also accounts for heterogeneity and time variation in risk aversion. For that, we estimate risk aversion for each consultant-year observation from regressions based on Equation 4 across asset classes. This estimation approach is analogous to how Couts et al. (2025) estimate risk aversion for their expected return decomposition analysis, and also parallels how risk aversion is estimated in cross-sectional asset pricing papers. We

continue to assume a constant b so that $\gamma^b_{c,t}$ is proportional to $\gamma_{c,t}$. As we detail in Internet Appendix D.4, separately estimating $\gamma_{c,t}$ and $\gamma^b_{c,t}$ that are heterogeneous and time-varying is econometrically impractical given the high collinearity of the two explanatory variables in Equation 4 ($\mathbb{C}ov_{c,t}[r, r_c]$ and $\mathbb{C}ov_{c,t}[r, r_c - r^b]$).

To understand our estimation approach, start from the simplest case of the MVE allocation with no constraints ($\ell^{cash}_{c,t} = 0$ and $\ell_{c,t} = \mathbf{0}$) and no benchmarking incentive ($b = 0$). In this case, Equation 4 reduces to

$$\mu_{c,n,t} \; = \alpha_{c,n,t} \; + \; \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r_c] \tag{8}$$

so that we can estimate $\gamma_{c,t}$ from the slope of a regression of $\mu_{c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ across asset classes (for each consultant-year observation). However, under constraints, the left hand side of Equation 8 becomes $\mu_{c,t} - \mathbf{1}\ell^{cash}_{c,t} + \ell_{c,t}$ so that we also need to measure $\ell^{cash}_{c,t}$ and $\ell_{c,t}$, which depend on $\gamma_{c,t}$, creating a circular problem. To solve this issue, we start from an arbitrary $\gamma_{c,t}$ value and update it through a sequence of regressions of $\mu_{c,t} - \mathbf{1}\ell^{cash}_{c,t} + \ell_{c,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ until $\gamma_{c,t}$ converges. Finally, further adding the benchmarking incentive implies the left hand side of Equation 8 becomes $\mu_{c,t} - \mathbf{1}\ell^{cash}_{c,t} + \ell_{c,t} + [b/(1-b)] \cdot \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r^b]$ and the slope coefficient becomes $\gamma_{c,t}/(1-b)$. We still rely on a sequence of regressions until $\gamma_{c,t}$ converges. However, in this case we perform the iterative procedure for multiple values of b and choose the one that matches the average tracking errors from the MVE and pension fund allocations (analogous to how we calibrate b with a constant γ). Further details about our $\gamma_{c,t}$ estimation are provided in Internet Appendix D.4.[15]

Figure 3 plots the histogram of the $\gamma_{c,t}$ estimates for our full specification, labeled "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." The overall risk aversion values are reasonable. The median and average

[15]Internet Appendix D.4 also discusses a potential structural endogeneity issue with our baseline γ estimation. Specifically, under our mean-variance structure in Subsection 2.1, the αs are functions of belief wedges, which correlate with pension fund weights through the ω^α component of Equation 5. Since pension fund weights are used in the calculation of $\mathbb{C}ov_{c,t}[r_n, r_c]$, the estimation is subject to structural endogeneity. However, we show that this structural endogeneity is not a major concern for two reasons. First, the correlation between $\mathbb{C}ov[r_n, r_c]$ and ϵ_μ across asset classes is small and statistically insignificant for the subset of consultant-years over which we observe ϵ_μ (see Internet Appendix D.6). Second, our main results are very similar if we use an instrumental variables approach to estimate γ that relies on $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$, which does not use portfolio weights, as an instrument for $\mathbb{C}ov_{c,t}[r_n, r_c]$ (see Internet Appendix D.5).

risk aversion values are 2.7 and 7.7, respectively, which are in line with risk aversion values estimated in the prior literature (and with the $\gamma = 3$ value we use in our constant risk aversion specifications). The distribution is also right skewed and has large variation. For instance, the 10% and 90% risk aversion quantiles are 1.3 and 23.3, respectively. Moreover, the minimum and maximum risk aversion values are 0.4 and 53.8. The latter value is quite high, but somewhat expected because these risk aversion estimates are subject to non-trivial estimation error since they are based on five observations each (i.e., there are five asset classes in each cross section).

The "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)" specification in Table 3 provides summary statistics for the resulting MVE weights (ω^*). As in the analogous specification with $\gamma = 3$, the variability in the MVE allocations is comparable to the variability in actual pension fund allocations. In addition, the mismatch between average allocations improves slightly. In particular, the MVE allocation to fixed income (hedge funds) is lower (higher) than the respective pension fund allocation by around four percentage points instead of six.

Figure 4 provides a visual representation of the results from this full specification with estimated $\gamma_{c,t}$. Specifically, for each asset class, it plots the empirical density function of ω against the empirical density function of ω^*. The density functions are relatively similar regardless of which asset class we consider.

Overall, once non-negative weight constraints and benchmarking incentives are incorporated, the distribution of MVE allocations closely resembles that of pension fund allocations. Nonetheless, on average, institutions appear to assign higher (lower) perceived benefits to fixed income (hedge funds) than implied by their consultants' beliefs, even under a mean-variance objective with benchmarking. Importantly, these patterns persist regardless of whether risk aversion is assumed constant or allowed to vary across funds and over time.

3 Regressing Pension Fund Portfolio Allocations on MVE Allocations

This section examines how closely pension fund portfolio allocations align with the mean-variance efficient (MVE) allocations implied by their consultants' beliefs. In particular, we

estimate regressions of consultant-level pension fund allocations (ω) onto their corresponding MVE allocations (ω^*):

$$\omega_{c,n,t} \ = \text{Fixed Effect} \ + \ \text{Slope} \cdot \omega^*_{c,n,t} \ + \ \epsilon_{c,n,t}, \tag{9}$$

The results are provided in Table 4. The columns differ in the source of variation explored: across funds (using $n \times t$ fixed effects), across asset classes (using $c \times t$ fixed effects), and over time (using $c \times n$ fixed effects). Columns also vary by whether the regressions focus on raw weights (ω) or active weights ($\widetilde{\omega}$). The $\widetilde{\omega}$ columns are based on an equation analogous to Equation 9, but that controls for benchmark allocations by replacing raw weights (ω and ω^*) with active weights ($\widetilde{\omega} = \omega - \omega^b$ and $\widetilde{\omega}^* = \omega^* - \omega^b$). For variation across funds, using raw weights or active weights leads to identical results (since ω^b does not vary across funds), but Table 4 still reports both columns to avoid confusion.

We structure the discussion of Table 4 in four subsections. Subsection 3.1 begins with the simple case of frictionless MVE allocations, showing that regressions of ω onto ω^* exhibit meaningful explanatory power even though the slope coefficients are small. Subsection 3.2 then adds our two realistic frictions (non-negative weight constraints and benchmarking incentives), finding that they substantially strengthen the explanatory power of ω^* and go a long way in resolving the issue of small slopes. Finally, Subsection 3.3 shows that reverse causality cannot explain the results we obtain from regressing ω onto ω^* and Subsection 3.4 addresses other potential concerns with these regressions.

3.1 Regressions of ω onto ω^* without Frictions

The first three specifications in Table 4 examine the link between pension fund allocations and MVE allocations without frictions. In particular, we consider unconstrained MVE allocations that have $b = 0$ and a common risk aversion of $\gamma = 3$, differing only based on the belief components that are allowed to vary.

The first specification holds correlations fixed at zero and sets $\sigma^2_{c,n,t} = \sigma^2_n$ (the asset class average), so that variation in MVE weights arises almost entirely from differences

in expected returns: $\omega^*_{c,n,t} = (1/\gamma) \cdot \mu_{c,n,t}/\sigma^2_n$.[16] In this simplistic case (which mirrors the literature approach of regressing ω onto μ), pension fund allocations exhibit only a weak connection to MVE allocations. For instance, the slope coefficients for raw weights across funds and over time are 0.03 ($t_{stat} = 1.52$) and 0.02 ($t_{stat} = 1.61$), respectively, with R^2 values below 2.0%. Across asset classes, the slope coefficient is even negative unless we control for benchmark weights by using active weights, in which case the slope coefficient becomes 0.03 ($t_{stat} = 1.64$) with $R^2 = 5.4\%$.

The second specification allows volatilities to vary across all dimensions: $\omega^*_{c,n,t} = (1/\gamma) \cdot \mu_{c,n,t}/\sigma^2_{c,n,t}$. This flexibility strengthens the overall alignment between pension fund and MVE allocations. The slope magnitudes remain similar, but statistical significance and explanatory power improve markedly. For the regression across asset classes using active weights, the t_{stat} and R^2 values increase to 2.14 and 9.1%. Similarly, for variation across funds, the t_{stat} and R^2 values rise to 2.42 and 5.8%. For variation over time, the slope coefficients become statistically significant (at least at the 10% level) whether we use raw or active weights, but the R^2 values remain very low.

The third specification, "UMVE ($b = 0$ & $\gamma = 3$)," accounts for the diversification benefit of different asset classes by allowing correlations to differ from zero: $\omega^*_{c,n,t} = (1/\gamma) \cdot [\Sigma^{-1}_{c,t}\mu_{c,t}]_n$. This version represents the fully unconstrained MVE allocation without benchmarking incentives. While slope magnitudes again remain similar, the link between pension fund and MVE allocations strengthens substantially. For variation across asset classes, the regression of $\widetilde{\omega}$ onto $\widetilde{\omega}^*$ has a slope t_{stat} of 2.54 and an R^2 of 22.2%, indicating that MVE active weights explain a meaningful share of the variation in pension fund active weights. Variation across funds shows a comparable improvement ($t_{stat} = 2.40$ and $R^2 = 17.3\%$). The slopes from regressions of pension fund allocations on unconstrained MVE allocations are also statistically significant ($t_{stat} = 2.52$ for raw weights and $t_{stat} = 1.82$ for active weights), but the R^2 values remain below 2% for both raw and active weights. Thus, accounting for diversification ben-

[16] We set $\sigma^2_{c,n,t} = \sigma^2_n$ instead of $\sigma^2_{c,n,t} = \sigma^2$ so that the units of ω^* are comparable across asset classes (otherwise, asset classes with low μ would have little effect in the results across funds and over time). However, the results are very similar if we use $\sigma^2_{c,n,t} = \sigma^2$ instead.

efits embodied by Σ drastically strengthens the pass-through of beliefs to allocations (even before accounting for frictions). This result represents a significant contribution to the literature which largely has ignored this effect (particularly when studying belief pass throughs for single assets classes).

3.2 Regressions of ω onto ω^* with Frictions

Despite the significant t_{stat} values and non-trivial R^2 levels for the link between pension fund allocations and unconstrained MVE allocations, the respective slope coefficients are small. In particular, the largest slope coefficient using unconstrained MVE allocations is 0.03, almost two orders of magnitude smaller than the theoretical ideal of 1.00. This result is related to the excessive variability of ω^* relative to ω in the absence of frictions (as highlighted in Table 3), yielding relatively low slope coefficients from regressions based on Equation 9. We now consider (institutional and informational) frictions as a potential resolution for this low pass through of MVE allocations to observed allocations.

We begin by introducing a common institutional friction: investors are prohibited from taking negative positions in any asset class. In particular, the fourth specification in Table 4, "CMVE ($b = 0$ & $\gamma = 3$)," uses our constrained MVE weights, which display more reasonable levels of variability (as shown in Table 3). There is little change in the R^2 values relative to the unconstrained MVE specification. However, the slope coefficients tend to be much larger when using constrained MVE weights. For instance, the slope coefficient increases to 0.16 ($t_{stat} = 2.45$) for variation across funds and to 0.11 ($t_{stat} = 2.02$) for variation in active weights across asset classes. The slope coefficients also increase for time variation, but the link between MVE allocations and pension fund allocations is generally much weaker there (e.g., even with constraints the slopes are still around 0.03).

While accounting for constraints leads to a stronger positive link between $\widetilde{\omega}$ and $\widetilde{\omega}^*$ in all dimensions, the relation between ω and ω^* across asset classes remains negative. That is, we only obtain a positive link between pension fund allocations and MVE allocations across asset classes after controlling for benchmark weights (through the use of active weights).

This result suggests that, in the absence of benchmarking incentives, our mean-variance framework cannot explain why a pension fund would allocate differently across asset classes at a given point in time. So, our fifth specification in Table 4, "CMVE ($b \neq 0$ & $\gamma = 3$)," adds the benchmarking incentive, which can reflect institutional or informational frictions (see Subsection 2.1). This addition has a large positive effect on the link between pension fund allocations and MVE allocations for all cases. For instance, for raw weights, the slope coefficients increase to 0.95 ($t_{stat} = 11.5$ and $R^2 = 75.6\%$) across asset classes, 0.47 ($t_{stat} = 3.51$ and $R^2 = 29.6\%$) across funds, and 0.32 ($t_{stat} = 3.70$ and $R^2 = 12.4\%$) over time.

As discussed in Subsection 2.1, part of the link between pension fund allocations and MVE allocations with a benchmarking incentive arises directly from the benchmark weights being part of the ω^* expression. However, this mechanical component is not the full driver of the improved fit we observe. Specifically, the mechanical component has no effect for variation across funds (since benchmark weights do not vary across funds). Moreover, replacing raw weights with active weights still leads to slope coefficients of 0.40 ($t_{stat} = 2.99$ and $R^2 = 32.9\%$) across asset classes and 0.07 ($t_{stat} = 2.54$ and $R^2 = 0.7\%$) over time, reflecting a non-trivial improvement relative to the analogous specifications without benchmarking incentives.

The sixth specification in Table 4, "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)," further adds variation in risk aversion by estimating it separately for each consultant-year observation (as discussed in Subsection 2.3). The overall results are even stronger than the ones obtained under $\gamma = 3$. The biggest difference is on variation over time. In that case, the regression of ω onto ω^* has a slope coefficient of 0.50 ($t_{stat} = 4.55$, $R^2 = 36.6\%$) and the analogous regression using active weights has a slope coefficient of 0.31 ($t_{stat} = 2.11$, $R^2 = 14.9\%$).

Figure 5 provides a visual representation of our results through scatterplots of pension fund allocations against MVE allocations from our full specification, CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$). Each of the six scatterplots reflects a column of Table 4, varying based on the weight type (raw or active) and source of variation (across asset classes, across funds, or over time). For instance, Figure 5(a) plots ω against ω^* (both demeaned using $c \times t$ fixed effects to focus on

variation across asset classes). We can visually see the strong connection between ω and ω^* in all dimensions for both raw weights and active weights.

In addition, Figure 6 provides barplots summarizing the results from regressions of pension fund allocations on MVE allocations from our full specification separately for each asset class (for both raw weights and active weights). Each bar reflects a slope coefficient, with the numbers on the top of each bar reporting the respective t_{stat} (in parenthesis) and within R^2 values. The key message is that, after accounting for frictions and risk aversion variation, the positive link between pension fund and MVE allocations tends to be strong for each individual asset class whether we consider variation across funds or over time. A caveat is that each asset-class regression is based on only 122 observations. As a result, slope estimates exhibit non-trivial dispersion across asset classes and sources of variation, and some specifications (particularly those exploiting time-series variation in active weights) yield comparatively small and statistically insignificant coefficients.

Overall, the link between pension fund allocations and the MVE allocations implied from the beliefs of their consultants is significant (statistically and economically). Moreover, the link is much stronger once we account for non-negative weight constraints and benchmarking incentives, which reflect institutional and informational frictions.

3.3 Addressing Reverse Causality Concerns with the Regressions of ω onto ω^*

A potential concern with our regression results is reverse causality: pension funds may choose their target allocations based on motives unrelated to their views about future asset class returns (e.g., distortionary incentives) and then rely on reported beliefs merely to rationalize those choices. We argue that this channel is unlikely to explain our findings because we proxy for pension fund beliefs using the beliefs of their consultants rather than beliefs reported by the pension funds themselves. To clarify why this mitigates reverse causality, we distinguish between two possible mechanisms.

The first mechanism is that consultants choose their reported beliefs to justify the portfolio allocations that their pension fund clients wish to implement. This is implausible for

two reasons. First, as discussed in Section 1.4, general consultants advise a very large and diverse client base, which limits their incentives to tailor CMAs to their public pension fund clients at the expense of their other clients. Specifically, the average consultant in our sample advises 275 clients, with public pension funds representing only 17.3% of their clients and 7.2% of their regulatory assets under advisory. Second, the beliefs data are inconsistent with consultants catering to the incentives of their public pension fund clients. For instance, consultant μs are on average lower than the μ_{AARR} values reported by pension funds (Figure 2), indicating that consultants do not inflate their expected returns to conform to the incentive of these public pension funds to aggressively discount their liabilities (see also Footnote 10). Moreover, consultant μs aggregated across consultants are unbiased predictors of future realized asset class returns (Couts, Gonçalves, and Loudis (2024)). If consultant beliefs were chosen to rationalize arbitrary allocations driven by motives unrelated to their views about future returns, this forecasting property would not obtain.

The second mechanism is that pension funds strategically select consultants not ex ante based on shared beliefs but rather ex post to justify their desired portfolio allocations that are based on non-belief factors. While selection based on factors other than beliefs may be present in the data, it cannot explain our results for two reasons. First, pension funds rarely change their general consultants: in our baseline sample, only 6.2% of fund-years correspond to the start of a new fund-consultant relationship, implying an expected relationship length of 16.1 years (see Internet Appendix Figure IA.1 for details). In contrast, beliefs implied by CMAs vary substantially over time (Couts et al. (2025) and Dahlquist and Ibert (2025)). If pension funds were dynamically selecting consultants mainly to validate their time varying desired allocations, we would observe much more frequent fund-consultant switching than what is present in the data. Second, if the link between ω onto ω^* we uncover was driven by the relatively few changes in fund-consultant relationships in our sample, we would expect this link to weaken as time since the start of the relationship increases. Internet Appendix C.2 shows that this is not the case. If anything, the strength of the link between ω onto ω^*

increases with the duration of the fund-consultant relationship.[17]

3.4 Addressing Other Potential Concerns with the Regressions of ω onto ω^*

Another potential worry with our regression results is that part of the link between pension fund allocations and MVE allocations is due to the inclusion of ω^b into the ω^* expression through the benchmarking incentive. Our baseline analysis deals with this issue through the use of active weights since $\widetilde{\omega}^*$ is proportional to $\widetilde{\omega}^{mv}$ so that it does not have the extra ω^b term (as shown in Equation 7). However, in this baseline approach, we explain variation in active weights, not raw weights. Internet Appendix C.1 considers an alternative approach that focuses on explaining raw weights. Specifically, we regress ω onto ω^* controlling for ω^b. The results are very similar to the ones we report in Table 4 using active weights.

More broadly, Internet Appendix C.1 also reports regressions that control for some important fund characteristics studied in the prior literature: pension fund size, a dummy for whether the pension fund has an internal team to manage its equity portfolio, the GASB funding ratio, and net cash flows over assets. We consider these characteristics in isolation as well as their interactions with asset class dummies to account for their heterogeneous effects across asset classes. The advantage of adding these controls is that they proxy for pension fund incentives. The disadvantage is that they may also reflect pension fund beliefs, and we do not want to control for variation in beliefs. Nevertheless, the results suggest that MVE allocations remain strongly linked to pension fund allocations after including these controls.

In addition, we consider a set of robustness checks on basic empirical decisions needed for our analysis. In particular, Internet Appendix C.3 reports results using three-way wild bootstrap clustering (by n, c, t) instead of two-way clustering (by c, t) as in our baseline analysis. The three-way clustering allows us to account for potential persistent shocks affecting the be-

[17] One possible explanation for the increasing strength of the link between ω onto ω^* over the duration of the fund-consultant relationship is that the views of pension funds are more influenced by the views of consultants they trust more, with trust plausibly increasing as the relationship matures. However, our argument does not require taking a stand on the specific economic mechanism underlying this pattern. What matters to us is that the link between ω onto ω^* does not decline over the duration of the fund-consultant relationship, as it would be the case if this link was driven by the reverse causality channel.

liefs of all consultants about individual asset classes, but requires a wild bootstrap approach since we only have five asset classes (see Cameron, Gelbach, and Miller (2008) for details on wild bootstrap clustering). In addition, Internet Appendix C.4 reports results using equal-weighted averages across funds instead of AUM-weighted averages (e.g., when constructing ω and ω^b). Similarly, Internet Appendix C.5 provides results that rely on actual pension fund allocations instead of target pension fund allocations. Finally, Internet Appendix C.6 covers results that replace consultant μs with pension fund μ_{AARR} values when available. In each of these four robustness checks, the results are similar to the ones we provide in the main text.

We also provide an analysis that expands our sample by including CMAs that do not cover Hedge Funds (in these cases, we remove Hedge Funds from ω and the mean-variance optimization we use to construct ω^*). The sample increases from 122 consultant-years (1,422 fund-years) to 189 consultant-years (2,060 fund-years). Yet, the overall results (reported in Internet Appendix C.7) remain similar to the ones we report in the main text.

4 Decomposing Variation in Pension Fund Allocations

The prior section shows that pension fund allocations have a strong connection to the MVE allocations implied from the beliefs of their consultants. However, the prior section is based on regressions of ω onto ω^*, which do not take full advantage of our structural mean-variance framework. Specifically, the mean-variance structure in Section 2.1 implies that each pension fund allocation vector is given by $\omega = \omega^\alpha + \omega^*$, where ω^* is the MVE allocation from the perspective of its consultant and ω^α is the allocation wedge arising from belief wedges between the pension fund and its consultant. Consequently, we have the variance decomposition

$$1 = \underbrace{\frac{\mathbb{C}ov[\omega\ ,\ \omega^\alpha]}{\mathbb{V}ar[\omega]}}_{\text{Fraction of } \mathbb{V}ar[\omega] \text{ driven by } \omega^\alpha} + \underbrace{\frac{\mathbb{C}ov[\omega\ ,\ \omega^*]}{\mathbb{V}ar[\omega]}}_{\text{Fraction of } \mathbb{V}ar[\omega] \text{ driven by } \omega^*}. \tag{10}$$

This mean-variance structure implies that the R^2 values from the regressions in the prior section understate the strength of the link between ω and ω^*. To see this, note that pension fund beliefs move less than one-to-one with consultant beliefs (as shown in Figure 2), and

thus an increase in ω^* is associated with a decline in $\omega^\alpha = \omega - \omega^*$. As such, ω^α and ω^* are negatively correlated so that a regression of ω onto ω^* as in Equation 9 leads to a downward bias in the slope estimate. That is, the regression $\omega = \text{FE} + \text{Slope} \cdot \omega^* + \epsilon$ has an omitted variable (ω^α) that correlates negatively with ω^*, inducing Slope < 1.[18] Since $R^2 = \frac{\mathbb{V}ar[\text{Slope}\cdot\omega^*]}{\mathbb{V}ar[\omega]} = \text{Slope} \cdot \frac{\mathbb{C}ov[\omega\,,\,\omega^*]}{\mathbb{V}ar[\omega]}$ and Slope < 1, we then have that the R^2 values reported in the prior section are lower than the fraction of $\mathbb{V}ar[\omega]$ driven by ω^* under our structural mean-variance framework.

So, this section explores the link between ω and ω^* directly from variance decompositions as in Equation 10. Importantly, since the motivation for these variance decompositions rely more directly on the structural interpretation of our mean-variance framework, we focus on the specifications that incorporate non-negative weight constraints and benchmarking incentives. This approach ensures that we can properly interpret the variance decompositions since actual pension funds face institutional and informational frictions that lead to non-negative weight constraints and benchmarking incentives.

Table 5 reports results from these variance decompositions. Specifically, we compute $\omega^\alpha = \omega - \omega^*$ and report the slope coefficients (in percentage terms) from regressions of ω^α and ω^* onto ω (and analogously for active weights). We include the same fixed effects as in Table 4 to separately capture variation across across funds, across asset classes, and over time. The results indicate that most of the variation in observed allocations across pension funds and asset classes can be explained by MVE allocations. For instance, MVE allocations explain around 60% of the variation in pension fund allocations across funds. Similarly, MVE allocations explain between 74% and 88% of the variation in pension fund allocations across asset classes depending on what type of weights we focus on (raw or active) and whether we incorporate variation in risk aversion or not. In terms of variation over time, incorporating variation in risk aversion is more consequential. With a $\gamma = 3$, MVE allocations explain 38% of $\mathbb{V}ar[\omega]$ and 10% of $\mathbb{V}ar[\widetilde{\omega}]$. In contrast, with $\gamma = \hat{\gamma}_{c,t}$, MVE allocations explain 73% of

[18] The correlation between ω^* and $\omega^\alpha = \omega - \omega^*$ is -0.42 across pension funds, -0.07 across asset classes, and -0.60 over time. These correlations are consistent with the Slope < 1 estimates from the prior section.

$\mathbb{V}ar[\omega]$ and 48% of $\mathbb{V}ar[\widetilde{\omega}]$.

Figure 7 provides barplots summarizing the results from variance decompositions of ω (or $\widetilde{\omega}$) for each asset class. Each bar reflects the fraction of $\mathbb{V}ar[\omega]$ (or $\mathbb{V}ar[\widetilde{\omega}]$) driven by the MVE allocations from our full specification, with the respective t_{stat} values on the top of the bars. The key message is that, after accounting for frictions and risk aversion variation, MVE allocations tend to explain a meaningful fraction of the variation in pension fund allocations across funds and over time for all asset classes. As with Figure 6, one caveat is that we only have 122 observations when focusing on a single asset class, and thus the fraction of weight variability explained by MVE weights is below 30% and statistically insignificant in a few instances.

Internet Appendices C.4 through C.6 report robustness analyses for Table 5. These analyses parallel the robustness checks in Section 3.3 for Table 4. The findings are highly consistent with those in Table 5, reinforcing the robustness of our main findings.

Overall, a large share of the variation in pension fund allocations (ω) is explained by the corresponding MVE allocations (ω^*). This relationship holds for each asset class and remains robust to empirical choices. However, it is important to highlight that these results rely more directly on the structural interpretation of our mean-variance framework in Subsection 2.1. As such, our interpretation of these findings relies on the plausibility of our non-negative weight constraints and benchmarking incentives as properly describing frictions faced by pension funds.

5 Studying the MVE Allocation Wedges

The prior sections focus on the link between pension fund allocations (ω) and the MVE allocations implied by the beliefs of their consultants (ω^*). In this section, we instead study the MVE allocation wedges: $\omega^\alpha = \omega - \omega^*$. Subsection 5.1 quantifies the economic magnitude of these MVE allocation wedges while Subsection 5.2 explores their link to fund characteristics as in general MVE allocation wedges can arise from belief wedges or other factors affecting pension fund allocations (e.g., distortionary incentives). As in the previous section, we restrict

attention to the two specifications that incorporate non-negative weight constraints and benchmarking incentives, as these frictions are essential for a realistic assessment of MVE allocation wedges.

5.1 MVE Allocation Wedges: Perceived Economic Magnitude

To start, Figure 8 plots the density functions of ω^{α} for each asset class. Clearly, the ω^{α} distributions center around zero, consistent with the results from Table 3. However, the distributions have non-trivial dispersion, indicating that values of ω^{α} far from zero are fairly common. This result highlights that MVE allocation wedges can be important even though ω^{*} explains the majority of the variation in ω.

Nevertheless, it is hard to quantify the economic importance of allocation wedges entirely from the magnitudes of ω^{α}. The reason is that large values of $|\omega^{\alpha}|$ may translate into small differences in utility from ω to ω^{*}. To address this issue, we follow an approach similar to the one used by Campbell and Thompson (2008) to evaluate the benefits of market timing. That is, we calculate the utility gain from moving from the suboptimal allocation ω to the optimal allocation ω^{*} (where optimality is with respect to the beliefs of consultants). Specifically, for each consultant-year, we evaluate the perceived utility obtained from the allocations ω and ω^{*} (under the objective function in Equation 1 evaluated under consultant beliefs). Then, we calculate $\Delta\text{Utility} = \text{Utility}(\omega^{*}) - \text{Utility}(\omega)$. This utility gain can be interpreted as the annual fee an investor would be willing to pay to move from the allocation ω to ω^{*} (if the investor had the utility function in Equation 1 and the same beliefs as its consultant).[19]

Figure 9 plots the density functions of ΔUtility for the two specifications that incorporate non-negative weight constraints and benchmarking incentives. The median annual fee is around 0.15% for both specifications. Whether this value is large or not is somewhat subjective. However, as a reference point, we calculate the aggregate cost pension funds incur

[19] An alternative approach would be to calculate the Sharpe ratio differential $\Delta\mathbb{SR} = \mathbb{SR}(\omega^{*}) - \mathbb{SR}(\omega)$. However, maximizing Sharpe ratio is not the objective function in Equation 1. As such, $\Delta\mathbb{SR}$ values can be misleading (they can even be negative). Empirically, the median $\Delta\mathbb{SR}$ is close to 0.01 (in annual units) for both specifications we focus on. However, in both cases we also have some negative $\Delta\mathbb{SR}$ values, which makes the center of the $\Delta\mathbb{SR}$ distribution less informative about the typical economic magnitude of ω^{α} .

to manage their investments.[20] The average and median annual aggregate costs are 0.46% and 0.51%, respectively. So, under the beliefs of consultants, the annual cost of holding the suboptimal allocation ω when allocation ω^* is available is roughly one third of the annual cost of managing the pension fund investments. We view this result as suggesting that MVE allocation wedges are economically relevant, but not as large as some readers might have thought. Nevertheless, Figure 9 also shows that ΔUtility has substantial dispersion and is quite large (above 1% per year) for a relatively small fraction of our sample, highlighting the large heterogeneity present in the economic magnitudes of ω^α across institution-years.

5.2 MVE Allocation Wedges: Link to Fund Characteristics

Under the structural interpretation of our mean-variance framework in Subsection 2.1, MVE allocation wedges arise entirely from belief wedges between pension funds and their consultants. More generally, however, pension fund allocations should respond to other factors beyond beliefs. For example, it is natural to expect pension fund allocations to be affected by the fund's liabilities and liquidity needs. We explore this aspect below.

Table 6 reports regressions of MVE allocation wedges (ω^α) onto lagged fund characteristics. As before, we focus on the specifications that incorporate non-negative weight constraints and benchmarking incentives. The fund characteristics are pension fund size, a dummy for whether the pension fund has an internal team to manage its equity allocation, the GASB funding ratio, and net cash flows over assets.[21] Pension fund size captures

[20] We obtain the aggregate cost pension funds incur to manage their investments using the CEM Benchmarking dataset (from 2001 to 2022). It captures the costs for internally managed assets and external investments (including funds-of-funds). Internal investment costs include compensation and benefits of employees managing internal portfolios as well as expenses for support staff, consulting, research, legal, trading services, and allocated overhead costs. External investment costs include the management fees paid to investment consultants and external asset managers as well as costs (compensation, travel, and education expenses) for internal staff whose sole responsibility is to select and monitor external managers. However, external costs do not include performance fees, carried interest, and rebates (as they are typically directly subtracted from the returns). As such, our cost estimates are lower than the true costs pension funds incur to manager their investments.

[21] Fund size is based on the fund AUM. The internal management team dummy is constructed using the Thomson-Reuters 13F reports on direct U.S. equity holdings. The GASB Funding Ratio (assets over liabilities) measures the funding status of a pension fund under GASB guidelines that use the expected rate

potential economies of scale in investment costs and superior access to external managers, which could be relevant for allocation decisions in alternative assets. Net cash flows capture differences in liquidity needs and investment horizons across funds, and financial theory suggests that investors should adopt more conservative portfolios when faced with cash outflows. Establishing an internal investment team could also relate to the allocation wedges, as it requires hiring and retaining investment staff and could reflect differences in long-term beliefs about an asset class. Finally, while the GASB funding status does not represent a comparable market measure of funding, it could affect actuarial, financial, and investment decisions. Depending on the political process and the governance structure, pension funds with different GASB funding ratios face heterogeneous incentives to reach for yield and to inflate the discount rate used for valuing liabilities (i.e., the expected rate of return under GASB guidelines).

The key finding from Table 6 is that these four simple fund characteristics explain a non-trivial fraction of the variation in MVE allocation wedges in both MVE specifications. However, the level of variation explained varies markedly across asset classes. For example, for public and private equity, fund characteristics explain between 19% to 26% of the variation in ω^{α}. In contrast fund characteristics explain a relatively small portion of the variation in allocation wedges within fixed income, real estate, and hedge funds.

Precisely interpreting the coefficients in these regressions is hard given the highly endogenous nature of the variables involved. Nevertheless, we find that larger pension funds tend to invest less in public equity and more in private equity than justified by the beliefs of their consultants. One potential reason is that variation in allocations to private equity (relative to public equity) is driven not only by beliefs but also by access, with larger funds having better access to private markets. We also find that funds that have an internal public equity team have a positive allocation wedge to public equity and a negative one to private equity. One possibility for this result is that having an internal public equity team proxies for how

of return as a discount rate. And the net cash flows over assets variable captures the difference between the inflows from contribution payments and the outflows for benefit payments and administrative expenditures, scaled by pension fund assets at the beginning of the year.

bullish the pension fund is on public equity (i.e., a positive public equity belief wedge), with public and private equity treated as substitutes. Finally, we find that a higher GASB funding ratio (i.e., a better funding position) is associated with a positive allocation wedge to private equity and a negative one to hedge funds. One potential explanation is that funds that have a better funding status have more flexibility to keep their allocation to alternatives invested in more illiquid assets (with private equity being more illiquid than hedge funds). Of course, these interpretations remain suggestive, as other mechanisms may also be at play. Further research could explore the underlying economic drivers of these allocation wedges in greater depth.

6 Conclusion

This paper examines how the beliefs of institutional investors shape their multi-asset portfolio allocations by linking the observed allocations of U.S. public pension funds to the long-term capital market assumptions of their general consultants. By embedding these beliefs in a flexible mean-variance framework that accounts for consultant-fund belief wedges, we show that expectations about returns, volatilities, and correlations explain a substantial share of the variation in pension fund allocations (across pension funds, across asset classes, and over time). Accounting for diversification benefits embodied in cross-asset correlations is an important innovation of our paper relative to prior work, and we show that it matters quantitatively for the link between MVE allocations and observed allocations. Even more importantly, we demonstrate that the pass through of beliefs to portfolio allocations is muted when frictionless MVE allocations are used to capture the effect of beliefs, but becomes strong after accounting for non-negative weight constraints and benchmarking incentives (arising from institutional or informational frictions), especially if we also incorporate heterogeneity in risk aversion.

Our findings highlight that beliefs play a central role in institutional portfolio choice, but their effects are mediated by the frictions that shape investment decisions. Accounting for these frictions reconciles the apparent weak pass through of MVE allocations to observed

allocations and clarifies how beliefs are translated into portfolio positions. More broadly, our results underscore the importance of modeling beliefs jointly with the institutional environment in which they are formed and acted upon. Finally, our framework can be applied more generally to study the belief-driven behavior of other classes of investors, opening the door to addressing a wider set of questions in financial economics.

References

Adam, K. and S. Nagel (2023). "Expectations Data in Asset Pricing". In: *Handbook of Economic Expectations*. Academic Press. Chap. 16, pp. 477–506.

Andolfatto, A. and F. Bastianello (2025). "From Numbers to Words: Breaking Down Institutional Beliefs". Working Paper.

Andonov, A., M. Bonetti, and I. Stefanescu (2025). "Choosing Pension Fund Investment Consultants". Working Paper.

Andonov, A. and J. D. Rauh (2022). "The Return Expectations of Public Pension Funds". In: *Review of Financial Studies* 35.8, pp. 3777–3822.

Basak, S. and A. Pavlova (Aug. 2013). "Asset Prices and Institutional Investors". In: *American Economic Review* 103.5, pp. 1728–58.

Bastianello, F. (2024). "Time-Series and Cross-Section of Risk Premia Expectations: A Bottom-Up Approach". Working Paper.

Bastianello, F., P. H. Décaire, and M. Guenzel (2025). "Mental Models and Financial Forecasts". Working Paper.

Bastianello, F. and C. Peng (2025). "Global Fund Managers' Beliefs, Perceived Mispricing, and Asset Allocation". Working Paper.

Begenau, J., P. Liang, and E. Siriwardane (2025). "The Rise in Alternatives". In: *Review of Financial Studies* Forthcoming.

Beutel, J. and M. Weber (2023). "Beliefs and Portfolios: Causal Evidence". Working Paper.

Black, F. (1972). "Capital Market Equilibrium with Restricted Borrowing". In: *Journal of Business* 45.3, pp. 444–455.

Black, F. and R. Litterman (1991). "Asset Allocation: combining investor views with market equilibrium". In: *Journal of Fixed Income* 1.2, pp. 7–18.

Black, F. and R. Litterman (1992). "Global Portfolio Optimization". In: *Financial Analyst Journal* 48.5, pp. 28–43.

Brennan, M. J. (1993). "Agency and Asset Pricing". Working Paper.

Buffa, A. M., D. Vayanos, and P. Woolley (2022). "Asset Management Contracts and Equilibrium Prices". In: *Journal of Political Economy* 130.12, pp. 3146–3201.

Cameron, C. A., J. B. Gelbach, and D. L. Miller (2008). "Bootstrap-based improvements for inference with clustered errors." In: *Review of Economics and Statistics* 90.3, pp. 414–427.

Campbell, J. Y. and S. B. Thompson (2008). "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" In: *Review of Financial Studies* 21.4, pp. 1509–1531.

Couts, S. J., A. S. Gonçalves, J. Loudis, and Y. Liu (2025). "Institutional Investors' Subjective Risk Premia: Time Variation and Disagreement". Working Paper.

Couts, S. J., A. S. Gonçalves, and J. A. Loudis (2024). "The Subjective Risk and Return Expectations of Institutional Investors". Working Paper.

Dahlquist, M. and M. Ibert (2024). "Equity Return Expectations and Portfolios: Evidence from Large Asset Managers". In: *Review of Financial Studies* 37.6, pp. 1887–1928.

Dahlquist, M. and M. Ibert (2025). "Institutions' Return Expectations across Assets and Time". Working Paper.

Décaire, P. H. and J. R. Graham (2024). "Valuation Fundamentals". Working Paper.

Delao, R., X. Han, and S. Myers (2024). "The Cross-section of Subjective Expectations: Understanding Prices and Anomalies". Working Paper.

Delao, R. and S. Myers (2021). "Subjective Cash Flow Expectations and Discount Rate Expectations". In: *Journal of Finance* 76.3, pp. 1339–1387.

Egan, M., A. MacKay, and H. Yang (2024). "What Drives Variation in Investor Portfolios? Estimating the Roles of Beliefs and Risk Preferences". Working Paper.

Evans, R., J.-P. Gomez, L. Ma, and Y. Tang (2024). "Peer Versus Pure Benchmarks in the Compensation of Mutual Fund Managers". In: *Journal of Financial and Quantitative Analysis* 59.7, pp. 3101–3138.

Gabaix, X. and R. S. J. Koijen (2022). "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis". Working Paper.

Giglio, S., M. Maggiori, J. Rillo, J. Stroebel, S. Utkus, and X. Xu (2025). "Investor Beliefs and Expectation Formation". In: *Annual Review of Financial Economics* Forthcoming.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). "Five Facts About Beliefs and Portfolios". In: *American Economic Review* 111.5, pp. 1481–1522.

Gromb, D. and D. Vayanos (2010). "Limits of Arbitrage". In: *Annual Review of Financial Economics* 2, pp. 251–275.

Ibert, M. (2025). "Tactical Asset Allocations of Large Asset Managers". Working Paper.

Koijen, R. S. J. and M. Yogo (2019). "A Demand System Approach to Asset Pricing". In: *Journal of Political Economy* 127.4, pp. 1475–1515.

Kremens, L., I. Martin, and L. Varela (2024). "Long-Horizon Exchange Rate Expectations". Working Paper.

Loudis, J. A. (2024). "Stock Price Reactions to the Information and Bias in Analyst-Expected Returns". In: *The Accounting Review* 99.4, pp. 281–313.

Miller, E. M. (1977). "Risk, Uncertainty, and Divergence of Opinion". In: *Journal of Finance* 32.4, pp. 1151–1168.

Novy-Marx, R. and J. Rauh (2011). "Public Pension Promises: How Big Are They and What Are They Worth?" In: *Journal of Finance* 66.4, pp. 1211–1249.

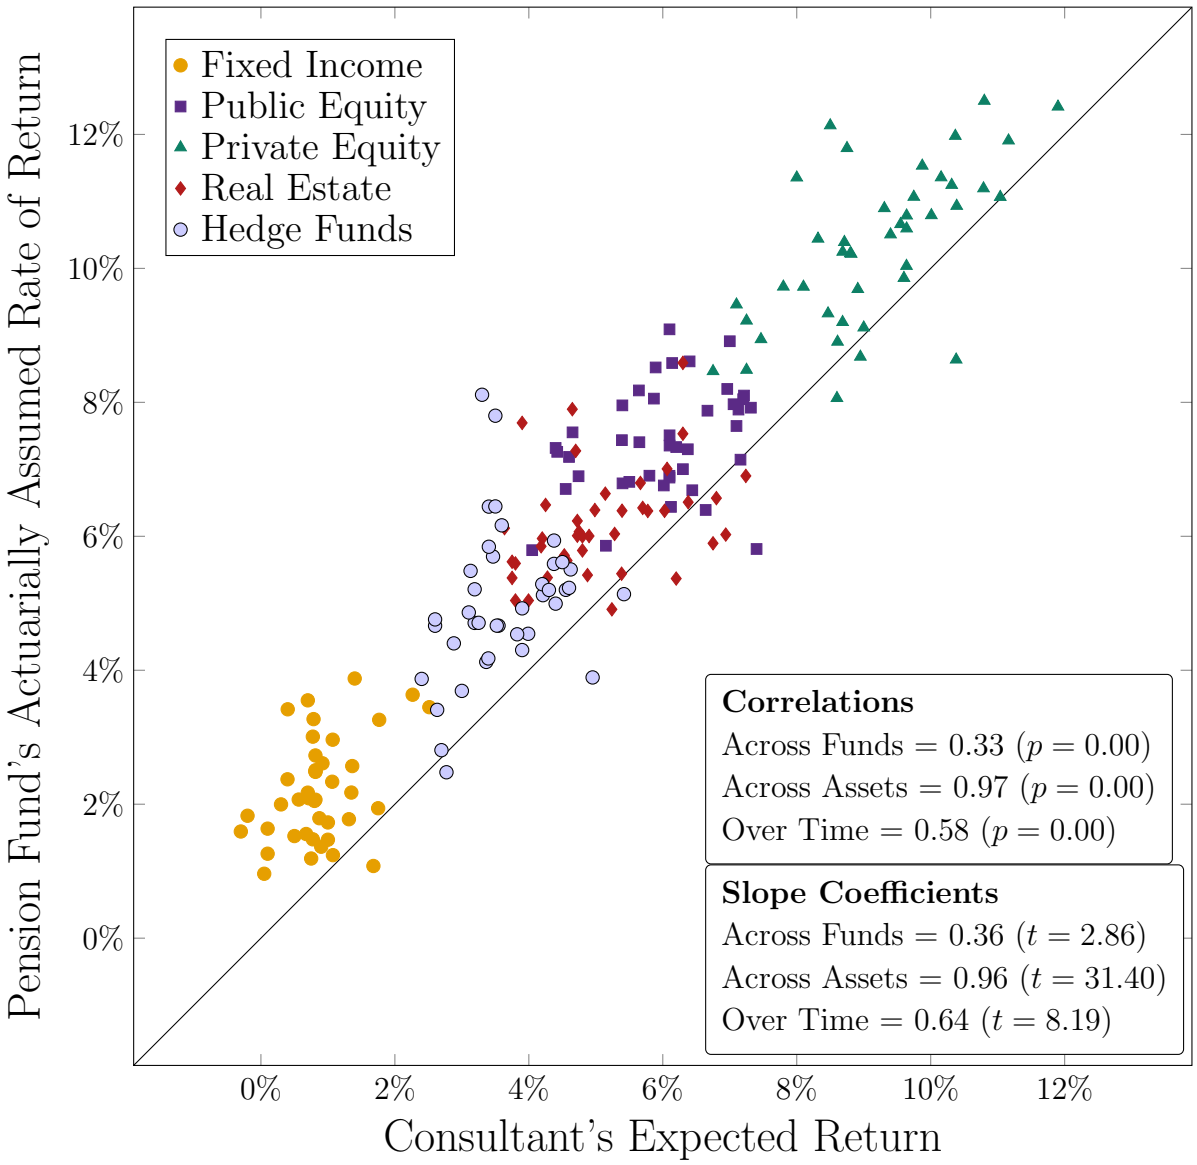


Figure 2
Pension Fund μ_{AARR} vs Consultant μ

This figure depicts pension fund actuarially assumed rates of return (μ_{AARR}) for five asset classes against the respective general consultant CMA-implied expected returns (μ) for the same five asset classes. Both μ_{AARR} and μ are in excess of their respective cash expected returns. Observations vary by consultant (c), asset class (n), and year (t). We obtain μ_{AARR} for consultant c in year t by aggregating the μ_{AARR} values of all pension funds matched to consultant c in year t. Section 1 describes the data on beliefs and Subsection 1.4 describes the results from this figure.

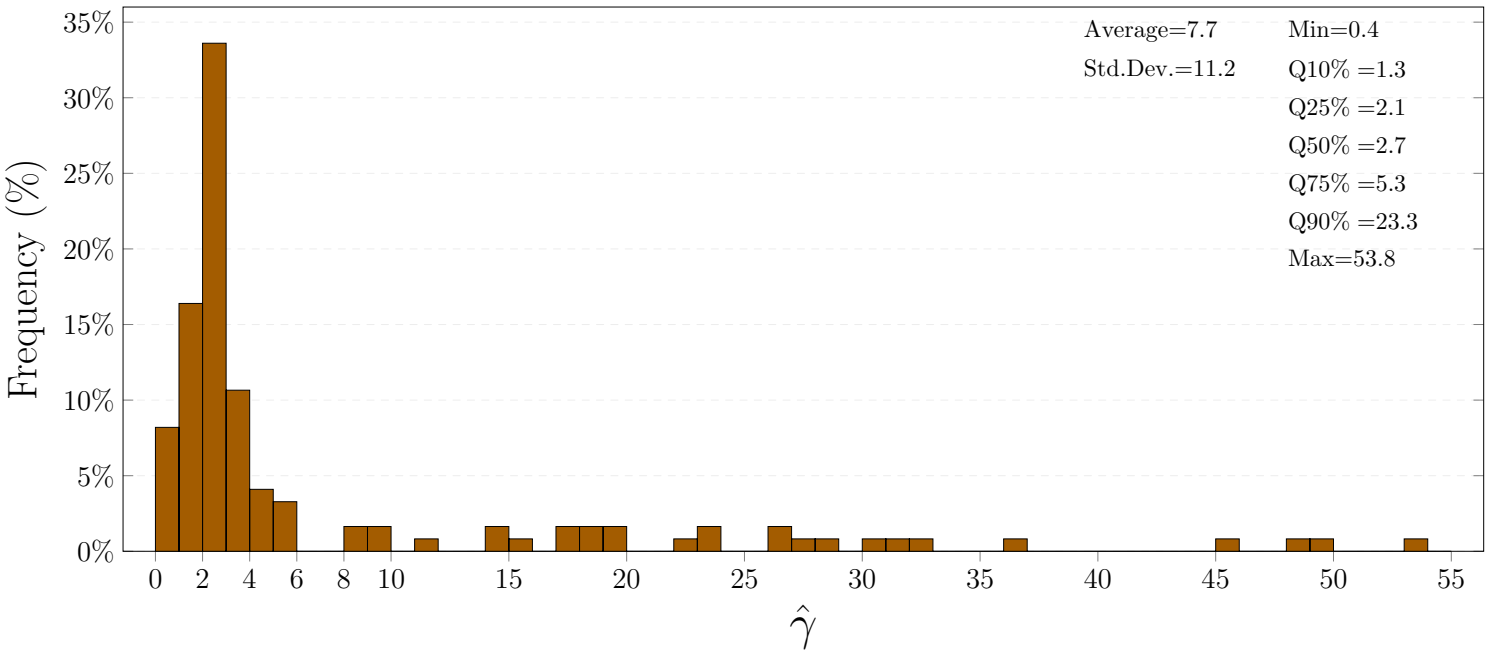


Figure 3
Distribution of Risk Aversion Estimates

This figure depicts the histogram of risk aversion estimates used for our MVE specification with non-negative weight constraints and a benchmarking incentive, "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." Observations vary by consultant (c) and year (t). Section 1 describes the data on beliefs and Subsection 2.3 describes the risk aversion estimation procedure and the results from this figure.

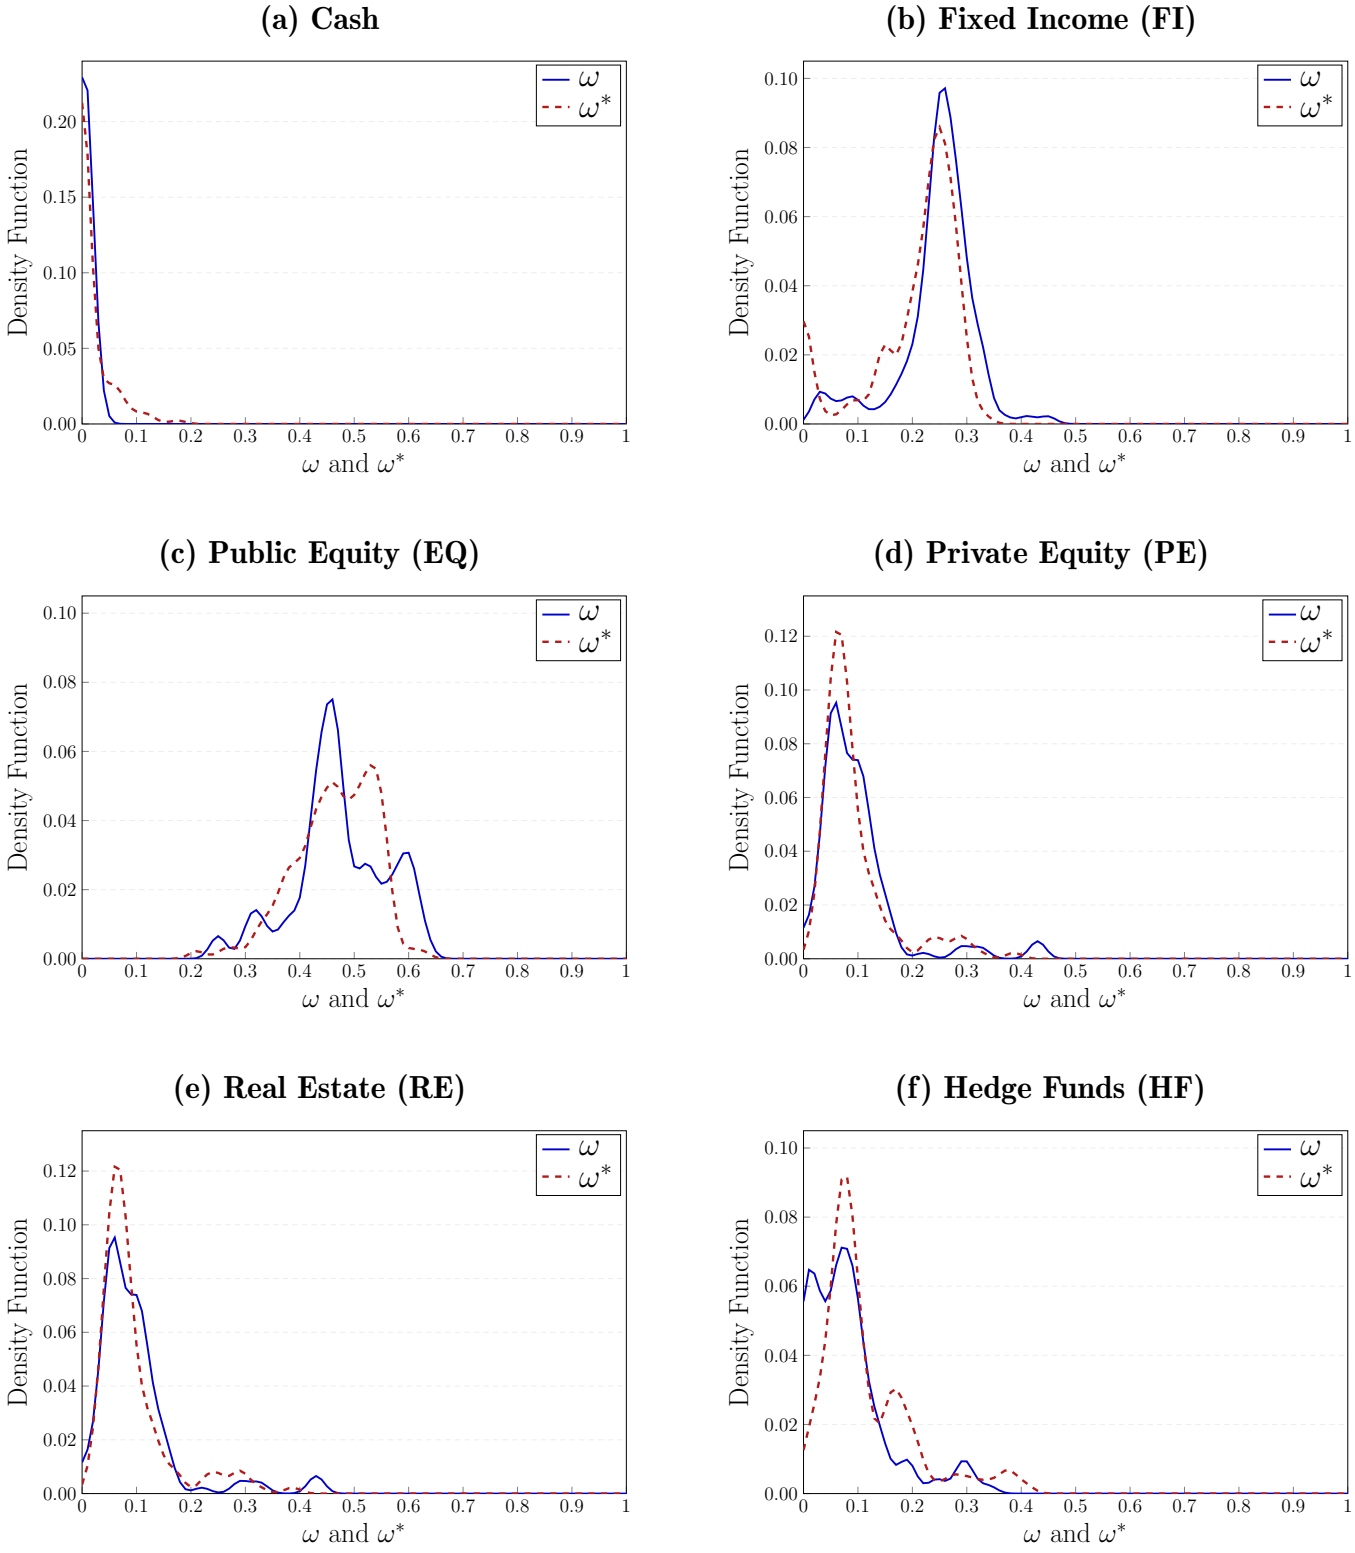


Figure 4
Distributions of Pension Fund Allocations and MVE Allocations

This figure depicts empirical density functions for pension fund portfolio allocations (ω) and the MVE allocations implied by the beliefs of their consultants (ω^*). The ω^* values are from our full MVE specification (with varying risk aversion, non-negative weight constraints, and a benchmarking incentive), labeled "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. Section 1 describes the data on beliefs and Subsection 2.3 describes the results from this figure.

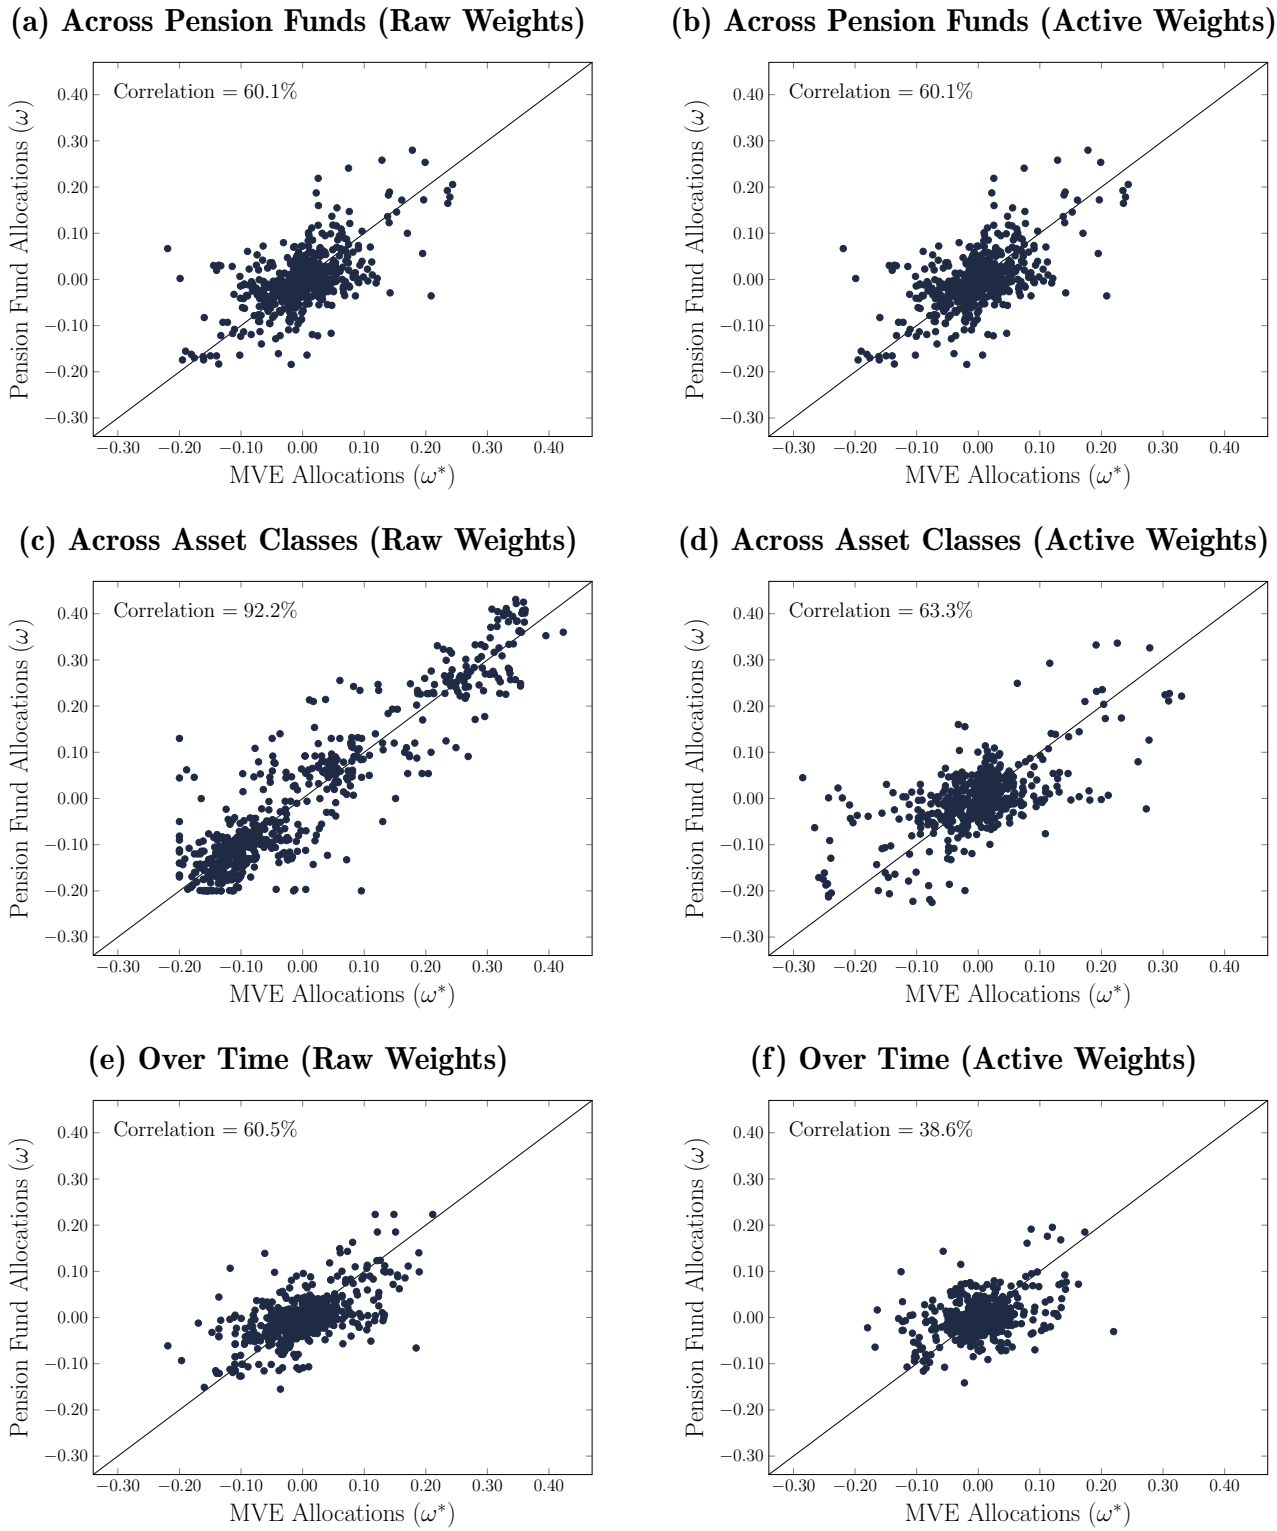


Figure 5
Pension Fund Allocations vs MVE Allocations

This figure depicts scatterplots of pension fund portfolio allocations (ω) against the MVE allocations implied by the beliefs of their consultants (ω^*), both demeaned by fixed effects to focus on different sources of variation. The ω^* values are from our full MVE specification (with varying risk aversion, non-negative weight constraints, and a benchmarking incentive), labeled "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. Section 1 describes the data on beliefs and Subsection 3.2 describes the results from this figure.



Figure 6

Regressing Pension Fund Portfolio Allocations on MVE Allocations (by Asset Class)

This figure reports slope coefficients (with t_{stat} and R^2_{within}) from regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of their consultants (ω^*). The ω^* values are from our full MVE specification (with varying risk aversion, non-negative weight constraints, and a benchmarking incentive), labeled "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. Bars capturing variation across pension funds (at the consultant level) have t fixed effects while bars capturing variation over time have c fixed effects. Section 1 describes the data on beliefs and Subsection 3.2 describes the results from this figure.

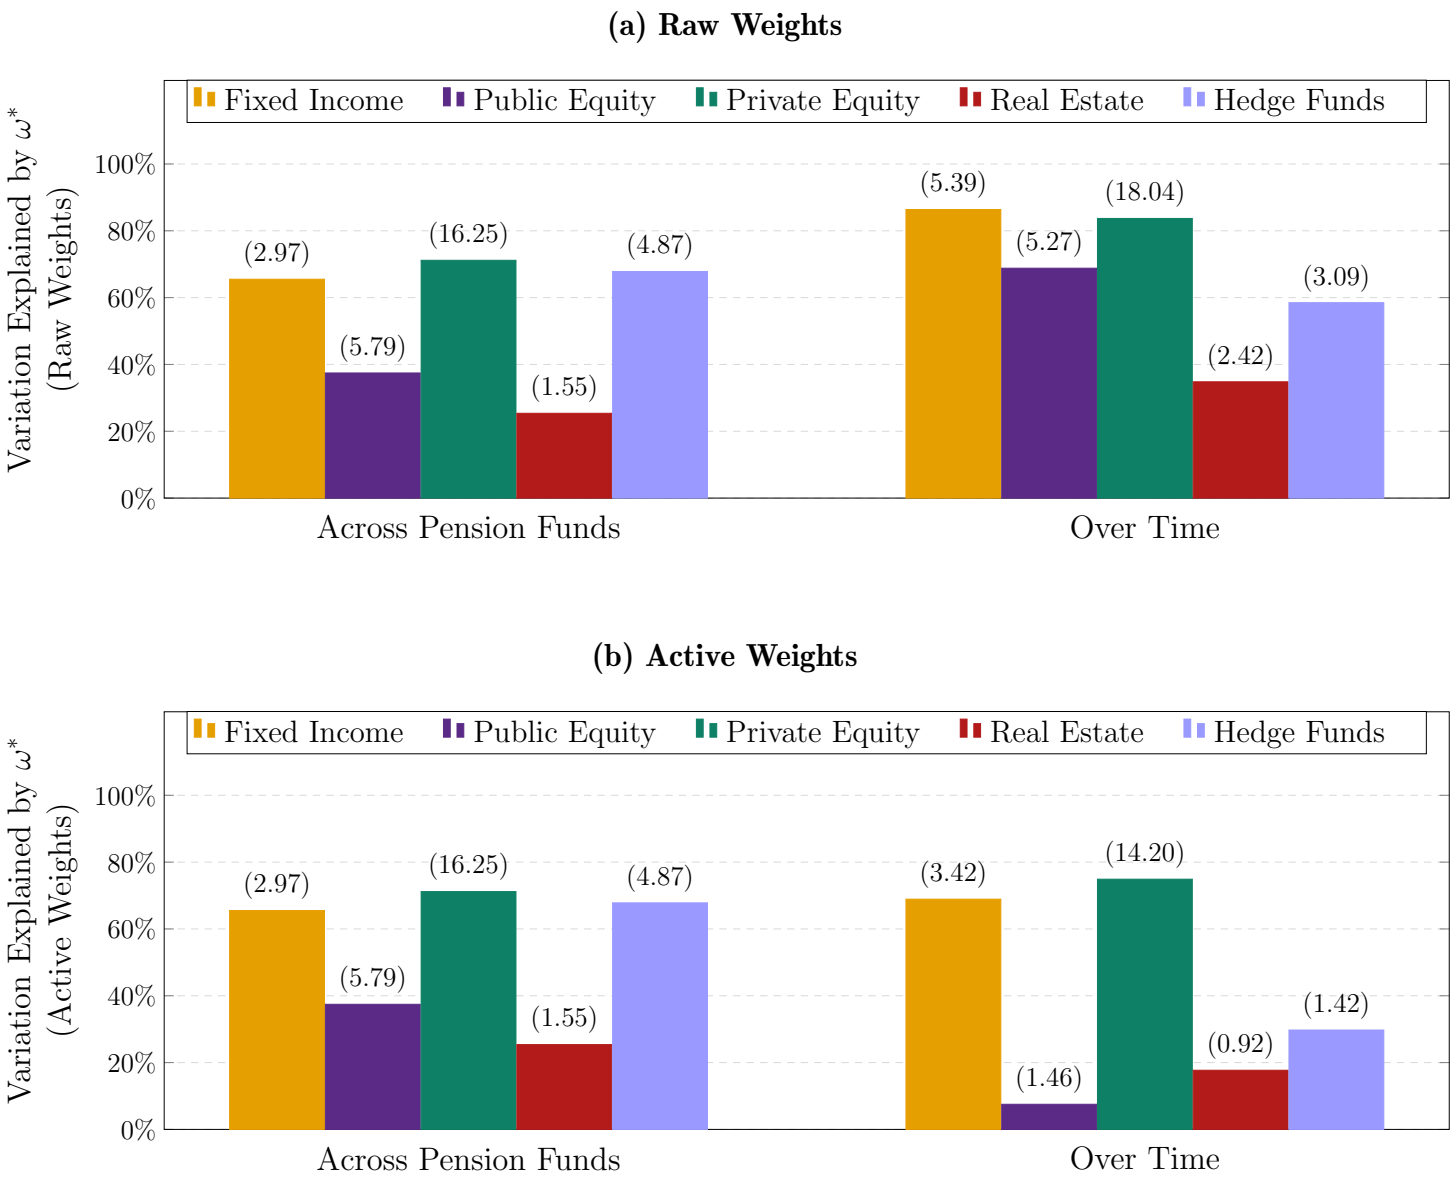
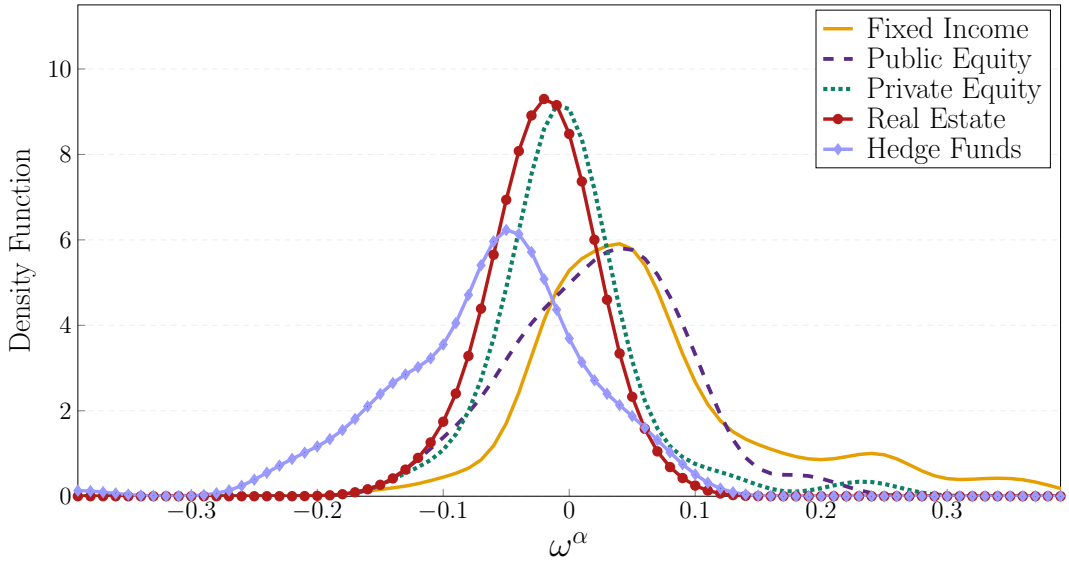


Figure 7
Decomposing Variation in Pension Fund Portfolio Allocations (by Asset Class)

This figure reports the fraction of variation in pension fund portfolio allocations (ω) explained by the MVE allocations implied by the beliefs of their consultants (ω^*), as well as the respective t_{stat} values. The ω^* values are from our full MVE specification (with varying risk aversion, non-negative weight constraints, and a benchmarking incentive), labeled "CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)." Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. We use Equation 10 for the decomposition of the variation in weights (ω) and an analogous equation for the decomposition of the variation in active weights ($\widetilde{\omega} = \omega - \omega^b$). Bars capturing variation across pension funds (at the consultant level) have t fixed effects while bars capturing variation over time have c fixed effects. Section 1 describes the data on beliefs and Section 4 describes the results from this figure.

(a) CMVE ($b \neq 0$ & $\gamma = 3$)



(b) CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)

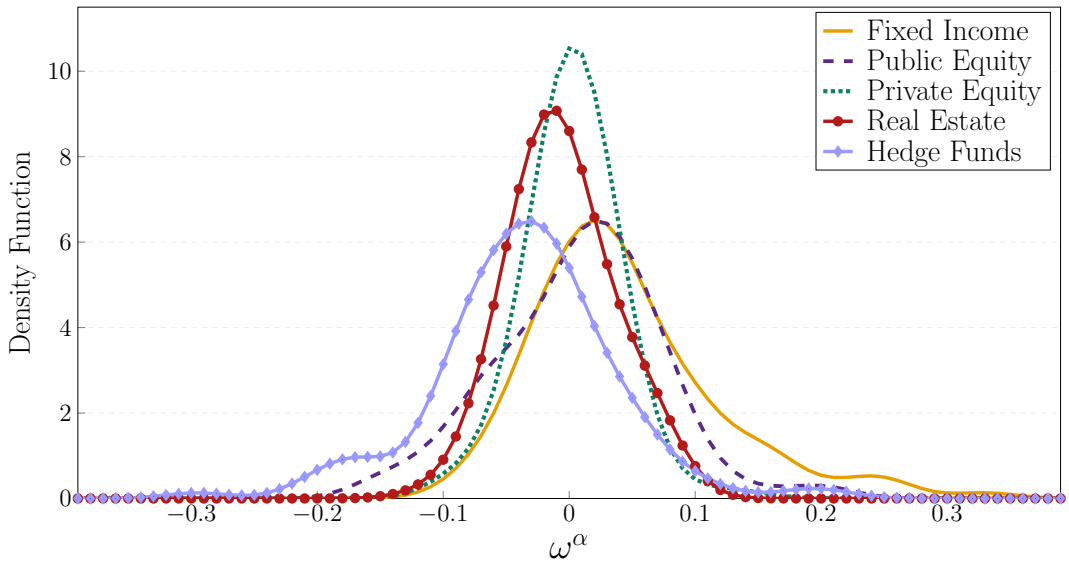


Figure 8
Distributions of MVE Allocation Wedges

This figure depicts empirical density functions for the allocation wedges ($\omega^{\alpha} = \omega - \omega^{*}$) between pension fund portfolio allocations (ω) and the MVE allocations implied by the beliefs of their consultants (ω^{*}). The ω^{*} values are from our MVE that consider non-negative weight constraints and a benchmarking incentive, with risk aversion calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. Section 1 describes the data on beliefs and Subsection 5.1 describes the results from this figure.

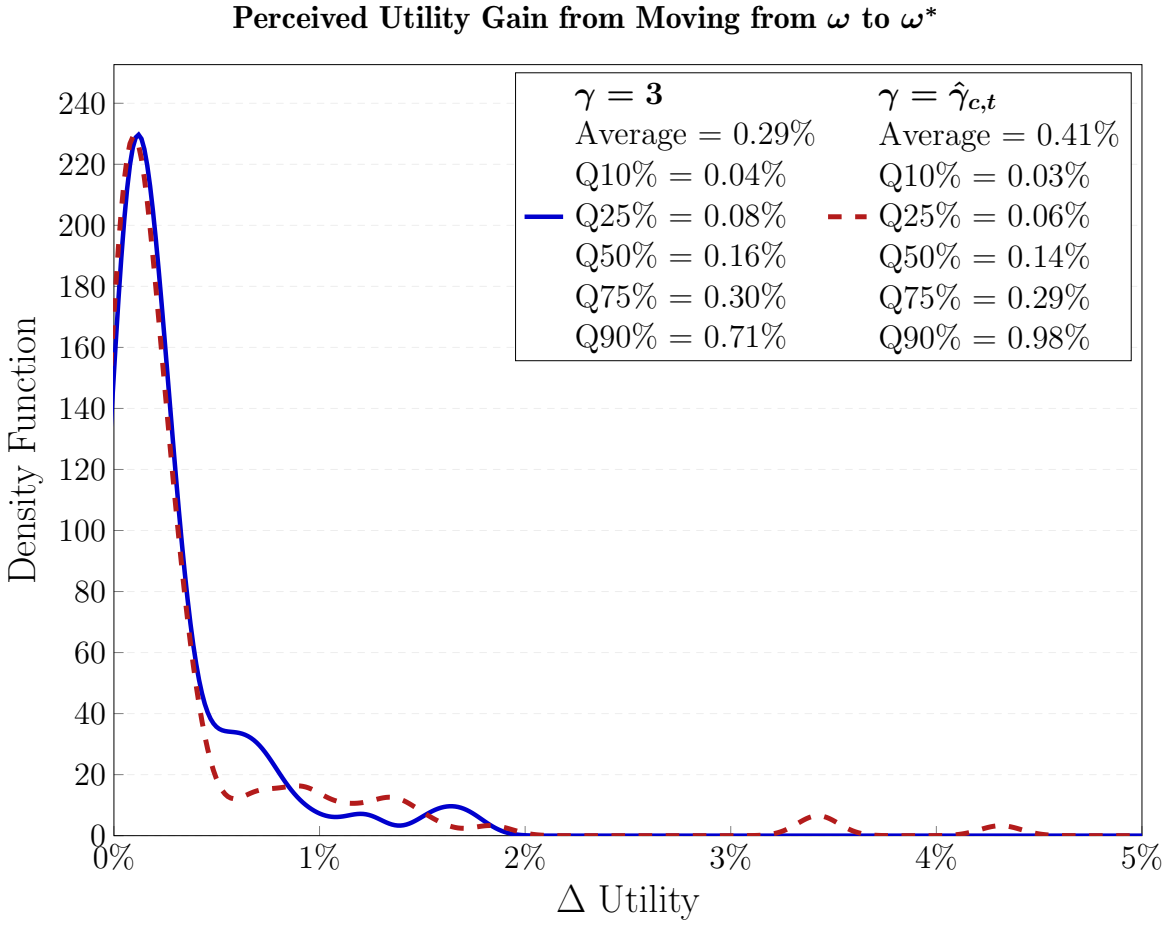


Figure 9
MVE Allocation Wedges: Perceived Economic Magnitudes

This figure depicts empirical density functions for the perceived utility gain from moving from the suboptimal pension fund allocation ω to the optimal MVE allocation ω^* (where optimality is with respect to the consultant beliefs). Specifically, for each consultant-year, we evaluate the perceived utility obtained from the allocations ω and ω^* (under the objective function in Equation 1 evaluated under consultant beliefs). Then, we calculate $\Delta\text{Utility} = \text{Utility}(\omega^*) - \text{Utility}(\omega)$. This utility gain can be interpreted as the annual fee an investor would be willing to pay to move from the allocation ω to ω^* (if the investor had the utility function in Equation 1 and the same beliefs as its consultant). The ω^* values are from our MVE that consider non-negative weight constraints and a benchmarking incentive, with risk aversion calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). Section 1 describes the data on beliefs and Subsection 5.1 describes the results from this figure.

Table 1

The Matching Between Pension Fund Dataset and Consultant CMAs

This table reports, for each allocation year in our sample, information on the raw pension fund dataset, the matching between that dataset and our dataset of CMAs, and the asset classes covered in the matched CMAs. In terms of the raw pension fund dataset, we report the total number of pension funds, the total number of general consultants to these pension funds, and the total dollar AUM of these pension funds. In terms of the matching to our dataset of CMAs, we report the fraction of pension funds, consultants, and dollar AUM we are able to match to consultant CMAs. In terms of the asset classes available in the CMAs, we report the number of CMAs covering fixed income (FI), public equity (EQ), private equity (PE), real estate (RE), and hedge funds (HF). Section 1.1 describes our dataset on pension funds, Section 1.2 describes our dataset on CMAs, and Section 1.3 describes the results from this table.

Allocation	In Pension Fund Dataset			% Matched to CMAs			# of Matched CMAs with*					
Year	Funds	Consultants	$AUM	Funds	Consultants	$AUM	FI	EQ	PE	RE	HF	All Asset Classes
2001	112	30	1.9T	30%	10%	30%	3	3	2	3	0	0
2002	119	31	1.8T	40%	19%	40%	6	6	4	5	1	1
2003	124	33	1.8T	35%	21%	32%	7	7	5	5	1	1
2004	132	33	2.1T	40%	18%	38%	6	6	4	5	1	1
2005	140	36	2.3T	39%	22%	39%	8	8	7	7	3	3
2006	144	38	2.6T	38%	21%	40%	8	8	8	8	5	5
2007	152	35	3.0T	41%	26%	40%	9	9	9	9	6	6
2008	157	37	2.7T	54%	30%	62%	11	11	11	11	7	7
2009	162	35	2.3T	44%	26%	43%	9	9	9	9	5	5
2010	163	33	2.5T	55%	33%	62%	11	11	11	11	6	6
2011	166	32	2.8T	57%	34%	55%	11	11	11	11	7	7
2012	164	30	2.8T	65%	37%	68%	11	11	11	11	8	8
2013	166	30	3.1T	57%	33%	56%	10	10	10	10	7	7
2014	175	29	3.5T	63%	34%	64%	10	10	10	10	8	8
2015	174	26	3.5T	58%	38%	60%	9	10	10	10	7	6
2016	175	25	3.5T	67%	36%	77%	9	9	9	9	6	6
2017	176	25	3.8T	65%	36%	71%	9	9	9	9	6	6
2018	177	24	4.0T	71%	38%	79%	9	9	9	9	7	7
2019	176	22	4.1T	71%	41%	82%	9	9	9	9	6	6
2020	176	21	4.2T	74%	38%	88%	8	8	8	8	6	6
2021	177	21	5.2T	78%	48%	90%	10	10	10	10	8	8
2022	177	22	4.8T	75%	36%	88%	8	8	8	8	6	6
2023	169	22	4.9T	67%	32%	75%	7	7	7	7	6	6
Total	3,653	670	73.1T	57%	30%	65%	198	199	191	194	123	122
Main Sample	3,653	670	73.1T	39%	18%	37%	122	122	122	122	122	122

*** FI:** Fixed Income, **EQ:** Public Equity, **PE:** Private Equity, **RE:** Real Estate **HF:** Hedge Funds

Table 2
Summary Statistics: Portfolio Allocations and Belief Elements

This table reports, for each asset class, summary statistics for portfolio allocations (ω) and belief elements (μ, σ, and ρ). In terms of allocations, ω_i reflects the pension fund allocations whereas ω_j captures the consultant-level allocations, which are given by the AUM weighted average of ω_i across all pension funds matched to the given consultant in the given year. In terms of belief elements, μ_c, σ_c, and ρ_c reflect the expected return, volatility, and correlation from the consultant CMA (all in excess of cash). The asset classes are fixed income (FI), public equity (EQ), private equity (PE), real estate (RE), and hedge funds (HF). The summary statistics reported are the number of observations, the pooled average and standard deviation of the given allocation or belief element, and the quantiles of the pooled distribution for the given allocation or belief element. Section 1 describes the data on portfolios allocations and beliefs and Subsection 1.3 describes the results from this table.

Statistic	Asset Class	# Obs	Average	Std. Dev.	Q1%	Q5%	Q10%	Q25%	Q50%	Q75%	Q90%	Q95%	Q99%
ω_i	Fixed Income	1,422	24.5%	7.2%	5.0%	13.7%	16.0%	20.0%	24.4%	29.0%	33.0%	36.0%	41.0%
	Public Equity	1,422	47.8%	10.3%	22.0%	30.0%	34.3%	40.9%	48.0%	55.0%	61.0%	64.5%	68.2%
	Private Equity	1,422	7.6%	6.3%	0.0%	0.0%	0.0%	2.0%	7.0%	11.0%	16.0%	18.0%	26.0%
	Real Estate	1,422	10.9%	5.8%	0.0%	0.0%	3.6%	8.0%	11.0%	15.0%	18.0%	20.0%	26.0%
	Hedge Funds	1,422	6.6%	7.6%	0.0%	0.0%	0.0%	0.0%	5.0%	10.0%	18.0%	21.0%	30.0%
ω_c	Fixed Income	122	24.5%	7.1%	3.0%	8.7%	16.7%	22.9%	25.5%	28.3%	31.1%	33.0%	41.0%
	Public Equity	122	47.1%	8.5%	25.0%	32.0%	36.9%	42.7%	46.4%	52.8%	59.3%	60.1%	62.5%
	Private Equity	122	9.7%	7.7%	0.0%	2.9%	4.0%	5.6%	8.0%	11.3%	15.0%	29.0%	43.0%
	Real Estate	122	10.5%	3.4%	0.0%	5.8%	6.5%	8.0%	10.7%	12.3%	15.1%	15.7%	17.5%
	Hedge Funds	122	8.0%	7.3%	0.0%	0.0%	0.2%	2.6%	6.8%	10.0%	17.6%	25.0%	30.8%
μ_c	Fixed Income	122	1.2%	0.7%	-0.2%	0.1%	0.3%	0.7%	1.1%	1.6%	2.0%	2.4%	3.3%
	Public Equity	122	6.1%	1.0%	3.2%	4.3%	4.6%	5.4%	6.1%	6.8%	7.3%	7.5%	7.7%
	Private Equity	122	9.5%	1.6%	6.2%	7.2%	7.5%	8.5%	9.5%	10.4%	11.4%	12.7%	13.4%
	Real Estate	122	4.8%	1.1%	2.5%	3.2%	3.4%	3.9%	4.8%	5.7%	6.4%	6.7%	6.9%
	Hedge Funds	122	3.8%	1.0%	2.1%	2.6%	2.7%	3.1%	3.5%	4.4%	5.4%	5.8%	6.4%
σ_c	Fixed Income	122	5.6%	1.6%	3.7%	3.8%	4.0%	4.3%	5.4%	6.9%	7.6%	8.7%	10.1%
	Public Equity	122	17.6%	1.0%	14.8%	16.0%	16.1%	16.9%	17.6%	18.1%	18.9%	19.0%	19.2%
	Private Equity	122	26.8%	4.6%	15.9%	20.1%	20.9%	23.9%	26.3%	30.0%	33.1%	34.0%	38.1%
	Real Estate	122	14.1%	2.7%	8.7%	9.1%	10.6%	12.5%	14.3%	16.4%	17.9%	18.0%	19.5%
	Hedge Funds	122	9.0%	2.3%	4.3%	5.0%	6.6%	7.8%	9.1%	9.9%	10.6%	11.9%	17.6%
ρ_c	FI-EQ	122	20.3%	20.6%	-9.5%	-9.0%	0.8%	5.8%	17.4%	27.9%	51.2%	70.6%	71.8%
	FI-PE	122	5.3%	19.1%	-33.7%	-22.6%	-18.8%	-8.2%	4.1%	19.2%	26.5%	33.0%	51.7%
	FI-RE	122	5.4%	20.1%	-56.1%	-22.3%	-20.3%	-6.5%	6.8%	17.4%	26.3%	33.6%	50.4%
	FI-HF	122	17.1%	21.2%	-19.7%	-15.8%	-11.1%	2.1%	15.3%	27.6%	51.7%	60.4%	72.1%
	EQ-PE	122	74.1%	11.6%	19.8%	61.8%	62.9%	70.0%	74.7%	80.3%	90.0%	91.6%	94.9%
	EQ-RE	122	38.6%	23.0%	-9.8%	-2.6%	7.3%	23.3%	40.6%	51.7%	70.1%	73.6%	74.8%
	EQ-HF	122	67.2%	15.3%	12.5%	47.5%	50.4%	59.4%	68.7%	80.3%	82.0%	86.5%	91.4%
	PE-RE	122	42.1%	20.2%	-0.6%	-0.1%	14.6%	29.4%	44.5%	55.2%	70.4%	70.9%	73.5%
	PE-HF	122	54.7%	18.9%	12.9%	28.7%	30.1%	33.9%	61.5%	68.2%	76.5%	81.1%	83.6%
	RE-HF	122	26.1%	22.9%	-26.6%	-25.2%	-4.9%	9.5%	24.7%	43.7%	53.0%	59.5%	64.5%

Table 3
Mean-Variance Efficient (MVE) Allocations

This table reports descriptive statistics for mean-variance efficient (MVE) allocations as implied by the beliefs of consultants, ω^*, for different specifications of our mean-variance solution. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The UMVE specification imposes no constraints (so that the lagrange multipliers underlying ω^{mv} are zero) whereas the CMVE specifications incorporate non-negative weight constraints. Specifications with $b = 0$ set the benchmarking incentive to zero whereas specifications with $b \neq 0$ calibrate the benchmarking incentive so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). Section 1 describes the data on pension fund portfolio allocations and consultant beliefs and Section 2 provides details about the MVE allocations and describes the results from this table.

AC	ω^* Specification	Average	Std. Dev.	Q1%	Q5%	Q10%	Q25%	Q50%	Q75%	Q90%	Q95%	Q99%
	UMVE ($b = 0$ & $\gamma = 3$)	107.7%	131.2%	-277.6%	-162.0%	-24.2%	45.3%	116.1%	188.3%	260.5%	308.7%	389.0%
	CMVE ($b = 0$ & $\gamma = 3$)	8.7%	13.6%	0.0%	0.0%	0.0%	0.0%	0.0%	15.6%	29.0%	38.3%	44.9%
FI	**CMVE ($b \neq 0$ & $\gamma = 3$)**	18.0%	9.9%	0.0%	0.0%	0.0%	12.4%	21.4%	25.7%	27.3%	28.6%	30.4%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	20.1%	8.8%	0.0%	0.0%	0.2%	17.3%	23.3%	25.8%	28.0%	28.9%	30.6%
	Pension Fund Allocations	24.5%	7.1%	3.0%	8.7%	16.8%	22.9%	25.6%	28.3%	31.1%	33.0%	40.1%
	UMVE ($b = 0$ & $\gamma = 3$)	-37.3%	96.2%	-411.8%	-197.6%	-147.4%	-68.8%	-14.4%	17.3%	49.0%	63.4%	93.2%
	CMVE ($b = 0$ & $\gamma = 3$)	11.7%	14.8%	0.0%	0.0%	0.0%	0.0%	1.4%	20.5%	35.9%	41.2%	46.5%
EQ	**CMVE ($b \neq 0$ & $\gamma = 3$)**	44.8%	7.6%	25.9%	30.4%	33.1%	41.2%	46.7%	50.5%	53.3%	54.1%	55.3%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	46.4%	7.3%	26.5%	33.9%	37.0%	42.0%	46.6%	52.3%	54.7%	55.7%	58.8%
	Pension Fund Allocations	47.2%	8.5%	25.0%	32.0%	36.9%	42.8%	46.5%	53.1%	59.4%	60.6%	62.4%
	UMVE ($b = 0$ & $\gamma = 3$)	26.6%	73.2%	-226.3%	-42.2%	-23.2%	0.6%	19.9%	45.2%	115.3%	128.2%	232.9%
	CMVE ($b = 0$ & $\gamma = 3$)	22.9%	16.8%	0.0%	0.0%	1.7%	12.5%	21.1%	29.9%	46.3%	56.4%	67.4%
PE	**CMVE ($b \neq 0$ & $\gamma = 3$)**	9.7%	5.2%	2.3%	3.6%	4.6%	6.1%	8.1%	11.9%	18.6%	20.0%	22.2%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	9.4%	6.8%	2.3%	3.6%	4.1%	5.7%	7.3%	10.1%	17.4%	25.4%	31.8%
	Pension Fund Allocations	9.7%	7.7%	0.0%	2.9%	4.0%	5.4%	8.0%	11.2%	15.0%	28.7%	43.0%
	UMVE ($b = 0$ & $\gamma = 3$)	77.3%	90.4%	-19.7%	2.3%	8.7%	18.9%	53.6%	101.3%	167.5%	247.5%	429.4%
	CMVE ($b = 0$ & $\gamma = 3$)	25.3%	18.0%	0.0%	0.0%	0.0%	9.1%	28.0%	38.7%	46.3%	55.3%	65.7%
RE	**CMVE ($b \neq 0$ & $\gamma = 3$)**	12.6%	3.6%	5.3%	6.9%	8.0%	10.3%	12.3%	15.1%	17.1%	17.9%	21.5%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	11.1%	3.7%	2.5%	5.9%	6.8%	8.8%	11.2%	13.6%	15.7%	17.2%	19.3%
	Pension Fund Allocations	10.5%	3.4%	0.9%	5.8%	6.5%	8.0%	10.7%	12.3%	15.1%	15.7%	17.5%
	UMVE ($b = 0$ & $\gamma = 3$)	216.4%	235.6%	-38.1%	20.6%	29.3%	49.2%	138.9%	248.7%	623.7%	686.7%	861.6%
	CMVE ($b = 0$ & $\gamma = 3$)	30.8%	21.3%	0.0%	0.0%	0.0%	11.5%	31.9%	47.2%	55.3%	60.9%	79.4%
HF	**CMVE ($b \neq 0$ & $\gamma = 3$)**	14.3%	11.0%	0.0%	2.5%	4.2%	7.0%	10.5%	17.5%	33.3%	39.1%	42.3%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	11.5%	8.5%	0.1%	2.2%	4.0%	6.5%	8.6%	15.8%	20.7%	30.8%	38.3%
	Pension Fund Allocations	8.0%	7.4%	0.0%	0.0%	0.2%	2.6%	6.7%	10.0%	17.3%	25.0%	30.6%

Table 4
Regressing Pension Fund Portfolio Allocations on MVE Allocations

This table reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of their consultants (ω^*). Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The UMVE specification imposes no constraints (so that the lagrange multipliers underlying ω^{mv} are zero) whereas the CMVE specifications incorporate non-negative weight constraints. Specifications with $b = 0$ set the benchmarking incentive to zero whereas specifications with $b \neq 0$ calibrate the benchmarking incentive so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection 3 describes the results from this table.

Variation =		**Across Pension Funds**		**Across Asset Classes**		**Over Time**	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
$\omega^*_{c,n,t} = \frac{1}{\gamma}\cdot\mu_{c,n,t}/\sigma^2_n$	**Slope**	0.03	0.03	-0.04	0.03	0.02	0.00
	(t_{stat})	(1.52)	(1.52)	(-3.48)	(1.64)	(1.61)	(0.47)
	R^2_{within}	1.9%	1.9%	1.4%	5.4%	1.0%	0.1%
$\omega^*_{c,n,t} = \frac{1}{\gamma}\cdot\mu_{c,n,t}/\sigma^2_{c,n,t}$	**Slope**	0.02	0.02	-0.02	0.02	0.00	0.00
	(t_{stat})	(2.42)	(2.42)	(-2.00)	(2.14)	(2.89)	(1.65)
	R^2_{within}	5.8%	5.8%	0.9%	9.1%	0.1%	0.1%
UMVE ($b = 0$ & $\gamma = 3$)	**Slope**	0.02	0.02	-0.02	0.02	0.01	0.00
	(t_{stat})	(2.40)	(2.40)	(-2.35)	(2.54)	(2.52)	(1.82)
	R^2_{within}	17.3%	17.3%	3.5%	22.2%	1.8%	1.4%
CMVE ($b = 0$ & $\gamma = 3$)	**Slope**	0.16	0.16	-0.17	0.11	0.03	0.03
	(t_{stat})	(2.45)	(2.45)	(-1.62)	(2.02)	(1.43)	(2.66)
	R^2_{within}	15.8%	15.8%	3.7%	24.0%	0.5%	0.7%
CMVE ($b \neq 0$ & $\gamma = 3$)	**Slope**	0.47	0.47	0.95	0.40	0.32	0.07
	(t_{stat})	(3.51)	(3.51)	(11.5)	(2.99)	(3.70)	(2.54)
	R^2_{within}	29.6%	29.6%	75.6%	32.9%	12.4%	0.7%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	**Slope**	0.62	0.62	0.97	0.54	0.50	0.31
	(t_{stat})	(3.00)	(3.00)	(16.6)	(2.71)	(4.55)	(2.11)
	R^2_{within}	36.1%	36.1%	85.0%	40.0%	36.6%	14.9%
	# Obs	610	610	610	610	610	610

Table 5
Decomposing Variation in Pension Fund Portfolio Allocations

This table reports decompositions of the variance of pension fund portfolio allocations (ω) into the effect of MVE allocations (ω^*) and deviations from MVE allocations ($\omega^\alpha = \omega - \omega^*$). Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The two CMVE specifications considered incorporate non-negative weight constraints (i.e., lagrange multipliers affect ω^{mv}) and a benchmarking incentive (i.e., $b \neq 0$). The benchmarking incentive parameter (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We use Equation 10 for the decomposition of the variation in weights (ω) and an analogous equation for the decomposition of the variation in active weights ($\widetilde{\omega} = \omega - \omega^b$). Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection 4 describes the results from this table.

Variation =		**Across Pension Funds**		**Across Asset Classes**		**Over Time**	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	**From ω^***	63%	63%	79%	83%	38%	10%
		(4.38)	(4.38)	(29.3)	(4.98)	(5.39)	(6.75)
	From ω^α	37%	37%	21%	17%	62%	90%
		(2.62)	(2.62)	(7.59)	(1.01)	(8.71)	(60.2)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	**From ω^***	58%	58%	88%	74%	73%	48%
		(4.81)	(4.81)	(35.5)	(5.56)	(6.49)	(2.75)
	From ω^α	42%	42%	12%	26%	27%	52%
		(3.43)	(3.43)	(4.96)	(2.00)	(2.44)	(3.01)
# Obs		610	610	610	610	610	610

Table 6

MVE Allocation Wedges: Link to Fund Characteristics

This table reports regressions of MVE allocation wedges ($\omega^{\alpha} = \omega - \omega^{*}$) onto a set of explanatory variables. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^{*} = (1-b) \cdot \omega^{mv} + b \cdot \omega^{b}$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^{b} reflecting benchmark allocations. We use our full specification, which incorporates non-negative weight constraints (i.e., lagrange multipliers that affect ω^{mv}) and a benchmarking incentive (i.e., $b \neq 0$). The benchmarking incentive parameter (b) is calibrated so that the average tracking error from ω^{*} is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). The explanatory variables (all aggregated to the level of the consultant) are the lagged log pension fund size, a dummy for whether the pension fund has an internal public equity team, the lagged GASB funding ratio, and the lagged net pension fund cash flows over assets (see Footnote 21 for measurement details). Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection 5.2 describes the results from this table.

Variation =	**CMVE ($b \neq 0$ & $\gamma = 3$)**					**CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)**				
	FI	**EQ**	**PE**	**RE**	**HF**	**FI**	**EQ**	**PE**	**RE**	**HF**
Log Pension Fund Size	0.01	-0.02	0.02	-0.00	-0.01	-0.00	-0.03	0.02	0.00	0.01
	(0.94)	(-2.61)	(3.29)	(-0.23)	(-1.15)	(-0.53)	(-4.96)	(5.06)	(0.25)	(1.79)
Dummy(Internal EQ Team)	-0.13	0.19	-0.12	0.02	0.05	-0.00	0.16	-0.09	-0.03	-0.06
	(-12.1)	(5.58)	(-3.13)	(0.81)	(1.07)	(-0.32)	(4.34)	(-5.54)	(-1.37)	(NaN)
GASB Funding Ratio	0.12	-0.02	0.20	-0.02	-0.25	0.26	-0.03	0.07	-0.02	-0.29
	(1.44)	(-0.21)	(2.14)	(-0.31)	(-12.9)	(3.00)	(-0.47)	(1.96)	(-0.33)	(-2.54)
Net Cash Flow / Assets	-0.57	0.89	-0.81	0.06	0.36	-1.20	-0.31	-0.34	0.23	1.25
	(-0.61)	(1.21)	(-1.20)	(0.11)	(0.83)	(-2.53)	(-0.62)	(-1.26)	(0.46)	(3.45)
$\boldsymbol{R^2}$	7.4%	21.1%	19.1%	0.7%	7.1%	6.9%	25.7%	24.4%	5.3%	13.9%
# Obs	122	122	122	122	122	122	122	122	122	122

Internet Appendix

"Subjective Beliefs and the Portfolio Allocations of Institutional Investors"

By Aleksandar Andonov, Spencer J. Couts, Andrei S. Gonçalves, Johnathan A. Loudis, and Andrea Rossi

Contents

A Matching Pension Fund Allocations to Consultant CMAs **IA.1**

B Matching Asset Classes from CMAs to Pension Fund Allocations **IA.3**

C Supplementary Empirical Results **IA.4**

C.1 Regressing Pension Fund Allocations on MVE Allocations (with Controls) . IA.4

C.2 Regressing Pension Fund Allocations on MVE Allocations (Fund-Level) . . . IA.5

C.3 Main Results using Three-Way Wild Bootstrap Clustering IA.7

C.4 Main Results using Equal-Weighted Pension Fund Allocations IA.7

C.5 Main Results using Actual Pension Fund Allocations IA.8

C.6 Main Results using Pension Fund μ_{AARR} when Available IA.8

C.7 Main Results including CMAs that do not cover Hedge Funds IA.8

D Mean-Variance Efficient (MVE) Allocations: Details **IA.10**

D.1 Providing Potential Economic Mechanisms for the Benchmarking Incentive . IA.10

D.2 Deriving the MVE Allocation Solution IA.14

D.3 Obtaining the Lagrange Multipliers given Risk Aversion Parameters IA.18

D.4 Estimating Risk Aversion IA.19

D.5 Main Results with IV Risk Aversion Estimates IA.23

D.6 Why Structural Endogeneity in Risk Aversion Estimation is not a Concern? IA.23

A Matching Pension Fund Allocations to Consultant CMAs

We match each fund-year observation to a consultant CMA as follows. For each pension fund allocation, we identify its general consultant and the corresponding allocation year. We then link that allocation to the consultant's CMA from the preceding year, when available. As we detail below, this procedure mitigates look-ahead bias (i.e., using beliefs formed after the allocation date) while ensuring that beliefs are not generated too far in advance of the fund's investment decisions.

To begin, Table IA.2 (Panel A) reports the distribution of the calendar months in which pension fund allocations are measured. The large majority of pension fund allocations are observed in June of each year. The remaining observations are typically from the months after June, and fewer than 1% of the allocations are dated prior to June. Because these allocations are matched to consultant CMAs from the preceding year, this timing structure leaves little scope for look-ahead bias.

Although we do not observe the exact dates on which consultants form their beliefs, we do observe the dates of the documents from which most consultant CMAs are extracted. Because beliefs must necessarily be formed before the documents that report them are created, these document dates provide an upper bound on the timing of belief formation.

Table IA.2 (Panel B) reports the distribution of the time gap, in months, between the pension fund allocation date and the corresponding CMA document date. This gap is negative in only 16 fund-year observations (2.18% of the sample), indicating that in the vast majority of cases we can verify that CMAs are produced no later than the allocations are formed. Even in these 16 cases, it is likely that the underlying CMAs were in fact produced before the allocations. The reason is that, for these observations, the CMAs are obtained from third-party documents that report consultant CMAs rather than from the original CMA reports themselves. As a result, the reporting dates we observe likely lag the true CMA production dates, which would again precede the allocation decisions.

Beyond minimizing concerns about look-ahead bias, our matching procedure also ensures

that consultant beliefs are not measured too far in advance of the fund's investment decisions. Specifically, Table IA.2 (Panel B) shows that more than 90% of the matches involve a CMA date that falls within 12 months of the corresponding pension fund allocation.

B Matching Asset Classes from CMAs to Pension Fund Allocations

Our pension fund allocations cover six broad asset classes: cash (Cash), fixed income (FI), public equity (EQ), private equity (PE), real estate (RE), and hedge funds (HF). In contrast, consultant-year CMA covers a range of asset classes with names that do not necessarily match the exact names of these six asset classes (which we refer to as the "broad asset classes"). Moreover, the asset classes covered in the CMAs vary across consultants and over time. So, we develop a detailed procedure to match the asset classes covered in the CMAs to our broad asset classes. First, we identify the asset classes in each CMA based on the asset class name used in the CMA report and/or the actual portfolio index stated in the CMA report. Second, we manually map each asset class in each institution-year CMA to an institution-specific asset class name (fixed over time) that reflects the underlying asset class well. Third, we map each institution-specific asset class name to a slightly more general asset class name (which we refer to as the master asset class) that reflects the institution-specific asset class name reasonably well while allowing for small mismatches to accommodate asset classes from different institutions under the same master asset class. Fourth, for each CMA, we match each of our broad asset classes to the most closely related master asset class available (with the possibility of no match).

Table IA.3 provides the results from our asset class matching procedure. The first column shows the broad asset classes we cover in the paper. The other columns provide the list of master asset classes that we match to these broad asset classes.

C Supplementary Empirical Results

This section provides results that supplement the findings we report in the main text.

C.1 Regressing Pension Fund Allocations on MVE Allocations (with Controls)

Our baseline regressions of pension fund allocations onto MVE allocations are performed without control variables, as suggested by the mean-variance framework of Subsection 2.1. This subsection considers regressions of ω onto ω^* while controlling for other variables.

C.1.1 Controlling for Benchmark Weights

As discussed in the main text, one worry with our regression results is that part of the link between pension fund allocations and MVE allocations is due to the inclusion of ω^b into the ω^* expression through the benchmarking incentive. Our baseline analysis deals with this issue through the use of active weights since $\widetilde{\omega}^*$ is proportional to $\widetilde{\omega}^{mv}$ so that it does not have the extra ω^b term (as shown in Equation 7). However, in this baseline approach, we explain variation in active weights, not raw weights. Table IA.5 considers an alternative approach that focuses on explaining raw weights. Specifically, we regress ω onto ω^* controlling for ω^b. We consider the same specifications provided in Table 4 and find that the results are very similar to the original ones based on active weights (reproduced in Table IA.5 for comparison).

C.1.2 Controlling for Pension Fund Characteristics

Table IA.6 further reports regressions that control for some important fund characteristics studied in the prior literature: pension fund size, a dummy for whether the pension fund has an internal team to manage its equity allocation, the GASB funding ratio, and net cash flows over assets. We aggregate these characteristics to the consultant-year level and consider them in isolation (in Panel A) as well as interacted with asset class dummies (in Panel B) to account for their heterogeneous effects across asset classes. The advantage of adding these controls is that they proxy for pension fund incentives. The disadvantage is that they may

also reflect pension fund beliefs, which we do not want to control for.

The results from Table IA.6 show that ω^* continues to be economically and statistically linked to ω in all three dimensions we study even after controlling for these characteristics. When considering the characteristics in isolation, the slope coefficients are very similar to the ones we report in the main text. When considering characteristics interacted with asset class dummies, the slope coefficients tend to somewhat decline (albeit they remain statistically and economically strong). The reason is that ω^* is a noisy proxy for the pension fund MVE allocation vector (as it is based on consultant beliefs) so that its effect decreases when we control for variables that also correlate with pension fund MVE allocations (i.e., fund-level characteristics interacted with asset class dummies).

C.2 Regressing Pension Fund Allocations on MVE Allocations (Fund-Level)

Our baseline regressions of pension fund allocations onto MVE allocations are performed at the consultant level. The reason is that MVE allocations are obtained from the beliefs of consultants, and thus do not vary across pension funds within a given consultant-year. This subsection considers regressions performed at the pension fund level. To be consistent with our consultant-level regressions, we use a weighted least squares with weights based on the AUM within the fund-years matched to each consultant-year.

C.2.1 Regressing $\omega_{i,n,t}$ onto $\omega^*_{c,n,t}$

Table IA.7 (Panel A) provides results from regressions of $\omega_{i,n,t}$ onto $\omega^*_{c,n,t}$. We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. All slope coefficients are statistically significant and numerically identical to the analogous ones in the main text (i.e., the last two specifications of Table 4). The within R^2 values decline as there is an extra source of variation in $\omega_{i,n,t}$ compared to the $\omega_{c,n,t}$ variable used in the main text (and $\omega^*_{c,n,t}$ cannot explain this extra source of variation by construction). Strikingly, however, the within R^2 values remain relatively large, suggesting much of the total variation in pension fund allocations can be explained by the MVE allocations implied by

the beliefs of their consultants.

C.2.2 Regressing $\omega_{i,n,t}$ onto $\omega^*_{c,n,t}$, $\tau_{i,t}$, and $\tau_{i,t} \cdot \omega^*_{c,n,t}$

Fund-level regressions are also useful to help rule out the hypothesis that our results are driven by a reverse causality channel. As discussed in Subsection 3.3, our results could be driven by pension funds selecting consultants whose beliefs justify their desired allocations. While only 6.2% of the fund-years in our baseline sample correspond to the start of a new fund-consultant relationship, it could be that the selection effect is so strong that it lasts for many years, thereby influencing our results in a non-trivial way. However, under this hypothesis, we would expect the link between $\omega_{i,n,t}$ and $\omega^*_{c,n,t}$ to weaken as time since the start of the relationship ($\tau_{i,t}$) increases. We find that this is not the case. If anything, the link between $\omega_{i,n,t}$ and $\omega^*_{c,n,t}$ strengthens as $\tau_{i,t}$ increases.

We demonstrate the above result using the following fund-level regressions

$$\omega_{i,n,t} = \text{Fixed Effect} + b_\tau \cdot \tau_{i,t} + \text{Slope}_{i,t} \cdot \omega^*_{c,n,t} + \epsilon_{i,n,t}, \tag{IA.1}$$

where $\text{Slope}_{i,t} = b_\omega + b_{\tau\cdot\omega} \cdot \tau_{i,t}$ so that

$$\omega_{i,n,t} = \text{Fixed Effect} + b_\tau \cdot \tau_{i,t} + b_\omega \cdot \omega^*_{c,n,t} + b_{\tau\cdot\omega} \cdot \tau_{i,t} \cdot \omega^*_{c,n,t} + \epsilon_{i,n,t}. \tag{IA.2}$$

Table IA.7 (Panel B) reports the estimates of b_ω and $b_{\tau\cdot\omega}$ from these regressions (we omit b_τ for brevity, but it tends to be small and statistically insignificant). In all cases, $b_{\tau\cdot\omega}$ is positive (and even statistically insignificant in most cases). This result is inconsistent with the hypothesis that the link between ω and ω^* is driven by the selection effect, which would imply a stronger link for smaller values of τ (i.e., the period right after the selection happens). We find consistent (untabulated) results when replacing $\tau_{i,t}$ with a dummy for $\tau_{i,t} = 0$ in these regressions (i.e., in this case, the $b_{\tau\cdot\omega}$ coefficient is always negative).

What explains the positive $b_{\tau\cdot\omega}$? One possible explanation is that the views of pension funds are more influenced by the views of consultants they trust more, with trust plausibly increasing as the relationship matures. Alternatively, a positive $b_{\tau\cdot\omega}$ can also arise if consultant CMAs have a causal effect on pension fund beliefs and there is a temporal lag in

the response of allocations to beliefs (e.g., due to institutional inertia or slow rebalancing in private markets). However, our argument does not require taking a stand on the specific economic mechanism underlying the positive $b_{\tau\cdot\omega}$. What matters to us is that $b_{\tau\cdot\omega}$ is not negative, as it would be the case if the link between ω and ω^* was driven by the reverse causality channel we discuss.

C.3 Main Results using Three-Way Wild Bootstrap Clustering

Our baseline analysis uses a two-way clustering (by consultant and year) for the standard errors. Table IA.8 reports results that instead rely on a three-way clustering (by consultant, year, and asset class). The three-way clustering allows us to account for potential persistent shocks affecting the beliefs of all consultants about individual asset classes. However, since we only have five asset classes (i.e., too few clusters in the asset class dimension), we perform the three-way clustering using the wild bootstrap algorithm from Cameron, Gelbach, and Miller (2008). Accordingly, we report finite sample p-values (instead of t_{stat} values), which can be interpreted without any asymptotic approximation. We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. The overall results are similar to the ones we report in the main text. The only result that becomes statistically weak (with p-values above 0.10) is the link between $\widetilde{\omega}$ and $\widetilde{\omega}^*$ over time for the model with $\gamma = 3$. All other results remain statistically strong.

C.4 Main Results using Equal-Weighted Pension Fund Allocations

Our baseline analysis uses AUM-weighted averages across funds when constructing ω and ω^b (which better reflect the economic importance of different pension funds). Table IA.9 reports results using equal-weighted averages across funds instead. We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. The overall results are similar to the ones we report in the main text.

C.5 Main Results using Actual Pension Fund Allocations

Our baseline analysis uses target pension fund allocations to construct ω (which better reflect pension fund decisions). Table IA.10 reports results that rely on actual pension fund allocations instead. We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. The overall results are similar to the ones we report in the main text.

C.6 Main Results using Pension Fund μ_{AARR} when Available

Our baseline analysis uses consultant beliefs to proxy for pension fund beliefs. As discussed in detail in the main text, this approach mitigates reverse causality concerns for two reasons. First, consultants are unlikely to tailor their CMAs to the incentives of their public pension fund clients, as public pension funds account for only a small fraction of consultants' clients and regulatory assets under advisory. Second, pension funds rarely switch consultants, making endogenous selection an unlikely explanation for our findings. Moreover, if anything, our results become stronger (not weaker) as the length of the fund-consultant relationship increases, which is inconsistent with the idea that selection drives ours findings.

However, for completeness, Table IA.11 reports results that use the actuarially assumed rates of return (μ_{AARR}) disclosed by pension funds in their financial reports whenever available, and consultant μs otherwise. Since pension funds do not report covariance matrices, we always construct Σ using consultant CMAs. We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. The overall results are similar to the ones we report in the main text.

C.7 Main Results including CMAs that do not cover Hedge Funds

Our baseline analysis only includes CMAs that cover Hedge Funds. This approach avoids misspecification since pension funds also consider Hedge Funds when making allocation decisions. However, as discussed in the main text, this approach also induces a non-trivial

reduction in our sample size (since many CMAs do not cover Hedge Funds). So, for completeness, Table IA.12 also reports results that include CMAs that do not cover Hedge Funds. For the respective fund-years, we remove Hedge Funds from ω and the mean-variance optimization we use to construct ω^*. We also rescale ω so that it adds to 100% across the asset classes included for each fund-year (but the results are similar without this rescaling). We apply an analogous rescaling to the ω^b used for each fund-year (so, in this analysis, ω^b varies across funds based on whether the matched CMA includes Hedge Funds or not). We focus on the two most relevant specifications, which incorporate non-negative weights and a benchmarking incentive. The qualitative results are similar to the ones we report in the main text. Quantitatively, the results are somewhat weaker, which is expected given that observations without Hedge Funds embed a misspecification in this analysis.

D Mean-Variance Efficient (MVE) Allocations: Details

This section provides details about the solution and estimation for the MVE allocations we use to explain actual pension fund allocations empirically. In the main text, all pension funds are assumed to have the same benchmark portfolio. The derivations in this section are slightly more general in that they also allow the benchmark portfolio to vary across pension funds.

D.1 Providing Potential Economic Mechanisms for the Benchmarking Incentive

The investor's mean-variance problem in the main text (Equation 1) has a benchmarking incentive through the $\mathbb{V}ar_{i,t}[r_i - r_{b,i}]$ term in the objective function. The prior literature has proposed multiple channels for this type of benchmarking incentive (e.g., Evans et al. (2024) and Buffa, Vayanos, and Woolley (2022)). So, we purposely set up the optimization problem in Equation 1 to be flexible on the economic mechanism inducing the benchmarking incentive (in the spirit of Brennan (1993) and Basak and Pavlova (2013)). However, this section formalizes two portfolio allocation problems that do not directly include a benchmarking incentive but still induce exactly the optimization problem in Equation 1 (thereby providing potential economic foundations for the benchmarking incentive we consider). Of course, other economic foundations for the same benchmark incentive exist. So, the purpose of this section is to illustrate how benchmark incentives can arise from reasonable economic mechanisms rather than to provide an exhaustive list of all possible economic foundations for the benchmarking incentive we rely on in our analysis.

D.1.1 Constrained Mean-Variance Problem with Extra Tracking Error Constraint

Consider an investor who chooses a vector of portfolio weights ($\omega_{i,t}$) each period to maximize the long-term mean-variance utility (without a direct benchmarking incentive)

$$\begin{aligned} \underset{\{\omega_{i,t}\}}{Max} \;\; \mathbb{E}_{i,t}[\omega_{i,t}^{'} r] \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[\omega_{i,t}^{'} r] \\ s.t. \quad (1 - \mathbf{1}^{'}\omega_{i,t}) \; \geq 0 \\ \omega_{i,t} \; \geq \mathbf{0} \\ \mathbb{V}ar_{i,t}[r_i - r_i^b] \; \leq TE_{i,t}^2 \end{aligned} \tag{IA.3}$$

The Lagrangian of this problem is given by

$$\mathcal{L}_{i,t} \; = \omega_{i,t}^{'}\mu_{i,t} - \frac{\gamma_{i,t}}{2} \cdot \omega_{i,t}^{'}\Sigma_{i,t}\omega_{i,t} + \ell_{i,t}^{c} \cdot (1 - \mathbf{1}^{'}\omega_{i,t}) + \ell_{i,t}^{'}\omega_{i,t} - \ell_{i,t}^{b} \cdot (\omega_{i,t} - \omega_{i,t}^{b})^{'}\Sigma_{i,t}(\omega_{i,t} - \omega_{i,t}^{b}) \tag{IA.4}$$

which is equivalent to the Lagrangian of the problem in the main text (in Equation 1) after replacing $\gamma_{i,t}^{b}/2$ (the benchmarking risk aversion parameter) with $\ell_{i,t}^{b}$ (the lagrange multiplier on the tracking error constraint).

Intuitively, even though the investor does not have a direct benchmarking incentive in its objective function, he faces the institutional constraint that its tracking error relative to the benchmark cannot exceed $TE_{i,t}$ (see Buffa, Vayanos, and Woolley (2022) for a similar constraint). So, the investor perceives an indirect benchmarking incentive. Moreover, the incentive is stronger when the constraint is more binding.

D.1.2 Constrained Mean-Variance Problem with Black-Litterman Bayesian Investor

Consider an investor who chooses a vector of portfolio weights ($\omega_{i,t}$) each period to maximize the long-term mean-variance utility (without a direct benchmarking incentive)

$$\begin{aligned} \underset{\{\omega_{i,t}\}}{Max} \;\; \mathbb{E}_{i,t}[\omega_{i,t}^{'} r] \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[\omega_{i,t}^{'} r] \\ s.t. \quad (1 - \mathbf{1}^{'} \omega_{i,t}) \; \geq 0 \\ \omega_{i,t} \; \geq \mathbf{0} \end{aligned} \tag{IA.5}$$

The investor estimates the covariance matrix of returns as $\Sigma_{i,t}$ and perceives no uncertainty in regards to this estimate such that $\mathbb{V}ar_{i,t}[r] = \Sigma_{i,t}$. In contrast, the investor collects a set of expected return signals and combines them into a single expected return signal that he perceives to be an unbiased but noisy estimate for the rational expected return:

$$\mu_{i,t} \; \sim \mathcal{N}\left(\mathbb{E}_t[r] \, , \, \kappa_{i,t} \cdot \Sigma_{i,t} \right) \tag{IA.6}$$

To form his final (or posterior) expectation, the investor uses the prior (in the spirit of Black and Litterman (1991, 1992)) that a given benchmark portfolio, $\omega_{i,t}^{b} \geq \mathbf{0}$, would be optimal for him to hold under rational expected returns:[IA.1,IA.2]

$$\mathbb{E}_t[r] \; \sim \mathcal{N}\left(\gamma_{i,t} \cdot \mathbb{C}ov_{i,t}[r, \omega_{i,t}^{b'} r] \, , \, \tau_{i,t} \cdot \Sigma_{i,t} \right) \tag{IA.7}$$

Given Equations IA.6 and IA.7, the investor's Bayesian posterior mean is given by

$$\mathbb{E}_{i,t}[r] \; = (1 - b_{i,t}) \cdot \mu_{i,t} \; + \; b_{i,t} \cdot \gamma_{i,t} \cdot \mathbb{C}ov_{i,t}[r, r_i^b] \tag{IA.8}$$

where $b_{i,t} = (1/\tau_{i,t})/(1/\tau_{i,t} + 1/\kappa_{i,t})$. Using Equation IA.8, we can rewrite the mean-variance

[IA.1] Note that the first order condition of the mean-variance problem in Equation IA.5 is $\mathbb{E}_{i,t}[r] = \gamma_{i,t} \cdot \mathbb{C}ov_{i,t}[r, \omega_{i,t} r]$. As such, $\mathbb{E}_t[r] = \gamma_{i,t} \cdot \mathbb{C}ov_{i,t}[r, \omega_{i,t}^{b} r]$ holds if and only if $\omega_{i,t}^{b}$ is optimal for the investor to hold under rational expected returns.

[IA.2] The prior in Black and Litterman (1991, 1992) is that a given portfolio $\omega_{i,t}^{m} \geq \mathbf{0}$ (which investor i perceives to be the market portfolio) is mean-variance efficient. In that case, the prior mean would be $\gamma_{i,t} \cdot \mathbb{C}ov[r, \omega_{i,t}^{m'} r]$ with $\gamma_{i,t} = \mathbb{SR}_{i,t}[r_i^m]$ reflecting the Sharpe ratio of the market portfolio as perceived by the investor. Instead of imposing a prior that assumes the market portfolio is efficient, we impose a prior that assumes a given benchmark portfolio is optimal for the investor. So, while the prior in Black and Litterman (1991, 1992) imposes the CAPM, our prior imposes a slightly more general model in which the investor's wealth portfolio prices the underlying assets (in the spirit of Bretscher, Lewis, and Santosh (2024)).

objective function as

$$(1 - b_{i,t}) \cdot \omega_{i,t}^{'} \mu_{i,t} \; + \; b_{i,t} \cdot \gamma_{i,t} \cdot \mathbb{C}ov_{i,t}[r_i, r_i^b] \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i] \tag{IA.9}$$

Then, substituting $\mathbb{C}ov_{i,t}[r_i, r_i^b] = 0.5 \cdot (\mathbb{V}ar_{i,t}[r_i] + \mathbb{V}ar_{i,t}[r_i^b] - \mathbb{V}ar_{i,t}[r_i - r_i^b])$, we can further rewrite the mean-variance objective function as

$$(1 - b_{i,t}) \cdot \omega_{i,t}^{'} \mu_{i,t} \; - \; (1 - b_{i,t}) \cdot \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i] \; - \; b_{i,t} \cdot \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i - r_i^b] \; + \; b_{i,t} \cdot \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i^b] \tag{IA.10}$$

Now, note that $(1 - b_{i,t})$ is positive and $\mathbb{V}ar_{i,t}[r_i^b]$ does not depend on $\omega_{i,t}$. As such, maximizing the objective function in Equation IA.10 is equivalent to maximizing the alternative objective function

$$\omega_{i,t}^{'} \mu_{i,t} \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i] \; - \; \frac{1}{2} \cdot \frac{b_{i,t}}{1 - b_{i,t}} \cdot \gamma_{i,t} \cdot \mathbb{V}ar_{i,t}[r_i - r_i^b] \tag{IA.11}$$

Finally, defining $\gamma_{i,t}^b = \frac{b_{i,t}}{1 - b_{i,t}} \cdot \gamma_{i,t}$ allows us use Equation IA.11 to rewrite the mean-variance problem in Equation IA.5 as

$$\begin{aligned} \underset{\{\omega_{i,t}\}}{Max} \quad & \omega_{i,t}^{'} \mu_{i,t} \; - \; \frac{\gamma_{i,t}}{2} \cdot \mathbb{V}ar_{i,t}[r_i] - \; \frac{\gamma_{i,t}^b}{2} \cdot \mathbb{V}ar_{i,t}[r_i - r_i^b] \\ s.t. \quad & (1 - \mathbf{1}^{'} \omega_{i,t}) \; \geq 0 \\ & \omega_{i,t} \; \geq \mathbf{0} \end{aligned} \tag{IA.12}$$

which is exactly the mean-variance problem with benchmarking incentive we have in the main text (see Equation 1).

Note also that the $\gamma_{i,t}^b$ definition above ensures $b_{i,t}$ satisfies

$$b_{i,t} \; = \frac{\gamma_{i,t}^b}{\gamma_{i,t} + \gamma_{i,t}^b} \; = \; \frac{1/\tau_{i,t}}{1/\tau_{i,t} + 1/\kappa_{i,t}} \tag{IA.13}$$

with the first equality being consistent with the $b_{i,t}$ definition in the main text and the second equality providing an alternative way to interpret $b_{i,t}$ (as a ratio of the prior precision to the posterior precision). If the prior (that the benchmark allocation is optimal for the investor) is very precise, then $b_{i,t}$ is close to one and the investor optimally allocates most of its portfolio

to the benchmark portfolio.

D.2 Deriving the MVE Allocation Solution

Consider an investor who chooses a vector of portfolio weights ($\omega_{i,t}$) each period to maximize the long-term mean-variance utility

$$\begin{aligned} \underset{\{\omega_{i,t}\}}{Max} \;\; \mathbb{E}_{i,t}[r_i] \;-\; \frac{\gamma_{i,t}}{2}\cdot \mathbb{V}ar_{i,t}[r_i] \;-\; \frac{\gamma^b_{i,t}}{2}\cdot \mathbb{V}ar_{i,t}[r_i - r^b_i] \qquad \text{(IA.14)} \\ s.t. \quad (1 - \mathbf{1}^{'}\omega_{i,t}) \;\geq 0 \\ \omega_{i,t} \;\geq \mathbf{0} \end{aligned}$$

where $r_{i,t+h_i} = \omega^{'}_{i,t} r_{t+h_i}$ is the long-term excess return on the wealth portfolio of investor i and $r^b_{i,t+h_i} = \omega^{b'}_{i,t} r_{t+h_i}$ is the long-term excess return on a benchmark investor i cares about.

As discussed in the main text, the two constraints reflect non-negative weight constraints. Moreover, while $\gamma_{i,t} \geq 0$ reflects the standard risk aversion parameter, $\gamma^b_{i,t} \geq 0$ reflects aversion to deviations from the benchmark. So, the $-\frac{\gamma^b_{i,t}}{2}\cdot \mathbb{V}ar_{i,t}[r_i - r^b_i]$ term in the objective function captures a benchmarking incentive (see Subsection D.1 for economic foundations).

The solution to the optimization problem in Equation IA.14 results in

$$\mu_{i,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t} \;= \gamma_{i,t}\cdot \mathbb{C}ov_{i,t}[r, r_i] \;+\; \gamma^b_{i,t}\cdot \mathbb{C}ov_{i,t}[r, r_i - r^b_i] \qquad \text{(IA.15)}$$

which can be alternatively written as

$$\omega_{i,t} \;= \frac{1}{\gamma_{i,t} + \gamma^b_{i,t}}\cdot \Sigma^{-1}_{i,t}(\mu_{i,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t}) \;+\; \frac{\gamma^b_{i,t}}{\gamma_{i,t} + \gamma^b_{i,t}}\cdot \omega^b_{i,t} \qquad \text{(IA.16)}$$

where $\mu_{i,t} = \mathbb{E}_{i,t}[r]$ and $\Sigma_{i,t} = \mathbb{V}ar_{i,t}[r]$, with the scalar $\ell^c_{i,t} \geq 0$ and the vector $\ell_{i,t} \geq \mathbf{0}$ (with $\ell_{i,n,t} = [\ell_{i,t}]_n$) reflecting the lagrange multipliers on the non-negative weight constraints for cash and the risky assets, respectively.[IA.3]

We do not observe the beliefs of investor i. Instead, we observe the beliefs of consultant

[IA.3] Equations IA.15 and IA.16 have a straightforward derivations. The first order condition w.r.t $\omega_{i,t}$ is

$$\partial_{\omega_{i,t}}\left(\omega^{'}_{i,t}\mu_{i,t} - \frac{\gamma_{i,t}}{2}\cdot\omega^{'}_{i,t}\Sigma_{i,t}\omega_{i,t} + \ell^c_{i,t}\cdot(1 - \mathbf{1}^{'}\omega_{i,t} - 0) + \ell^{'}_{i,t}\omega_{i,t} - \frac{\gamma^b_{i,t}}{2}\cdot(\omega_{i,t} - \omega^b_{i,t})^{'}\Sigma_{i,t}(\omega_{i,t} - \omega^b_{i,t})\right) = 0$$

c, which are directly related to (but not perfectly revealing of) the beliefs of investor i. So, we define the belief gaps

$$\epsilon_{\mu,i,t} = \mu_{i,t} - \mu_{c,t} \tag{IA.17}$$

$$\epsilon_{\Sigma,i,t} = \Sigma_{i,t} - \Sigma_{c,t} \tag{IA.18}$$

Substituting Equations IA.17 and IA.18 into Equation IA.15, we have

$$\mu_{c,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t} = \alpha_{i,t} + \gamma_{i,t} \cdot \mathbb{C}ov_{c,t}[r, r_i] + \gamma^b_{i,t} \cdot \mathbb{C}ov_{c,t}[r, r_i - r^b_i] \tag{IA.19}$$

which yields[IA.4]

$$\omega_{i,t} = \omega^\alpha_{i,t} + \frac{1}{\gamma_{i,t} + \gamma^b_{i,t}} \cdot \Sigma^{-1}_{c,t}(\mu_{c,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t}) + \frac{\gamma^b_{i,t}}{\gamma_{i,t} + \gamma^b_{i,t}} \cdot \omega^b_{i,t} \tag{IA.20}$$

where $\omega^\alpha_{i,t} = -\Sigma^{-1}_{c,t}\alpha_{i,t}/(\gamma_{i,t} + \gamma^b_{i,t})$, with $\alpha_{i,t} = \gamma_{i,t} \cdot \epsilon_{\Sigma,i,t}\omega_{i,t} + \gamma^b_{i,t} \cdot \epsilon_{\Sigma,i,t}(\omega_{i,t} - \omega^b_{i,t}) - \epsilon_{\mu,i,t}$.

The expression for $\omega^\alpha_{i,t}$ may look strange in that (ceteris paribus) higher alphas imply lower weights. In particular, one may think that if an investor perceives a positive alpha for a given asset then the investor must hold an abnormally high weight on that asset to take advantage of the opportunity. That interpretation is not correct since alphas are always zero in this model from the perspective of investors (as shown in Equation IA.15). However, from the perspective of econometricians (who do not observe investors' beliefs), alphas deviate from zero. So, the correct interpretation is that a positive alpha for a given asset under the beliefs of consultant c is equivalent to investor i holding an abnormally low weight in that asset from the perspective of consultant c. That is, Equations IA.19 and IA.20 both reflect the investor's portfolio allocation decisions, just in different format.

---

which can be written as

$$\mu_{i,t} - \gamma_{i,t} \cdot \Sigma_{i,t}\omega_{i,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t} - \gamma^b_{i,t} \cdot \Sigma_{i,t}\omega_{i,t} + \gamma^b_{i,t} \cdot \Sigma_{i,t}\omega^b_{i,t} = 0.$$

Then, isolating $\omega_{i,t}$ leads to Equation IA.16 and isolating $\mu_{i,t} - \mathbf{1}\ell^c_{i,t} + \ell_{i,t}$ leads to

$$\mu_{i,t} - \mathbf{1}\ell^c_{i,t} + \Omega\ell_{i,t} = \gamma_{i,t} \cdot \Sigma_{i,t}\omega_{i,t} + \gamma^b_{i,t} \cdot \Sigma_{i,t}(\omega_{i,t} - \omega^b_{i,t})$$

which is the same as Equation IA.15.

[IA.4] To derive Equation IA.20, note that $\mathbb{C}ov_{c,t}[r, r_i] = \Sigma_{c,t}\omega_{i,t}$ and $\mathbb{C}ov_{c,t}[r, r_i - r^b_i] = \Sigma_{c,t}(\omega_{i,t} - \omega^b_{i,t})$ so that isolating $\omega_{i,t}$ in Equation IA.19 yields Equation IA.20.

Empirically, alphas capture both belief gaps between i and c as well as deviations of observed allocations from mean-variance allocations. So, the magnitude of (and variation in) α and ω^{α} values provide an upper bound on the effect of deviations from mean-variance allocations, which is what we try to measure in this paper. To make the upper bound as tight as possible (by averaging out belief gaps across i matched to the same c), we aggregate Equation IA.19 across the set of investors ($I_{c,t}$) matched to consultant c at time t, which implies

$$\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t} \; = \alpha_{c,t} \; + \; \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r, r_c] \; + \; \gamma_{c,t}^{b} \cdot \mathbb{C}ov_{c,t}[r, r_c - r_c^b] \tag{IA.21}$$

so that defining $\omega_{c,t}^{\alpha} = -\Sigma_{c,t}^{-1}\alpha_{c,t}/(\gamma_{c,t} + \gamma_{c,t}^{b})$, we have

$$\omega_{c,t} \; = \omega_{c,t}^{\alpha} \; + \; \frac{1}{\gamma_{c,t} + \gamma_{c,t}^{b}} \cdot \Sigma_{c,t}^{-1}(\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t}) \; + \; \frac{\gamma_{c,t}^{b}}{\gamma_{c,t} + \gamma_{c,t}^{b}} \cdot \omega_{c,t}^{b} \tag{IA.22}$$

where $r_c = \omega_{c,t}^{'} r$, with $\omega_{c,t} = \sum_{i \in I_{c,t}} w_{i,t} \cdot \omega_{i,t}$ for arbitrary $w_{i,t}$ weights. Equations IA.21 and IA.22 are derived below in Subsection D.2.1, which also shows that all other c-level terms in these equations are weighted averages of the respective i-level terms. The only exceptions are the risk aversion parameters, which are implicitly defined from $1/(\gamma_{c,t}+\gamma_{c,t}^{b}) = \sum_{i \in I_{c,t}} w_{i,t} \cdot 1/(\gamma_{i,t} + \gamma_{i,t}^{b})$ and $\gamma_{c,t}^{b} = b_{c,t} \cdot (\gamma_{c,t} + \gamma_{c,t}^{b})$ where $b_{c,t} = \sum_{i \in I_{c,t}} w_{i,t} \cdot \gamma_{i,t}^{b}/(\gamma_{i,t} + \gamma_{i,t}^{b})$. Empirically, we set $w_{i,t}$ to be proportional to the AUM of the set of investors ($I_{c,t}$) matched to consultant c at time t.

Note that we have $b_{c,t} = \gamma_{c,t}^{b}/(\gamma_{c,t} + \gamma_{c,t}^{b})$, which captures the relative strength of the benchmarking incentive in the investor's objective function and allows us to write Equation IA.22 in a more intuitive form:

$$\omega_{c,t} \; = \omega_{c,t}^{\alpha} \; + \; \omega_{c,t}^{*} \; = \; \omega_{c,t}^{\alpha} \; + \; (1 - b_{c,t}) \cdot \omega_{c,t}^{mv} \; + \; b_{c,t} \cdot \omega_{c,t}^{b} \tag{IA.23}$$

where

$$\omega_{c,t}^{mv} \; = \frac{1}{\gamma_{c,t}} \cdot \Sigma_{c,t}^{-1}(\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t}) \tag{IA.24}$$

so that the investor's portfolio allocation is given by the optimal allocation under the consultant c beliefs, $\omega_{c,t}^{*} = (1 - b_{c,t}) \cdot \omega_{c,t}^{mv} \; + \; b_{c,t} \cdot \omega_{c,t}^{b}$, plus a component associated with belief gaps

between investors and their consultants, $\omega^{\alpha}_{c,t}$. Moreover, optimal allocation has a standard mean-variance component (ω^{mv}) and a benchmarking component (ω^{b}).

Equations IA.21 and IA.23 are Equations 4 and 5 in the main text.

D.2.1 Deriving Equations IA.21 and IA.22

Consider $I_{c,t}$ investors matched to consultant c at time t and define arbitrary weights $(w_{i,t})$ that satisfy $\sum_{i\in I_{c,t}} w_{i,t} = 1$.

We start by defining two sets of adjusted weights. First, $w^{\gamma}_{i,t} = w_{i,t}\cdot(\gamma_{c,t}+\gamma^{b}_{c,t})/(\gamma_{i,t}+\gamma^{b}_{i,t})$, where $1/(\gamma_{c,t}+\gamma^{b}_{c,t}) = \sum_{i\in I_{c,t}} w_{i,t}\cdot 1/(\gamma_{i,t}+\gamma^{b}_{i,t})$. Second, $w^{b}_{i,t} = w_{i,t}\cdot b_{i,t}/b_{c,t}$, where $b_{i,t} = \gamma^{b}_{i,t}/(\gamma_{i,t}+\gamma^{b}_{i,t})$ and $b_{c,t} = \sum_{i\in I_{c,t}} w_{i,t}\cdot b_{i,t}$. Note that $w^{\gamma}_{i,t}$ and $w^{b}_{i,t}$ are weights since they satisfy $\sum_{i\in I_{c,t}} w^{\gamma}_{i,t} = 1$ and $\sum_{i\in I_{c,t}} w^{b}_{i,t} = 1$.

Then, take the weighted sum (with $w^{\gamma}_{i,t}$ weights) on both sides of Equation IA.19 to obtain

$$\begin{aligned}\sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot \text{LHS}_{i,t} &= \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\left(\mu_{c,t} - \mathbf{1}\ell^{cash}_{i,t} + \ell_{i,t}\right)\\ &= \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\mu_{c,t} \;-\; \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\mathbf{1}\ell^{cash}_{i,t} \;+\; \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\ell_{i,t}\\ &= \mu_{c,t} \;-\; \mathbf{1}\ell^{cash}_{c,t} \;+\; \ell_{c,t} \qquad \text{(IA.25)}\end{aligned}$$

and

$$\begin{aligned}\sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot \text{RHS}_{i,t} =& \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\alpha_{i,t} + \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot(\gamma_{i,t}+\gamma^{b}_{i,t})\cdot\mathbb{C}ov_{c,t}[r,r_i] - \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\gamma^{b}_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r^{b}_{i}]\\ =&\alpha_{c,t} + (\gamma_{c,t}+\gamma^{b}_{c,t})\cdot\sum_{i\in I_{c,t}} w_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r_i] \;-(\gamma_{c,t}+\gamma^{b}_{c,t})\cdot\sum_{i\in I_{c,t}} w_{i,t}\cdot b_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r^{b}_{i}]\\ =&\alpha_{c,t} + (\gamma_{c,t}+\gamma^{b}_{c,t})\cdot\sum_{i\in I_{c,t}} w_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r_i] \;-(\gamma_{c,t}+\gamma^{b}_{c,t})\cdot b_{c,t}\cdot\sum_{i\in I_{c,t}} w^{b}_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r^{b}_{i}]\\ =&\alpha_{c,t} + (\gamma_{c,t}+\gamma^{b}_{c,t})\cdot\sum_{i\in I_{c,t}} w_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r_i] \;-\gamma^{b}_{c,t}\cdot\sum_{i\in I_{c,t}} w^{b}_{i,t}\cdot\mathbb{C}ov_{c,t}[r,r^{b}_{i}]\\ =&\alpha_{c,t} + (\gamma_{c,t}+\gamma^{b}_{c,t})\cdot\mathbb{C}ov_{c,t}[r,r_c] - \gamma^{b}_{c,t}\cdot\mathbb{C}ov_{c,t}[r,r^{b}_{c}]\\ =&\alpha_{c,t} + \gamma_{c,t}\cdot\mathbb{C}ov_{c,t}[r,r_c] + \gamma^{b}_{c,t}\cdot\mathbb{C}ov_{c,t}[r,r_c - r^{b}_{c}] \qquad \text{(IA.26)}\end{aligned}$$

where $r_c = \omega'_{c,t}r$ and $r^{b}_{c} = \omega^{b\prime}_{c,t}r$, with $\omega_{c,t} = \sum_{i\in I_{c,t}} w_{i,t}\cdot\omega_{i,t}$ and $\omega^{b}_{c,t} = \sum_{i\in I_{c,t}} w^{b}_{i,t}\cdot\omega^{b}_{i,t}$. Moreover, we have $x_{c,t} = \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot x_{i,t}$ for $x = \ell^{cash}, \ell, \alpha$. Finally, note that the risk

aversion parameters are $1/(\gamma_{c,t}+\gamma^b_{c,t}) = \sum_{i\in I_{c,t}} w_{i,t}\cdot 1/(\gamma_{i,t}+\gamma^b_{i,t})$ and $\gamma^b_{c,t} = b_{c,t}\cdot(\gamma_{c,t}+\gamma^b_{c,t})$ where $b_{c,t} = \sum_{i\in I_{c,t}} w_{i,t}\cdot b_{i,t}$.

Combining Equations IA.25 and IA.26 yields Equation IA.21. Then, note that $\mathbb{C}ov_{c,t}[r, r_c] = \Sigma_{c,t}\omega_{c,t}$ and $\mathbb{C}ov_{c,t}[r, r_c - r^b_c] = \Sigma_{c,t}(\omega_{c,t} - \omega^b_{c,t})$ so that isolating $\omega_{c,t}$ in Equation IA.21 yields Equation IA.22.

D.3 Obtaining the Lagrange Multipliers given Risk Aversion Parameters

Equations IA.21 and IA.22 are closed form solutions for our portfolio allocation problem given the lagrange multipliers. However, there are no closed form solutions for the lagrange multipliers themselves. As in any optimization problem, the lagrange multipliers are indirect functions of the problem inputs. Defining $w^{\gamma}_{i,t} = w_{i,t}\cdot(\gamma_{c,t}+\gamma^b_{c,t})/(\gamma_{i,t}+\gamma^b_{i,t})$, Subsection D.2.1 shows that $\ell^{cash}_{c,t} = \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\ell^{cash}_{i,t}$ and $\ell_{c,t} = \sum_{i\in I_{c,t}} w^{\gamma}_{i,t}\cdot\ell_{i,t}$, with $\ell^{cash}_{i,t} = \ell^{cash}(\mu_{i,t}, \Sigma_{i,t}, \gamma_{i,t}, \gamma^b_{i,t})$ and $\ell_{i,t} = \ell(\mu_{i,t}, \Sigma_{i,t}, \gamma_{i,t}, \gamma^b_{i,t})$ reflecting the lagrange multipliers of the investor-level optimization problem. Since we do not observe $\mu_{i,t}$ and $\Sigma_{i,t}$, we cannot recover $\ell^{cash}_{i,t}$ and $\ell_{i,t}$ to aggregate to $\ell^{cash}_{c,t}$ and $\ell_{c,t}$. However, we can obtain an approximation for the consultant-level lagrange multipliers through $\ell_{c,t} = \ell(\mu_{c,t}, \Sigma_{c,t}, \gamma_{c,t}, \gamma^b_{c,t}, \omega^b_{c,t})$ and $\ell^{cash}_{c,t} = \ell^{cash}(\mu_{c,t}, \Sigma_{c,t}, \gamma_{c,t}, \gamma^b_{c,t}, \omega^b_{c,t})$. These functions do not have closed form solutions, but are straightforward to obtain numerically. In particular, given $\mu_{c,t}$, $\Sigma_{c,t}$, $\gamma_{c,t}$, $\gamma^b_{c,t}$, and $\omega^b_{c,t}$, we obtain the lagrange multipliers as follows:

1. Solve (numerically) the following quadratic optimization problem with linear inequality constraints:

$$\omega^*_{c,t} = \underset{\{\omega_{c,t}\}}{argmax}\ \omega'_{c,t}\mu_{c,t} - \frac{\gamma_{c,t}}{2}\cdot\omega'_{c,t}\Sigma_{c,t}\omega_{c,t} - \frac{\gamma^b_{c,t}}{2}\cdot(\omega_{c,t}-\omega^b_{c,t})'\Sigma_{c,t}(\omega_{c,t}-\omega^b_{c,t}) \quad \text{(IA.27)}$$

$$s.t. \quad (1 - \mathbf{1}'\omega_{c,t}) \geq 0$$

$$\omega_{c,t} \geq \mathbf{0}$$

   which we know has solution of the form:

$$\mu_{c,n,t} - \ell^{cash}_{c,t} + \ell_{c,n,t} = \gamma_{c,t}\cdot 1'_n\Sigma_{c,t}\omega^*_{c,t} + \gamma^b_{c,t}\cdot 1'_n\Sigma_{c,t}(\omega^*_{c,t} - \omega^b_{c,t}) \quad \text{(IA.28)}$$

2. Select one risky asset with $\omega^*_{c,n,t} > 0$. For this asset, we have $\ell_{c,n,t} = 0$. So, we can recover

$$\ell^{cash}_{c,t} = \mu_{c,n,t} - \gamma_{c,t} \cdot 1_n^{'} \Sigma_{c,t} \omega^*_{c,t} - \gamma^b_{c,t} \cdot 1_n^{'} \Sigma_{c,t} (\omega^*_{c,t} - \omega^b_{c,t}) \tag{IA.29}$$

3. For all assets with $\omega^*_{c,n,t} > 0$, we have $\ell_{c,n,t} = [\ell_{c,t}]_n = 0$. For all other assets, recover $\ell_{c,n,t}$ from

$$\ell_{c,n,t} = \gamma_{c,t} \cdot 1_n^{'} \Sigma_{c,t} \omega^*_{c,t} + \gamma^b_{c,t} \cdot 1_n^{'} \Sigma_{c,t} (\omega^*_{c,t} - \omega^b_{c,t}) - (\mu_{c,n,t} - \ell^{cash}_{c,t}) \tag{IA.30}$$

As discussed in Section 1, we measure $\mu_{c,t}$ and $\Sigma_{c,t}$ directly from CMAs. Also, as discussed in Section 2, we obtain $\omega^b_{c,t}$ by aggregating the $t-1$ weights of all pension funds. The next subsection explains how we estimate $\gamma_{c,t}$ and $\gamma^b_{c,t}$ while using the above procedure to recover $\ell^{cash}_{c,t} = \ell^{cash}(\mu_{c,t}, \Sigma_{c,t}, \gamma_{c,t}, \gamma^b_{c,t}, \omega^b_{c,t})$, and $\ell_{c,t} = \ell(\mu_{c,t}, \Sigma_{c,t}, \gamma_{c,t}, \gamma^b_{c,t}, \omega^b_{c,t})$.

D.4 Estimating Risk Aversion

We now explain how we estimate risk aversion ($\gamma_{c,t}$) allowing for heterogeneity and time variation. We also provide some descriptive statistics for the different risk aversion estimates we obtain.

D.4.1 Unconstrained Mean-Variance Problem without Benchmarking Incentive

It is instructive to start from the mean-variance solution with no constraints ($\ell^{cash}_{c,t} = 0$ and $\ell_{c,t} = \mathbf{0}$) and no benchmarking incentive ($b = 0$). In this case, Equation 4 reduces to

$$\mu_{c,n,t} = \eta_{c,t} + \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r_c] + \epsilon_{c,n,t} \tag{IA.31}$$

where $\alpha_{c,n,t} = \eta_{c,t} + \epsilon_{c,n,t}$ holds by construction.

So, we estimate $\gamma_{c,t}$ from an ordinary least squares (OLS) regression of $\mu_{c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ across asset classes for each institution-year. This regression provides an unbiased estimate of $\gamma_{c,t}$ so long as $\mathbb{C}ov_{c,t}[r_n, r_c]$ is uncorrelated with $\epsilon_{c,n,t}$ across asset classes. This condition holds in the mean-variance framework of Subsection 2.1 so long as belief gaps are uncorrelated with $\mathbb{C}ov_{c,t}[r_n, r_c]$ across asset classes.

As detailed in Section D.6, the correlation between $\mathbb{C}ov_{c,t}[r_n, r_c]$ and $\epsilon_{\mu,c,n,t}$ across asset classes is very small in the subsample over which we can calculate ϵ_μ. Nevertheless, we do not observe $\epsilon_{\Sigma,c,n,t}$, which should also affect $\epsilon_{c,n,t}$. Moreover, we expect at least some correlation between $\mathbb{C}ov_{c,t}[r_n, r_c]$ and $\epsilon_{c,n,t}$ since $\mathbb{C}ov_{c,t}[r_n, r_c] = \mathbb{C}ov_{c,t}[r_n, \omega_{c,t}' r]$, with $\omega_{c,t} = \omega_{c,t}^{\alpha} + \omega_{c,t}^{*}$ where $\omega_{c,t}^{\alpha} = -\Sigma_{c,t}^{-1}\alpha_{c,t}/\gamma_{c,t}$ is a function of alpha (i.e., we have a structural endogeneity issue). As such, we also consider risk aversion estimates from an instrumental variable (IV) regression of $\mu_{c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ with instrument $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$ (which does not depend on $\omega_{c,t}$). Subsection D.4.4 shows that the resulting risk aversion estimates are very similar across the two estimation approaches. In turn, Section D.5 shows that our main results are also similar if we use the IV risk aversion estimates instead of the OLS ones. For completeness, Section D.6 then discusses in more details the structural endogeneity issue that motivates the IV estimation and clarifies why this issue is not a major concern in our analysis.

D.4.2 Constrained Mean-Variance without Benchmarking Incentive

Incorporating non-negative weight constraints (i.e., allowing for $\ell_{c,t}^{cash} \neq 0$ and $\ell_{c,t} \neq \mathbf{0}$) but keeping no benchmarking incentives ($b = 0$), Equation 4 becomes

$$\mu_{c,n,t} - \ell_{c,t}^{cash} + \ell_{c,n,t} \; = \eta_{c,t} \; + \; \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r_c] \; + \; \epsilon_{c,n,t} \tag{IA.32}$$

so that we can use the regression approach described in the prior subsection. The only difference is that we do not know $\ell_{c,t}^{cash}$ and $\ell_{c,t}$ before we have a $\gamma_{c,t}$ estimate. As such, we follow an iterative procedure for each institution-year:

1. Start with $\ell_{c,t}^{cash} = 0$ and $\ell_{c,t} = \mathbf{0}$
2. Estimate $\gamma_{c,t}$ from a regression of $\mu_{c,n,t} - \ell_{c,t}^{cash} + \ell_{c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$
3. Given the $\gamma_{c,t}$ estimate, update $\ell_{c,t}^{cash}$ and $\ell_{c,t}$ following the procedure described in Section D.3
4. Repeat Steps 2 and 3 until $\gamma_{c,t}$, $\ell_{c,t}^{cash}$, and $\ell_{c,t}$ converge

with Step 2 considering both the OLS and IV regressions discussed in Subsection D.4.1.

The resulting estimates are summarized in Subsection D.4.4.

D.4.3 Constrained Mean-Variance with Benchmarking Incentive

Further incorporating a benchmarking incentive (through $b_{c,t} = \gamma^b_{c,t}/(\gamma_{c,t} + \gamma^b_{c,t}) \neq 0$), Equation 4 becomes

$$\mu_{c,n,t} - \ell^{cash}_{c,t} + \ell_{c,n,t} \;= \eta_{c,t} \;+\; (\gamma_{c,t} + \gamma^b_{c,t}) \cdot \mathbb{C}ov_{c,t}[r_n, r_c] \;-\; \gamma^b_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r^b] \;+\; \epsilon_{c,n,t} \tag{IA.33}$$

so that we need to modify our approach to account for the $\mathbb{C}ov_{c,t}[r_n, r^b]$ term.

In principle, we could estimate $\gamma_{c,t}$ and $\gamma^b_{c,t}$ jointly by replacing the sequence of univariate regressions of Subsection D.4.2 with bivariate regressions that include both $\mathbb{C}ov_{c,t}[r_n, r_c]$ and $\mathbb{C}ov_{c,t}[r_n, r^b]$. In practice, however, this approach would be econometrically unreliable. The reason is that $\mathbb{C}ov_{c,t}[r_n, r_c]$ and $\mathbb{C}ov_{c,t}[r_n, r^b]$ are highly correlated across asset classes. In particular, the correlation is higher than 0.97 for all institution-years in our baseline sample. Consequently, we instead assume $\gamma^b_{c,t}$ is proportional to $\gamma_{c,t}$ so that $b_{c,t} = b$. In this case, Equation IA.33 simplifies to

$$\mu_{c,n,t} - \ell^{cash}_{c,t} + \ell_{c,n,t} + \gamma^b_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r^b] \;= \eta_{c,t} \;+\; \frac{\gamma_{c,t}}{1-b} \cdot \mathbb{C}ov_{c,t}[r_n, r_c] \;+\; \epsilon_{c,n,t} \tag{IA.34}$$

so that we can estimate risk aversion by combining the estimation framework from Subsection D.4.2 with the b calibration we use in the constant risk aversion specification. In particular, we apply the following iterative procedure to estimate $\gamma_{c,t}$ given b:

1. Start with $\ell^{cash}_{c,t} = 0$, $\ell_{c,t} = \mathbf{0}$, and $\gamma^b_{c,t} = 0$
2. Estimate $\gamma_{c,t} = slope_{c,t} \cdot (1-b)$ from a regression of $\mu_{c,t} - \ell^{cash}_{c,t} + \ell_{c,n,t} + \gamma^b_{c,t} \cdot \mathbb{C}ov_{c,t}[r_n, r^b]$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$
3. Given the $\gamma_{c,t}$ estimate, update $\ell^{cash}_{c,t}$ and $\ell_{c,t}$ following the procedure described in Section D.3 as well as $\gamma^b_{c,t}$ from $\gamma^b_{c,t} = \gamma_{c,t} \cdot b/(1-b)$

4. Repeat Steps 2 and 3 until $\gamma_{c,t}$, $\ell_{c,t}^{cash}$, and $\ell_{c,t}$ converge

with Step 2 considering both the OLS and IV regressions discussed in Subsection D.4.1.

We apply the iterative procedure above for a b grid and choose the b value that leads to the distribution of MVE allocation tracking errors,

$$TE_{c,t}^{*} = \sqrt{\mathbb{V}ar_{c,t}[\omega_{c,t}^{*\prime} r - \omega_t^{b\prime} r]} = \sqrt{\widetilde{\omega}_{c,t}^{*\prime} \Sigma_{c,t} \widetilde{\omega}_{c,t}^{*}} = (1-b) \cdot \sqrt{\widetilde{\omega}_{c,t}^{mv\prime} \Sigma_{c,t} \widetilde{\omega}_{c,t}^{mv}}, \tag{IA.35}$$

that has the same average as the distribution of pension fund allocation tracking errors,[IA.5]

$$TE_{c,t} = \sqrt{\mathbb{V}ar_{c,t}[\omega_{c,t}^{\prime} r - \omega_t^{b\prime} r]} = \sqrt{\widetilde{\omega}_{c,t}^{\prime} \Sigma_{c,t} \widetilde{\omega}_{c,t}}. \tag{IA.36}$$

The resulting estimates are summarized in Subsection D.4.4.

D.4.4 Descriptive Statistics for Risk Aversion Estimates

Table IA.13 provides descriptive statistics for risk aversion estimates for three model specifications, with the last one being the full specification considered in the main text. The γ distributions have almost the same median across specifications. However, γ estimates under the specification with $b \neq 0$ display substantially more variability and right skewness, leading to much higher average and standard deviation. Within each specification, the table also considers two risk aversion estimations, one by OLS and one by IV (as discussed in Subsection D.4.1). The two are highly correlated and lead to similar risk aversion distributions. The main difference is that the distribution based on IV estimates has a little more dispersion. This is not surprising since IV estimates are less precise than OLS estimates, so the additional estimation error should add some variability to the distribution of risk aversion estimates. Nevertheless, the differences are not large.

Next, Section D.5 shows that our main results are similar if we replace the OLS risk aversion estimates with IV risk aversion estimates in our full specification. For completeness, Section D.6 then discusses in more details the structural endogeneity issue that motivates

[IA.5] The resulting b is 0.93 when estimating γ^{OLS} and 0.94 when estimating γ^{IV} (compared to the $b = 0.89$ obtained under $\gamma = 3$). However, for some fund-years, the γ estimation fails to converge at these b values. In those cases, we select the nearest lower b grid point at which the corresponding γ estimation converges.

the IV estimation and clarifies why this issue is not a major concern in our analysis.

D.5 Main Results with IV Risk Aversion Estimates

Our full specification in the main text estimates time-varying and heterogeneous risk aversion through OLS (applied to Equation 4). Table IA.14 reproduces the results from our full specification estimated by OLS (originally reported in Tables 4 and 5) together with analogous results that rely on the IV estimation discussed in the prior section. Panel A focuses on regressions of pension fund allocations on MVE allocations and Panel B performs decompositions of the variation in pension fund allocations. The overall results are qualitatively similar between OLS and IV. Quantitatively, however, there tends to be a small decline (going from OLS to IV) in the slope coefficients, t-statistics, R^2 values, and share of variation attributed to MVE weights (on all three dimensions of variation). In principle, the relatively small quantitative differences between OLS and IV can be due to a potential bias in the OLS estimates or the extra noise present in the IV estimates. Empirically, however, the center of the OLS and IV risk aversion distributions are almost identical (see Subsection D.4.4). So, the small differences likely steams from the extra noise in the IV estimates.

D.6 Why Structural Endogeneity in Risk Aversion Estimation is not a Concern?

Consider again the mean-variance solution with no constraints (i.e., $\ell_{c,t}^{cash} = 0$ and $\ell_{c,t} = \mathbf{0}$) and no benchmarking incentive (i.e., $b = 0$). In this case, Equations 4 reduces to

$$\mu_{c,t} - \mathbf{1}\ell_{c,t}^{cash} + \ell_{c,t} \;= \eta_{c,t} \;+\; \gamma_{c,t} \cdot \mathbb{C}ov_{c,t}[r, r_c] \;+\; \epsilon_{c,n,t} \tag{IA.37}$$

where $\alpha_{c,n,t} = \eta_{c,t} + \epsilon_{c,n,t}$ holds by construction.

We estimate $\gamma_{c,t}$ as the slope coefficient from an OLS regression of $\mu_{c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ across asset classes for each institution-year. However, our mean-variance framework of Section 2.1 implies $\alpha_{i,t} = \gamma_{i,t} \cdot \epsilon_{\Sigma,i,t}\omega_{i,t} - \epsilon_{\mu,i,t}$ is a function of belief wedges. As such, our estimation strategy only leads to unbiased risk aversion estimates if $\mathbb{C}ov_{c,t}[r_n, r_c]$ is uncorrelated with belief wedges. However, we have $\mathbb{C}ov_{c,t}[r_n, r_c] = \mathbb{C}ov_{c,t}[r_n, \omega_{c,t}' r]$, with $\omega_{c,t} = \omega_{c,t}^{\alpha} + \omega_{c,t}^{*}$

where $\omega^{\alpha}_{c,t} = -\Sigma^{-1}_{c,t}\alpha_{c,t}/\gamma_{c,t}$ is a function of alpha. So, $\mathbb{C}ov_{c,t}[r_n, r_c]$ should correlate with belief wedges through $\omega_{c,t}$.

Empirically, the correlation between $\mathbb{C}ov_{c,t}[r_n, r_c]$ and $\epsilon_{\mu,c,t}$ is small and statistically insignificant, which suggests that this structural endogeneity issue is not a major concern. Specifically, using the subsample over which we observe ϵ_{μ}, we regress $-\epsilon_{\mu,c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_c]$ across asset classes (i.e., using consultant-by-year fixed effects). The first column in Table IA.15 provides the results, which show that the correlation between $-\epsilon_{\mu,c,n,t}$ and $\mathbb{C}ov_{c,t}[r_n, r_c]$ is economically and statistically insignificant.

However, $\mathbb{C}ov_{c,t}[r_n, r_c]$ can also correlate with $\alpha_{i,t}$ through $\epsilon_{\Sigma,i,t}$, which we do not observe. As such, we further consider an IV estimation strategy. Specifically, we use $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$ as an instrument for $\mathbb{C}ov_{c,t}[r_n, \omega^{'}_{c,t}r]$ so that any endogeneity through $\omega_{c,t}$ is entirely removed. The second column in Table IA.15 shows results from a regression of $-\epsilon_{\mu,c,n,t}$ onto $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$, which demonstrate that the two variables are uncorrelated across asset classes, and thus this aspect of the exogeneity condition is satisfied for the proposed instrument. In addition, the third column in Table IA.15 provides results from regressing $\mathbb{C}ov_{c,t}[r_n, r_c]$ onto $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$ across asset classes, showing that the instrument proposed is highly relevant. Since the prior section shows that our main results are very similar whether we use our baseline OLS risk aversion estimation of this alternative IV risk aversion estimation, we can conclude that structural endogeneity is not the driver of our results.

References for Internet Appendix

Basak, S. and A. Pavlova (Aug. 2013). "Asset Prices and Institutional Investors". In: *American Economic Review* 103.5, pp. 1728–58.

Black, F. and R. Litterman (1991). "Asset Allocation: combining investor views with market equilibrium". In: *Journal of Fixed Income* 1.2, pp. 7–18.

Black, F. and R. Litterman (1992). "Global Portfolio Optimization". In: *Financial Analyst Journal* 48.5, pp. 28–43.

Brennan, M. J. (1993). "Agency and Asset Pricing". Working Paper.

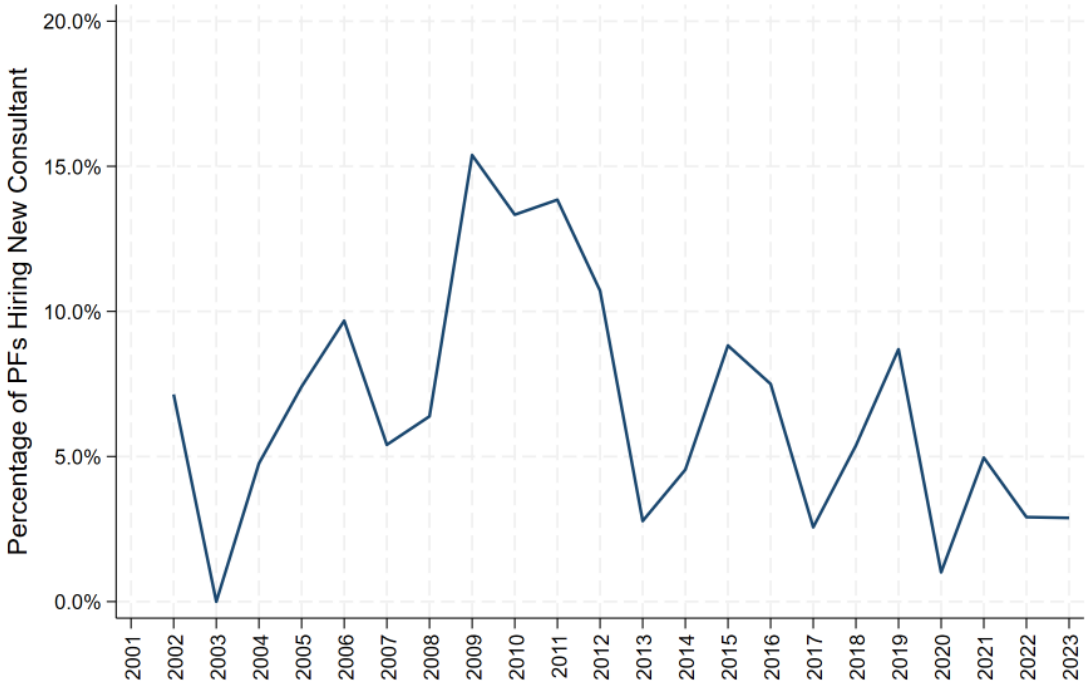
Bretscher, L., R. Lewis, and S. Santosh (2024). "Investor Betas". Working Paper.

Buffa, A. M., D. Vayanos, and P. Woolley (2022). "Asset Management Contracts and Equilibrium Prices". In: *Journal of Political Economy* 130.12, pp. 3146–3201.

Cameron, C. A., J. B. Gelbach, and D. L. Miller (2008). "Bootstrap-based improvements for inference with clustered errors." In: *Review of Economics and Statistics* 90.3, pp. 414–427.

Evans, R., J.-P. Gomez, L. Ma, and Y. Tang (2024). "Peer Versus Pure Benchmarks in the Compensation of Mutual Fund Managers". In: *Journal of Financial and Quantitative Analysis* 59.7, pp. 3101–3138.

(a) Baseline Sample (1,422 Fund-Years)



(b) Total Sample (3,653 Fund-Years)

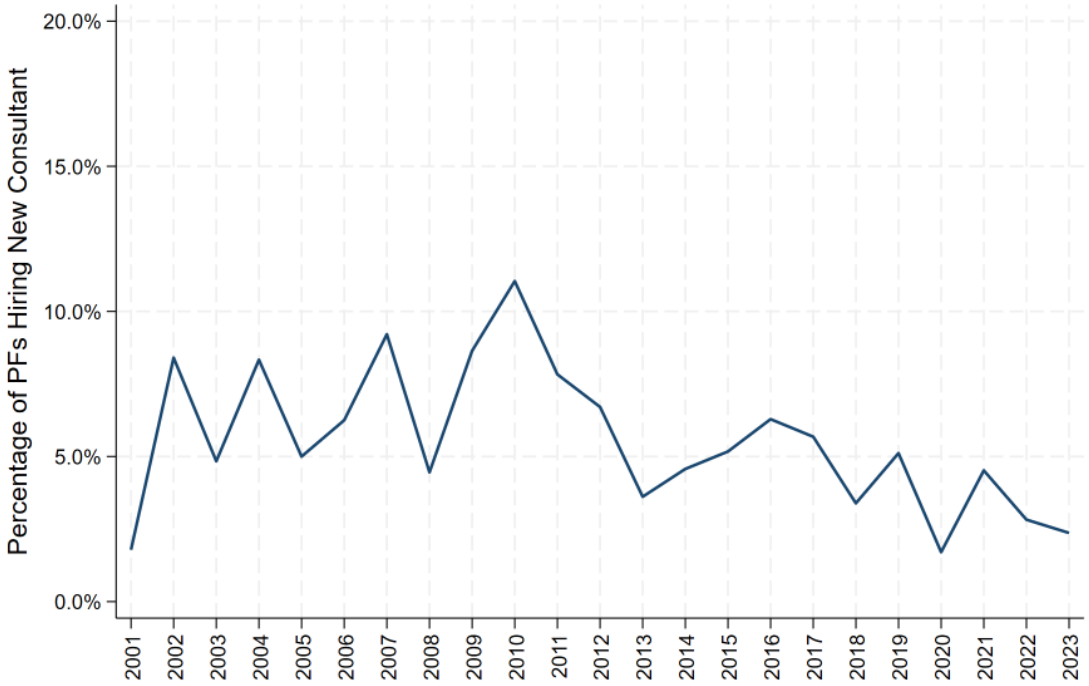


Figure IA.1
Fraction of New Fund-Consultant Relationships (by Year)

This figure plots the fraction of pension funds that hire a new general consultant each year (note that pension funds can have at most one general consultant at a given point in time). Panel A covers our main sample (of fund-years matched to CMAs that cover all asset classes in our analysis) while Panel B covers the total sample in the first column of Table 1. Overall, the fraction of fund-years with a new fund-consultant relationship is 6.2% in Panel A and 5.5% in Panel B. Section 1 describes our data on pension funds and their consultants.

Table IA.1
Institutions that are General Pension Fund Consultants at Least Once

This table reports the list of all institutions classified as general consultants in our pension fund dataset for at least one fund-year. We report the general consultant name, for how many fund-years it is classified as a general consultant (column Fund-Years), the fraction of those fund-years for which we have the consultant CMA (column %Matched), and the fraction of pension fund AUM associated with the consultant CMAs we cover (%Matched AUM). We order consultants by the number of Fund-Years in the pension fund dataset. Section 1.1 describes our dataset on pension funds, Section 1.2 describes our dataset on CMAs, and Section 1.3 describes the results from this table.

General Consultant	Fund-Years	%Matched	%Matched AUM	General Consultant	Fund-Years	%Matched	%Matched AUM
NEPC	457	100%	100%	**Angeles Investments**	14	0%	0%
Callan	453	100%	100%	**Cambridge**	14	0%	0%
Wilshire	314	95%	93%	**Wainwright**	12	0%	0%
AonHewitt	226	57%	65%	**Jeffrey Slocum & Associates**	11	0%	0%
Meketa	193	100%	100%	**Richards & Tierney**	11	0%	0%
Strategic Inv. Group	167	49%	45%	**CSG**	10	0%	0%
RVKuhns	163	58%	63%	**Gregory W. Group**	10	0%	0%
Verus	160	96%	99%	**North Point**	9	0%	0%
PCA	147	81%	78%	**The Bogdahn Group**	9	0%	0%
Summit	147	0%	0%	**Hammond Associates**	8	0%	0%
Marquette	137	0%	0%	**Holbein Associates**	8	0%	0%
Mercer	127	28%	36%	**Fiduciary Investment Adv.**	7	0%	0%
Ennis Knupp	105	0%	0%	**Prime Buchholz**	6	0%	0%
Asset Consulting Group	67	0%	0%	**Standard & Poor**	6	0%	0%
Milliman	44	100%	100%	**AmSouth Bank**	5	0%	0%
Segal	40	10%	15%	**Nuveen**	5	0%	0%
Windmark	38	0%	0%	**Monroe Vos Consulting**	4	0%	0%
Rocaton	35	0%	0%	**Dorn, Helliesen & Cottle**	3	0%	0%
AndCo	34	0%	0%	**Evaluation Associates**	3	0%	0%
Russell	34	29%	31%	**Fiducient Advisors**	3	0%	0%
Dahab Associates	31	0%	0%	**Graystone Consulting**	3	0%	0%
Cliffwater	27	100%	100%	**Bank of New York**	2	0%	0%
Becker Burke	24	0%	0%	**Mellon Consultants**	2	0%	0%
Buck Consultants	21	0%	0%	**Standard Valuations**	2	0%	0%
Gray & Co.	21	0%	0%	**Peter Chan & Co.**	1	0%	0%
RogersCasey	19	0%	0%	**Wellington Management**	1	0%	0%
Regions Bank	18	0%	0%	**No General Consultant**	186	-	-
Clearstead	17	0%	0%				
Stephens Inc	17	0%	0%				
Towers Watson	15	7%	1%				
				Total	3,653	57%	65%

Table IA.2
Matching Pension Fund Allocations to Consultant CMAs

This table reports information on our matching between pension fund allocations and their corresponding consultant CMAs. For each pension fund allocation, we identify its general consultant and the corresponding allocation year. We then link that allocation to the consultant's CMA from the preceding year, when available. Panel A reports the distribution of the calendar months in which pension fund allocations are measured. Panel B reports the distribution of the time gap, in months, between the pension fund allocation date and the corresponding CMA document date. Because beliefs must necessarily be formed before the documents that report them are created, these document dates provide an upper bound on the timing of belief formation. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs and Section A describes the results from this table.

PANEL A - Distribution of Calendar Month for Pension Fund Portfolio Allocation

Month	Frequency	Percent	Cumulative Percent
3	6	0.42%	0.42%
6	926	65.12%	65.54%
8	12	0.84%	66.39%
9	45	3.16%	69.55%
12	433	30.45%	100%

PANEL B - Distribution of Gap Between Allocation Month and CMA Month

Allocation Month - CMA Month	Frequency	Percent	Cumulative Percent
-3	15	1.05%	1.05%
-1	1	0.07%	1.13%
0	178	12.52%	13.64%
2	5	0.35%	13.99%
3	24	1.69%	15.68%
4	48	3.38%	19.06%
5	103	7.24%	26.30%
6	407	28.62%	54.92%
7	224	15.75%	70.68%
8	4	0.28%	70.96%
9	54	3.80%	74.75%
10	24	1.69%	76.44%
11	67	4.71%	81.15%
12	148	10.41%	91.56%
13	115	8.09%	99.65%
15	2	0.14%	99.79%
17	3	0.21%	100%

Table IA.3
Constructing our CMA Asset Classes

This table provides the results from our procedure to match the broad asset classes covered in pension fund portfolio allocation data to the asset classes covered in the CMAs. The procedure is as follows. First, we identify the asset classes in each CMA based on the asset class names used in the CMA report and/or the actual portfolio index stated in the CMA report. Second, we manually map each asset class in each institution-year CMA to an institution-specific asset class name (fixed over time) that reflects the underlying asset class well. Third, we map each institution-specific asset class name to a slightly more general asset class name (which we refer to as the master asset class) that reflects the institution-specific asset class name reasonably well while allowing for small mismatches to accommodate asset classes from different institution under the same master asset class. Fourth, for each CMA, we match each of our broad asset classes to the most closely related master asset class available (with the possibility of no match). The first column shows the broad asset classes we cover in the paper. The other columns provide the list of master asset classes that we match to these broad asset classes. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs and Section B describes the results from this table.

	Master Asset Classes Used (in order)				
Broad Asset Class	**Option #1**	**Option #2**	**Option #3**	**Option #4**	**Option #5**
Cash	US Cash	3-Month Libor	US Treasuries ST	US Inflation	
Fixed Income	Global Bonds	Global Credit	Government and Corporate Bonds	US Bonds	US Bonds Credit
Equity	Global Equities	Global Dev. Equities	US Equities	US Equities Large Cap	Global Equities Large Cap Ex CAD
Private Equity	Global Private Equity	Private Equity+Venture Capital	Private Equity Fund of Funds	US Private Equity	Private Investments
Real Estate	Global Real Estate	Real Estate Public+Private	Private Real Estate	US Core Real Estate	Global REITs
Hedge Funds	Hedge Funds	Equal-Weighted Average of Different Hedge Fund Strategies*			

	Master Asset Classes Used (in order)				
Broad Asset Class	**Option #6**	**Option #7**	**Option #8**	**Option #9**	**Option #10**
Cash					
Fixed Income	US Fixed Income Core+	US Corporate Fixed Income	US Corporate Inv. Grade	Intermediate Bonds	Global Fixed Income Ex US
Equity					
Private Equity	Buyouts				
Real Estate	REITs	Real Assets	Global Real Estate Hedged		
Hedge Funds					

* For Hedge Funds, we take the average of all available hedge fund strategies when the primary and substitute 1 categories are missing.

The hedge fund strategies are (they vary by institution-year observation) "Funds of Funds", "Multi Strategy", "Discretionary", "Open Mandate", "Directional", "Non-Directional", "Event Driven", "Market Neutral", "Relative Value", "Long-Short", "Long Bias", "Macro", "CTA", "Equity Style", "Credit Style Bonds Hedged", "Asymmetric Style", "Equity Hedged", "Convertible", "Moderate Aggregate Risk", "Absolute Returns", "Managed Futures"

Table IA.4
Lagrange Multipliers

This table reports the distributions of lagrange multipliers on the non-negative weight constraints for different specifications of our MVE allocation. Specifications with $b = 0$ set the benchmarking incentive to zero whereas specifications with $b \neq 0$ calibrate the benchmarking incentive so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). The lagrange multipliers cover cash ($\ell_{c,t}^{cash}$) and risky asset classes ($\ell_{c,t}$), which are fixed income (FI), public equity (EQ), private equity (PE), real estate (RE), and hedge funds (HF). They are based on the solution described in D.3. For each asset class, the lagrange multipliers vary by consultant (c) and year (t). Section 1 describes the data on pension fund portfolio allocations and consultant beliefs and Section 2 provides details about the MVE allocations.

Multiplier	ω^* Specification	# Obs	Average	Std. Dev.	Q1%	Q5%	Q10%	Q25%	Q50%	Q75%	Q90%	Q95%	Q99%
	CMVE ($b = 0$ & $\gamma = 3$)	196	1.7%	1.2%	0.0%	0.3%	0.5%	0.8%	1.5%	2.3%	3.4%	4.1%	4.6%
ℓ^{Cash}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	1.3%	0.9%	0.0%	0.0%	0.3%	0.7%	1.2%	1.8%	2.5%	2.9%	4.2%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	1.1%	0.9%	0.0%	0.0%	0.0%	0.2%	0.9%	1.7%	2.4%	2.5%	3.5%
	CMVE ($b = 0$ & $\gamma = 3$)	196	0.9%	1.1%	0.0%	0.0%	0.0%	0.0%	0.3%	1.6%	2.5%	3.4%	3.9%
ℓ_{FI}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	0.1%	0.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	1.2%	2.4%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	0.1%	0.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.2%	1.2%	2.2%
	CMVE ($b = 0$ & $\gamma = 3$)	196	0.5%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.7%	1.6%	2.0%	2.4%
ℓ_{EQ}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b = 0$ & $\gamma = 3$)	196	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	1.1%
ℓ_{PE}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b = 0$ & $\gamma = 3$)	196	0.1%	0.3%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.2%	0.8%	1.5%
ℓ_{RE}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	CMVE ($b = 0$ & $\gamma = 3$)	196	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.2%	0.8%
ℓ_{HF}	**CMVE ($b \neq 0$ & $\gamma = 3$)**	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	196	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table IA.5
Regressing Pension Fund Portfolio Allocations on MVE Allocations (Controlling for Benchmark Weights)

This table reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of their consultants (ω^*). Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The UMVE specification imposes no constraints (so that the lagrange multipliers underlying ω^{mv} are zero) whereas the CMVE specifications incorporate non-negative weight constraints. Specifications with $b = 0$ set the benchmarking incentive to zero whereas specifications with $b \neq 0$ calibrate the benchmarking incentive so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). The $\omega|\omega^b$ columns use raw weights (for the dependent and independent variables) but control for benchmark weights in the regressions. The $\widetilde{\omega} = \omega - \omega^b$ columns directly use active weights for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.1.1 describes the results from this table.

Variation =		**Across Pension Funds**		**Across Asset Classes**		**Over Time**	
ω^* Specification		$\omega \mid \omega^b$	$\widetilde{\omega} = \omega - \omega^b$	$\omega \mid \omega^b$	$\widetilde{\omega} = \omega - \omega^b$	$\omega \mid \omega^b$	$\widetilde{\omega} = \omega - \omega^b$
$\omega^*_{c,n,t} = \frac{1}{\gamma}\cdot\mu_{c,n,t}/\sigma^2_n$	**Slope**	0.03	0.03	0.01	0.03	0.00	0.00
	(t_{stat})	(1.52)	(1.52)	(1.13)	(1.64)	(0.42)	(0.47)
	R^2_{within}	1.9%	1.9%	0.5%	5.4%	0.1%	0.1%
$\omega^*_{c,n,t} = \frac{1}{\gamma}\cdot\mu_{c,n,t}/\sigma^2_{c,n,t}$	**Slope**	0.02	0.02	0.02	0.02	0.00	0.00
	(t_{stat})	(2.42)	(2.42)	(1.94)	(2.14)	(1.41)	(1.65)
	R^2_{within}	5.8%	5.8%	3.6%	9.1%	0.1%	0.1%
UMVE ($b = 0$ & $\gamma = 3$)	**Slope**	0.02	0.02	0.02	0.02	0.00	0.00
	(t_{stat})	(2.40)	(2.40)	(2.33)	(2.54)	(1.81)	(1.82)
	R^2_{within}	17.3%	17.3%	13.9%	22.2%	1.2%	1.4%
CMVE ($b = 0$ & $\gamma = 3$)	**Slope**	0.16	0.16	0.13	0.11	0.03	0.03
	(t_{stat})	(2.45)	(2.45)	(1.80)	(2.02)	(2.73)	(2.66)
	R^2_{within}	15.8%	15.8%	13.0%	24.0%	1.2%	0.7%
CMVE ($b \neq 0$ & $\gamma = 3$)	**Slope**	0.47	0.47	0.36	0.40	0.07	0.07
	(t_{stat})	(3.51)	(3.51)	(2.67)	(2.99)	(2.40)	(2.54)
	R^2_{within}	29.6%	29.6%	23.4%	32.9%	0.7%	0.7%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	**Slope**	0.62	0.62	0.51	0.54	0.31	0.31
	(t_{stat})	(3.00)	(3.00)	(2.46)	(2.71)	(2.16)	(2.11)
	R^2_{within}	36.1%	36.1%	31.7%	40.0%	14.0%	14.9%
	# Obs	610	610	610	610	610	610

Table IA.6
Regressing Pension Fund Portfolio Allocations on MVE Allocations (Controlling for Fund Characteristics)

This table reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of their consultants (ω^*).Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. All columns control for the same four fund characteristics used in Table 6 (lagged log pension fund size, a dummy for whether the pension fund has an internal public equity team, the lagged GASB funding ratio, and the lagged net pension fund cash flows). Panel A considers these characteristics in isolation (aggregated to the consultant-year level) and Panel B interacts these characteristics with asset class dummies. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.1.2 describes the results from this table.

PANEL A - Controlling for Fund Characteristics

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.47	0.47	0.95	0.40	0.32	0.07
	(t_{stat})	(3.51)	(3.51)	(11.5)	(2.99)	(3.70)	(2.51)
	R^2_{within}	29.6%	29.6%	75.6%	32.9%	12.4%	0.7%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.62	0.62	0.97	0.54	0.50	0.31
	(t_{stat})	(3.01)	(3.01)	(16.6)	(2.71)	(4.56)	(2.11)
	R^2_{within}	36.2%	36.2%	85.0%	40.0%	36.6%	14.9%
	# Obs	610	610	610	610	610	610

PANEL B - Controlling for Fund Characteristics Interacted with Asset Class Dummies

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.29	0.29	0.47	0.31	0.23	0.08
	(t_{stat})	(4.07)	(4.07)	(4.99)	(3.58)	(2.04)	(2.23)
	R^2_{within}	13.0%	13.0%	27.9%	14.4%	7.3%	0.9%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.44	0.44	0.61	0.47	0.44	0.32
	(t_{stat})	(4.21)	(4.21)	(5.57)	(3.56)	(3.30)	(2.48)
	R^2_{within}	25.3%	25.3%	45.0%	29.0%	31.6%	16.5%
	# Obs	610	610	610	610	610	610

Table IA.7
Regressing Pension Fund Portfolio Allocations on MVE Allocations (Fund-Level Regressions and Interaction with Time since Fund-Consultant Match, τ)

This table reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of their consultants (ω^*). Observations vary by pension fund (i), asset class (n), and year (t). To be consistent with our consultant-level (c) regressions, we use a weighted least squares with weights based on the AUM within the fund-years matched to each consultant-year. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Panel A considers regressions of ω onto ω^* whereas Panel B reports regressions of ω onto ω^*, τ, and $\tau\cdot\omega^*$ (analogous for active weights), where τ reflects the number of years since the given fund-consultant match was formed. For brevity, we omit from the table the slope coefficient for τ (but it tends to be small and statistically insignificant). Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.1.2 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	ω^*	0.47	0.47	0.95	0.40	0.32	0.07
	(t_{stat})	(3.36)	(3.36)	(11.0)	(2.86)	(3.54)	(2.43)
	R^2_{within}	16.8%	16.8%	68.5%	20.5%	5.7%	0.2%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	ω^*	0.62	0.62	0.97	0.54	0.50	0.31
	(t_{stat})	(2.88)	(2.88)	(15.9)	(2.60)	(4.36)	(2.02)
	R^2_{within}	20.5%	20.5%	77.1%	24.9%	16.7%	5.2%
	# Obs	7,100	7,100	7,100	7,100	7,100	7,100

PANEL B - Regressing Pension Fund Portfolio Allocations on MVE Allocations (Interacted with τ)

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	ω^*	0.35	0.19	0.92	0.13	0.32	-0.07
	(t_{stat})	(2.58)	(2.12)	(11.5)	(1.59)	(3.78)	(-1.28)
	$\tau\cdot\omega^*$	0.02	0.04	0.00	0.04	0.00	0.02
	(t_{stat})	(3.53)	(3.17)	(0.60)	(3.65)	(0.21)	(2.19)
	R^2_{within}	19.0%	21.3%	68.6%	25.1%	5.7%	3.3%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	ω^*	0.47	0.22	0.93	0.17	0.50	0.02
	(t_{stat})	(2.15)	(1.07)	(18.8)	(0.97)	(4.18)	(0.40)
	$\tau\cdot\omega^*$	0.02	0.04	0.01	0.04	0.00	0.03
	(t_{stat})	(4.46)	(3.24)	(1.07)	(4.11)	(0.35)	(3.01)
	R^2_{within}	22.8%	23.8%	77.1%	28.8%	16.8%	7.6%
	# Obs	7,100	7,100	7,100	7,100	7,100	7,100

Table IA.8
Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (using Three-Way Wild Boostrap Clustering)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c, n, and t using a wild cluster bootstrap method, with p-values reported in brackets (see Subsection C.3 for details). Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.3 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

Variation =		Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.47	0.47	0.95	0.40	0.32	0.07
	[*p-val*]	[0.03]	[0.03]	[0.02]	[0.04]	[0.06]	[0.31]
	R^2_{within}	29.6%	29.6%	75.6%	32.9%	12.4%	0.7%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.62	0.62	0.97	0.54	0.50	0.31
	[*p-val*]	[0.04]	[0.04]	[0.01]	[0.06]	[0.03]	[0.08]
	R^2_{within}	36.1%	36.1%	85.0%	40.0%	36.6%	14.9%
	# Obs	610	610	610	610	610	610

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

Variation =		Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	From ω^*	63%	63%	79%	83%	38%	10%
		[0.04]	[0.04]	[0.03]	[0.02]	[0.06]	[0.26]
	From ω^α	37%	37%	21%	17%	62%	90%
		[0.16]	[0.16]	[0.04]	[0.46]	[0.03]	[0.01]
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	From ω^*	58%	58%	88%	74%	73%	48%
		[0.03]	[0.03]	[0.02]	[0.03]	[0.02]	[0.06]
	From ω^α	42%	42%	12%	26%	27%	52%
		[0.05]	[0.05]	[0.05]	[0.13]	[0.03]	[0.02]
	# Obs	610	610	610	610	610	610

Table IA.9
Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (using Equal Weights)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t (using equal-weights). MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.4 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.45	0.45	0.98	0.37	0.30	0.02
	(t_{stat})	(3.88)	(3.88)	(11.9)	(2.98)	(3.25)	(0.61)
	R^2_{within}	27.8%	27.8%	77.6%	29.6%	11.0%	0.1%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.63	0.63	1.01	0.55	0.50	0.31
	(t_{stat})	(3.50)	(3.50)	(18.9)	(3.08)	(4.33)	(1.97)
	R^2_{within}	38.0%	38.0%	87.7%	41.1%	38.3%	14.9%
	# Obs	610	610	610	610	610	610

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	From ω^*	61%	61%	80%	80%	37%	3%
		(5.05)	(5.05)	(34.6)	(4.99)	(4.88)	(0.74)
	From ω^α	39%	39%	20%	20%	63%	97%
		(3.22)	(3.22)	(8.90)	(1.24)	(8.42)	(23.2)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	From ω^*	60%	60%	87%	75%	76%	48%
		(5.19)	(5.19)	(40.7)	(5.43)	(9.36)	(2.89)
	From ω^α	40%	40%	13%	25%	24%	52%
		(3.47)	(3.47)	(6.24)	(1.86)	(2.94)	(3.08)
	# Obs	610	610	610	610	610	610

Table IA.10
Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (using Actual Pension Fund Allocations)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.5 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.50	0.50	1.02	0.38	0.39	0.12
	(t_{stat})	(3.99)	(3.99)	(14.3)	(2.69)	(3.67)	(2.42)
	R^2_{within}	35.0%	35.0%	77.3%	30.4%	15.0%	1.6%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.66	0.66	1.03	0.53	0.58	0.38
	(t_{stat})	(2.89)	(2.89)	(24.5)	(2.31)	(4.07)	(2.30)
	R^2_{within}	38.3%	38.3%	86.4%	36.0%	38.7%	18.2%
	# Obs	610	610	610	610	610	610

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	From ω^*	70%	70%	76%	81%	38%	14%
		(5.77)	(5.77)	(33.5)	(4.21)	(5.64)	(4.12)
	From ω^α	30%	30%	24%	19%	62%	86%
		(2.51)	(2.51)	(10.8)	(1.01)	(9.02)	(25.2)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	From ω^*	58%	58%	84%	68%	67%	48%
		(8.00)	(8.00)	(72.3)	(5.35)	(9.35)	(4.70)
	From ω^α	42%	42%	16%	32%	33%	52%
		(5.84)	(5.84)	(14.2)	(2.53)	(4.58)	(5.12)
	# Obs	610	610	610	610	610	610

Table IA.11
Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (using Pension Fund Expected Returns when Available)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*), with consultant μ replaced by pension fund μ (when available). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.6 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.46	0.46	0.94	0.40	0.30	0.08
	(t_{stat})	(4.18)	(4.18)	(12.6)	(3.38)	(3.01)	(1.87)
	R^2_{within}	32.5%	32.5%	75.9%	34.7%	9.7%	1.0%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.63	0.63	0.96	0.58	0.53	0.34
	(t_{stat})	(3.68)	(3.68)	(19.7)	(3.18)	(5.38)	(2.60)
	R^2_{within}	39.5%	39.5%	86.6%	43.2%	37.1%	17.2%
	# Obs	610	610	610	610	610	610

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	From ω^*	70%	70%	81%	87%	32%	13%
		(4.61)	(4.61)	(29.5)	(5.08)	(5.91)	(1.71)
	From ω^α	30%	30%	19%	13%	68%	87%
		(1.98)	(1.98)	(6.86)	(0.78)	(12.3)	(11.9)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	From ω^*	63%	63%	90%	74%	70%	50%
		(4.85)	(4.85)	(44.1)	(5.60)	(5.72)	(3.00)
	From ω^α	37%	37%	10%	26%	30%	50%
		(2.86)	(2.86)	(4.78)	(1.95)	(2.50)	(3.00)
	# Obs	610	610	610	610	610	610

Table IA.12

Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (including CMAs that do not cover Hedge Funds)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is calibrated (to $\gamma = 3$) or estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider weights ($\omega$) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. We remove Hedge Funds from ω and ω^* for the consultant-years with CMAs that do not cover Hedge Funds. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection C.7 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	Slope	0.37	0.35	1.01	0.35	0.31	0.03
	(t_{stat})	(2.49)	(2.35)	(16.6)	(2.76)	(4.03)	(0.78)
	R^2_{within}	19.3%	17.8%	79.2%	27.1%	12.0%	0.1%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	Slope	0.45	0.44	0.95	0.44	0.42	0.20
	(t_{stat})	(2.20)	(2.09)	(20.4)	(2.39)	(3.73)	(1.48)
	R^2_{within}	25.3%	23.6%	85.1%	31.8%	28.2%	7.6%
	# Obs	878	878	878	878	878	878

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

	Variation =	Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = 3$)	From ω^*	52%	51%	79%	77%	39%	4%
		(3.02)	(2.84)	(33.7)	(4.30)	(5.22)	(0.91)
	From ω^α	48%	50%	21%	23%	61%	96%
		(2.76)	(2.79)	(9.08)	(1.28)	(8.21)	(21.7)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	From ω^*	56%	54%	89%	72%	66%	37%
		(4.16)	(3.85)	(37.9)	(5.66)	(6.06)	(1.97)
	From ω^α	44%	46%	11%	28%	34%	63%
		(3.29)	(3.30)	(4.55)	(2.15)	(3.06)	(3.31)
	# Obs	878	878	878	878	878	878

Table IA.13
Risk Aversion Estimates

This table reports the distributions of risk aversion estimates for different specifications of our mean-variance solution. The UMVE specification imposes no constraints (so that the lagrange multipliers underlying ω^{mv} are zero) whereas the CMVE specifications incorporate non-negative weight constraints. Specifications with $b = 0$ set the benchmarking incentive to zero whereas specifications with $b \neq 0$ calibrate the benchmarking incentive so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$). We consider two estimation methods: ordinary least squares (OLS) and instrumental variable (IV). Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection D.4 describes the risk aversion estimation methods we consider as well as the results from this table.

Est.	ω^* Specification	$\mathbb{C}or[\gamma^{OLS},\gamma^{IV}]$	Average	Std. Dev.	Q1%	Q5%	Q10%	Q25%	Q50%	Q75%	Q90%	Q95%	Q99%
OLS	UMVE ($b = 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.88	2.9	0.6	1.9	2.0	2.2	2.4	2.8	3.3	3.9	4.1	4.6
	CMVE ($b = 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.92	2.8	0.6	1.6	1.8	2.1	2.4	2.7	3.1	3.7	3.9	4.3
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.94	7.7	11.2	0.5	0.9	1.3	2.1	2.7	5.3	23.3	31.1	49.3
IV	UMVE ($b = 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.88	2.8	0.7	0.9	1.9	2.1	2.4	2.7	3.2	3.7	3.9	4.4
	CMVE ($b = 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.92	2.6	0.7	0.4	1.7	1.9	2.3	2.6	3.0	3.5	3.7	4.2
	CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}$)	0.94	8.1	12.3	0.4	0.8	1.2	1.8	2.6	6.0	23.5	32.2	55.0

Table IA.14

Explaining Variation in Pension Fund Portfolio Allocations with MVE Allocations (OLS vs IV)

Panel A reports regressions of pension fund portfolio allocations (ω) onto the MVE allocations implied by the beliefs of consultants (ω^*). Panel B reports decompositions of the variance of ω into the effect of ω^* and $\omega^\alpha = \omega - \omega^*$. Observations vary by consultant (c), asset class (n), and year (t). Pension fund allocations for consultant c in year t reflect the aggregate allocations of all pension funds matched to consultant c in year t. MVE allocations are given by $\omega^* = (1-b)\cdot\omega^{mv} + b\cdot\omega^b$ (as in Equation 5), with ω^{mv} reflecting pure MVE allocations (from Equation 6) and ω^b reflecting benchmark allocations. The benchmarking incentive (b) is calibrated so that the average tracking error from ω^* is the same as that from ω. The risk aversion parameter is estimated allowing for time variation and heterogeneity ($\gamma = \widehat{\gamma}_{c,t}$) using two alternative estimation methods: ordinary least squares (OLS) and instrumental variable (IV) (see Subsection D.6 for details). We consider weights (ω) and active weights ($\widetilde{\omega} = \omega - \omega^b$) for both the dependend and independent variables. Columns capturing variation across pension funds (at the consultant level) have $t \times n$ fixed effects, columns capturing variation across asset classes have $t \times c$ fixed effects, and columns capturing variation over time have $n \times c$ fixed effects. Standard errors are clustered by c and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection D.5 describes the results from this table.

PANEL A - Regressing Pension Fund Portfolio Allocations on MVE Allocations

Variation =		Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}^{OLS}$)	Slope	0.62	0.62	0.97	0.54	0.50	0.31
	(t_{stat})	(3.00)	(3.00)	(16.6)	(2.71)	(4.55)	(2.11)
	R^2_{within}	36.1%	36.1%	85.0%	40.0%	36.6%	14.9%
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}^{IV}$)	Slope	0.54	0.54	0.95	0.50	0.44	0.23
	(t_{stat})	(2.10)	(2.10)	(14.2)	(2.11)	(2.93)	(1.33)
	R^2_{within}	28.5%	28.5%	83.4%	33.8%	29.1%	8.7%
	# Obs	610	610	610	610	610	610

PANEL B - Decomposing Variation in Pension Fund Portfolio Allocations

Variation =		Across Pension Funds		Across Asset Classes		Over Time	
ω^* Specification		ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$	ω	$\widetilde{\omega} = \omega - \omega^b$
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}^{OLS}$)	From ω^*	58%	58%	88%	74%	73%	48%
		(4.81)	(4.81)	(35.5)	(5.56)	(6.49)	(2.75)
	From ω^α	42%	42%	12%	26%	27%	52%
		(3.43)	(3.43)	(4.96)	(2.00)	(2.44)	(3.01)
CMVE ($b \neq 0$ & $\gamma = \widehat{\gamma}_{c,t}^{IV}$)	From ω^*	53%	53%	88%	67%	66%	37%
		(3.94)	(3.94)	(37.5)	(4.59)	(5.22)	(1.93)
	From ω^α	47%	47%	12%	33%	34%	63%
		(3.55)	(3.55)	(5.08)	(2.25)	(2.73)	(3.26)
	# Obs	610	610	610	610	610	610

Table IA.15

Regressions Evaluating the Effect of Structural Endogeneity on Risk Aversion Estimation

This table reports regressions evaluating the effect of structural endogeneity on risk aversion estimates, which is a consequence of the fact that the beliefs of pension funds differ from the beliefs of their consultants. The structural endogeneity induces correlation (across asset classes) between the explanatory variables and the error terms of the regressions we can use to estimate risk aversion (Equations 4 general case and Equation IA.37 in the simplified specification we study here). Let $\epsilon_{\mu,c,n,t} = \Sigma_{i \in I_{c,t}} w_{i,t} \cdot (\mu_{i,n,t} - \mu_{j,n,t})$ reflect the AUM-weighted expected return gap from pension funds matched to consultant c at time t. Assuming that risk aversion is homogenous across pension funds matched to the same consultant and that pension funds and their consultants agree on their subjective covariance matrices, we have $\alpha_{c,n,t} = -\epsilon_{\mu,c,n,t}$, which allow us to evaluate the effect of structural endogeneity on our risk aversion estimation. The first column shows that structural endogeneity has a weak effect (economically and statistically) on our baseline risk aversion estimates under the stated homogeneity assumptions. Still, the last two columns show that $\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$ serves as a good instrument for $\mathbb{C}ov_{c,t}[r_n, r_c]$ when estimating risk aversion by instrumental variable (IV) regressions applied to Equation IA.37. So, we also explore this approach to estimate risk aversion and obtain results that are similar to the ones we report in the main text (see Table IA.14). Standard errors are clustered by j and t. Section 1 describes the data on pension fund portfolio allocations and consultant beliefs, Section 2 provides details about the MVE allocations, and Subsection D.6 describes the results from this table.

	Exogeneity Condition		**Relevance Condition**
Dependent Variable (y) =	$-\epsilon_{\mu,c,n,t}$	$-\epsilon_{\mu,c,n,t}$	$\mathbb{C}ov_{c,t}[r_n, r_c]$
Independent Variable (x) =	$\mathbb{C}ov_{c,t}[r_n, r_c]$	$\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$	$\mathbb{C}ov_{c,t}[r_n, r_{EQ}]$
Slope	0.038	0.016	0.687
(t_{stat})	(0.41)	(0.24)	(47.1)
R^2_{within}	0.3%	0.1%	97.1%
# Obs	210	210	610
Year×Consultant FE	x	x	x