

NBER WORKING PAPER SERIES

WHAT DRIVES VARIATION IN INVESTOR PORTFOLIOS? ESTIMATING THE
ROLES OF BELIEFS AND RISK PREFERENCES

Mark L. Egan
Alexander MacKay
Hanbin Yang

Working Paper 29604
<http://www.nber.org/papers/w29604>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2021, Revised September 2025

We thank Vivek Bhattacharya, John Beshears, John Campbell, Xavier Gabaix, Sam Hanson, David Laibson, Erik Loualiche, Maarten Meeuwis, Stavros Panageas, Andrei Shleifer, Paolo Sodini, Adi Sunderam, Kairong Xiao, Motohiro Yogo, and seminar and conference participants at Harvard Business School, Indiana University, Johns Hopkins Carey Business School, Micro and Macro Implications of Household Behaviour and Financial Decision-Making Seminar, NBER Asset Pricing Meetings, Oklahoma State University, Southern Methodist University, the Einaudi Institute for Economics and Finance, the Stanford Institute for Theoretical Economics, the University of Mannheim, the University of Minnesota Corporate Finance Conference, the University of North Carolina Roundtable for Junior Faculty in Finance, the University of Texas, the University of Toulouse, the Western Finance Association Meetings, and the Workshop on Financial Intermediation and Regulation at Queens University. The authors gratefully acknowledge financial support from the Division of Research of the Harvard Business School. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Mark L. Egan, Alexander MacKay, and Hanbin Yang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Drives Variation in Investor Portfolios? Estimating the Roles of Beliefs and Risk Preferences

Mark L. Egan, Alexander MacKay, and Hanbin Yang

NBER Working Paper No. 29604

December 2021, Revised September 2025

JEL No. G0, G11, G12, G40, G5, G51, J32

ABSTRACT

We present a portfolio choice demand model that allows for the nonparametric estimation of investors' (subjective) expectations and risk preferences. Using comprehensive 401(k)-plan-level data from 2009 through 2019, we explore heterogeneity in asset allocations using our empirical framework. We recover investors' beliefs about each asset and examine the implications and potential sources of those beliefs. Heterogeneity in expectations across investors accounts for twice as much variation in portfolio holdings as heterogeneity in risk aversion. Belief heterogeneity is partly driven by investors' characteristics and experiences, reflecting local sources of information such as county-level GDP and employers' past performance.

Mark L. Egan

Harvard University

Department of Economics

and NBER

megan@hbs.edu

Hanbin Yang

London Business School

hyang@london.edu

Alexander MacKay

University of Virginia

Department of Economics

mackay@virginia.edu

1 Introduction

Investors make portfolio decisions by allocating among multiple assets. As is well understood, heterogeneity in (subjective) beliefs about expected investment performance and preferences for risk can lead to different allocations across investors. Understanding these dimensions of heterogeneity can provide insight into investor decisions in the cross section and help illuminate how allocations across individual assets evolve over time. Further, the distribution of beliefs and risk preferences may have important implications for asset prices and investor welfare.

In this paper, we develop a portfolio choice demand model in which investors can have heterogeneous and subjective beliefs about individual funds. The model allows us to rationalize observed portfolios in terms of differences in beliefs and risk preferences, and we provide an empirical strategy that can nonparametrically recover these two components based on investors' revealed preferences. The model can also quantify the impacts of heterogeneity by constructing counterfactual portfolios that would be chosen if investors were identical. We apply the approach to aggregated portfolio data from 401(k) plans. In this context, we find that heterogeneous beliefs about expected returns play a first-order role in determining portfolio allocations. We exploit the panel structure of our data to shed new light on the drivers of investor expectations and how changes in expectations affect allocations. Heterogeneity across investors can be partially explained by exposure to local information—through demographic, geographic, and employment channels—which can drive substantial differences in beliefs, investment decisions, and ultimately investor welfare.

We estimate our portfolio choice demand model using data encompassing the allocation of funds within 401(k) plans from 2009 through 2019. Our data comes from BrightScope Beacon, which provides annual plan-level details about investment menus and fund allocations for 70,000 different 401(k) plans. The data cover 85 percent of assets in employer-sponsored investment accounts (defined contribution plans) that are subject to the Employee Retirement Income Security Act of 1974 (ERISA). 401(k) plans are a key component of investor wealth and a significant source of exposure to equity markets. As of 2021, Americans held roughly \$7 trillion in 401(k) assets¹ and, over our sample period, plan participation rates are high—74 percent on average. Approximately half of Americans participate in the stock market, and for 60 percent of those participants, defined contribution plans are their sole source of equity exposure (Badarinza et al., 2016).²

There is substantial heterogeneity in allocations across 401(k) plans. We show that, on average, 44 percent of assets are allocated to US equity funds, but this ranges from 17 percent to 64 percent for the 10th and 90th percentiles. These differences cannot be explained by differences in retirement-plan participation among these groups, as we study the within-plan

¹https://www.ici.org/faqs/faq/401k/faqs_401k

²Defined contribution plans account for the bulk of equity participation in the US and roughly one-third of retirement assets. https://www.ici.org/system/files/2021-06/21_rpt_recsurveyq1.pdf

allocation decisions conditional on participation. Nor can they be explained by differences in available investment options across plans. While some earlier work indicates that the choices of investors are driven by the menu of funds (Benartzi and Thaler, 2001, 2007), we find a substantially weaker relationship. Instead, our analysis suggests that investors make conscious (and different) allocation decisions. We estimate that a 10 basis point (bp) increase in fund expense ratios is associated with a 6.7% decrease in demand, which suggests that fees play an important role in allocation decisions.³

Consequently, we focus on how, conditional on available investment options and fees, differences in risk preferences and beliefs across investors explain variation in holdings. To interpret the decisions of investors, we model an investor’s portfolio decision as a myopic mean-variance optimization problem a la Markowitz (1952). When forming a portfolio, an investor trades off her subjective expectations with the corresponding additional risk, according to her risk preference. We interpret the risk preference parameter as risk aversion, though we discuss other possible interpretations. Our baseline model is stylized—we do not directly account for outside assets, dynamic allocations over time, or trading frictions such as inattention. We use a series of robustness checks to help assess the potential impact of these limitations on our estimates. Overall, we view our model as a natural starting point for considering how heterogeneity in investor beliefs affects allocations over multiple risky assets.

We implement a new identification strategy to address the challenge of recovering beliefs and risk aversion using holdings data. We show how exogenous variation in fees, such as variation induced by an instrument, can be used to recover the joint distributions of beliefs and risk preferences across investors. Fees affect net returns; by understanding how investors would re-allocate in response to a change in fees, we can measure how investors trade off expected returns and risk. Conceptually, our method allows us to nonparametrically recover each investor’s risk preferences and beliefs about each asset in their portfolio. Our approach can be implemented in portfolio choice problems where investors face plausibly exogenous investment costs, such as varying fees. The model can be applied to data with portfolios of individual investors, but we show how it can also be applied to aggregate allocations, which we consider a useful feature.

We use the model to estimate the time-varying distributions of risk aversion and expected returns for each investment option, which may vary arbitrarily across plans. When applied to aggregate portfolios—like the retirement plans in our data—our estimates reflect the average preferences and beliefs of investors within each plan. For example, consider the electric truck manufacturer Rivian. Our estimates suggest that in 2019 these participants, on average, had a risk aversion of 3.5 and were relatively optimistic about the return of the market. They expected the return of Vanguard’s large cap equity fund (VFIAX) to be 8.09% and expected the return of Vanguard’s small cap equity fund (VSMAX) to be 10.23%. In contrast, employees participating

³This is consistent with the evidence documented in Kronlund et al. (2021).

in the candy company Jelly Belly’s 401(k) plan exhibited similar risk aversion (3.3) but were more pessimistic about the market. These participants expected the return of VFIA, which was also an option in their 401(k) plan, to be 4.45% in 2019. Importantly, our framework captures this heterogeneity across investors, which we show has critical implications for holdings and welfare.

Overall, we recover time-varying distributions of both risk aversion and beliefs that are consistent with previous research and realized returns. Across funds and categories, investors have higher expected returns for riskier assets, and expected returns have reasonable magnitudes. In our baseline specification, the average investor in our sample behaved as if she expected the excess return of the market to be 9.6% over the period 2009-2019. To put this in perspective, the compound annual excess return of the S&P 500 was 10.7% over the same period. Across investors, the standard deviation of expected market returns is 2.3%, reflecting large heterogeneity in beliefs. Because our framework recovers investors’ asset-level beliefs, we can also use it to infer beliefs about specific funds and factors. We find that the average expected return of the value and size factors was roughly 2% each, although expected returns on the value factor varied over time. We estimate an average constant relative risk aversion parameter close to 4, which is comparable to what other researchers have found in the literature.⁴

To provide some external validation for our estimates, we also show that, for individual funds, the first and second moments of our estimated distributions of beliefs are positively and significantly correlated with the corresponding moments of the (implied) beliefs of equity analysts. We also find that our model performs well out of sample: our estimates of beliefs and risk aversion predict investors’ holdings at the fund level, generating an in-sample R^2 of 23% and a comparable out-of-sample R^2 of 17%. Furthermore, by leveraging the mean-variance structure of the decision problem in conjunction with our preference and belief estimates, we further improve the out-of-sample predictability of our model.

Investor behavior may not perfectly align with the assumptions of our model. In general, investors’ beliefs about expected returns and risk might be either distorted and/or uninformed. Our model accounts for this possibility by allowing beliefs to vary arbitrarily across investors.⁵ If investors do not trade off risk and returns when choosing portfolios and/or solve different portfolio choice problems, our results should then be viewed as investors acting as if their preferences and beliefs are as such. We still believe these “as if” preferences can be useful because they may be used to understand differences in portfolios across investors even if the underlying decisions are more or less complicated. A second issue is that investors may not make active decisions in every period. Our findings are robust to such inertia in that we obtain

⁴For example, using life cycle models, Fagereng et al. (2017) estimate relative risk aversion of 7.3, Calvet et al. (2019) estimate relative risk aversion of 5.8, Meeuwis (2019) estimate relative risk aversion of 5.4, and Choukhmane and de Sliva (2022) estimate relative risk aversion of 3.1.

⁵In our baseline framework we assume that investors have heterogeneous beliefs about expected returns but agree on the covariance structure of returns. If investors have heterogeneous beliefs about the covariance structure, our methodology will recover investors’ risk-adjusted expected returns.

similar estimates of risk aversion and expected returns if we estimate the model using only new plans in the year of inception and no default options, where allocation decisions are more likely to reflect investors' active choices. Despite these potential limitations, we obtain plausible estimates of beliefs and preferences, and we therefore view our approach as a reasonable way to capture investor portfolio decisions.

Accounting for heterogeneity in both risk aversion and beliefs is important for fitting the investment patterns in the data. A simple two-parameter model with risk aversion and beliefs explains over 50% of the reduced-form variation in equity holdings across plans. To more precisely evaluate the extent to which heterogeneity in beliefs and risk aversion shape investment behavior, we use our model to calculate counterfactual allocations where investors have identical beliefs, identical risk aversion, or both. We find that heterogeneity in beliefs contributes to the majority of variation in across-plan allocations. We also explore, to the extent that this belief heterogeneity results from incomplete information and biases, the potential costs of such distortions. We calculate the ex-ante loss in risk-adjusted returns associated with holding inefficient portfolios due to these belief distortions. This cost is modest for most investors, yet it can be significant for those with more extreme beliefs. The median (mean) investor's ex-ante risk-adjusted expected return is 0.17 (0.72) percentage points lower due to these belief distortions. Although belief distortions have a meaningful impact on portfolio allocations, their welfare effects are mitigated in the 401(k) setting, where investors choose from a curated menu of diversified mutual funds that are often already near the efficient frontier.

We use our estimates to explore the determinants of heterogeneity in beliefs and risk preferences. We find that wealthier and more educated investors tend to have more optimistic market expectations, which is consistent with previous experimental and survey evidence documenting that households with lower socioeconomic status are more pessimistic about future stock returns and macroeconomic conditions (Kuhnen and Miu, 2017; Das et al., 2020).⁶ Conversely, older and minority investors tend to have more pessimistic market expectations. We also find that investors' beliefs are correlated with their work experience. For example, investors working in the real estate sector are 27% (2.3 pp) more optimistic about the expected return of the market than investors working in the construction sector, despite both sectors having potentially similar risk exposures.

Risk aversion also varies with demographics and employment. Older and more educated investors behave as if they are more risk averse, while wealthier investors, as measured by income, appear more risk tolerant. The variation in risk aversion and beliefs provides insight into why equity exposure varies with investor demographics. For example, our results suggest that beliefs, rather than risk aversion, explain why educated investors tend to tilt their portfolios

⁶This finding is also consistent with the prior literature that finds that equity allocation is positively correlated with wealth (Heaton and Lucas, 2000; Wachter and Yogo, 2010; Bach et al., 2020; Fagereng et al., 2020) and education (Black et al., 2018). Bekaert et al. (2017) document how both menu design and investor characteristics are related to international equity exposure in 401(k) plans.

towards equities. Conversely, both risk aversion and beliefs help explain why older investors tend to have lower equity exposure.

Lastly, we explore the dynamic factors driving heterogeneity in beliefs. Consistent with a long literature documenting that investors extrapolate their beliefs across a number of settings,⁷ we find that investors form extrapolative beliefs. We also present new evidence that investors extrapolate from local information. We find that local economic conditions, such as county-level GDP, population, and home price growth, are positively correlated with beliefs about market returns, above and beyond what is available from aggregate, macro-level information. For the subset of publicly traded employers, we also find that investors' expectations are positively correlated with the past performance of their employer, as measured by returns, investment, employment growth, and sales growth, even after controlling for industry-by-year fixed effects. Moreover, investors appear to hold more optimistic beliefs about their own sector in general. These key findings—that expectations demonstrate systematic and predictable cross-sectional differences and, in the time series, are influenced by local factors—point to the importance of personal experiences in the formation of heterogeneous beliefs (e.g., Malmendier and Nagel, 2011, 2015; Kuchler and Zafar, 2019). Our finding that investors respond to local information indicates that even potentially irrelevant information helps shape beliefs that have real stakes (Bordalo et al., 2022), which can be costly for investors.

The paper proceeds as follows: In Section 2, we describe the data used in our analysis and present some basic facts about how portfolio allocations differ across investors and over time. We introduce our model and estimation procedure in Section 3. In Section 4, we present our baseline estimates and show how risk aversion and beliefs vary in the cross section. In Section 5, we assess the implications of belief heterogeneity for allocations and investor welfare. We explore the dynamic factors that explain the formation of investor expectations in Section 6. Section 7 concludes.

Related Literature

A primary contribution of our paper is to develop an empirical model of portfolio demand that allows for arbitrary beliefs about individual funds. Our approach fits within the growing demand-based asset pricing literature that has been heavily influenced Koijen and Yogo (2019a). Koijen and Yogo (2019a) develop and estimate an asset pricing model based on investor portfolio choices for which the optimal portfolio can be reduced to a characteristics-based logit demand function that is often used in the industrial organization literature (Berry, 1994; Berry et al., 1995). This methodology has been extended to study various settings, including exchange rates (Koijen and Yogo, 2019b), cryptocurrencies (Benetton and Compani,

⁷For example, previous work documents extrapolation in the stock market (Benartzi, 2001; Greenwood and Shleifer, 2014), the housing market (Case et al., 2012), risk taking (Malmendier and Nagel, 2011), investment decisions (Gennaioli et al., 2016), and inflation markets (Malmendier and Nagel, 2015).

2021), bonds (Bretscher et al., 2020; Darmouni et al., 2022), competition in the stock market (Haddad et al., 2021), and global equities (Koijen et al., 2019).

Similar to this literature, we focus on the holdings and asset demand of investors. Our alternative framework and identification strategy emphasize separately estimating investors' subjective, heterogeneous, and potentially irrational beliefs alongside their heterogeneous risk preferences. Our approach reflects a portfolio choice model where investors optimize potentially complementary allocations across individual funds. Rather than using the underlying discrete-choice framework of Berry (1994), we incorporate arbitrary investor heterogeneity directly into a classic (Markowitz, 1952) portfolio problem. Our framework relates to Shumway et al. (2009), who use a revealed-preference approach to understand risk-adjusted beliefs of fund managers. A key distinction between our work and that of Shumway et al. (2009) is our focus on separately identifying risk aversion and beliefs.

Our model allows for the nonparametric identification of beliefs and risk preferences, which we identify by exploiting exogenous variation in fund expense ratios. In previous work, Egan et al. (2022) use variation in the cost of leverage to recover investor beliefs in the context of a discrete-choice problem for a particular asset. By contrast, here we develop a new methodology to address an investor's entire portfolio problem. At a high level, our approach relates to Barseghyan et al. (2013), who estimate beliefs and risk aversion in the context of insurance.

We study heterogeneity across investors using large-scale data on holdings. Alternatively, data from surveys that elicit beliefs from a sample of investors can also be used to inform the relationships between beliefs and allocations.⁸ Our work complements the recent work from Giglio et al. (2021) that follows this latter approach. Using novel survey and account-level data from Vanguard, Giglio et al. (2021) find evidence that beliefs are reflected in the portfolios of investors. Like Giglio et al. (2021), we document substantial and persistent heterogeneity in beliefs across retail investors and find similar relationships between beliefs and equity exposure. When looking at tax-advantaged retail accounts, which is most similar to our sample, Giglio et al. (2021) estimate that a 1pp increase in beliefs about stock market returns is correlated with a 1.34-3.55pp increase in equity share, depending on the investor's characteristics. Our estimates imply that a 1pp increase in beliefs about the stock market returns is correlated with a 3.68pp increase in equity share, close to the upper end this range.⁹ More recent field experiment/survey evidence from Beutel and Weber (2022) indicates that beliefs and portfolio decisions are causally linked and that individuals' portfolio choices are consistent with a standard Merton model of portfolio choice, which supports our empirical approach. Our approach provides insight into both investor risk aversion and beliefs based on aggregate portfolio data,

⁸See, e.g., Vissing-Jorgensen, 2003; Ben-David et al., 2013; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Nagel and Xu, 2019.

⁹See column 7 of Table 4 in Giglio et al., 2021. Unlike our approach, Giglio et al. (2021) do not control for risk aversion. When we do not control for risk aversion, we find that a 1pp increase in beliefs about stock market returns is correlated with a 2.77pp increase in equity share. See the discussion in Section 5.1.

allowing for insights when survey data is unavailable.¹⁰

A second primary contribution is to apply the model to provide insights about investor decisions in the context of 401(k) plans. Thus, our paper relates to the literature on retirement savings and household finance more generally (see Campbell, 2006, Benartzi and Thaler, 2007, Choi, 2015 and Gomes et al., 2021 for discussions of the literatures).¹¹ While we find that participation is high conditional on eligibility, Yogo et al. (2021) documents that many households, especially low-income households, do not have access to employer-sponsored retirement plans and that providing access could increase retirement account participation by upwards of 10pp. Another strand of literature focuses on menu design and fees (Pool et al., 2016; Pool et al., 2020; Bhattacharya and Illanes, 2021). For example, Bhattacharya and Illanes (2021) develop a Nash bargaining model of plan design and show how imperfect competition and agency frictions can lead to sub-optimal plans. By contrast, we focus on the asset allocation decisions conditional on both participation and the 401(k) menu.

Lastly, our paper also relates to the growing literature at the intersection of industrial organization and finance. In addition to the demand-based asset pricing literature, a increasing number of papers have used methods from the industrial organization literature to estimate demand for financial products, including deposits (Dick, 2008; Egan, Hortaçsu, and Matvos, 2017; Xiao, 2020; Wang, Whited, Wu, and Xiao, 2022; Egan, Lewellen, and Sunderam, 2022), digital currencies (Whited et al., 2022), insurance (Koijen and Yogo, 2016, 2022), and mortgages (Aguirregabiria et al., 2019; Robles-Garcia, 2019; Benetton, 2021; Benetton et al., 2021; Wang et al., 2022; Allen et al., 2023)

2 Data

2.1 Sources

Our primary data set comes from BrightScope Beacon. BrightScope Beacon provides detailed plan and fund level information for ERISA defined contribution plans, covering 85% of plan assets. BrightScope collects the data either directly from plan sponsors, or from publicly available sources ranging from The United States Department of Labor (DOL) to the Securities and Exchange Commission (SEC). We focus on 401(k) defined contribution plans. The data set

¹⁰Another alternative method uses data on asset prices to recover the distribution of beliefs of a single representative investor (Ross, 2015). By contrast, (i) we use data on allocations, and (ii) we recover the distribution of beliefs and risk aversion across investors. An advantage in our setting is that we observe plausibly exogenous variation in investment costs, which allows us to recover the distribution of both beliefs and risk aversion without making any assumptions about the structure of asset prices or beliefs.

¹¹A subset of this literature focuses on 401(k) enrollment and contributions and studies the effects of plan design such as automatic enrollment (e.g., Madrian and Shea (2001); Choi et al. (2007); Beshears et al. (2009); Carroll et al., 2009; and Choukhmane 2019) and firm matching (e.g., Choi et al., 2002; Duflo et al., 2006; Dworak-Fisher, 2011). Due perhaps in part to the impact of this earlier literature, we find that plan participation is relatively high (74% on average) in our sample.

covers 70,000 different 401(k) plans over the period 2009-2019, resulting in roughly 450k plan-by-year observations. For each 401(k) plan, BrightScope reports annual data on the specific investment options available to participants and the total amount invested (across all plan participants) in each investment option. BrightScope does not provide individual investor level holdings data but provides holdings at the plan level. The data also includes details on the investment options in terms of the fee structure and type of funds. Because each 401(k) plan offers, on average, 26 different investment options, we have 11 million observations at the investment option-by-plan-by-year level, which is the unit of observation in our baseline analysis.

We merge our investment menu level data from BrightScope with additional data from the DOL Form 5500. The DOL Form 5500 data provides additional plan-by-year level details on plan participants, including the number of plan participants, the plan participation rate, employer contributions, and the share of participants that are retired.

We supplement our 401(k) data with mutual fund and stock return data from CRSP. CRSP provides daily level return data for stocks and open-end funds and quarterly level expense data for open-end funds. We merge the investment option-by-plan-by-year data in BrightScope with data from CRSP at the ticker-by-year level.

2.2 Summary Statistics

Plan Characteristics: Table 1a displays plan level summary statistics for the BrightScope data. The average plan has \$85 million in assets and the average participant balance is \$66 thousand. Employers accounted for 34% of all contributions with the remaining coming from plan participants. Participants, on average, can choose from 26 different investment options in the plan menu. The average plan has 1,261 participants.

The results also indicate that participation rates are quite high and that most eligible employees participate in 401(k) plans. At the median (mean) plan in our sample, 83% (74%) of eligible employees participate. Participation rates remained relatively high and constant over our sample period of 2009-2019. While there has been concern about the lack of retirement savings in the US, these summary statistics suggest that 401(k) plan eligibility may be a more important factor in explaining low retirement savings than participation by eligible employees.¹²

Portfolio Holdings: We document substantial heterogeneity in 401(k) holdings across plans and over time. We group investment options into six major asset classes: US equities, bonds,

¹²In Appendix Table A11, we examine how participation rates vary with the demographics of eligible participants. We find that participation is positively correlated with age and negatively correlated with minority status. However, consistent with the evidence in Yogo et al. (2021) we find no relationship between minority status and participation once we condition income and wealth.

cash, target date funds, alternatives, and international equities.¹³ Figure 1 displays the portfolio weights for the major asset classes across plan-by-year observations. The average plan holds 44% of the 401(k) assets in US equities, but there is substantial heterogeneity across plans. The standard deviation of US equity allocations across plans is 19% with some plans having almost no money allocated to equities and others having 100% allocated to equities.¹⁴ Similarly, there is substantial heterogeneity in cash holdings across 401(k) plans. The average plan holds 11% in cash, but the standard deviation across plans is 13%.

Differences in allocations are predictable based on sectors of employment and demographic characteristics. Figure 2 displays the distribution of equity exposure by the 2-digit NAICS of the employer, where we determine industry based on the NAICS codes provided by BrightScope. Median equity exposure varies across sectors, ranging from 53.1 percent in Public Administration to 62.6 percent in Information. Some differences may be attributed to background risk, such as shocks to labor income. However, the pattern across sectors suggests that risk is not the only factor driving allocation decisions. For example, it is not obvious that employment in the Public Administration sector would be substantially riskier than employment in the Utilities sector. Differences in risk aversion and beliefs may play an important role in explaining variation across sectors, in addition to underlying risk.

We explore the relationships between beliefs and risk aversion to sectors and demographics in Section 4.¹⁵ Although we do not directly observe investor demographics, we merge our 401(k) data with industry-by-county-by-year level demographic proxies from American Community Survey (ACS). Appendix Table A2 shows that plans with wealthier and more highly educated participants tend to have higher equity exposure, while plans with a greater share of older, retired, and minority participants tend to have lower equity exposures.

Over the sample period, target date funds have become more popular, rising from 10 percent to over 30 percent of holdings. We exclude target date funds in our baseline estimates because they tend to be the default option in most plans and may not reflect active investor choices. Excluding target date funds, holdings across major asset classes have remained fairly stable over our sample period. Modest changes occurred around the time of the financial crisis, when investors increased the weight held in cash and bonds at the expense of US equities and international assets. Figure 3 displays the average portfolio weights for each of the major asset classes over time.

¹³When calculating US equity and bond shares for our summary statistics presented in this section, we assume that non-target date multi-asset class funds (i.e., allocation funds) hold sixty percent of their assets in US equities and forty percent in bonds. When estimating our quantitative model in Section 3, we calculate the equity/factor exposure of each fund using historical data.

¹⁴We find similar dispersion in equity exposures when we compute the equity beta for each portfolio (Appendix Figure A1a). We also find similar patterns if we focus on the cross section in 2019, restrict to plans with larger menus over at least 25 options, or examine 401(k) plans that were created after the Pension Protection Act of 2006, which changed how 401(k) plans were designed.

¹⁵Figure 2 displays the distribution of equity exposure by the 2-digit NAICS of the employer.

Fund/Investment Fees: BrightScope Beacon provides the latest expense ratios for each investment option, and we obtain historical expense ratio data for those investment options structured as mutual funds using data from CRSP. We obtain 6,596,581 fund-year observations for expense ratios. The mean and median values are 61 bps, and the standard deviation is 43 bps (Table 1b).

The data provide some reduced-form indicators that investors make active choices based on investment fees. Assets are disproportionately invested in funds with lower expense ratios. For example, in 2019 the average fund appearing on an investor's 401(k) menu charges an expense ratio of 57 bps; however, when weighted by dollars invested, the average expense ratio paid by investors is 26 bps (see Figure 4). In Appendix A.3, we formally estimate investors' sensitivity with respect to fees using the standard discrete-choice demand model developed in Berry (1994) and estimate an elasticity of demand equal to -0.40.¹⁶ A limitation of this demand system is that it imposes restrictive substitution patterns across assets, which may not be ideal for the 401(k) setting. A unique feature of asset demand, as opposed to demand for other goods, is that asset pricing theory provides specific insights into how investor substitution patterns should behave. These limitations and insights motivate our consideration of a portfolio choice model that we develop in Section 3.

Overall, these patterns indicate that there is substantial variation in portfolio holdings. These differences cannot be explained by differences in retirement-plan participation among these groups, as we study the within-plan allocation decisions conditional on participation. Nor can they be explained by the composition of the menu, as plan menus are largely uncorrelated with participant demographics.¹⁷ Instead, these results suggest that investors make conscious (and different) allocation decisions. The variation appears to be partially driven by investor characteristics and fees. In the next section, we develop and estimate a model of portfolio choice that allows us to explore the drivers of plan-level heterogeneity in terms of differences in investor beliefs and risk aversion.

¹⁶We compute the demand elasticity assuming a market share of 1/26 and fee of 0.61pp. It is also useful to compare our estimates with those found in other financial markets. For example, recent studies have found that demand is inelastic in bank deposit markets (0.20-0.75; Dick, 2008; Egan et al., 2017; and Xiao, 2020), privatized social security markets (0.3-1; Hastings et al., 2017), and equity brokerage markets (0.47; Di Maggio et al., 2021). In contrast, other researchers have found that demand is more elastic in life insurance markets (2.18; Koijen and Yogo, 2016) and mortgage markets (2-6; Buchak et al., 2018; Robles-Garcia, 2019; Benetton, 2021). This is intuitive and these results suggest that within a 401(k) plan, the available funds are less substitutable than mortgage providers and life insurers.

¹⁷See Appendix A.2. We also find that menu composition explains a smaller fraction of variation in investment allocations compared to that measured in Benartzi and Thaler (2001). This may be due to a greater number of options in investment menus and a higher level of investor sophistication in our sample.

3 Model

We model each investor's 401(k) portfolio allocation as a mean-variance decision problem to understand what drives the variation in investor portfolios. Each investor trades off her subjective and potentially biased expectation of the return associated with investing in one of the available 401(k) investment options with the additional risk scaled by risk aversion. Using this framework, we show how to separately identify an investor's beliefs about the expected returns of each asset and risk aversion.

Theoretically, our model can be applied to the portfolio decisions of individual investors. In practice, data on individual portfolios may not be widely available. We show how our empirical framework can accommodate aggregation across individual investors. Thus, our approach is similar in spirit to random coefficient demand models that employ aggregate data on shares (e.g., Berry et al., 1995), though our empirical model reflects a portfolio problem rather than a discrete-choice decision framework.

We use our estimates of beliefs and risk aversion to better understand the portfolio allocations of investors. Without this structural framework, an analysis of portfolio allocations provides limited insight into investors' decisions. That is because portfolio allocations are a function of both 401(k) plan design and investor preferences/beliefs. For example, if we were to observe an investor with a relatively small equity allocation, it could be because: (i) the investor is risk averse, (ii) the investor is pessimistic about the return of the market, and/or (iii) the equity investment options in the investor's 401(k) plan are expensive. Unlike portfolio allocations, our estimates of beliefs and risk aversion adjust for the menu of funds available in each investor's 401(k) plan. If two sets of investors with identical beliefs and preferences faced different plan menus, they may have different portfolio allocations. Our methodology allows us to recover the same set of beliefs despite the different observed allocations.

3.1 Investor's Problem

Each investor i forms a portfolio from the set of investment options (i.e. funds) $k = 1, \dots, K_i$ and a risk-free asset, where the set of investment options can vary across investors. We assume investors have mean-variance preferences with risk aversion λ_i . Investors choose the $K_i \times 1$ vector of weights ω_i to maximize

$$\max_{\omega} \omega_i'(\mu_i - p) + (1 - \omega_i' \mathbf{1})R_F - \frac{\lambda_i}{2} \omega_i' \Sigma_i \omega_i,$$

where μ_i is a vector of investor i 's expectations of fund returns, p is a vector of fund expenses, R_F is the risk-free return, Σ_i is the $K_i \times K_i$ covariance matrix of expected fund returns, and λ_i is risk aversion. The corresponding set of first order conditions is

$$\boldsymbol{\mu}_i - \boldsymbol{p} - \mathbf{1}R_F = \lambda_i \boldsymbol{\Sigma}_i \boldsymbol{\omega}_i. \quad (1)$$

For each investor, we have K_i first order conditions.

3.2 Empirical Framework

We assume that the return of each asset k follows a factor structure with L orthogonal factors f_{lt} and idiosyncratic component ϵ_{kt} . By construction the factors and idiosyncratic component each have a variance of one. We can then write returns as:

$$R_{kt} = \sum_{l=1}^L b_{klt} f_{lt} + \sigma_{kt} \epsilon_{kt},$$

yielding a covariance matrix

$$\boldsymbol{\Sigma}_{it} = \boldsymbol{b}_t \mathbf{I}_L \boldsymbol{b}_t' + \boldsymbol{\sigma}_t \mathbf{I}_K \boldsymbol{\sigma}_t'.$$

The factors are orthogonal by construction. We assume that the idiosyncratic component is uncorrelated across securities.

We assume investors agree on the factor structure and the loadings $(\boldsymbol{b}_t, \boldsymbol{\sigma}_t)$. Thus, differences in beliefs about returns for an asset k arise from differences in expected realizations of factors and the idiosyncratic component, $\mu_{ikt} = E_i[R_{kt}] = \sum_{l=1}^L b_{klt} E_i[f_{lt}] + \sigma_{kt} E_i[\epsilon_{kt}]$.

We can then rewrite the above first order condition for each asset k as

$$\mu_{ikt} - p_{kt} - R_F = \lambda_{it} \left(\sum_{l=1}^L b_{klt} \left(\sum_{j=1}^{K_i} b_{jlt} \omega_{ijt} \right) + \sigma_{kt}^2 \omega_{ikt} \right). \quad (2)$$

The term on the left hand side reflects the expected return net of fees associated with investing an additional dollar in fund k , and the term on the right hand side reflects the additional risk of investing an additional dollar in fund k . Note that the first order conditions are additively separable in expected returns and risk. This is convenient for considering data that is aggregated across investors. For example, in our empirical application, we observe the aggregated portfolio for all investors participating in the same defined contribution retirement plan m , and we employ related aggregate first order conditions in estimation.

Let \mathcal{I}_m denote the set of individuals participating in defined contribution plan m and A_i denote investor i 's total portfolio value. We can then write the value-weighted average of the first order conditions (eq. 2) across all individuals participating in defined contribution plan m (and facing the same set of options, $K_i = K_m$) as

$$\left(\frac{1}{\sum_{i \in \mathcal{I}_m} A_i} \right) \sum_{i \in \mathcal{I}_m} A_i (\mu_{ikt} - p_{kt} - R_F) = \lambda_{mt} \left(\frac{1}{\sum_{i \in \mathcal{I}_m} A_i} \right) \sum_{i \in \mathcal{I}_m} A_i \left(\sum_{l=1}^L b_{klt} \left(\sum_{j=1}^{K_m} b_{jlt} \omega_{ijt} \right) + \sigma_{kt}^2 \omega_{ikt} \right),$$

where we assume that all investors in plan m have the same risk aversion λ_{mt} . We make

the assumption that risk aversion is heterogeneous across plans but common within a plan for simplicity, but this is not a necessary assumption for recovering investor beliefs and preferences. We discuss and relax this assumption in Appendix E.2. Under this assumption, we obtain

$$\bar{\mu}_{mkt} - p_{kt} - R_F = \lambda_{mt} \left(\sum_{l=1}^L b_{kl} \left(\sum_{j=1}^{K_m} b_{jlt} \bar{\omega}_{mjt} \right) + \sigma_{kt}^2 \bar{\omega}_{mkt} \right),$$

where $\bar{\mu}_{mkt}$ is the average (dollar-weighted) expected return of asset k at time t across investors participating in defined contribution plan m that purchase asset k . The weight $\bar{\omega}_{mkt}$ is the average (dollar-weighted) portfolio weight.

Given the factor structure \mathbf{b}_t and the idiosyncratic variance σ_t , we can compute the risk associated with each fund k . We can then estimate the linear regression equation:

$$\left(\sum_{l=1}^L b_{kl} \left(\sum_{j=1}^{K_m} b_{jlt} \bar{\omega}_{mjt} \right) + \sigma_{kt}^2 \bar{\omega}_{mkt} \right) = \theta_{mt} p_{kt} + \epsilon_{mkt}, \quad (3)$$

where the parameter θ is the negative inverse of risk aversion (i.e., $\theta = \frac{-1}{\lambda}$) and ϵ_{mkt} is equal to average investor beliefs divided by risk aversion (i.e., $\epsilon_{mkt} = (\bar{\mu}_{mkt} - R_F)/\lambda_{mt}$). Eq. (3) is the heart of our estimation strategy. Identification requires exogenous variation in the fees investors pay for each investment option that is orthogonal to average investor beliefs (ϵ_{mkt}). With an appropriate instrumental variables strategy, we are able to recover the parameter θ_{mt} and consequently risk aversion λ_{mt} , regardless of the potential correlation between p_{kt} and ϵ_{mkt} . In principle, with a sufficient number of funds per plan, we could nonparametrically identify separate values of risk aversion for each plan and year combination. Given risk aversion, we can recover average beliefs as $\lambda_{mt} \epsilon_{mkt} = \bar{\mu}_{mkt} - R_F$.

3.3 Implementation

3.3.1 Risk

To estimate risk aversion and recover investor beliefs, we need to estimate the factor structure of fund returns (\mathbf{b}_t, σ_t). We estimate the factor structure using a 6-factor model where we include the Fama-French 3 factors and three bond factors: the excess return of long term government bonds, the excess return of investment-grade bonds, and the excess return of high-yield bonds.¹⁸

We estimate factor loadings for each mutual fund and equity in CRSP using weekly return

¹⁸We calculate long term government bond returns using Vanguard's Long-Term Treasury Fund (VUSTX) returns, the investment grade bond returns using Vanguard's Long-Term Investment-Grade Fund (VWESX) returns, and high yield bond returns using Vanguard High-Yield Corporate Fund (VWEHX) returns. We calculate excess returns relative to the risk free rate as reported in the Fama and French database.

data over the previous ten years where we allow factor loadings to vary year-to-year. We then merge the estimated factor loadings with our BrightScope data at the fund-by-year level using mutual fund and stock tickers. Our data also contains non-mutual fund and stock options, such as separate accounts. For these investment options, we do not observe high-frequency data, but we do observe their category classifications. We calculate the risk associated with these investment options based on the average risk of all other funds that belong to the same Morningstar category in the same year.¹⁹

As a robustness check, we also consider a simpler factor structure where we construct the factors by forming equal weighted portfolios based on the broad BrightScope categories reported in Table 1, with the idea that investors think of risk in terms of broad asset classes (e.g., bonds, international stocks, cash, etc.). We also estimate a 55-factor model following Shumway et al. (2009). Estimates of beliefs and risk aversion using these alternative methodologies are highly correlated with our baseline estimates. We provide comparison statistics in Table A9. In Appendix E.3, we also explore the case where investors account for labor income risk and find that investors behave as if they neglect risks related to labor income.

3.3.2 Expenses

We determine fund expenses using data from CRSP. One concern is that fund fees may be endogenously related to investor beliefs. For example, if a mutual fund provider anticipated that investors were optimistic about the returns of a particular fund, the fund provider might find it optimal to increase its expense ratio. This endogeneity would result in an upward bias in θ in eq. (3).²⁰

To help address this concern, we include plan-by-year fixed effects and fund classification-by-year²¹ fixed effects in our main specification. Thus, we allow fees to rise endogenously in response to the expectations of investors in specific plans or for specific fund categories in specific years, and we identify model parameters based on variation in expenses within plan-by-year and within classification-by-year. After including these fixed effects, the potential endogeneity concern would then be that, conditional on a 401(k) plan and fund classification, the residual variation in expenses is correlated with the residual variation in investor beliefs for specific funds.²²

¹⁹We do not observe the Morningstar category for a handful of options. For these funds we calculate risk based on the average risk of all other funds that belong to the same BrightScope category in the same year.

²⁰We do not model the selection of funds into menus on individual plans. With certain models of fund selection, this could introduce a distinct form of endogeneity bias in fees. To assess this, we have estimated our model while also including each fund's unconditional probability of menu inclusion as an independent variable, which we treat as a control function for selection. This specification yields a coefficient on fees that is 4 percent smaller in magnitude than that of our baseline specification. This difference is small and statistically insignificant, and it gives us some reassurance that endogeneity from menu selection is not a first-order concern in our setting.

²¹Fund classification categories include, e.g., US Equity Large Cap Value Equity, Real Estate Equity, etc.

²²For example, suppose that (i) Fidelity anticipates that participants in IBM's 401(k) plan have more optimistic beliefs about Fidelity's Large Cap Growth Index Fund relative to the other investment options in IBM's 401(k) plan

We instrument for fees to account for the potential endogeneity. In our baseline specification, we use Hausman-type instruments. Specifically, we use the average fee charged by the same mutual fund provider in other Lipper investment objective categories.²³ This instrument will be relevant (correlated with fees) when a provider's cost of operating a mutual fund is correlated with its costs of operating its other mutual funds, perhaps as a result of the provider's scale and technology. The instrument meets the exclusion restriction (provides exogenous variation) when participants' beliefs about the idiosyncratic expected returns of a given fund (after controlling for plan-by-year and category-by-year fixed effects) are, on average, uncorrelated with fees a provider charges on its funds from different Lipper investment objective categories. We consider both of these conditions to be plausible in our setting. A threat to exogeneity would be that, for example, an investor's belief about the expected return of Fidelity's Large Cap Growth Index Fund is correlated with the expenses Fidelity charges on its bond funds.

As robustness checks, we consider a suite of other instruments, including using fund-level variation in custodian fees (Janssen and Thiel, 2025), instruments based on the characteristics of other 401(k) plans offering the same funds (Waldfogel, 2003; Fan, 2013), and variation from provider switches that generates plausibly exogenous variation in fund pricing (Yang, 2025). We also consider two alternative Hausman instruments. We use (a) the average fees charged by the same provider for funds not offered in any 401(k) plan in our sample, in a different investment category and year, and (b) the average fees charged by the same provider in its separate accounts. These instruments draw on distinct market or investor segments and help mitigate a potential concern with our baseline strategy about joint fee-setting across categories.

We report full details of the instrument construction and corresponding results in Appendix C. Overall, we find estimates that are consistent with our baseline estimates across the different specifications.

3.3.3 Portfolio Weights

We construct portfolio weights using total assets (across all participants in the plan) for each investment option and year reported in BrightScope. When constructing portfolio weights we treat all investment options categorized in BrightScope as "Cash/Stable Value" as risk-free assets. We also exclude funds classified in BrightScope as target date funds because these funds are often the default option and tend to be held by passive investors. However, as reported in Appendix Table A9, we find qualitatively similar estimates if we include target date funds in

(average absorbed by plan-by-year fixed effects) and relative to average beliefs about other large cap growth funds (average absorbed by classification-by-year fixed effects) and that, as a result, (ii) Fidelity increases the expense ratio it charges on its Large Cap Growth Index Fund. While certainly possible, the fact that mutual fund fees are infrequently updated and set uniformly helps alleviate these endogeneity concerns.

²³When forming the instrument for fund k in plan m , we exclude all funds appearing on the menu for plan m when calculating the average fee charged by the mutual provider who manages fund k . The standard deviation of fees in our data is 0.43. The residual variation, after controlling for plan-by-year and category-by-year fixed effects, is 0.20. The residual standard deviation of our instrument, after including the same fixed effects, is 0.23.

our analysis.

3.4 Estimation

We estimate the empirical analog of the investor's first order conditions for choosing an optimal portfolio (eq. 3) in the following regression specification:

$$\varsigma_{mkt}^2 = \theta_{mt} p_{kt} + \phi_{mt} + \phi_{j(k)t} + \epsilon_{mkt}, \quad (4)$$

where

$$\varsigma_{mkt}^2 = \left(\sum_{l=1}^L b_{klt} \left(\sum_{j=1}^{K_m} b_{jlt} \bar{\omega}_{mjt} \right) + \sigma_{kt}^2 \bar{\omega}_{mkt} \right)$$

and ϕ_{mt} and $\phi_{j(k)t}$ are plan-by-year and fund type-by-year fixed effects. Here, subscript m denotes specific 401(k) plans, and $j(k)$ denotes fund type based on the fund's classification in both Morningstar and BrightScope as well as whether the fund is an index/passive fund. Thus, the fixed effect $\phi_{j(k)t}$ is a quadruple interaction term (i.e., Morningstar Category \times BrightScope Category \times Passive \times Year). Observations are at the investment option-by-plan-by-year level. Because each observation reflects the average behavior of plan participants, we weight each observation by the total assets of the 401(k) plan when estimating eq. (4). Our estimates allow us to recover risk aversion as $\hat{\lambda}_{mt} = -\frac{1}{\hat{\theta}_{mt}}$. In principle, risk aversion is nonparametrically identified for each plan-year, provided a sufficient number of funds per plan. In practice, we parameterize θ_{mt} to allow for some flexibility.

Our empirical framework also allows us to recover the average expected returns within investors in a 401(k) plan for each investment option available in the plan. We recover the average beliefs for each investment option based on our estimate of θ_{mt} , our estimated fixed effects, and the residual from eq. (4) after netting out fees:

$$\widehat{\bar{\mu}_{mkt} - R_F} = -\frac{1}{\hat{\theta}_{mt}} \left(\hat{\phi}_{mt} + \hat{\phi}_{j(k)t} + \hat{\epsilon}_{mkt} \right). \quad (5)$$

Given each investor's beliefs about the expected return and the factor loadings for each investment option/fund, we can use the estimated distribution of beliefs to recover investors' expectations of the market return. We estimate the plan-by-year average expected market return at time t for each plan m based on the regression:

$$\widehat{\bar{\mu}_{mkt} - R_F} = \delta_{mt} b_{1kt} + \eta_{mkt}, \quad (6)$$

where b_{1kt} is the loading for fund k on the market factor at time t . Observations are at the fund-by-plan-by-year level. The parameter δ_{mt} , which varies at the plan-by-year level, reflects the average expected return of the market across participants in plan m at time t . Note that because

the other factors are orthogonal to the market by construction, we do not need to control for the other factors in eq. (6).

3.5 Identification and Interpretation

We estimate risk aversion by measuring how investors trade off risk and expected returns in eq. (4). Specifically, we estimate risk aversion by examining how investors adjust their portfolio risk exposure in response to changes in expense ratios. Changes in expense ratios are equivalent to shifts in expected returns, allowing us to calculate risk-return tradeoffs. An investor optimally sets the expected return of an investment equal to the marginal risk, scaled by risk aversion. Our approach relies on the following assumptions: we can correctly measure investors' beliefs about risk; investors make allocation decisions considering their retirement accounts only; investors solve a myopic portfolio problem; and investors only trade off risk and expected returns.

These assumptions impose some relatively strong restrictions on investor behavior; however, our estimates may still be interpretable if these assumptions are violated. Here, we briefly discuss how the interpretation of our estimates of beliefs and risk aversion would change if our baseline assumptions are violated. We provide a more detailed analysis in Appendix D. It is important to emphasize that while our methodology places some strong constraints on investor behavior, we do not impose rational beliefs in our analysis. Our framework allows for behavioral biases and mistakes in investor beliefs about expected returns.

- **Outside Assets:** By focusing on 401k investments, we only observe part of an investor's overall portfolio. According to 2019 Survey of Consumer Finance (SCF), working-age individuals who have retirement accounts allocate on average 64% of their financial assets in retirement accounts, compared to 7% in other investment funds, directly owned stocks, or bonds. Thus, retirement assets represent the vast majority of risky financial assets for most individuals with access to retirement accounts. In Appendix D, we provide more details on how ignoring outside cash holdings and human capital may influence our estimates.
- **Dynamic Allocation Across Multiple Periods:** We model an investor's allocation decision as a myopic portfolio choice problem. It is well known that when investors have power utility and return distributions are independent over time, long-term portfolio choice is equivalent to myopic portfolio choice. More general time-varying returns could introduce intertemporal hedging demand, which would not impact our estimates of risk aversion but would potentially be captured in our estimates of beliefs.
- **Optimization Error:** We assume that investors actively trade off and equate marginal risk with returns when making investment decisions. However, there are a few reasons this could be violated in the data. Departures from the optimization condition in eq.

(1) can be represented as optimization errors, and these errors can be interpreted as true errors or as unobserved preferences of consumers for particular funds. Provided the optimization error is additively separable, it will appear in the residual of eq. (4). This type of estimation error does not bias our estimate of risk preferences, even if it is correlated with beliefs or expense ratios, because we instrument for expense ratios. However, the idiosyncratic fund-specific beliefs we recover from the data will be measured with noise as a result of the optimization error. While this affects our measurement of these idiosyncratic beliefs, it does not bias our estimates of an investor's beliefs about the market or other factor returns, provided that the optimization error is orthogonal to a fund's factor loadings.

- **Inattention:** Relatedly, one might be concerned that investors do not actively trade off expected returns with risk. With inattentive investors, our estimate of risk aversion could be biased upwards because investors would appear as if they are insensitive to expected returns/fees. To assess this possibility, we exploit the fact that, when a 401(k) plan is introduced, any allocation into non-target date funds reflects an active choice of the participant. Thus, we use the allocation decisions in the first year of each plan as a robustness check, which we discuss in Section 4. We find that investors in existing plans behave as if their risk aversion is slightly higher than that of investors in new plans, which is consistent with the presence of investor inattention. However, inattention does not appear to be the driving factor behind our estimate of risk aversion.
- **Neglecting Fees:** In our framework we assume that investors respond to fees and changes in expected returns one-for-one. To the extent that investors are less sensitive to fees, perhaps as a result of salience or some other factor, relative to expected returns, our estimates of risk aversion will be biased upward. This is because, in our context, an investor's sensitivity to fees corresponds to the inverse of risk aversion. If investors appear artificially insensitive to fees, this will cause us to underestimate the inverse of risk aversion and thus overestimate risk aversion.
- **Heterogeneous Beliefs about Risk:** We assume that investors understand and agree on the risk of each investment options, and that we, as the econometricians, assess risk in the same way. Heterogeneous beliefs about risk will introduce an measurement error in our dependent variable. Because our dependent variable is the product of risk and allocation, this measurement error is not additively separable and can introduce biases in our risk aversion estimates. Interestingly, if investors disagree about risks but agree on the expected returns of assets, we still recover unbiased estimates of their expected returns on average. In Appendix D, we discuss further how heterogeneous beliefs about risk affect our estimates.

To summarize this section, we introduce a new methodology that independently identifies beliefs and preferences, which may be useful when survey data is not available. Naturally, this approach comes with certain trade-offs and imposes notable constraints on investor behavior. Some of these limitations stem from data-related challenges (for instance, the lack of access to individual data or the full scope of investors' portfolios), while others are more intrinsic to our method itself (such as the assumption that investors have mean-variance preferences and solve a myopic portfolio choice problem). As discussed in the following section, we show that we recover reasonable distributions of beliefs and preferences, and the first and second moments of these belief distributions align with survey estimates, lending some support to our approach.

4 Estimates of Risk Preferences and Beliefs

Here, we present our baseline estimates of risk aversion and beliefs and examine how they vary across investor demographics and characteristics. We find substantial heterogeneity across investors and find that this heterogeneity is highly correlated with investor demographics. To provide some external validation for our estimates, we also show that, for individual funds, the first and second moments of our estimated distributions of beliefs are positively and significantly correlated with the corresponding moments of the (implied) beliefs of equity analysts. We also find that our model performs well out of sample; our estimates of beliefs and risk aversion predict investors' holdings.

4.1 Risk Preferences

We report our baseline model estimates corresponding to eq. (4) in Table 2. We estimate each specification using two-stage least squares, where we instrument for expenses and the corresponding interaction terms using Hausman-type instruments as described in Section 3.3.²⁴ The independent variables, other than the dummy variables *Unionized* and *Existing 401(k) Plan*, are all standardized such that they are in units of standard deviations. We weight each observation by the total assets of the 401(k) plan.

In specification (1), we keep the parameter θ_{mt} and consequently risk aversion fixed across 401(k) plans. In specifications (2)-(5), we allow θ_{mt} and risk aversion to vary across plans based on plan characteristics/demographics. In specifications (3) and (5), we also allow for arbitrary year-by-year variation in the mean level of θ by interacting fees with time dummy variables. For each specification, the left column reports the model estimates and standard errors. Recall that the parameter θ_{mt} corresponds to the negative inverse of risk aversion ($\theta_{mt} = -\frac{1}{\lambda_{mt}}$). For ease of interpretation, we report the corresponding estimates in terms of risk aversion and demographic interactions in the right column (λ).

²⁴The first-stage F-statistics for specifications (1)-(5) in Table 2 are 3082, 734, 758, 683, and 726, respectively.

We estimate mean risk aversion ranging from 3.6 to 5.2 across our main specifications. We also find that accounting for heterogeneity in risk aversion, as discussed further below, is important for explaining investment decisions. The interaction terms in Table 2 indicate how demographics are correlated with the parameter θ_{mt} . We find evidence that older plan participants behave as if they are more risk averse. The results in specification (2) of Table 2 indicate that a one standard deviation increase in age is associated with a 0.38 (8%) increase in risk aversion. Education is positively correlated with risk aversion. A one standard deviation increase in fraction with some college education is correlated with a 0.56 (12%) increase in risk aversion (specification 2, Table 2). Wealthier investors, as measured by median family income, tend to behave as if they are less risk averse, such that a one standard deviation increase in log income is correlated with a 0.47 (10%) decrease in risk aversion (specification 2, Table 2).

Lastly, in specifications (4) and (5) of Table 2 we allow risk aversion to vary in the year the 401(k) plan was first introduced. As discussed in Section 3.5, if investors are inattentive, they may appear more risk averse in the data than they actually are. Consistent with this, we find that investors in existing plans behave as if their risk aversion is 0.89 (25%) higher compared to investors in the year of plan inception (specification 4). In Appendix Table A9, we also show that we get similar estimates of beliefs and risk aversion if we restrict our sample to the first year each 401(k) plan was introduced.

The remainder of our analysis employs estimates based on specification (5) of Table 2, where we allow the average level of risk aversion to vary over time. When constructing our estimates of risk aversion and beliefs, we set the variable *Existing 401(k) Plan* equal to zero to adjust for potential effects of inattention, and we winsorize both risk aversion and beliefs at the 1 percent level.

Figure 5a displays the estimated distribution of risk aversion over time. The solid red line displays the average risk aversion across plans and the dashed/dotted lines correspond to different quantiles of the distribution. The results suggest that risk aversion fell in 2010 as the economy was coming out of the global financial crisis and then peaked again in 2012 around the time of the European sovereign debt crisis. Figure 5a also illustrates the range of heterogeneity in risk aversion across plans/investors. Plans in the 90th percentile of the risk aversion distribution behave as if they are more than 25% more risk averse than plans in the 10th percentile of the risk aversion distribution. We find that this dispersion in risk aversion helps explain investors' portfolio allocations in Section 5.1.

Though risk aversion is an important factor in explaining allocations, investor beliefs are also important. Figure 5b plots the expected market returns versus risk aversion across plans. We find a modest positive correlation.²⁵ A substantial portion of the variation in beliefs cannot be predicted by risk aversion ($R^2 = 0.20$). We now turn to our estimates of investor beliefs.

²⁵As a caveat, if there is estimation error in our measure of risk aversion, this will translate into some positive correlation between beliefs and risk aversion given the way we construct beliefs.

4.2 Investor Beliefs

Our approach yields estimates of beliefs for each investment option in every plan. Figure 6a displays the estimated distributions of expectations of returns across investors for each investment category. For every plan in every year, we compute category-level expected returns by averaging expected returns across all investment options available in each category, and we plot the distribution of category-level returns across plans and years. Consistent with our priors, investors' expectations of returns are highest for small cap stock funds and lowest for bond funds. Similarly, Figure 6b displays the distribution of investors' expectations for three popular Vanguard index funds. Investors expect higher returns from Vanguard's small cap funds and lower returns from large cap funds. We find meaningful variation in expected returns at the level of individual funds.²⁶

4.2.1 Market Beliefs

We now focus on expectations about market returns. Figure 7a displays the distribution of beliefs about the market return (δ_{mt}) over the period 2009-2019, where we allow risk aversion to vary across plans (corresponding to specification 5 in Table 2). The bright red solid line displays the average belief across plans over time. The results suggest that optimism remained relatively constant over the early part of our sample as the average investor expected the market return to be roughly 11%. Investors remained optimistic through 2017 and then the average expected return fell to roughly 7.4% in 2019. The average expected return over our sample is 9.6%, which is close to the realized excess return of the S&P 500 over this period. The compound annual growth rate (CAGR) of the excess return of the S&P 500 over the period 2009-2019 was roughly 10.7%.

There is also substantial heterogeneity in beliefs across plans. In Figure 7a, we plot the 10th, 25th, 50th, 75th, and 90th percentile of expected returns in addition to the mean. Moving from the 10th to the 90th percentile of the distribution implies an increase in expected returns of roughly 5 percentage points in most years. For example, in 2011, the 10th percentile expected return is 8 pp and the 90th percentile is 14 pp. The standard deviation in expected market returns across plans within a year is 2.30 pp on average.

The differences in expected returns across plans are persistent. To demonstrate this, we calculate the average deviation from the within-year mean for each plan over time. Figure 7b displays the average plan-level deviation from the mean, i.e., the persistent cross-sectional variation in expected returns across plans. The standard deviation is 1.8pp, which is close to the plan-year standard deviation of 2.3. Thus, our estimates imply that relative pessimism and relative optimism about market returns are persistent features of investor beliefs.

Note that our analysis examines the cross-sectional dispersion in the average plan beliefs,

²⁶We report the portfolio-level implied Sharpe ratios in Figure A1b of the Appendix.

where each plan is a collection of individuals with the median plan having more than 200 participants. To the extent that there is variation in investor beliefs within plans, the dispersion shown in Figure 7a and Figure 7b could underestimate the individual-level dispersion in beliefs.

To better understand what drives heterogeneity in investor beliefs, we regress market beliefs (δ_{mt}) on a vector of plan characteristics (X_{mt}). Because risk aversion and beliefs tend to be positively correlated in the data and risk aversion is a deterministic function of the covariates X_{mt} , we examine how the variation in market beliefs that is orthogonal to risk aversion (δ_{mt}^*) varies with plan characteristics according to the following regression:

$$\delta_{mt}^* = X'_{mt}\Gamma + \nu_{mt}. \quad (7)$$

Observations are at the plan-by-year level. The dependent variable δ_{mt}^* measures the residualized variation in expected market returns averaged across investors participating in plan m at time t that is orthogonal to risk aversion.²⁷

Table 3 displays the estimates corresponding to eq. (7). We include year fixed effects in each specification. Columns (1)-(11) display univariate regressions and column (12) includes the full set of variables. In general, we find that wealthier and more educated investors tend to have more optimistic expectations about the market. This helps explain why wealthier investors have higher equity participation rates. The results in column (2) indicate that a one standard deviation increase in log income is associated with a 0.17 pp increase in expected market returns. Similarly, a one standard deviation increase in the fraction of college educated individuals is associated with a 0.23 pp increase in expected market returns (column 4, Table 3).

In contrast, we find that older investors, retirees, and minorities tend to have more pessimistic expectations about market returns. The results in column (12) indicate that a one standard deviation increase in the fraction of Hispanic (black) individuals is correlated with a 0.08 pp (0.10 pp) decrease in expected returns. One potential explanation is that market expectations could be affected by trust (Guiso et al., 2008; Gennaioli et al., 2015) which may differ across ethnicities (Chiteji and Stafford, 2000).

We also find some evidence that participants' beliefs are shaped by their industry. The results in column (9) and (12) indicate that investors who work in riskier sectors, as measured by the equity beta of their sector, tend to have more optimistic beliefs. We look at this further by examining how beliefs about the market vary across sectors in Figure 8. The results suggest that there is substantial heterogeneity across sectors. At the median, investors from the most optimistic sector, Real Estate, expect the market return to be roughly 30 percent higher than investors from the least optimistic sector, Accommodation and Food Services (10.8% versus 8.3%). Investors in Real Estate also have meaningfully higher expected returns than those

²⁷We calculate δ_{mt}^* as the residual from the regression of δ_{mt} on the parameter θ_{mt} , which corresponds to the negative inverse of risk aversion.

in Construction (10.8% versus 8.5%), despite having arguably similar risk profiles. We also find evidence that there is substantial heterogeneity in beliefs within a sector. The average interquartile range of beliefs within a sector is 3.1 pp. In other words, within a sector those investors in the 75th percentile of the beliefs distribution expect the market return to be roughly 40% higher than investors in the 25th percentile of the beliefs distribution.

It is interesting to compare equity allocations with beliefs across sectors as shown in Figures 2 and 8. To facilitate the comparison, Figure 8 sorts the sectors by median share allocated to U.S. equities. Though equity allocation and expected market returns are correlated, the correlation is far from perfect, suggesting the important role of variation in risk aversion across sectors as well.

4.2.2 Beliefs about Fama-French Factors

Similar to beliefs about market returns, we can extend equation (6) to estimate plan m 's expected return (δ_{mlt}) for factor l at time t :

$$\bar{\mu}_{mkt} - R_F = \sum_{l=1}^L \delta_{mlt} b_{lkt} + \eta_{mkt}$$

For brevity, we focus our discussion and analysis on the Fama French factors, size and value ($\delta_{m, size, t}$, $\delta_{m, value, t}$), because they are two of the most commonly studied factors in asset pricing and many mutual fund providers offer funds explicitly based on these two factors.

Figure 9 panels (a) and (b) display the distribution of beliefs about the size and value factors over the period 2009–2019. Overall, expected returns of the size and value factors are lower compared to the beliefs about market, which is consistent with lower realized returns of size and value during our sample period. Panel (a) illustrates that beliefs about the expected returns of the size factor are relatively stable, with the mean expected return averaging roughly two percent over the sample period. Similarly, the average belief about the value factor is approximately 2 percentage points (Figure 9, panel (b)).

As with beliefs about market returns, investors became more optimistic about the expected returns of value stocks coming out of the Great Recession, although beliefs declined slightly thereafter. The time series of beliefs about the value factor appears to track the realized returns of the value factor, which performed poorly during the financial crisis, recovered after the crisis, and again performed poorly toward the end of the sample period. This pattern suggests that beliefs about value tend to be extrapolative, which is consistent with prior literature on belief formation and with our findings in Section 6. Although investors appear to be extrapolative, investors still believe in a modestly positive value premium, despite negative realized return of the value factor during our sample period. Our results also indicate that similar to market return beliefs, beliefs about other factors are heterogeneous. The interdecile range of beliefs

about the expected returns of the size and value factors is approximately 2 and 3 percentage points, respectively. In Appendix Tables A12 and A13, we explore how expected returns about factors vary with investor demographics. We find that older and more educated investors appear more optimistic about the returns of both the size and value factors.

4.3 Assessing Beliefs

Our framework provides a new methodology for recovering investors' expected returns and risk aversion in the absence of survey data. While our estimates of risk aversion are broadly consistent with those in the literature, a natural question arises: How do our estimates of expectations align with other measures? Furthermore, even if the model-implied estimates of beliefs are reasonable, how well does our model explain the actual investment decisions of investors?

To assess the validity of our model estimates, we proceed in two steps. First, we evaluate how well the model predicts investors' portfolios out-of-sample. Second, we compare our estimates of beliefs with analyst expectations, focusing on both average expected returns and the dispersion in expected returns across investors.

4.3.1 Out-of Sample Predictions

We evaluate our model using a simple out-of-sample prediction exercise. The data structure we use for estimation naturally lends itself to such an exercise. When estimating our model (eq. (4)), we restrict our observations to plan-year-investment option pairs corresponding to mutual funds, as opposed to other investment vehicles such as collective investment trusts (CITs) and separate accounts. This is because we only observe panel data on expense ratios for mutual funds and are able to construct valid instruments for them. We refer to these mutual fund observations as "in-sample" and the remaining investment options as "out-of-sample." It is worth noting that when constructing the dependent variable in our estimation of equation (4), we use holdings of these out-of-sample investment options to measure asset covariances. However, these out-of-sample options are not directly used in the estimation of equation (4).

For the out-of-sample prediction exercise, we examine how well our estimated beliefs and risk aversion explain holdings of the out-of-sample investment options, which are excluded from our estimation. To form predictions for each out-of-sample observation, we use the estimated risk aversion for each plan. For beliefs, we rely on the factor structure of these out-of-sample options and the factor beliefs estimated using in-sample mutual funds. Specifically, following our analysis in Section 4.2.2, we measure beliefs about out-of-sample investment option k for investors in plan m at time t as:

$$\mu_{mkt} = \sum_{l=1}^L \delta_{mlt} b_{lkt}.$$

Given our estimates of beliefs and preferences, we estimate a simple regression to assess the in-sample and out-of-sample predictive power of the primitives we recover. Specifically, we estimate the regression:

$$\ln \omega_{mk} - \ln \omega_{m,cash} = \beta^0 + \beta^\mu \mu_{mk} + \beta^\lambda \lambda_m + \beta^\sigma \sigma_k^2 + \epsilon_{mk} \quad (8)$$

Observations are at the plan m by investment option k level, as of 2019. We focus on the 2019 cross-section because it is the only year for which we observe fee data for the out-of-sample investment options, allowing us to compute net-of-fee beliefs. The dependent variable is the log difference between allocation to investment option k and cash. We control for plan-investment-option-specific expected returns (based on factors), plan-specific risk aversion, and the variance of investment option k .

We report the corresponding estimates in Table 4. We find that our factor-based beliefs and risk aversion measure explain substantial variation in portfolio weights. At the bottom of the table, we report the in-sample and out-of-sample R^2 values. Our estimates exhibit substantial predictive power in both contexts. The regression yields an in-sample R^2 of 23% and an out-of-sample R^2 of 17%. We use this simple reduced-form regression to illustrate that our model has roughly similar predictive power out of sample. Moreover, we show in Appendix F that if we further leverage the mean-variance structure of the decision problem, in conjunction with our preference and belief estimates, we can generate an out-of-sample R^2 that is 25 percentage points higher than that obtained using naive estimates.

4.3.2 External Validity: Comparison with Analyst Expected Returns

We compare our estimates of investor beliefs with analyst expectations. For each 401(k) plan, we recover the average participant's belief for each mutual fund in their 401(k) menu. While survey data on individual mutual fund expected returns are unavailable, we can calculate expected mutual fund returns using data from equity analysts. Equity analysts provide price targets, and subsequently, implicit expected return expectations for individual stocks. Using data on analyst expected returns and security-level mutual fund holdings for equity funds, we can calculate expected mutual fund returns based on return expectations for the underlying securities. This allows us to calculate both the average expected return and the dispersion in expected returns for each equity mutual fund.²⁸

Figure 10 panel (a) displays a binned scatter plot of the average analyst's expected mutual fund returns versus the average 401(k) investor's expected mutual fund returns across equity mutual funds. Observations are at the mutual fund-by-year level. There is a strong positive and statistically significant relationship between equity analysts' expected mutual fund returns and

²⁸Other papers such as Han et al. (2022) have used analyst price targets to calculate expected returns. See Appendix E.4 for further details on how we construct analyst implied expected mutual fund returns.

401(k) investors' expected mutual fund returns. We also find evidence that the second moment of the distribution of expectations is positively and significantly correlated across equity analysts and 401(k) investors. Panel (b) displays a binned scatter plot of the dispersion in expected mutual fund returns across analysts versus the dispersion in expected mutual fund returns across 401(k) investors. The results suggest that belief disagreement among 401(k) investors is correlated with belief disagreement among equity analysts. Though there are substantial differences between the two populations, we would expect to find positive correlations if there were fundamental properties of beliefs that were shared across investors. Overall, we view the results displayed in Figure 10 as providing additional credibility to our estimates of beliefs.

4.4 Alternative Specifications and Robustness

We consider several alternative specifications to assess the robustness of the estimated parameters. First, we re-estimate the model to include target date funds, which are excluded from our baseline analysis. By excluding target date funds in our baseline analysis, we are implicitly assuming that target date investors only invest in target date funds. When we include target date funds, we are assuming that target date fund investors also hold non-target date funds. Second, to account for potential inertia in investor behavior, we estimate the model using only new plans. For all of these specifications, we find very similar estimates of risk aversion and expected returns. The mean risk aversion we estimate in these alternative models is nearly identical to our baseline estimate (3.55 and 3.56 vs. 3.55). The mean expected return ranges from 9.7 to 9.9, similar to our baseline estimate of 9.6. In Appendix E.1, we present results for these and other robustness checks, including the alternative measures of risk described in Section 3.3.

5 Implications of Preference and Belief Heterogeneity

Our results above indicate substantial heterogeneity in beliefs and risk aversion across investors. We explore the corresponding implications of this heterogeneity for investors' portfolios. We start by quantifying how both risk aversion and beliefs impact investors' portfolios by examining the reduced-form correlations between investor portfolios and preferences and beliefs. We then calculate counterfactual portfolios while imposing that investors have the same beliefs and preferences. Although both beliefs and preferences matter, our reduced-form and counterfactual results suggest that heterogeneity in beliefs explains more of the variation in portfolio holdings than risk preferences.

We then explore, to the extent that this belief heterogeneity results from incomplete information and biases, the potential costs of such distortions. These belief distortions can cause investors to hold inefficient and suboptimal portfolios. This cost is modest for most investors, yet it can be significant for those with more extreme beliefs.

5.1 What Explains Holdings? Beliefs vs. Risk Aversion

We examine how dispersion in beliefs and risk aversion explain variation in equity exposure across plans in the following regression:

$$Equity Share_{mt} = \beta \lambda_{mt} + \gamma \delta_{mt} + \epsilon_{mt}. \quad (9)$$

Observations are at the plan-by-year level. The dependent variable $Equity Share_{mt}$ measures the share of assets in plan m that are invested in US equities. The dependent variables λ_{mt} and δ_{mt} measure the risk aversion and average market expectations of investors in plan m at time t .

Table 5 displays how dispersion in risk aversion and expectations explain 401(k) portfolios. The dependent variable in the regression specification displayed in columns (1) and (2) is the share of the portfolio held in equities (US and international equities), the dependent variable in columns (3) and (4) is the share held in US equities, and the dependent variable in column (5) and (6) is the share held in cash. To aid interpretation we also normalize risk aversion and investor beliefs such that each is mean zero and has a variance equal to one.

The results are intuitive and suggest that variation in beliefs and risk aversion both play important roles in explaining investor equity and cash allocations. The results in column (2) indicate that a one standard deviation increase in expected returns is correlated with a 13.7 pp (20% relative to the mean allocation) increase in an investor's equity allocation and a one standard deviation increase in risk aversion is correlated with a 7.4 pp (11% relative to the mean allocation) decrease in an investor's equity allocation. Conversely, an increase in expected market returns is negatively correlated with an investor's cash allocation, while an increase in risk aversion is positively correlated with her cash allocation. The results also indicate that our simple two-parameter model explains a fair amount of the variation in equity and portfolio holdings. Variation in beliefs and risk aversion explain 51% of the reduced-form variation in equity exposure.

Despite utilizing a different methodology, our point estimates of the relationship between an investor's equity share and expected market returns closely align with those reported by Giglio et al. (2021). Giglio et al. (2021) employ novel survey and account-level data from Vanguard to examine the relationship between investor beliefs and their portfolio decisions. Specifically, when analyzing tax-advantaged retail accounts, which closely resemble our sample, Giglio et al. (2021) estimate that a 1 percentage point (pp) increase in beliefs about stock market returns is correlated with a 1.34-3.55pp increase in equity share, depending on the investor's characteristics (see column 7 of Table 4 in Giglio et al. (2021)). Our baseline estimates in column (1) of Table 5 suggest that a 1pp increase in beliefs about the stock market returns is correlated with a 3.68pp increase in equity share. Part of the reason we estimate a stronger relationship is because we simultaneously control for risk aversion. In untabulated results, we find that a 1pp increase in beliefs about stock market returns is correlated with a 2.77pp increase in equity

share if we do not simultaneously control for risk aversion, as in Giglio et al. (2021). This comparison with Giglio et al. (2021) lends additional credence to our methodology.

These findings, in conjunction with our findings from Section 4.1 and 4.2 provide a useful lens for understanding why portfolio allocations vary across investors. For example, older investors have lower equity exposure because they are both more risk averse and more pessimistic. Individuals with more education allocate more towards equity because they have optimistic beliefs despite being more risk averse. Beliefs rather than risk aversion explain why equity allocation varies across ethnicities.

As an alternative way to illustrate the relative importance of heterogeneity in beliefs and risk aversion, we simulate allocations under counterfactual environments in which investors have identical beliefs, identical risk aversion, or both. For these counterfactuals, we use our method to calculate a single “average” expected return for each fund and an average risk aversion parameter, separately by year. We then calculate the optimal portfolio such that equation (2) is satisfied when replacing our estimated beliefs/risk aversion with the average values.

For the risk aversion parameter, we use the mean estimated value across plans, weighted by total plan assets. For expected returns, we aggregate fund balances across all plans and calculate the implied beliefs for each fund that would rationalize this aggregate portfolio under the average risk aversion parameter.²⁹ For the purposes of these counterfactuals, we only focus on plans with more than three investment options for one year (2016).

Figure 11 plots the densities of equity allocations across plans. The dashed line indicates the distribution of assets held in U.S. equity funds in our data. The solid line indicates the counterfactual distribution when removing heterogeneity in beliefs, i.e., assigning all investors identical fund-specific expected returns. The dotted line indicates the counterfactual distribution when assigning all investors identical values for risk aversion. Finally, the counterfactual distribution where we assign investors identical beliefs and risk aversion is given by the gray shaded area. To show the different counterfactuals on a more reasonable scale, the top of the density is visually cropped in the figure.

These counterfactual allocations indicate the importance of heterogeneity in beliefs and risk aversion in investor portfolio choice. Assigning investors identical beliefs greatly reduces the variation in equity allocations across plans. By comparison, assigning all investors the same risk aversion slightly increases the variation in equity allocations across plans. In this sample, the across-plan standard deviation in equity allocations is 0.132. With identical beliefs, the standard deviation falls to 0.072, but with average risk aversion, it increases to 0.141. Removing heterogeneity in beliefs and risk aversion together further reduces variation across plans, lowering the standard deviation in equity allocations to 0.043. The residual variation when investors have identical beliefs and risk aversion is due to differences in menus across plans. Our

²⁹ Alternatively we could calculate beliefs using the average estimated belief (across investors) for each asset using our estimates from Section 4.2. The correlation between this measure of implied beliefs and the average estimated belief (across investors) in 2016 is 0.91.

estimates indicate that both heterogeneity in beliefs and risk aversion are important; however, these simulations suggest that variation in beliefs plays a bigger role in driving variation across plans.³⁰

In Appendix Figure A2, we also show counterfactual allocations to cash funds. Assigning investors identical beliefs leads to a modest decrease in across-plan standard deviation but substantially reduces the right tails. Some investors can be overly pessimistic and allocate a significant portion of their assets to cash. Assigning all investors the same risk aversion also slightly increases the variation in cash allocation across plans. Removing both sources of heterogeneity reduces the standard deviation of cash holdings from 0.13 to 0.11. Again, the residual variation with identical beliefs and risk aversion is due to plan menus.³¹

5.2 Welfare Costs of Belief Distortions

To the extent that heterogeneity in beliefs stems from incomplete information and investor biases, it could prove costly for investors, particularly if it leads to distorted portfolio decisions. We use our framework to calculate the potential costs of such distortions. As a benchmark, we consider a counterfactual scenario in which all investors share the same unbiased beliefs about returns, denoted as $\bar{\mu}_t$. Using these unbiased beliefs, we compute the optimal portfolio for each plan m , denoted $\bar{\omega}_{mt}$, and compare the optimal portfolio weights and associated welfare against the actual portfolio weights observed in the data, ω_{mt} . Based on the mean-variance utility model, we calculate the change in consumer surplus as:

$$\Delta Consumer Surplus_{mt} = (\bar{\omega}_{mt} - \omega_{mt})'(\bar{\mu}_t - \mathbf{p}_t - R_{Ft}) - \frac{\lambda_m}{2} (\bar{\omega}'_{mt} \Sigma_t \bar{\omega}_{mt} - \omega'_{mt} \Sigma_t \omega_{mt}). \quad (10)$$

This change in consumer surplus quantifies the ex ante cost of harboring heterogeneous and biased beliefs in terms of annual forgone risk-adjusted returns.

We must take a stance on the appropriate beliefs to assign to $\bar{\mu}_t$. We use the average beliefs as described above in Section 5.1, which allows us to assess the welfare costs of heterogeneity independently from the potential costs arising from inaccurate average beliefs. Though average beliefs may be biased, employing them as a benchmark is a logical starting point for assessing the costs of heterogeneity.

Figure 12 illustrates the distribution of $\Delta Consumer Surplus_{mt}$ at the plan level for the year 2016. Ex ante, incorrect and heterogeneous beliefs reduce the average investor's risk-adjusted return by 0.72 percentage points per annum; however, this distribution is skewed. The cost of biased beliefs for the median investor is relatively modest, lowering the median investor's expected risk-adjusted returns by 0.17 percentage points per annum, but it is significantly higher

³⁰This relates to the finding in Kojien and Yogo (2019a) that latent demand plays an important role in asset demand.

³¹One reason plan menus play a larger role in the variation in cash holdings is that some plans offer low-risk bond funds that are close substitutes for cash funds, whereas other funds do not.

for investors with extreme beliefs. In the Appendix, we further explore the heterogeneity. Appendix Table A6 shows that costs of belief distortions are lower for wealthier and more educated investors, while they are higher for retirees and those who work in riskier sectors. These results suggest that the impacts of divergent beliefs are greater among the latter two groups.

Despite investors suffering from sizable belief distortions and these distortions having large impacts on investors' portfolio allocations (e.g., Figure 11), our results suggest that most investors construct portfolios that are relatively efficient in terms of maximizing the portfolio Sharpe ratio. For the average plan, the annual Sharpe ratio of the optimal counterfactual portfolio, $\bar{\omega}_{mt}$, is only 0.01 higher than that of the actual portfolio, ω_{mt} . The relatively modest average cost of belief distortions reflects the constrained choice architecture of 401(k) plans. These menus typically consist of professionally managed and well diversified funds that are reasonably close to the efficient frontier. As a result, even if investors have distorted beliefs, the scope for making costly mistakes is mitigated. By contrast, if individuals were selecting from individual stocks or bonds, the potential for loss would be far greater. In Appendix G, we demonstrate that the welfare costs of belief heterogeneity when investors allocate across individual stocks could be an order of magnitude larger, with a median (mean) of 4.7% (7.5%).

6 Evidence on the Formation of Beliefs

Investor beliefs are heterogeneous and play a critical role in determining investor portfolios, with potentially important welfare implications. Here, we provide insight into how beliefs are formed across investors. A large previous literature documents that investors extrapolate beliefs from past returns and experiences. Our unique setting provides additional insight into the types of information that investors use to form extrapolative beliefs.

As one might expect, investors respond to contemporaneous, commonly available information. We show in Appendix B.1 that investor beliefs are positively related with fund-specific past returns, even when those funds are new to a particular plan.³²

We present new evidence that investors rely on different information sets when forming expectations. In particular, we find that personal experience shapes beliefs above and beyond aggregate market information. Specifically, we examine how expected market returns are influenced by local economic conditions and the past performance of an investor's employer, while controlling for time-varying macroeconomic conditions using time fixed effects. Both local conditions and employer performance are positively correlated with expectations, suggesting that investors extrapolate broader market returns from idiosyncratic, individualized experience.

We also find that an investor's sector of employment shapes not only their beliefs about the overall market, but also their sector-specific return expectations. Investors tend to be more

³²We look at newly added plans to rule out the possibility that extrapolation reflects investors' prior-year holdings in their 401(k) plans.

optimistic about sectors they work in, especially in industries like real estate, tech services, and finance, indicating that sector affiliation serves as another lens through which individualized experience shapes expectations.

Our findings, which document systematic and predictable drivers of heterogeneity of beliefs, suggest that a standard rational expectations model may not capture the investment behavior across households. Investor beliefs are correlated with observable characteristics such as wealth and income, and appear to depend on past market returns as well as recent employer performance. Such findings suggest a violation of rational expectations; consistent with this, we show in Appendix B.2 that forecast errors are predictable.

6.1 Extrapolation from Local Economic Conditions

We examine the role of experience in shaping extrapolation *above and beyond* what is available in aggregate information. One source of individualized experience comes from local economic conditions. We examine the relationship between investors' beliefs and local economic conditions in the following regression:

$$\delta_{mt} = \text{Local Economic Conditions}'_{mt} \Gamma + \phi_t + \phi_m + \varepsilon_{mt}. \quad (11)$$

Observations are at the plan-by-year level. The dependent variable δ_{mt} measures the expected market return averaged across investors in plan m at time t . The term $\text{Local Economic Conditions}_{mt}$ is a vector of county-by-year level measures of economic conditions including: GDP growth, business establishment growth, annual home price growth, and population growth.³³ We also control for year (ϕ_t) and plan (ϕ_m) fixed effects. Thus we measure how, conditional on aggregate macroeconomic conditions, changes in local economic conditions are correlated with changes in investors' beliefs.

We report the estimates corresponding to eq. (11) in Table 6. In each specification, we find a strong positive relationship between local economic conditions and investors' beliefs about stock market returns. The results in column (1) indicate that a 1% increase in county population is correlated with a 0.13pp increase in expected returns. Similarly, the results in column (3) indicate that a 10% increase in county home prices is associated with a 0.22pp increase in expected returns. We find a positive relationship between each measure of local economic activity and market expectations, even when we use within plan variation. In Appendix Table A14, we show that these effects translate to equity holdings as well. Overall, these results suggest that idiosyncratic experiences may drive differences in expected returns across investors,

³³We measure home price growth using data from the FHFA, GDP growth from the Bureau of Economic Analysis, establishment growth from the County Business Patterns, and population growth from the Census. We determine the location based on the employer's headquarters that is reported in BrightScope. While the median firm in our sample has only 223 401(k) participants, measuring local economic conditions based on the employer's headquarters may introduce measurement error for larger firms with establishments in multiple counties.

potentially through how they shape forecasts of future returns.

6.2 Extrapolation from Employer Performance

An advantage in our setting is that we observe details on the investors' employers, the 401(k) plan sponsors. This allows us to explore how investors' beliefs depend on their employment. Using the sponsor's EIN, we link our BrightScope 401(k) data with balance sheet, income statement, and market return data from Compustat and CRSP.

We examine the relationship between the financial performance of an investor's employer and the investor's beliefs in the following regression:

$$\delta_{mt} = \varphi Performance_{mt} + \phi_t + \phi_m + \eta_{mt}. \quad (12)$$

Observations are at the plan-by-year level where we restrict the data set to those plans where the sponsor is publicly traded. The dependent variable δ_{mt} measures the expected market return averaged across investors in plan m at time t . The independent variable $Performance_{mt}$ measures the financial performance of plan sponsor m at time t . We measure firm performance in terms of last year's annual stock market return, sales growth, investment, and employment growth.

We report the estimates corresponding to eq. (12) in Table 7. Consistent with our previous results, we find that beliefs are highly correlated with local conditions. In each specification, we document a positive and significant relationship between sponsor performance and participants' expectations about the market. The results are robust to the inclusion of both plan fixed effects (odd columns) and industry-by-year fixed effects (even columns). Including industry-by-year fixed effects allows us to effectively compare the beliefs of two investors working in the same industry at the same time but for different firms. In columns (1) and (2), we find that investors become more optimistic about the market following strong performance of their employer. The effect is marginally stronger when we include industry-by-year fixed effects, which suggests that investors are more sensitive to industry or risk-adjusted returns than absolute returns. The results in column (4) indicate that investors become 0.18 pp more optimistic about the expected return of the market following a 10% increase in investment. Similarly, we estimate that a one standard deviation increase in sales growth (24%) is associated with a 0.10 pp increase in the expected return of the market (column 6). In Appendix Table A15, we document that we find a similar positive relationship between equity holdings and employer performance. Overall, this suggests that investors may misattribute the performance of their employer to the performance of the economy more generally, or that they use this more local experience to form an idiosyncratic forecast of future market returns.

Extrapolations based on local economic conditions and employer performance point to the importance of idiosyncratic experiences in belief formation (e.g., Malmendier and Nagel, 2011,

2015; Kuchler and Zafar, 2019) and indicate that even potentially irrelevant information helps shape beliefs that have real stakes (Bordalo et al., 2022). These findings demonstrate the value of our empirical methodology in understanding beliefs and risk aversion using large-scale data on portfolio allocations.

6.3 Sector-Specific Beliefs

Finally, we explore whether investors' employment sectors shape not only their beliefs about the overall market but also their expectations for sector-specific returns. Specifically, we ask: do investors believe their own sector will outperform?

We examine how an investor's expected returns for a fund vary with the fund's sectoral exposure and the investor's employment sector. To measure a fund's sectoral exposure, denoted ϖ_{kst} , we decompose its holdings at the NAICS 2-digit sector s -level by computing the fund's (k) share allocated to stocks in sector s . We study investor beliefs about the expected returns across different sectors and how beliefs vary with whether the investor is employed in that sector in the following regression:

$$\bar{\mu}_{mkt} - R_F = \sum_{s \in \mathcal{S}} \varpi_{s k t} (\beta_s + 1\{s = s(m)\} \beta_s^{own}) + \phi_{mt} + \phi_{j(k)t} + \epsilon_{mkt}. \quad (13)$$

The dependent variable is the expected return of fund k for investors in plan m at time t . The coefficients β_s capture sector-wide return expectations, and β_s^{own} captures how those expectations differ for investors employed in sector s . ϕ_{mt} and $\phi_{j(k)t}$ are plan-by-year and fund type-by-year fixed effects. This specification allows sector-specific optimism to vary by sector: for example, real estate employees may be particularly optimistic about real estate, while manufacturing workers may be more skeptical about their own sector's prospects.

Figure 13 panels (a) and (b) display the corresponding regression coefficients. All sectoral exposures ϖ_{kst} are standardized. Panel (a) displays how the sector exposure contributes to investor return expectations, β_s , ignoring the own sector interaction term. Investors are most optimistic about the Information sector, which likely reflects the boom in tech stocks during this sample period. A one standard deviation (8 percent exposure) increase in information sector exposure leads to 40 bps higher expected return. Investors are also optimistic about construction, manufacturing, and finance, while pessimistic about utilities.

Panel (b) shows how sector beliefs shift when the investor works in the sector (β_s^{own}). On average, investors are modestly more optimistic about their own sector. A one standard deviation increase in exposure to their employment sector is associated with a 1.5 basis point increase in expected returns. However, the pattern varies across sectors. For example, real estate employees are more optimistic about their own sector than others. Workers in science/technology services and finance/insurance also display positive own-sector bias. In contrast, construction workers are less optimistic about their own sector than others.

Overall, these results underscore the importance of local and sector-specific experience in shaping beliefs, contributing to a growing understanding of how personal circumstances influence return expectations.

7 Conclusion

We examine how households allocate their 401(k) portfolios. Allocations vary dramatically across plans and vary in systematic ways with participant and employer characteristics. For example, plans with more educated participants tend to hold more of their portfolio in equities and less in cash. In contrast, the investment options available to plan participants do not vary systematically with participant characteristics.

To understand these patterns, we develop a framework for estimating investor beliefs and risk aversion. By measuring how investors re-optimize their portfolios in response to exogenous changes in investment fees, we are able to separately identify risk aversion from beliefs. Studying 401(k) plan allocations, where investors choose from a preset menu of investment options with variations in expense ratios, offers an ideal setting for our approach.

We find that there is substantial heterogeneity in both risk aversion and beliefs across investors. The differences in expectations and risk aversion are correlated with observable investor characteristics and help explain the heterogeneity in asset allocation across plans. For example, our results suggest that differences in beliefs, rather than risk aversion, help explain why educated investors tend to hold more equities and less cash. Counterfactual simulations suggest that heterogeneity in beliefs drives the majority of variation in equity allocations. The potential welfare costs of this belief heterogeneity appear modest for the median investor but can be quite large for investors with extreme beliefs.

An important feature of our model is that we do not impose any restrictions on the rationality of beliefs. In fact, we find that investor beliefs violate rational expectations. Investors extrapolate their beliefs from both past fund returns and from individualized experience based on local economic conditions and employer performance.

Our results also highlight the importance of accounting for and understanding heterogeneity in both beliefs and risk aversion. We show that both sources of heterogeneity play important roles in explaining equity participation rates across investors, and they could potentially have important implications for asset prices and other macroeconomic phenomena. Our framework can also be applied in other settings to provide insight about investor beliefs and risk aversion, which could be particularly valuable when survey data is unavailable.

References

Aguirregabiria, V., R. Clark, and H. Wang (2019). The geographic flow of bank funding and access to credit: Branch networks, local synergies, and competition. *Working Paper*.

Allen, J., R. Clark, J.-F. Houde, S. Li, and A. V. Trubnikova (2023). The role of intermediaries in selection markets: Evidence from mortgage lending. *Working Paper*.

Ameriks, J. and S. P. Zeldes (2004). How do household portfolio shares vary with age. *Working paper*.

Amromin, G. and S. A. Sharpe (2014). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science* 60(4), 845–866.

Bach, L., Laurent E. Calvet, and Paolo Sodini (2020). Rich pickings? Risk, return, and skill in household wealth. *American Economic Review* 110(9), 46.

Badarinza, C., J. Y. Campbell, and T. Ramadorai (2016). International comparative household finance. *Annual Review of Economics* 8, 111–144.

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review* 103(6), 2499–2529.

Bekaert, G., K. Hoyem, W.-Y. Hu, and E. Ravina (2017). Who is internationally diversified? evidence from the 401 (k) plans of 296 firms. *Journal of Financial Economics* 124(1), 86–112.

Ben-David, I., J. R. Graham, and C. R. Harvey (2013). Managerial miscalibration. *The Quarterly Journal of Economics* 128(4), 1547–1584.

Benartzi, S. (2001). Excessive extrapolation and the allocation of 401 (k) accounts to company stock. *The Journal of Finance* 56(5), 1747–1764.

Benartzi, S. and R. Thaler (2007). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives* 21(3), 81–104.

Benartzi, S. and R. H. Thaler (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review* 91(1), 79–98.

Benetton, M. (2021). Leverage regulation and market structure: A structural model of the uk mortgage market. *The Journal of Finance* 76(6), 2997–3053.

Benetton, M. and G. Compiani (2021). Investors beliefs and cryptocurrency prices. *Working paper*.

Benetton, M., A. Gavazza, and P. Surico (2021). Mortgage pricing and monetary policy. *Working Paper*.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2009). 5. *The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States*. University of Chicago Press.

Beutel, J. and M. Weber (2022). Beliefs and portfolios: Causal evidence. *Working Paper*.

Bhattacharya, V. and G. Illanes (2021). The design of defined contribution plans. *Working paper*.

Black, S. E., P. J. Devereux, P. Lundborg, and K. Majlesi (2018). Learning to take risks? the effect of education on risk-taking in financial markets. *Review of Finance* 22(3), 951–975.

Bordalo, P., G. Burro, K. Coffman, N. Gennaioli, and A. Shleifer (2022). Imagining the future: Memory, simulation and beliefs about covid. *Working Paper*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Over-reaction in macroeconomic expectations. *Working paper*.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma (2020). Institutional corporate bond pricing. *Working paper*.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. *Working paper*.

Calvet, L. E., J. Y. Campbell, F. J. Gomes, and P. Sodini (2019). The cross-section of household preferences. *Working paper*.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2007). Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy* 115(5), 707–747.

Campbell, J. Y. (2006). Household finance. *The journal of finance* 61(4), 1553–1604.

Campbell, J. Y. and L. Viceira (2001). Who should buy long-term bonds? *American Economic Review* 91(1), 99–127.

Campbell, J. Y. and L. M. Viceira (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics* 114(2), 433–495.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The quarterly journal of economics* 124(4), 1639–1674.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012). What have they been thinking? Homebuyer behavior in hot and cold markets. *Brookings Papers on Economic Activity*, 265.

Catherine, S. (2021). Countercyclical labor income risk and portfolio choices over the life cycle. *The Review of Financial Studies* 35(9), 4016–4054.

Catherine, S., P. Sodini, and Y. Zhang (2022). Countercyclical income risk and portfolio choices: Evidence from sweden. *Swedish House of Finance Research Paper* (20-20).

Chiteji, N. S. and F. P. Stafford (2000). Asset ownership across generations. *Working paper*.

Choi, J. J. (2015). Contributions to defined contribution pension plans. *Annual Review of Financial Economics* 7, 161–178.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax Policy and the Economy* 16, 67–113.

Choi, J. J., D. Laibson, B. C. Madrian, A. Metrick, and J. M. Poterba (2007). 2. *For Better or for Worse: Default Effects and 401(k) Savings Behavior*, pp. 81–126. University of Chicago Press.

Choi, J. J. and A. Z. Robertson (2020). What matters to individual investors? evidence from the horse's mouth. *The Journal of Finance* 75(4), 1965–2020.

Choukhmane, T. (2019). Default options and retirement saving dynamics. *Working Paper*.

Choukhmane, T. and T. de Sliva (2022). What drives investors portfolio choices? separating risk preferences from frictions. *Working paper*.

Cocco, J. F. (2005). Portfolio choice in the presence of housing. *The Review of Financial Studies* 18(2), 535–567.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies* 18(2), 491–533.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105(8), 2644–78.

Darmouni, O., K. Siani, and K. Xiao (2022). Nonbank fragility in credit markets: Evidence from a two-layer asset demand system. *Working Paper*.

Das, S., C. M. Kuhnen, and S. Nagel (2020). Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies* 33(1), 395–432.

Di Maggio, M., M. Egan, and F. Franzoni (2021). The value of intermediation in the stock market. *Journal of Financial Economics*.

Dick, A. A. (2008). Demand estimation and consumer welfare in the banking industry. *Journal of Banking & Finance* 32(8), 1661–1676.

Du, W., S. Gadgil, M. B. Gordy, and C. Vega (2019). Counterparty risk and counterparty choice in the credit default swap market. *Working paper*.

Duflo, E., W. Gale, J. Liebman, P. Orszag, and E. Saez (2006). Saving incentives for low-and middle-income families: Evidence from a field experiment with H&R Block. *The Quarterly Journal of Economics* 121(4), 1311–1346.

Dworak-Fisher, K. (2011). Matching matters in 401 (k) plan participation. *Industrial Relations: A Journal of Economy and Society* 50(4), 713–737.

Egan, M. (2019). Brokers vs. retail investors: Conflicting interests and dominated products. *The Journal of Finance* 74(3), 1217–1260.

Egan, M., A. Hortaçsu, and G. Matvos (2017). Deposit competition and financial fragility: Evidence from the us banking sector. *American Economic Review* 107(1), 169–216.

Egan, M., S. Lewellen, and A. Sunderam (2022). The cross-section of bank value. *The Review of Financial Studies* 35(5), 2101–2143.

Egan, M. L., A. MacKay, and H. Yang (2022). Recovering investor expectations from demand for index funds. *The Review of Economic Studies* 89(5), 2559–2599.

Fagereng, A., C. Gottlieb, and L. Guiso (2017). Asset market participation and portfolio choice over the life-cycle. *The Journal of Finance* 72(2), 705–750.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). Heterogeneity and persistence in returns to wealth. *Econometrica* 88(1), 115–170.

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review* 103(5), 1598–1628.

Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations* 1, Volume 2, pp. 261–343. Elsevier.

Gennaioli, N., Y. Ma, and A. Shleifer (2016). Expectations and investment. *NBER Macroeconomics Annual* 30(1), 379–431.

Gennaioli, N., A. Shleifer, and R. Vishny (2015). Money doctors. *The Journal of Finance* 70(1), 91–114.

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

Gomes, F., M. Haliassos, and T. Ramadorai (2021). Household finance. *Journal of Economic Literature* 59(3), 919–1000.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *The Review of Financial Studies* 27(3), 714–746.

Guiso, L., T. Jappelli, and D. Terlizzese (1996). Income risk, borrowing constraints, and portfolio choice. *The American Economic Review*, 158–172.

Guiso, L., P. Sapienza, and L. Zingales (2008). Trusting the stock market. *The Journal of Finance* 63(6), 2557–2600.

Haddad, V., P. Huebner, and E. Loualiche (2021). How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing. *Working paper*.

Han, Y., D. Huang, D. Huang, and G. Zhou (2022). Expected return, volume, and mispricing. *Journal of Financial Economics* 143(3), 1295–1315.

Hastings, J., A. Hortaçsu, and C. Syverson (2017). Sales force and competition in financial product markets: the case of mexico's social security privatization. *Econometrica* 85(6), 1723–1761.

Heaton, J. and D. Lucas (2000). Portfolio choice and asset prices: The importance of entrepreneurial risk. *The Journal of Finance* 55(3), 1163–1198.

Janssen, A. and J. Thiel (2025). Do search costs explain persistent investment in active mutual funds? *Journal of Economic Dynamics and Control* 176, 105099.

Koijen, R. S., R. J. Richmond, and M. Yogo (2019). Which investors matter for global equity valuations and expected returns? *Working paper*.

Koijen, R. S. and M. Yogo (2016). Shadow insurance. *Econometrica* 84(3), 1265–1287.

Koijen, R. S. and M. Yogo (2019a). A demand system approach to asset pricing. *Journal of Political Economy* 127(4), 1475–1515.

Koijen, R. S. and M. Yogo (2019b). Exchange rates and asset prices in a global demand system. *Working paper*.

Koijen, R. S. and M. Yogo (2022). The fragility of market risk insurance. *The Journal of Finance* 77(2), 815–862.

Kronlund, M., V. K. Pool, C. Sialm, and I. Stefanescu (2021). Out of sight no more? the effect of fee disclosures on 401 (k) investment allocations. *Journal of Financial Economics* 141(2), 644–668.

Kuchler, T. and B. Zafar (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance* 74(5), 2491–2542.

Kuhnen, C. M. and A. C. Miu (2017). Socioeconomic status and learning from financial information. *Journal of Financial Economics* 124(2), 349–372.

Lynch, A. W. and S. Tan (2011). Labor income dynamics at business-cycle frequencies: Implications for portfolio choice. *Journal of Financial Economics* 101(2), 333–359.

Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics* 116(4), 1149–1187.

Malmendier, U. and S. Nagel (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126(1), 373–416.

Malmendier, U. and S. Nagel (2015). Learning from inflation experiences. *The Quarterly Journal of Economics* 131(1), 53–87.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance* 7(1), 77–91.

Meeuwis, M. (2019). Wealth fluctuations and risk preferences: Evidence from U.S. investor portfolios. *Working Paper*.

Nagel, S. and Z. Xu (2019). Asset pricing with fading memory. *Working paper*.

Pool, V. K., C. Sialm, and I. Stefanescu (2016). It pays to set the menu: Mutual fund investment options in 401 (k) plans. *The Journal of Finance* 71(4), 1779–1812.

Pool, V. K., C. Sialm, and I. Stefanescu (2020). Mutual fund revenue sharing in 401 (k) plans. *Working paper*.

Robles-Garcia, C. (2019). Competition and incentives in mortgage markets: The role of brokers. *Working paper*.

Ross, S. (2015). The recovery theorem. *The Journal of Finance* 70(2), 615–648.

Shumway, T., M. Szeffler, and K. Yuan (2009). The information content of revealed beliefs in portfolio holdings. *Working paper*.

Sinai, T. and N. S. Souleles (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics* 120(2), 763–789.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy* 112(3), 695–717.

Van Rooij, M., A. Lusardi, and R. Alessie (2011). Financial literacy and stock market participation. *Journal of Financial Economics* 101(2), 449–472.

Vissing-Jorgensen, A. (2003). Perspectives on behavioral finance: Does “irrationality” disappear with wealth? evidence from expectations and actions. *NBER Macroeconomics Annual* 18, 139–194.

Wachter, J. A. and M. Yogo (2010). Why do household portfolio shares rise in wealth? *The Review of Financial Studies* 23(11), 3929–3965.

Waldfogel, J. (2003). Preference externalities: An empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics* 34(3), 557–568.

Wang, Y., T. M. Whited, Y. Wu, and K. Xiao (2022). Bank market power and monetary policy transmission: Evidence from a structural estimation. *The Journal of Finance* 77(4), 2093–2141.

Whited, T. M., Y. Wu, and K. Xiao (2022). Will central bank digital currency disintermediate banks? *Working Paper*.

Xiao, K. (2020). Monetary transmission through shadow banks. *The Review of Financial Studies* 33(6), 2379–2420.

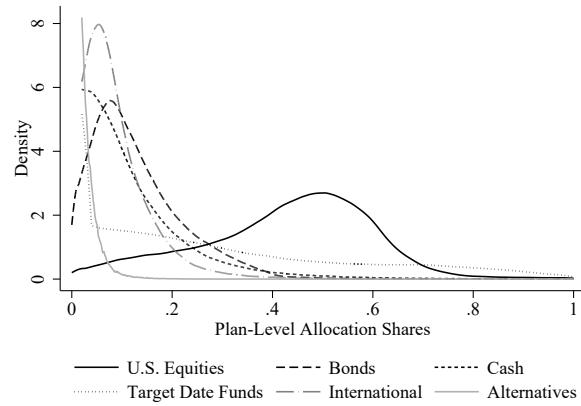
Yang, H. (2025). What determines 401 (k) plan fees? a dynamic model of transaction costs and markups. *Working Paper*.

Yao, R. and H. H. Zhang (2005). Optimal consumption and portfolio choices with risky housing and borrowing constraints. *The Review of Financial Studies* 18(1), 197–239.

Yogo, M., A. Whitten, and N. Cox (2021). Financial inclusion across the united states. *Working paper*.

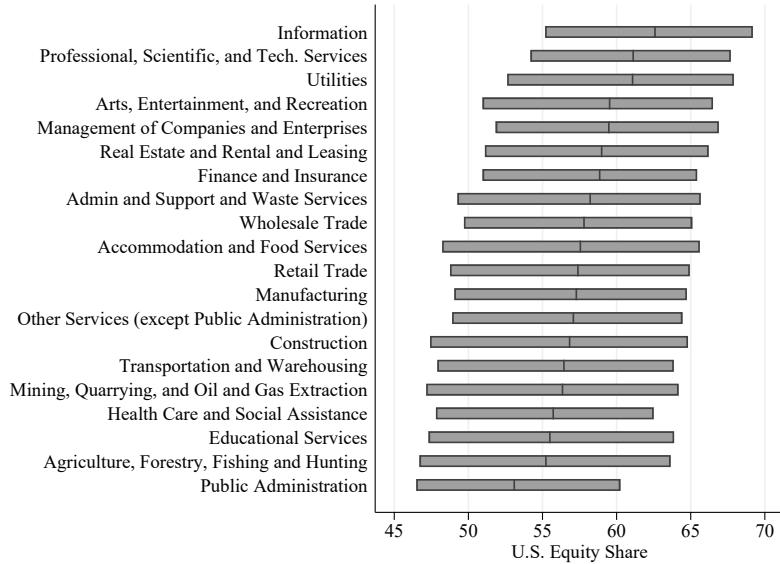
Tables and Figures

Figure 1: Distribution of Holdings



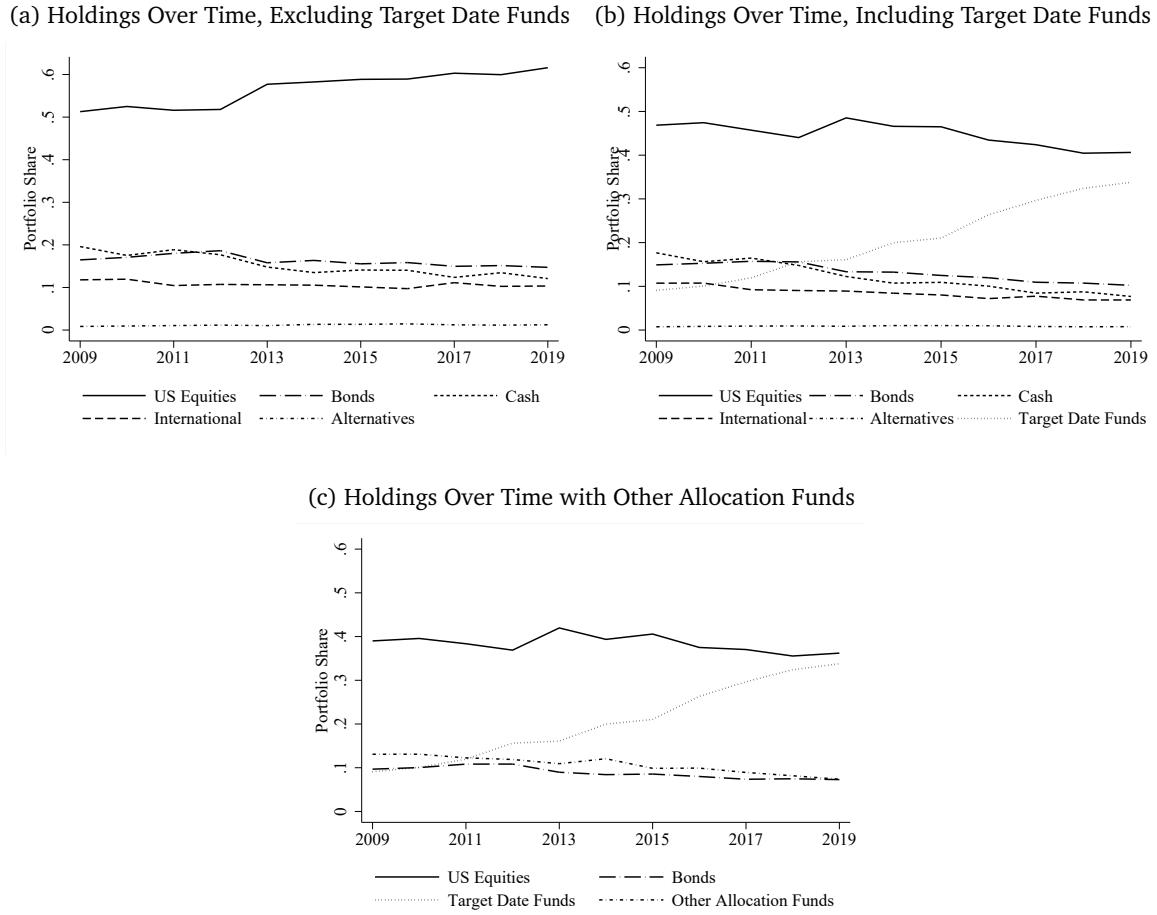
Notes: Figure 1 displays the distribution of holdings across 401(k) plans. To show the densities on the same scale, the figure censors observations with less than 2 percent share for cash, target date funds, international equities, and alternatives. Observations are at the plan-by-year level over the period 2009-2019 for plans with at least five investment options. **Alt Text:** Distribution of holdings in different asset classes across 401(k) plans.

Figure 2: Equity Allocation by Sector of Employment



Notes: Figure 2 displays the distribution of US equity allocations (i.e., share of plan assets held in US equities) across sectors (2-digit NAICS). The horizontal gray bars cover the 25th to 75th percentiles, and the short vertical lines indicate medians. When computing the share held in US equities, we drop all target date fund assets and assume that remaining non-target date allocation funds hold 60% in US equities. Observations are at the plan-by-year level over the period 2009-2019. **Alt Text:** Distribution of US equity allocations across 2-digit NAICS sectors.

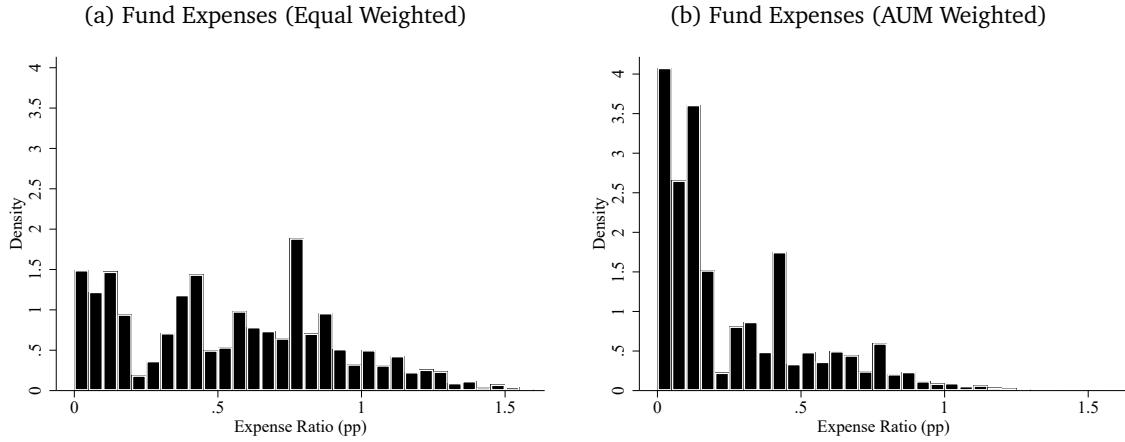
Figure 3: Holdings Over Time



Notes: Figure 3 displays the equal-weighted average holdings across plans over the period 2009-2019. In panel (a) we calculate portfolio shares excluding target date funds. In panel (b) we calculate portfolio shares including target date funds. In panel (c), we calculate shares for target date and non target date allocation funds, as well as US equity and bond assets without considering allocation funds.

Alt Text: Time series of average holdings across plans from 2009 to 2019, with and without target date funds.

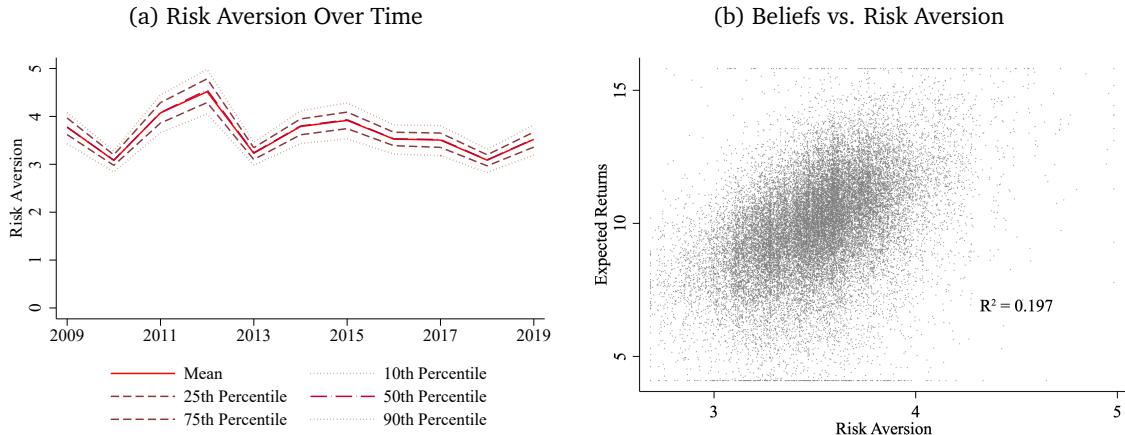
Figure 4: Fund Expenses



Notes: Figure 4 displays the distribution of fund expenses. Observations are at the fund-by-plan level as of 2019 as reported by BrightScope. Panel (a) displays the equal weighted distribution of fund expenses. Panel (b) displays the asset weighted distribution of fund expenses.

Alt Text: Distribution of fund expenses, equal weighted and asset weighted.

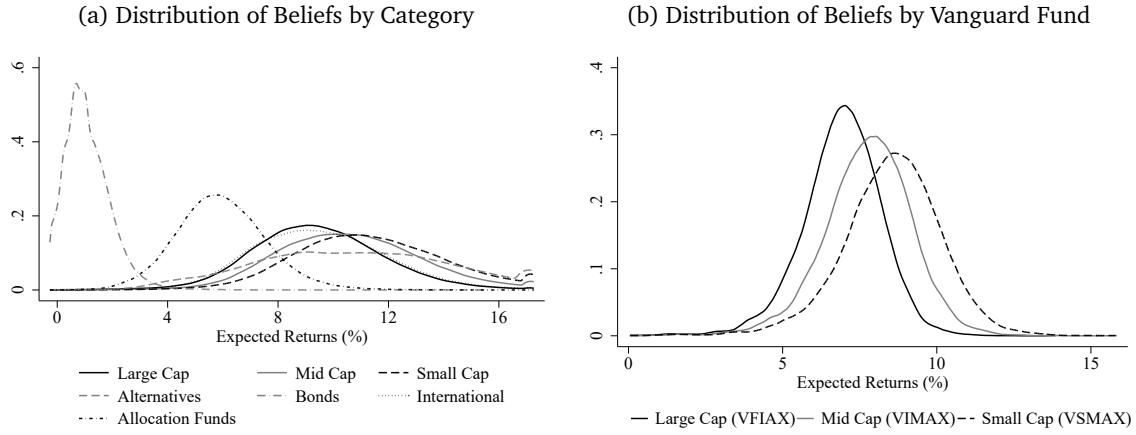
Figure 5: Estimates of Risk Aversion



Notes: Figure 5a displays estimated risk aversion over time. Figure 5b displays a scatter plot of the cross section of expected returns versus risk aversion as of 2016. The estimates correspond to the specification (5) of Table 2.

Alt Text: Estimated risk aversion over time and scatter plot of estimated risk aversion versus expected returns as of 2016.

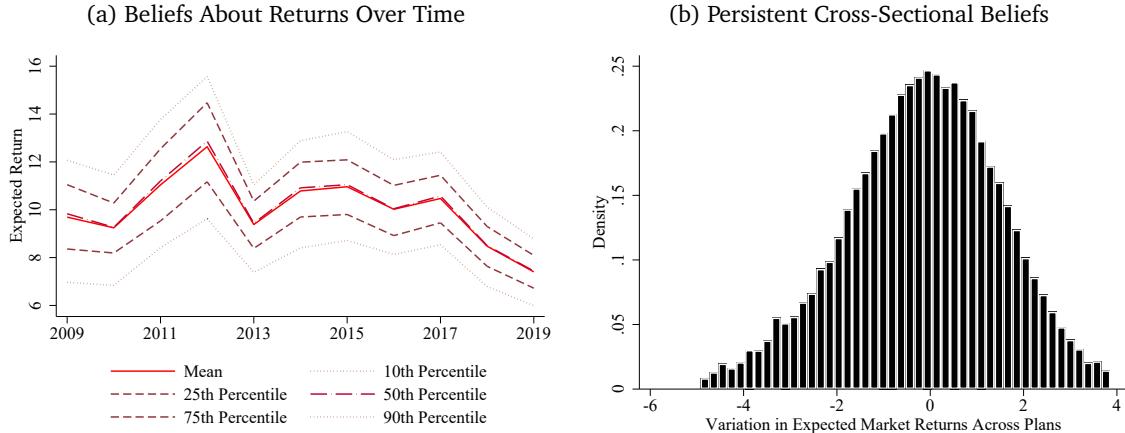
Figure 6: Distribution of Investor Beliefs by Investment Category and Fund



Notes: Figure 6a displays the estimated distributions of investors' expectations of returns across investors for each investment category. Figure 6b displays the estimated distributions of investors' expectations of returns for three popular Vanguard funds in the large, mid, and small cap categories: Vanguard 500 Index Fund (VFIAX), Vanguard Mid-Cap Index Fund (VIMAX), Vanguard Small-Cap Index Fund (VSMAX).

Alt Text: Estimated distributions of expected returns across investment categories and individual funds.

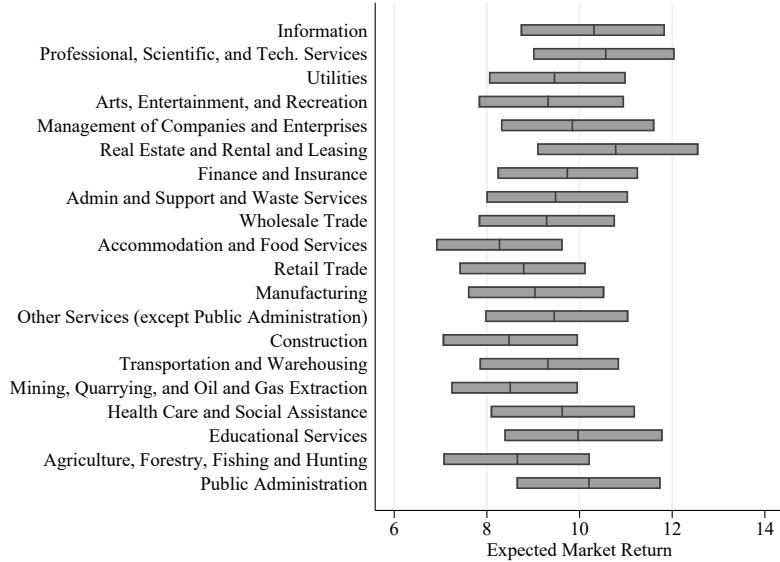
Figure 7: Estimates of Beliefs



Notes: Figure 7a displays the estimated distribution of expected market returns over time. Figure 7b displays the estimated cross-sectional distribution of expected market returns. In Figure 7b, expectations are de-meaned across investors within each year, and each observation reflects the average deviation from the yearly mean over the period 2009-2019. Negative values indicate plans with investors that have persistently pessimistic expectations relative to the mean.

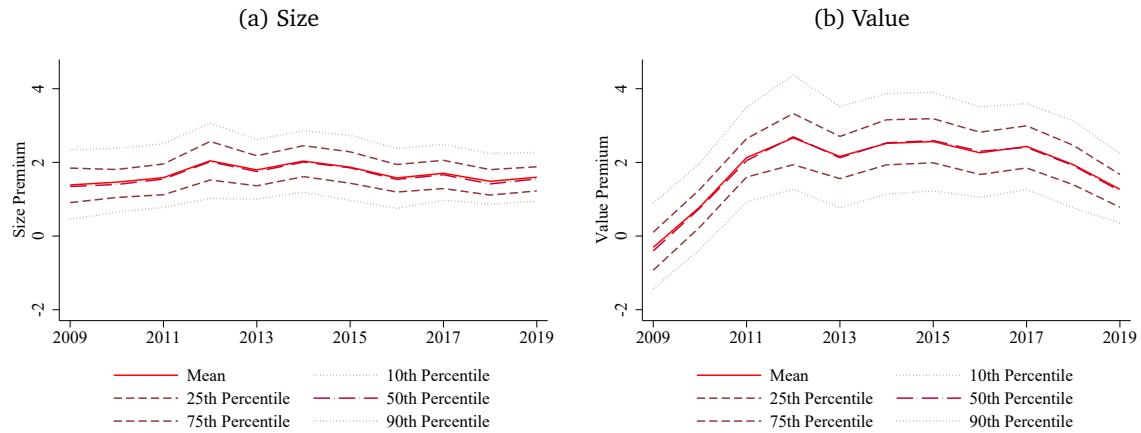
Alt Text: Estimated distribution of expected market returns over time and estimated cross-sectional distribution of expected market returns.

Figure 8: Expected Returns Across and Within Sectors



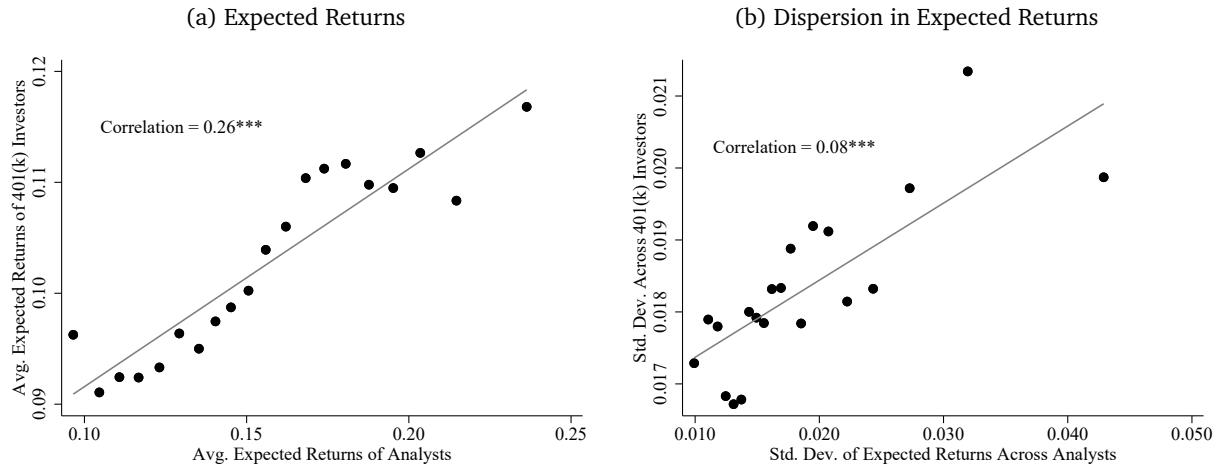
Notes: Figure 8 displays the distribution of expected market returns across sectors. The horizontal gray bars cover the 25th to 75th percentiles, and the short vertical lines indicate medians. Panel (a) is sorted by median U.S. equity allocation (see Figure 2). **Alt Text:** Estimated distribution of expected market returns across sectors.

Figure 9: Factor beliefs



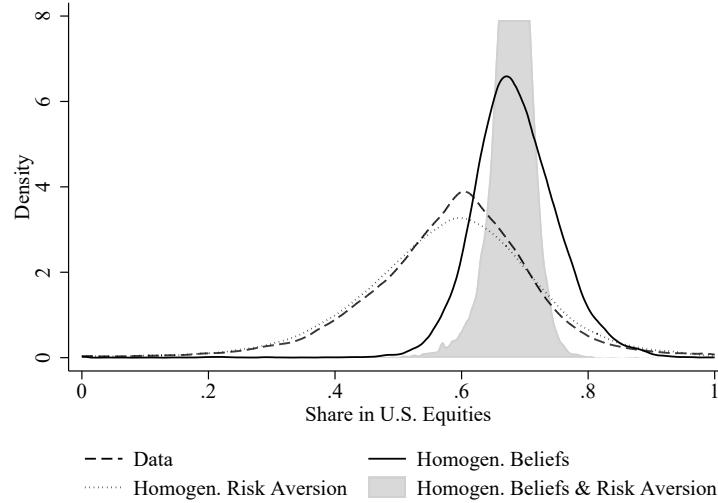
Notes: Figure 9 displays the estimated distribution of expected size and value factor returns over time. **Alt Text:** Estimated distribution of expected size and value factor returns.

Figure 10: Expected Returns Across Mutual Funds: Analysts vs. 401(k) Investors



Notes: Figure 10 panels (a) and (b) display the average and dispersion in expected returns across equity mutual funds for 401(k) investors and equity analysts. Observations are at the mutual fund-by-year level in both panels. Panel (a) displays a binned scatter plot of the average analyst's expected mutual fund returns versus the average 401(k) investor's expected mutual fund returns. Panel (b) displays a binned scatter plot of the standard deviation of expected mutual fund returns across analysts versus the standard deviation of expected mutual fund returns across 401(k) investors. **Alt Text:** Comparison of the average and dispersion of estimated expected returns relative to equity analysts.

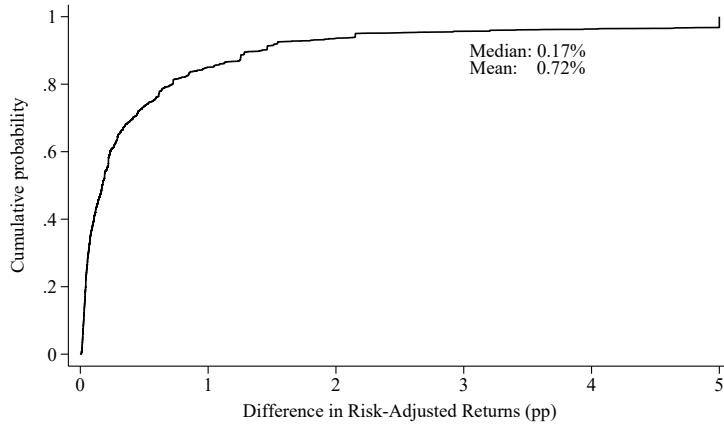
Figure 11: Counterfactual Allocations without Heterogeneity in Beliefs or Risk Aversion



Notes: Figure 11 displays actual and counterfactual densities of equity allocations by plan in 2016. The dashed line indicates the actual distribution of equity allocations across plans. The solid line indicates the counterfactual (optimal) allocations under the assumption that every investor has identical beliefs about each fund. The dotted line indicates the counterfactual allocations when investors have identical risk aversion parameters. The gray shaded area indicates allocations when investors share identical beliefs and risk aversion.

Alt Text: Counterfactual distributions of equity allocations without heterogeneity in beliefs or risk aversion.

Figure 12: Counterfactual: Welfare Costs of Belief Distortions

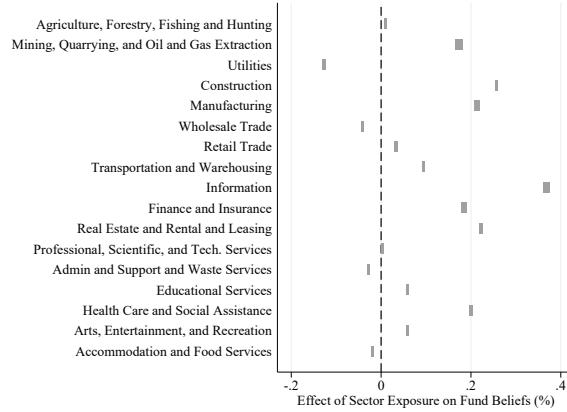


Notes: Figure 12 displays the distribution of $\Delta Consumer Surplus_{mt}$ at the plan level for the year 2016. The variable $\Delta Consumer Surplus_{mt}$ is calculated as the difference in consumer surplus if the investor held the optimal portfolio (given average beliefs in the economy) relative to the investors chosen portfolio (eq. 10). For reporting purposes, the distribution is censored at 5 percentage points.

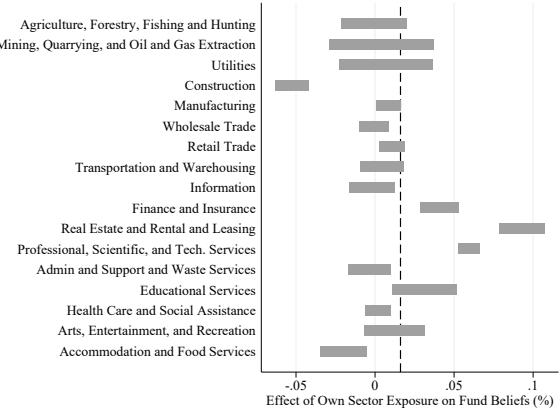
Alt Text: Cumulative probability of counterfactual welfare changes.

Figure 13: Sector Exposure

(a) Sector Belief



(b) Own Sector Belief



Notes: Figure 13 displays how investors' sectors of employment affect their beliefs based on regression estimates in eq 13. Panel (a) plots 95% confidence intervals of β_s without own sector interaction, capturing how fund allocation to different sector affect beliefs. Panel (b) shows estimated β^{own} on average and plots 95% confidence intervals of β^{own} for each sector. β^{own} captures investors' belief about their own sector.

Alt Text: confidence interval of sector exposure and own sector belief.

Table 1: Summary Statistics

(a) Plan Summary Statistics

	Obs	Mean	Std. Dev.	Median
Total Assets (millions)	442,631	84.749	689.657	10.722
Number of Plan Participants	425,075	1,261.304	92,360.288	223.000
Number of Investment Options	442,631	26.297	13.835	26.000
Average Account Balance	424,136	66,082.215	532,846.346	45,323.926
Plan Participation Rate	404,869	0.745	0.252	0.834
Employer Contribution Rate	392,401	0.337	0.245	0.290
Share Retired	406,258	0.008	0.014	0.001
Investment Category:				
US Equities	442,631	0.441	0.192	0.455
Target Date Funds	442,631	0.230	0.260	0.137
Bond Fund	442,631	0.126	0.096	0.106
Cash	442,631	0.113	0.127	0.078
International Stock	442,631	0.082	0.072	0.067
Alternatives	442,631	0.009	0.019	0.000
Investment Vehicle Type:				
Mutual Fund	442,631	0.612	0.407	0.823
Separate Account	442,631	0.191	0.356	0.000
Guaranteed Investment Contract	442,631	0.080	0.114	0.038
Collective Trust	442,631	0.053	0.169	0.000
Company Stock	442,631	0.030	0.154	0.000
Common Stock	442,631	0.010	0.076	0.000
Brokerage	442,631	0.009	0.054	0.000
Other	442,631	0.014	0.084	0.000

(b) Investment Option Summary Statistics

	Obs	Mean	Std. Dev.	Median
Volatility	10,781,851	0.137	0.043	0.148
Expense Ratio (pp; BrightScope)	1,856,108	0.569	0.383	0.590
Expense Ratio (pp; CRSP)	6,596,581	0.606	0.432	0.610

Notes: Table 1a displays plan level summary statistics. Observations are reported at the plan-by-year level over the period 2009-2019. Table 1b displays investment option-by-plan-by-year level summary statistics. Observations for *Expense Ratio (BrightScope)* are at the investment option-by-plan level as of 2019. Observations for all other variables are at the investment option-by-plan-by-year level over the period 2009-2019. *Volatility* corresponds to the dependent variable in eq. (4) and is annualized.

Alt Text: Summary statistics at the plan-year level and plan-investment option-year level.

Table 2: Estimated Model Parameters and Risk Aversion

	(1)		(2)		(3)		(4)		(5)	
	θ	λ								
Fee	-0.193*** (0.017)	5.171	-0.212*** (0.016)	4.718	-0.191*** (0.030)	5.246	-0.281*** (0.028)	3.558	-0.262*** (0.037)	3.823
\times Age			0.017** (0.007)	0.380	0.018** (0.008)	0.500	0.017** (0.007)	0.217	0.018** (0.008)	0.267
\times Frac Black			0.005 (0.005)	0.104	0.005 (0.005)	0.125	0.005 (0.005)	0.061	0.005 (0.005)	0.068
\times Frac Hispanic			-0.004 (0.011)	-0.096	-0.004 (0.011)	-0.111	-0.004 (0.011)	-0.053	-0.004 (0.011)	-0.057
\times Frac College			0.025** (0.010)	0.556	0.025** (0.010)	0.675	0.025** (0.010)	0.314	0.024** (0.010)	0.356
\times ln(Median Family Income)			-0.021* (0.012)	-0.465	-0.022* (0.012)	-0.601	-0.021* (0.012)	-0.261	-0.022* (0.012)	-0.316
\times ln(Median House Value)			0.009 (0.011)	0.193	0.009 (0.011)	0.242	0.009 (0.011)	0.109	0.009 (0.011)	0.127
\times Frac Employed			-0.005 (0.006)	-0.118	-0.004 (0.006)	-0.103	-0.005 (0.006)	-0.068	-0.004 (0.006)	-0.055
\times Unionized			0.016 (0.016)	0.352	0.015 (0.016)	0.403	0.016 (0.016)	0.203	0.015 (0.016)	0.218
\times Share Retired			-0.003 (0.006)	-0.068	-0.003 (0.006)	-0.082	-0.003 (0.006)	-0.040	-0.003 (0.006)	-0.046
\times ln(Avg. 401(k) Balance)			-0.000 (0.005)	-0.007	0.000 (0.005)	0.005	-0.001 (0.005)	-0.008	-0.000 (0.005)	-0.002
\times Existing 401(k) Plan							0.070*** (0.023)	0.885	0.072*** (0.023)	1.054
Observations	4,932,059		4,528,147		4,528,147		4,528,147		4,528,147	
Plan-Year FE	X		X		X		X		X	
Category-Year FE	X		X		X		X		X	
Year-Fee Interactions					X				X	
Estimated Risk Aversion										
Mean	5.171		4.781		5.236		3.584		3.546	
Std. Dev.	0.000		0.628		0.730		0.292		0.479	
Median	5.171		4.769		5.266		3.587		3.496	
Observations	442,631		402,497		402,497		402,497		402,497	

Notes: Table 2 displays two-stage least squares estimates corresponding to eq. (4). For each specification, the left column (θ) reports the linear regression estimates and standard errors, and the right column translates the coefficients in terms of in terms of risk aversion (λ) and the marginal effects for the average plan in 2009. Observations are at the investment option-by-plan-by-year level over the period 2009-2019. Standard errors are clustered at the plan level and are in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Alt Text: Two-stage least squares estimates of risk aversion corresponding to eq. (4).

Table 3: Residualized Variation in Expected Market Returns vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	-0.140*** (0.033)											-0.181*** (0.040)
ln(Income)		0.173** (0.064)										0.050 (0.056)
ln(Home Value)			0.154*** (0.048)									-0.001 (0.037)
College				0.226*** (0.073)								0.147** (0.064)
Employed					0.123* (0.061)							0.024 (0.033)
Black						-0.113*** (0.037)						-0.090*** (0.022)
Hispanic							-0.101* (0.051)					-0.085*** (0.029)
Unionized								-0.553*** (0.135)				-0.410*** (0.113)
Sector Equity Beta								0.133 (0.085)				0.133** (0.053)
Share Retired									-0.118*** (0.019)			-0.111*** (0.018)
ln(Avg. Acct. Bal.)										0.099** (0.041)		0.063 (0.049)
Observations	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268
R-squared	0.006	0.010	0.008	0.017	0.005	0.004	0.003	0.005	0.005	0.004	0.003	0.044
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Notes: Table 3 displays the results from linear regressions of residualized expected market returns on standardized demographic variables and a dummy variable for *Unionized*. Observations are at the investment plan-by-year level over the period 2009 through 2019. Standard errors are clustered 2-digit NAICs level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. Alt Text: Regressions of residualized expected market returns on standardized demographic variables.

Table 4: Predictive Allocations

VARIABLES	(1)	(2)
Factor-based Beliefs	12.80*** (0.09)	45.15*** (0.17)
Risk Aversion		-0.80*** (0.01)
Variance		-82.33*** (0.37)
Constant	-1.91*** (0.01)	1.17*** (0.04)
Observations	557,641	557,641
In sample R2	0.072	0.227
Out of sample MSE-R2	0.067	0.166
Out of sample OLS-R2	0.068	0.169

Notes: Table 4 displays the regression results corresponding to the predictive regression in eq. (8). We estimate the predictive regression using in sample mutual funds and display regression coefficients as well as in-sample R^2 . We then use this model to predict allocations to out-of-sample investment options, and report out-of-sample R^2 . Standard errors are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. Alt Text: Predictive regressions of allocations.

Table 5: Equity Holdings vs. Beliefs and Risk Aversion

VARIABLES	(1) All Equities	(2) All Equities	(3) US Equities	(4) US Equities	(5) Cash	(6) Cash
Risk Aversion (Std.)	-6.511*** (0.192)	-7.449*** (0.153)	-5.702*** (0.135)	-6.303*** (0.156)	4.400*** (0.262)	5.006*** (0.341)
Expected Returns (Std.)	9.974*** (0.367)	13.692*** (0.245)	8.561*** (0.280)	12.031*** (0.176)	-7.140*** (0.425)	-9.738*** (0.392)
Observations	243,268	243,268	243,268	243,268	243,268	243,268
R-squared	0.507	0.788	0.348	0.595	0.286	0.440
Year FE	X		X	X	X	X

Notes: Table 5 displays the regression results corresponding to a linear regression model. Observations are at the plan-by-year level over the period 2009-2019. Standard errors are clustered 2-digit NAICs level by year level and the plan level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. **Alt Text:** Regressions of equity holdings on estimated risk aversion and expected returns.

Table 6: Expected Market Returns vs. Local Economic Conditions

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Pop. Growth	0.125*** (0.016)	0.041*** (0.012)							0.117*** (0.021)	0.032** (0.013)
Home Price Growth		0.022*** (0.006)	0.005*** (0.002)						0.008 (0.007)	0.003 (0.002)
Establishment Growth			0.039*** (0.011)	0.016*** (0.004)					-0.022* (0.012)	0.006 (0.004)
GDP Growth				0.036*** (0.004)	0.005*** (0.001)	0.024*** (0.004)			0.003** (0.002)	
Observations	232,877	225,188	239,199	231,551	243,268	235,577	239,313	231,731	225,022	217,483
R-squared	0.357	0.871	0.344	0.865	0.343	0.864	0.344	0.864	0.359	0.872
Year FE	X	X	X	X	X	X	X	X	X	X
Plan FE	X		X		X		X		X	X

Notes: Table 6 displays results from linear regressions of expected market returns on local economic conditions. Observations are at the investment plan-by-year level over the period 2009 through 2019. Standard errors are in parenthesis and are clustered at the county-by-year level. *** p<0.01, ** p<0.05, * p<0.10. **Alt Text:** Regressions of expected market returns on local economic conditions.

Table 7: Expected Market Returns vs. Employer Returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
VARIABLES										
Firm Return (1 years)	0.001** (0.000)	0.001*** (0.000)							0.001** (0.000)	0.001*** (0.000)
Firm Investment			0.004** (0.002)	0.018*** (0.003)					0.005** (0.002)	0.016*** (0.004)
Sales Growth					0.000 (0.000)	0.004*** (0.001)			-0.000 (0.000)	0.001** (0.001)
Employment Growth							0.001 (0.001)	0.005*** (0.001)	0.000 (0.001)	0.002* (0.001)
Observations	11,495	11,738	10,262	10,474	11,233	11,452	11,216	11,441	9,889	10,081
R-squared	0.886	0.510	0.889	0.521	0.886	0.510	0.887	0.510	0.890	0.519
Year FE	X		X		X		X		X	
Plan FE	X		X		X		X		X	
NAICS-Year FE		X		X		X		X		X

Notes: Table 7 displays results from linear regressions of expected market returns on variables corresponding to employer performance. Observations are at the investment plan-by-year level over the period 2009 through 2019. Standard errors are in parenthesis and are clustered at the plan level. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of expected market returns on employer performance.

Appendix

A Additional Reduced-Form Results

A.1 Allocations and Investor Characteristics

Here, we present reduced-form relationships between investor characteristics and equity allocations, as well as the relationships for other asset classes. We examine how investment allocations vary across investor demographics in the following regression:

$$\text{Share in US Equities}_{mt} = X_{mt}\beta + \epsilon_{mt}. \quad (14)$$

Observations are at the 401(k) plan-by-year level. The dependent variable $\text{Share in US Equities}_{mt}$ reflects the share of assets held in equities in plan m at time t . When computing the share of assets held in US equities we exclude target date funds because they tend to be the default option in 401(k) plans.

We consider demographics, industry, and plan variables in X_{kt} . Following the literature, we focus on age, income, housing wealth, and race using county-by-industry-by-year level demographics information from the ACS. Since we do not perfectly observe participant demographics, this may introduce measurement error in our demographic covariates and could attenuate some of our results. We also include several plan-level characteristics using Form 5500 data. The Form 5500 data includes plan-by-year level information on the average account balance of plan participants, the share of participants that are retired, and plan age.

We present the corresponding estimates in Appendix Table A2. Column (1) displays the results with all covariates and time fixed effects. Column (2) also incorporates county and industry fixed effects. For ease of interpretation the independent variables are in units of standard deviation.

Wealth and Income: Plans with wealthier participants, measured by average account balances, allocate more towards equities. The results in column (2) indicate that a one standard deviation increase in the average account balance is correlated with a 0.89 pp increase in equity exposure. Previous research based on data from the Survey of Consumer Finances in the US (Heaton and Lucas, 2000; Campbell, 2006; Wachter and Yogo, 2010) and administrative data in Sweden and Norway (Bach et al., 2020; Fagereng et al., 2020) document a similar positive relationship between wealth and equity allocation. Because we study 401(k) portfolios conditional on participation, our results indicate that the positive relationship between wealth and equity allocation is not solely driven by participation costs along the extensive margin.

Similarly, we find some modest evidence that income and home wealth are positively correlated with equity exposure. The existing theoretical predictions regarding how equity exposure

vary by income and home value are mixed. Cocco et al. (2005) shows that labor income is a close substitute for risk-free asset and increase the demand for equity; however, other theoretical works highlight how income risk can also crowd out equity allocation (Catherine, 2021; Lynch and Tan, 2011; Storesletten et al., 2004). Using Administrative data in Sweden, Catherine et al. (2022) show that individuals with higher counterfactual income risk are less likely to participate and have lower equity share conditional on participation. Housing can also be considered as a long-term safe asset and hedges against rental prices (Sinai and Souleles, 2005). Housing also provides collateral for borrowing, and can increase equity holding thanks to lower borrowing constraints (Guiso et al., 1996). On the other hand, housing is illiquid. In life cycle models with housing decisions, Cocco (2005) and Yao and Zhang (2005) show that individuals with a higher fraction of total wealth in real estate invest less in risky assets.

Age and Retirement: We find that age and share of retired participants are negatively correlated with equity exposure. One standard deviation increases in participant age and the share of participants retired are associated with a 0.71 and 0.46 pp decline in US equity holdings, respectively (column 1). The decreasing age profile is consistent with standard life cycle models (Cocco et al., 2005) which consider the present value of future income as safe assets. Using novel survey data, Choi and Robertson (2020) find that years left until retirement is one of the most commonly cited factors for determining equity allocations. Empirical estimates tend to be mixed due to the identification challenge of collinearity among cohort, time and age effects. Using Norwegian administrative data, Fagereng et al. (2017) find that risky asset share of stock market participants is a decreasing function of age. However, Ameriks and Zeldes (2004) find evidence of hump-shaped patterns based on US data.

Other Demographics: We also find that more educated households have higher equity allocation. The results in column (2) indicate that a one standard deviation increase in the share of college educated individuals is correlated with a 0.86 pp increase in equity allocation. This relationship is consistent with the findings in Campbell (2006) and Black et al. (2018), and could potentially be driven by financial literacy (Calvet et al., 2007; Van Rooij et al., 2011).

We find that minorities invest less in equity. A one standard deviation increase in the fraction of Hispanic and black populations are correlated with 0.62 and 0.20 percentage point decreases in equity exposure. Campbell (2006) and Chiteji and Stafford (2000) also find that minorities have lower equity shares.

Other Asset Classes

The differences in equity allocation across plans documented above extends to other asset classes as well. Appendix Table A2 displays the regression results where we replicate eq. (14) for the other main asset classes. The dependent variable in columns (3)-(4) is the share in

bonds, in columns (5)-(6) is the share in cash, and in columns (7)-(8) is the share in international equities. A couple of interesting patterns emerge in Appendix Table A2. In general, the demographics that are positively (negatively) correlated with US equity ownership are also positively (negatively) correlated with international equity ownership with a few notable exceptions. For example, education is positively correlated with both US equity ownership and international equity ownership. However, wealth, as measured by account balances, is positively correlated with US equity ownership but negatively correlated with international equity ownership. These findings regarding international exposure are consistent with the evidence in Bekaert et al. (2017). Plans with a greater share of retirees and older participants tend to have higher bond and cash exposures and lower US and international equity exposures. Union membership and minority status are correlated with higher cash allocations but are negatively correlated with equity and bond allocations.

A.2 The Role of Investment Menus

A prior literature literature has indicated that the composition of the investment menu drives investors' allocation decisions (Benartzi and Thaler (2001)). Based on the previous results, it is fair to wonder to what extent the previous results that link demographics to allocations depends on the menu of options presented to employees.

In Appendix Table A3, we regress the number of funds available for each asset category on the demographic covariates. The estimates indicate that the menus themselves are largely uncorrelated with participant demographics. The point estimates are small and precisely estimated, and the overall explanatory of the regressions is small ($R^2 < 0.09$). This suggests that investor choices related to demographics—potentially arising from differences in beliefs and risk preferences—are driving the demographic differences in allocation decisions, not differences in menus.³⁴

In addition, we replicate the findings of Benartzi and Thaler (2001) in Appendix Table A4. Although we find that the menu composition is correlated with investment allocations, we find much weaker magnitudes than in Benartzi and Thaler (2001). Part of this may be due to the sample composition. Benartzi and Thaler (2001) study a cross-section of 170 plans in 1996 where the average plan has 6.8 different investment options. We study a much larger and more recent sample where the average plan has 26 options. Investors may have become more sophisticated in the past twenty years, and it is possible that investor behavior changes when facing a menu with more options, holding fixed investor sophistication.

The aggregate time series trends also point to the importance of investor choice, rather than naive allocations. Over our sample period, the share allocated to U.S. equities has grown (Figure 3a). However, the share of U.S. equity investment options has fallen over the same

³⁴In untabulated results, we also replicate Appendix Table A2 where we control for the composition of the 401(k) menu. We find that controlling for the composition of the menu has little impact on our estimates.

period, from 63 percent of fund menu options (excluding target date funds) in 2009 to 56 percent in 2019.³⁵ Thus, growth of non-equity options in the time series is not correlated with a shift to non-equity investment allocations.

A.3 Sensitivity of Investor Demand to Fees

To assess how investors respond to fees, we estimate the following demand specification:

$$\ln Share_{kmt} = \alpha p_{kt} + \phi_{mt} + \phi_{\tau(k)t} + \xi_{kmt}. \quad (15)$$

Observations are at the fund-by-plan-by-year level where we exclude target date funds. The dependent variable $Share_{kmt}$ measures the share of assets held in fund k in plan m at time t relative to the total assets in plan m at time t . Fund expense p_{kt} is the independent variable of interest. We include plan-by-year (ϕ_{mt}) and fund-type-by-year ($\phi_{\tau(k)t}$) fixed effects. We define fund type $\tau(k)$ based on the fund's classification in both Morningstar and BrightScope and whether it is a index/passive fund (i.e., Morningstar Category x BrightScope Category x Passive). Including plan-by-year fixed effects is important because it allows us to measure how investors trade off relative differences in expenses among the funds available in the investor's 401(k) menu rather than differences across 401(k) menus, which may be correlated with plan size. While we present eq. (15) as a simple linear specification, by including plan-by-year fixed effects, eq. (15) directly corresponds to the workhorse discrete-choice demand model developed in Berry (1994) that is commonly used in the industrial organization literature.³⁶ We use the same Hausman-type instruments as the estimation in our main model to address concerns for endogeneity of fund fees.

We report our demand estimates in Appendix Table A5. Column (1) displays the OLS results and column (2) displays the corresponding IV results. Note that OLS and IV estimates are quite similar, so the potential endogeneity concern appears minimal. The results indicate that, as expected, investors are sensitive to expenses. The results in column (2) indicate that a 10 bp

³⁵Consistent with our treatment in our main analysis, these numbers treat non-target date allocation funds as equity funds. The share in non-allocation U.S. equity funds has also declined slightly over this period.

³⁶Following the setup in (Berry, 1994) , the market share of product k in market m can be written in logs as

$$\ln share_{kmt} = \alpha p_{kt} + \xi_{kmt} - \ln \left(\sum_{k' \in \mathcal{K}_{mt}} \exp(\alpha p_{k't} + \xi_{k't}) \right),$$

where ξ_{kmt} captures unobserved product characteristics and \mathcal{K}_{mt} is the set of available products available in market m at time t . In the context of 401(k) choice, k refers to the fund and markets are defined based on the 401(k) plan menu. The plan-by-year fixed effect in eq. 15 absorbs the non-linear term $\ln \left(\sum_{k' \in \mathcal{K}_{mt}} \exp(\alpha p_{k't} + \xi_{k't}) \right)$ which is constant within a plan-year. This type of demand system has been used in a number of other financial applications such as demand for bank deposits (Dick, 2008; Egan, Hortaçsu, and Matvos, 2017; Wang, Whited, Wu, and Xiao, 2022), convertible bonds (Egan, 2019), credit default swaps (Du et al., 2019), insurance (Koijen and Yogo, 2016, 2022), mortgages (Benetton et al., 2021; Benetton, 2021) and investments more generally (Koijen and Yogo, 2019a,b; Koijen et al., 2019).

increase in fees is associated with a 6.7% decrease in demand. In the context of the discrete choice demand system developed in Berry (1994), the estimates in column (2) correspond to a demand elasticity of -0.40.³⁷

B Additional Results on the Formation of Beliefs

In the main text, we present evidence that investors use different information sets when forming beliefs. Specifically, we examine how expected market returns reflect local economic conditions and past performance of an investor’s employer while controlling for macroeconomic aggregate information, which can be captured with time fixed effects. We find that local economic conditions and employer past performance are positively correlated with beliefs about market returns, suggesting that investors form broader beliefs about market returns based on individualized experiences.

Here, we show that investors respond to contemporaneous, commonly available information. We focus on past fund returns, which are often a salient feature of 401(k) brochures. Consistent with extrapolation, we find that investor beliefs are positively related with fund-specific past returns. Importantly, we are able to use variation in 401(k) menus over time to show that the relationship between past returns and beliefs holds for funds that are newly added to investors’ 401(k) menus. Thus, the extrapolation we document cannot be explained by investors prior-year holdings in their 401(k) plans.

We also assess the rationality of investor beliefs. The above evidence, which documents systematic and predictable drivers of heterogeneity of beliefs, suggests that a standard rational expectations model may not capture the investment behavior across households. Investor beliefs are correlated with observable characteristics such as wealth and income, and appear to depend on past market returns as well as on recent employer performance. We find evidence consistent with the vast prior literature suggesting that investor forecasts violate full information rational expectations. Forecast errors are predictable and forecast revisions, measured by changes in investor expectations, are correlated with future forecast errors.

B.1 Extrapolation from Fund Returns

We examine how investors form their beliefs for a particular fund based on the fund’s return over the previous year. We estimate the regression:

$$\bar{\mu}_{mkt} = \rho Ret_{kt-1} + v_{kt}. \quad (16)$$

Observations are at the investment option-by-plan-by-year level. The dependent variable measures the average participant in plan m ’s expected return of fund k ($\bar{\mu}_{mkt}$). The independent

³⁷To compute the demand elasticity, we assume a market share of 1/26 and fee of 0.61pp.

variable Ret_{kt+1} measures the past monthly return of investment option k averaged over year $t - 1$ to t and is annualized. Appendix Table A7 displays the estimates corresponding to eq. (16). We examine extrapolation across three different subsets of the data: (i) the full data set in columns (1), and (4); (ii) fund-by-plan observations in the first year the fund was added to the plan in column (2);³⁸ and (iii) fund-by-plan observations corresponding to the first year a 401(k) plan was introduced in column (3). Samples (ii) and (iii) allow us to examine how investors extrapolate their beliefs about funds they have not previously held in their 401(k).

We find evidence that investors extrapolate their beliefs from past returns. The results in columns (2)-(4) indicate that investors extrapolate their beliefs from past returns for funds they did not hold in the past. The results in column (2) indicate that a ten percentage point increase in last year's return is correlated with a 0.16 pp increase in expected returns. In column (4) we interact past returns with the dummy variable $New Investment_{kt}$, which indicates whether the fund was added to the 401(k) menu in year t . We find a small statistically insignificant coefficient which indicates that investors extrapolate in the same way for both new and existing funds. The results in columns (2)-(4) show that the extrapolation we observe is not simply a function of investor inattention or inertia in portfolio rebalancing.

B.2 Are Beliefs Rational?

Lastly, we examine the rationality of investor beliefs by examining forecast errors. The previous results already provide suggestive evidence that investor beliefs are irrational. The unpredictability of forecast errors is a necessary condition for rational forecasts. We construct forecast errors at the plan-by-investment-by-year level as:

$$\varepsilon_{mkt+1} = Ret_{.kt+1} - \bar{\mu}_{mkt} \quad (17)$$

where $Ret_{.kt+1}$ measures the monthly return of investment option k averaged over year t to $t + 1$ and is annualized. The term $\bar{\mu}_{mkt}$ is the average participant in plan m 's expected return of fund k . We test the predictability of forecast errors in the following regression model:

$$\varepsilon_{mkt+1} = \alpha_0 + \alpha_1 X_{mkt} + \eta_{mkt+1}. \quad (18)$$

Observations are at the investment option-by-plan-by-year level. The vector X_{mkt} consists of lag forecast errors, lag fund returns, or changes in beliefs. We examine how forecast errors vary with past forecast errors, past fund returns, and changes in investor expectations.

Appendix Table A8 displays the estimation results corresponding to eq. (18). In short, we find overwhelming evidence that forecast errors are predictable. The results in column (1) indicate that forecast errors are persistent. We also find that investors tend to over predict fund

³⁸To keep the sample distinct from sample (iii), we exclude all fund-by-year observations when the 401(k) plan is introduced.

returns following past positive fund returns (columns 3 and 4). This is consistent with our previous finding that investors extrapolate from previous returns. We also find that changes in beliefs are negatively correlated with future forecast errors. This test is in a similar vein as the test developed in (Coibion and Gorodnichenko, 2015) and employed in (Bordalo et al., 2018) where the researchers examine how forecast errors correlate with forecast revisions. The negative relationship between changes in beliefs and future forecast errors suggests that investors overreact to news.

One might expect that these patterns are driven by inexperience in financial markets. In untabulated results, we replicate our analysis where we restrict our analysis to those plan sponsors in the finance and insurance sector (NAICS 52). Similar to our baseline results, we find that the beliefs of investors working in the financial sector are extrapolative, violate full information rational expectations, and tend to overreact to news.

C Alternative Instruments

One potential threat to the exclusion restriction is that providers may jointly set fees across fund categories, causing idiosyncratic beliefs of one fund to correlate with fees charged on other funds. We construct two alternative Hausman instruments. First, we use average fees charged by the same provider for funds not offered in any 401(k) plans in our sample and in a different investment category. These funds are likely subject to distinct pricing strategies. Second, we use 2019 BrightScope data on separate account expense ratios, which are typically negotiated individually with plan sponsors. These instruments draw on distinct market or investor segments and help mitigate concerns about joint fee-setting across categories.

As additional robustness checks, we consider three different sets of instruments that exploit different sources of variation in the data. First, we use fund-specific cost-shifters as instruments for fund expense ratios. These shifters are relevant to the extent that a fund's expense ratio at least partially reflects the cost of operating the fund, and the instrument is exogenous if these cost components are orthogonal to investors' beliefs. Following Janssen and Thiel (2025), we use fund-level variation in custodian fees as a cost-shifter instrument. Each fund must pay a custodian, which is a financial institution that safeguards fund assets, ensures regulatory compliance, and maintains transaction records, for its services. We find that the instrument is relevant and believe it is unlikely that custodial fees, conditional on other fund characteristics, are correlated with investors' beliefs.

Second, we use Waldfogel instruments, which are based on the exogenous characteristics of other 401(k) plans offering the same funds, following Waldfogel (2003) and Fan (2013). The intuition is that fund providers set a single expense ratio for a mutual fund (or share class) while targeting a broad investor base across plans with heterogeneous characteristics. Since we include plan-year fixed effects, variation from the plan's own characteristics is absorbed,

and identification comes from differences in average characteristics across other plans using the same fund. These instruments capture exogenous variation in demand from other plans, which influences how providers set prices, thereby making them relevant. The instruments are plausibly exogenous, as the average characteristics of other plans (e.g., plan age measured by the number of years since plan inception, total number of plan participants, and total plan assets) are likely orthogonal to investors' beliefs about a specific fund after controlling for own plan characteristics. Importantly, Waldfogel instruments do not rely on provider-level variation and thus help address concerns related to provider reputation.

Third, we exploit variation arising from provider switches, whether intentional or due to mergers and acquisitions. Such switches often result in the addition or removal of funds with varying expense ratios. One major type of addition is index funds, which became increasingly popular during our sample period (Yang, 2025). When plans switch providers, menus are frequently updated with more index funds, which tend to have lower fees than the previously offered funds. To capture this variation, we construct an instrument defined as a dummy equal to one for index funds in years following a provider switch. We use the interaction of the post-switch year and the index fund indicator, rather than a provider switch dummy, because we include plan-by-year fixed effects. This instrument is relevant if index funds introduced by new providers generally have lower fees. The exclusion restriction would be violated if, for example, investors were educated by the new provider about the benefits of index funds and subsequently formed higher expected gross returns for those funds, which coincidentally have lower fees.

We report estimates using different instruments in Table A1. OLS estimate in column (1) shows that without any fee instruments, our estimate is closer to zero, implying a larger risk aversion. We report our baseline specification in Column (2). The corresponding fee coefficient is -0.193. In columns (3) and (4), we report the estimates corresponding to alternative Hausman instruments. The estimates range from -0.198 to -0.232. We report the results corresponding to our cost-shifter and Waldfogel instruments in columns (5) and (6) and estimate coefficients of -0.275 and -0.166, respectively. Lastly, we report the results using variation in fund providers in columns (7) and (8). In column (7) we use variation across all provider switches and in column (8) we use variation in provider switches that only come from mergers. Overall, we find estimates that are consistent with our baseline estimates across the different specifications. In the last two specifications, we find that demand is slightly more sensitive to expense ratios (-0.40); however, the estimates are statistically indistinguishable from our baseline results.

D Considering Alternative Model Assumptions

We estimate risk aversion by measuring how investors trade off risk and expected returns in eq. (4). Specifically, we estimate risk aversion by examining how investors adjust their portfolio risk exposure in response to changes in expense ratios. Changes in expense ratios are equivalent to shifts in expected returns, allowing us to calculate risk-return tradeoffs. An investor optimally sets the expected return of an investment equal to the marginal risk, scaled by risk aversion. Our approach relies on the following assumptions: we can correctly measure investors' beliefs about risk; investors make allocation decisions considering their retirement accounts only; investors solve a myopic portfolio problem; and investors only trade off risk and expected returns.

These assumptions impose some relatively strong restrictions on investor behavior; however, our estimates may still be interpretable if these assumptions are violated. We discuss how the interpretation of our estimates of beliefs and risk aversion would change if our baseline assumptions are violated. It is important to emphasize that while our methodology places some strong constraints on investor behavior, we do not impose rational beliefs in our analysis and our framework allows for behavioral biases and mistakes in investor beliefs.

Outside Assets: By focusing on 401k investments, we only observe part of an investor's overall portfolio. According to 2019 Survey of Consumer Finance (SCF), working age individuals who have retirement accounts allocate on average 64% of their financial assets in retirement accounts, compared to 7% in other investment funds, directly owned stocks, or bonds. On average, these individuals hold 26% in cash (deposit, money market, etc.). Thus, retirement assets represent the vast majority of risky financial assets for most individuals with access to retirement accounts. In addition, human capital measured by net present value of future income approximates the risk and return profile of bonds for most individuals, and hence constitutes another important financial asset. Although retirement accounts are the primary source of risky equity for most investors, we could potentially over-estimate an investor's equity share by ignoring outside cash holdings and human capital.

To see how this would impact our estimates of beliefs and risk aversion, suppose that the true measure of risk is $\varsigma_{mkt}^2 = h_{mt}\tilde{\varsigma}_{mkt}^2 + \nu_{mt}$, where h_{mt} is the fraction of assets held in retirement accounts and ν_{mt} captures the additional risk from non-retirement account holdings. Rather than recover the true risk aversion, we will recover risk aversion scaled by h_{mt} . With additional information on h_{mt} , we can adjust our risk aversion accordingly. Our estimates of investor beliefs will capture an investor's true beliefs plus the term coming from additional risky asset holdings ν_{mt}/θ :

$$\widehat{\bar{\mu} - R_F} = -\frac{\hat{\epsilon}}{\hat{\theta}} = -\frac{\left(\frac{\varsigma^2 - \nu}{h} - \frac{\theta}{h}p\right)}{\frac{\theta}{h}} = -\frac{\varsigma^2 - \theta p - \nu}{\theta} = -\frac{\epsilon}{\theta} + \frac{\nu}{\theta} = \mu - R_F + \frac{\nu}{\theta}.$$

Since $\nu > 0, \theta < 0$, this will cause our estimates of beliefs to be biased downwards if investors have other equity assets outside of their 401k. Since the SCF data shows that 401(k) accounts are the primary source of risky assets for most households, the associated bias may be small.

Dynamic Allocation Across Multiple Periods: We model an investor's allocation decision as a myopic portfolio choice problem. It is well known that when investors have power utility and return distributions are independent over time, long-term portfolio choice is equivalent to myopic portfolio choice. More general time-varying returns could introduce intertemporal hedging demand, which would not impact our estimates of risk aversion but would potentially be captured in our estimates of beliefs. To illustrate, let ρ denote the vector of covariance of risky asset's excess return with the quality of future investment opportunities (e.g., the risk free rate in Campbell and Viceira (1999) or risk premium in Campbell and Viceira (2001)). The investor's first order condition would then be:

$$\lambda \Sigma \omega_i = \mu_i - \mathbf{p} - \mathbf{1} R_F - \psi \rho,$$

where the term $\psi \rho$ would be captured in our estimates of μ_i . For example, if equity is mean reverting, lower unexpected return today is correlated with better investment opportunities in the future. This would result in a positive hedging term and would potentially bias our estimates of beliefs upwards.

Optimization Error: We assume that investors actively trade off and equate marginal risk with returns when making investment decisions. However, there are a few reasons this could be violated in the data. Suppose that marginal risk is equal to expected returns plus some vector of optimization errors ζ_i :

$$\lambda \Sigma \omega_i = \mu_i - \mathbf{p} - \mathbf{1} R_F + \zeta_i. \quad (19)$$

One could interpret this optimization error as either a true error term or a variable that captures unobserved preferences of consumers. For example, it could be the case that even conditional on the risk and expected returns of a fund, investors have preferences for one fund over another. This type of optimization error would impact our estimation in the exact same way as a noisy measure of risk; if the optimization error is either not mean zero and/or correlated across investors within a plan, the risk aversion estimate would still be consistent but the beliefs estimates would reflect this preference (and potentially be biased).

Inattention: Relatedly, one might be concerned that investors do not actively trade off expected returns with risk. For example, investors may be inattentive such that only a fraction of investors actively update their portfolios every period (Gabaix, 2019). Generally speaking, our estimate of risk aversion could be biased upwards because investors would appear as if they

are insensitive to expected returns/fees. As such, investors would appear to be unwilling to take on additional risk after an increase in expected fund returns. If we were to systematically over-estimate risk aversion, this would result in over-estimating investor optimism regarding fund returns because investors equate expected returns scaled by risk aversion to risk.

Suppose a fraction π of investors rebalance each period following a Calvo model. Our observed holding at plan level $\bar{\omega}_t$ is a weighted average of current and past allocation. For simplicity, we omit investor and plan subscripts and only keep track of time subscript.

$$\bar{\omega}_t = \pi \omega_t + \sum_{l=1}^{\infty} \pi (1 - \pi)^l \omega_{t-l}$$

Our estimation Equation 4 becomes

$$\zeta_t^2 = \theta \left[\pi (\mu_t - p_t) + \sum_{l=1}^{\infty} \pi (1 - \pi)^l (\mu_{t-l} - p_{t-l}) \right]$$

Next, for simplicity we assume fund expense ratios are independent from beliefs in the past, present or future. Also, let expense ratios follow an AR(1) process when moving backward in time with $p_{t-1} = \phi p_t + \epsilon_{t-1}$ for $\phi \in (0, 1)$. Under these assumptions, OLS estimate of $\hat{\theta}$ can be expressed as the following:

$$\hat{\theta} = \theta \left[\sum_{l=0}^{\infty} \pi (1 - \pi)^l \phi^l \right] = \theta \left(\frac{\pi}{1 - (1 - \pi) \phi} \right) < \theta.$$

We underestimate θ and over estimate risk aversion $\lambda = 1/\theta$.

While some investors may be inattentive, survey evidence suggests that investors rebalance portfolios over time. In 2020 (2009) roughly 17% (15%) of DC participants changed the asset allocation of their account balance and 10% (19%) changed the asset allocation of their contribution.³⁹ Consistent with this survey evidence, we estimate that the one-year autocorrelation in 401(k) fund holdings is 0.77-0.89 after adjusting for returns. This autocorrelation indicates that investors actively rebalance their portfolios.

To help address potential concerns, we separately examine the investment allocation decisions of participants in the year the 401(k) plan was first introduced. When a 401(k) plan is introduced, any allocation into non-target date funds reflects an active choice of the participant. We discuss this robustness check in Section 4 and note that the estimated risk aversion appears roughly 20% lower in the year when the 401(k) plan was introduced. This suggests that some investors may be inattentive, but it does not appear to be the driving factor of our estimate of risk aversion. Moreover, our estimates are in line with previous estimates of risk aversion.

³⁹See https://www.ici.org/system/files/2021-09/21_rpt_recsurveyq2.pdf. ICI reports rebalancing activity for the first half of 2009 and 2020, which we annualize by multiplying them by two.

Heterogeneous Beliefs about Risk: We assume that investors agree on the covariance structure of returns, but potentially disagree about expected returns. However, investors can also disagree about factor loadings that lead them to form different beliefs about risk of assets. Such heterogeneous beliefs about risk will introduce an measurement error in our dependent variable. Because our dependent variable is the product of risk and allocation, measurement errors in risk is not additively separable and can introduce biases in our risk aversion estimates.

Consider a simple example where each investor i chooses from k assets. Here, by investor, we refer to the representative investor in each retirement plan. Investors perceive risk following a factor structure. For simplicity, we assume there are k factors. Each asset loads on one factor only, so the returns of these assets are uncorrelated, which simplifies notation. For this example, we also assume expense ratios are independent from beliefs. Let ν_{ik} denote investor i 's belief about the variance of asset k

$$\nu_{ik} = \nu + \beta_{ik}^2 \nu_k$$

Investor i 's holding in asset k is ω_{ik} . We re-arrange to obtain our estimation eq. (4)

$$\begin{aligned}\omega_{ik} &= \frac{\mu_{ik} - p_k}{\lambda \nu_{ik}} \\ \nu_{ik} \omega_{ik} &= \theta p_k - \theta \mu_{ik}\end{aligned}$$

Suppose the average belief about risk is equal to the variance of past return, which is our measurement of risk $E[\nu_{ik}]$. Then, our dependent variable is $E[\nu_{ik}] \omega_{ik}$, and we estimate the following:

$$\begin{aligned}\hat{\theta} &= \frac{COV(E[\nu_{ik}] \omega_{ik}, p_k)}{VAR(p_k)} \\ &= \frac{COV\left(E[\nu_{ik}] \frac{\theta}{\nu_{ik}} (p_k - \mu_{ik}), p_k\right)}{VAR(p_k)} \\ &= \frac{\theta E[\nu_{ik}] E\left[\frac{1}{\nu_{ik}}\right] (E[p_k^2] - E[p_k]^2)}{VAR(p_k)} \\ &= \theta E[\nu_{ik}] E\left[\frac{1}{\nu_{ik}}\right] \neq \theta.\end{aligned}$$

Hence, our estimated $\hat{\theta}$ would be biased. By convexity of $1/\nu$, $E\left[\frac{1}{\nu_{ik}}\right] > \frac{1}{E[\nu_{ik}]}$ and we overestimate $\hat{\theta}$. Since fee coefficient $\hat{\theta} < 0$, over estimating $\hat{\theta}$ suggests that we underestimate risk aversion $\hat{\lambda} < \lambda$. Intuitively, investor allocation ω is a function of $1/\nu$. The convexity suggests that investors who perceive small risk have very large holdings. Rationalizing such

large holdings with average risks requires lower risk aversion.

Next, we recover beliefs of expected return μ_{ik} in the following.

$$\begin{aligned}
\widehat{\mu_{ik}} &= \frac{E[\nu_{ik}]\omega_{ik} + \frac{1}{\lambda}E[\nu_{ik}]E\left[\frac{1}{\nu_{ik}}\right]p_k}{\frac{1}{\lambda}E[\nu_{ik}]E\left[\frac{1}{\nu_{ik}}\right]} \\
&= \frac{\frac{1}{\lambda}E[\nu_{ik}]\frac{1}{\nu_{ik}}(\mu_{ik} - p_k)}{\frac{1}{\lambda}E[\nu_{ik}]E\left[\frac{1}{\nu_{ik}}\right]} + p_k \\
&= \frac{\mu_{ik}}{\nu_{ik}E\left[\frac{1}{\nu_{ik}}\right]} + p_k \frac{\nu_{ik}E\left[\frac{1}{\nu_{ik}}\right] - 1}{\nu_{ik}E\left[\frac{1}{\nu_{ik}}\right]} \neq \mu_{ik}
\end{aligned}$$

The recovered beliefs for individual investors $\widehat{\mu_{ik}}$ are also biased. The first component in the final line of the equations above gives us a risk-adjusted expected return. If investor i believes that risk is high $\nu_{ik}E\left[\frac{1}{\nu_{ik}}\right] > 1$, the investor's actual belief about return is higher than our estimate. The second component is a function of expense ratio. Because we estimate a smaller fee coefficient (in absolute terms), we do not sufficiently incorporate expense ratio when recovering gross-of-fee return, so we need to add back this term.

Lastly, we consider the average recovered belief across investors. If beliefs about expected returns and risks are positively correlated, the covariance between the inverse of risk and belief in the second term is negative, $Cov\left(\frac{1}{\nu_{ik}}, \mu_{ik}\right) < 0$, and we underestimate beliefs. If beliefs and risks are orthogonal or if investors agree on expected returns μ_k , $Cov\left(\frac{1}{\nu_{ik}}, \mu_{ik}\right) = 0$ and we estimate unbiased average beliefs.

$$\begin{aligned}
E[\widehat{\mu_{ik}}] &= E\left[\frac{\mu_{ik} - p_k}{\nu_{ik}}\right] \frac{1}{E\left[\frac{1}{\nu_{ik}}\right]} + p_k \\
&= \frac{E\left[\frac{1}{\nu_{ik}}\mu_{ik}\right]}{E\left[\frac{1}{\nu_{ik}}\right]} \\
&= E[\mu_{ik}] + \frac{Cov\left(\frac{1}{\nu_{ik}}, \mu_{ik}\right)}{E\left[\frac{1}{\nu_{ik}}\right]}
\end{aligned}$$

E Robustness Checks

E.1 Alternative Model Specifications

We consider several alternative specifications to assess the robustness of the estimated parameters. First, we re-estimate the model to include target date funds, which are excluded from our baseline analysis. Second, to account for potential inertia in investor behavior, we estimate the model using only new plans. For all of these specifications, we find very similar estimates of risk aversion and expected returns. Results are reported in Appendix Table A9. The mean risk aversion we estimate in these alternative models is nearly identical to our baseline estimate (3.55 and 3.56 vs. 3.55). The mean expected return ranges from 9.7 to 9.9, similar to our baseline estimate of 9.6. As shown in panel (b), individual estimates of expectations and risk aversion are positively and significantly correlated with the baseline specification.

In addition, we consider alternative measures of risk based on both simplified and more extensive factor structures, as we describe in Section 3.3. As in the above specifications, we find that the estimates of risk aversion and beliefs are highly correlated with our baseline estimates. We do estimate a somewhat higher average risk aversion (7.63) with our simplified measure of risk relative to our baseline estimates. This can be explained in part because the standard deviation of our simplified measure of risk is roughly 40% smaller than the standard deviation of our baseline measure of risk, though the two measures are highly correlated (0.94).

E.2 Within-Plan Correlation in Beliefs and Risk Preferences

In our main analysis, we assume that all investors in the same plan share the same risk aversion. We would obtain identical initial estimates under a weaker assumption that the within-plan (e.g., across participants) heterogeneity in risk aversion is uncorrelated with the within-plan heterogeneity in beliefs. However, this has the minor inconvenience that we cannot directly translate the mean inverse risk aversion to mean risk aversion because of Jensen's inequality.

If beliefs and risk aversion are correlated among individuals within a plan, our framework still allows us to recover the average inverse of plan risk aversion. However, our measure of average plan beliefs will be a function of average beliefs and the within-plan covariance of risk aversion and beliefs. This would bias our estimates of beliefs, similar to the discussion of measurement error of risk in Section 3.5. As we show in Section 4, our estimated expected market returns seem reasonable and in line with realized returns over the period (approximately 10 pp), which suggests that within-plan correlation of beliefs and risk aversion is not a first-order concern.

Nonetheless, we perform a robustness check on the sensitivity of our results to within-plan correlation, focusing on expected market returns. Denote the expected returns for individual i on fund k as $\mu_{ik} = \delta_i b_{1k} + \eta_{ik}$, where b_{1k} is the loading on the market factor and δ_i is the

individual's beliefs about market returns. For simplicity, we omit market and time subscripts in this section. Plugging into the first order condition and taking the mean across investors for each k (in a given plan), we have the following:

$$\varsigma_k^2 = -E_k[\theta_{ik}\delta_{ik}]b_{1k} + E_k[\theta_{ik}](p_k - \eta_{ik}).$$

With an appropriate an instrumental variables strategy, we obtain a consistent estimate $\hat{\theta}_k = E_k[\theta_{ik}]$. The market belief we recover is

$$\hat{\delta}_k = \frac{E_k[\theta_{ik}\delta_{ik}]}{\hat{\theta}_k} = \bar{\delta}_k + \frac{Cov_k(\theta_{ik}, \delta_{ik})}{\hat{\theta}_k}.$$

If $Cov(\theta_{ik}, \delta_{ik}) \neq 0$, our estimated market belief will be biased.

To assess the sensitivity to this bias, we use the covariance across plans to approximate within plan covariance $Cov(\theta_{ik}, \delta_{ik})$. Within a plan, we assume that individuals' beliefs are positively correlated.⁴⁰ The across-plan covariance ρ can be expressed as

$$\rho = \frac{1}{n^2} \sum_{ij} Cov(\delta_i, \theta_j) = \frac{1}{n} Cov(\delta_i, \theta_i) + \frac{n-1}{n} \alpha Cov(\delta_i, \theta_i)$$

where n denotes the number of participants in a plan, and $\alpha \in (0, 1)$ governs how correlated beliefs are across individuals within a plan, i.e., $Cov(\delta_i, \theta_j) = \alpha Cov(\delta_i, \theta_i)$ for $i \neq j$. Then our estimate for within-plan correlation of beliefs and idiosyncratic risk aversion is:

$$Cov(\delta_i, \theta_i) = \frac{n}{1 + \alpha(n-1)} \rho$$

For each year, we compute the across-plan covariance ρ and generate adjusted beliefs with $\alpha = 0.5$ and $\alpha = 0.1$. For simplicity, we do not make different adjustments for different plan sizes but instead use $n = 1,261$ based on average number of plan participants. We estimate positive across-plan covariance, consistent with the positive correlation between belief and risk aversion in 5b. Since $\hat{\theta}_k < 0$, we underestimate expected returns.

In Figure A3, we plot our estimated year-by-year expected market returns while adjusting for potential within-plan correlation of beliefs and idiosyncratic risk aversion. Overall, the adjustment yields modest changes to the expected market return, which remains around 10 percent on average.

⁴⁰Our results demonstrating that individuals extrapolate belief from local information and employer performance provide support for positive correlation within a plan.

E.3 Accounting for Labor Income Risk

We also consider the case when investors account for labor income risk. Specifically, we model an investor's labor income risk as an additional asset with a fixed relative weight ϖ (relative to the value of the investor's 401(k) portfolio) and factor loadings b_{wlt} for each factor l . We can then rewrite an investor's first order condition as:

$$\mu_{ikt} - p_{kt} - R_F = \lambda \left(\sum_{l=1}^L b_{klt} \left(b_{wlt} \varpi + \sum_{j=1}^{K_i} b_{jlt} \omega_{ijt} \right) + \sigma_k^2 \omega_{ikt} \right).$$

Rearranging the terms yields:

$$\left(\sum_{l=1}^L b_{klt} \left(\sum_{j=1}^{K_i} b_{jlt} \omega_{ijt} \right) + \sigma_k^2 \omega_{ikt} \right) = \theta p_{kt} + \psi \left(\sum_{l=1}^L b_{klt} b_{wlt} \right) + \epsilon_{mkt}, \quad (20)$$

where the parameter θ is the negative inverse of risk aversion (i.e., $\theta = \frac{-1}{\lambda}$), ϵ_{kt} is equal to average investor beliefs divided by risk aversion (i.e., $\epsilon_{mkt} = (\bar{\mu}_{mkt} - R_F)/\lambda$), and ψ is equal to $-\varpi$.

We estimate the empirical equivalent of eq. (20) as

$$\varsigma_{mkt}^2 = \theta p_{mkt} + \psi \xi_{mkt}^2 + \phi_{mt} + \phi_{j(k)t} + \epsilon_{mkt}, \quad (21)$$

where:

$$\varsigma_{mkt}^2 = \left(\sum_{l=1}^L b_{klt} \left(\sum_{j=1}^{K_m} b_{jlt} \bar{\omega}_{mjt} \right) + \sigma_k^2 \bar{\omega}_{mkt} \right),$$

and

$$\xi_{mkt}^2 = \left(\sum_{l=1}^L b_{klt} b_{wlt} \right).$$

The term ξ_{mkt}^2 captures the additional risk of investing in asset k due to labor income risk.

Appendix Table A10 displays the corresponding estimates. In columns (1) and (2), we measure labor income risk using factor loadings of the sponsor's stock, and so we restrict to plans where the sponsor is publicly traded. We find similar estimate of θ as in our baseline specification in column (1). Our estimates imply that investors behave as if they are risk seeking with respect to their labor income risk ($\psi > 0 \implies \varpi < 0$). In column (2), we include fund-by-year fixed effects, which absorbs the term θ . The object of interest is the parameter $\psi = -\varpi$, where we find again that investors appear to be risk seeking with respect to their labor income. In columns (3) and (4), we include the full sample, where we proxy for the factor loadings for labor income risk using the equity factor loadings corresponding to the industry of the plan sponsor m . We find similar risk seeking patterns. One caveat is that the additional risk due to

labor income $\xi_{mkt}^2 = \left(\sum_{l=1}^L b_{klt} b_{wlt} \right)$ could be correlated with investor beliefs μ , which would make it endogenous in eq. (21). Directly addressing this endogeneity issue is challenging because it requires variation in the additional risk due to labor income that is orthogonal to investor beliefs.

One concern is that stock returns of the company or the industry do not accurately account for labor income variation. Hence, we also construct measures of labor income risks following Catherine et al. (2022). Without access to administrative data such as Swedish Wealth and Income Registry or US Social Security Master Earnings File (MEF), we use the Survey of Income and Program Participation (SIPP) which is a comprehensive survey that tracks the incomes of American individuals and households through several multiyear samples since 1984. We use SIPP samples from 1988 to 2021, which allows us to measure 24 yearly income changes.⁴¹

We measure the covariance between income shock moments and market returns following Catherine et al. (2022) closely, while acknowledging that we are unable to exactly replicate their methodology due to our limited sample size. We first group individuals by 2-digit NAICS codes. Second, we aggregate monthly personal earned income to the annual level, and measure income shocks as the residual from regressing changes in log yearly income on a third-order polynomial of age, 2-digit NAICS fixed effects, and an indicator for whether the individual has some college education. Third, we compute mean, variance, and skewness of income shocks across individuals in each industry and year, and estimate the covariance with contemporaneous and lag annual market return to measure the covariance, countercyclical variance, and cyclical skewness of labor income risk. To be consistent with measures of factor loading in our main analysis, we standardize market returns to have a standard deviation of one. We take the sum of the coefficients on both contemporaneous and lag market returns as measures of covariance between labor income shocks and market returns.

In columns (5)-(8), we construct ξ_{mkt}^2 using the market factor only $\xi_{mkt}^2 = b_{k1t} b_{w1}$, where we multiply covariance, countercyclical variance, and cyclical skewness by the market factor loading of each investment option. We want to test if individuals with high countercyclical labor income risk have low demand for funds with higher market exposures $\psi < 0$. We find risk seeking patterns with covariance and cyclical skewness, and a modest risk averse pattern with countercyclical covariance. The magnitudes of coefficients are much smaller compared to columns (1)-(4), which could be driven by our imprecise measures of labor income risks.

E.4 Calculating Analysts Beliefs about Fund Returns

We calculate analysts' expectations of fund returns based on fund holdings and analysts' price targets for individual equities. To begin, we construct expected returns at the individual share level. This is done by dividing the analyst's 12-month price target by the current share price

⁴¹SIPP samples prior to 1988 have limited industry information. We measure income change within each SIPP multiyear sample but not across samples.

and subtracting one. For each stock, we calculate the mean and standard deviation of all estimates in a given quarter, resulting in a dataset organized by ticker and quarter. To address the presence of outliers in analysts' estimates, we exclude observations where the analyst's implied expected returns exceed 100%.

Next, we merge our ticker-by-quarter level analyst expectations data with quarterly data on equity mutual fund holdings, as reported on Form 12S. We calculate fund-level implied expected returns by taking weighted averages of the ticker-level estimates. The weight on share s for fund k in quarter q is determined by the fraction of k 's holdings invested in share s for which we have analyst estimates in quarter q . We calculate fund-level implied standard deviations of expected returns, assuming that expected returns are independent across fund holdings.

Lastly, we compare our fund-level analyst expectations with our estimates of investor beliefs. Since our investor data is annual and our analyst beliefs data is quarterly, we utilize the analysts' beliefs as of the second quarter, which is roughly halfway through the year. Because many funds in our sample hold bonds and other non-equity securities, and our analyst beliefs data only pertains to equities, we limit our focus to those fund-quarter observations where we have expected returns data for at least 80% of portfolio holdings (weighted by AUM). We remove outliers by excluding observations in both the bottom and top 5% of the distributions of the mean and the standard deviation of fund-level analyst expectations.

F Out-of Sample Predictions

In this Appendix section, we expand the reduced-form regression in Section 4.3.1 to further leverage the mean-variance structure of the portfolio allocation problem. Within a plan, the allocations toward different investment options are not independent, and so we develop the following procedure to consider such interdependence when predicting allocation to out of sample investment options.

Let ω_k denote the allocation to fund k and μ_k the (net of fee) expected return of fund k . Let Σ denote the variance-covariance matrix across assets, and λ the risk aversion. Mean variance optimal allocations are:

$$\omega = \frac{1}{\lambda} \Sigma^{-1} \mu$$

From our estimation, we recovered μ_k for “in-sample” mutual funds, including the idiosyncratic belief component not explained by factors. For “out-of-sample” investment options, we can predict belief based on factor beliefs μ_k^f similar to the regression in Section 4.3.1. We then introduce a regularization procedure to add an idiosyncratic component ϵ_k to the belief of each “out-of-sample” investment option, such that we satisfy the following three constraints.

1. Mean-variance allocations of in-sample funds $\hat{\omega}_k^{in}$ match allocations in the data ω_k^{in}
2. Total mean-variance allocation matches total allocation we observe in the data. This is an adding up constraint, where we also use observed cash allocation.
3. Mean-variance allocation to each asset (both in and out of sample) is between 0 and 1

With this regularization, beliefs for each investment options are as follows:

$$\hat{\mu}_k = \begin{cases} \hat{\mu}_k & \text{if in sample} \\ \mu_k^f + \epsilon_k & \text{if out of sample} \end{cases}$$

where $\epsilon = \operatorname{argmin}_\epsilon \sum_k \epsilon_k^2$

s.t. $\hat{\omega}_k^{in} = \omega_k^{in}$

$$\sum_k \hat{\omega}_k = \sum_k \omega_k$$

$$\hat{\omega}_k \in (0, 1)$$

Then, we predict allocation using $\hat{\omega} = \frac{1}{\lambda} \Sigma^{-1} \hat{\mu}$. Figure A4 plots R^2 computed by regressing out of sample allocation on prediction in the black dashed line. The black line with markers illustrates how R^2 vary with different number of out of sample funds. For plans with a larger

number of out of sample funds, our procedure is less effective. Note that when predicting mean-variance portfolio allocation, the minimum R^2 is not zero. A more relevant benchmark is R^2 from a naive prediction based on 1/n allocations. The gray lines show R^2 by regressing on 1/n allocation to out of sample plans and fund category fixed effects. As a comparison, our estimated beliefs allow us to improve out of sample R^2 by 25 percentage points from 45% to 70%. Again, this exercise suggests that our recovered beliefs are informative of portfolio allocations.

G Stock-Level Counterfactual Allocations

We find in Section 5 that, while there is a large dispersion in beliefs across investors, the welfare costs associated with this dispersion are modest. One reason is that we focus on 401(k) holdings, which consist primarily of well-diversified mutual funds. This investment setting reduces the scope for large welfare losses from suboptimal asset allocations. To illustrate this point, we consider a counterfactual exercise where investors allocate across individual stocks. We compare allocations based on heterogeneous beliefs and common beliefs, and find that the welfare cost of belief distortions in this context are an order of magnitude larger than in 401(k) settings.

Specifically, instead of funds, we assume investors allocated across individual stocks held by mutual funds on their plan menus. To form stock allocation under heterogeneous belief ω^s_{mt} , we use beliefs about stock returns based on the factor structure of stocks and the factor beliefs estimated in Section 4.2.2. We measure beliefs about stock k for investors in plan m at time t as:

$$\mu_{mkt} = \sum_{l=1}^L \delta_{mlt} b_{lkt}.$$

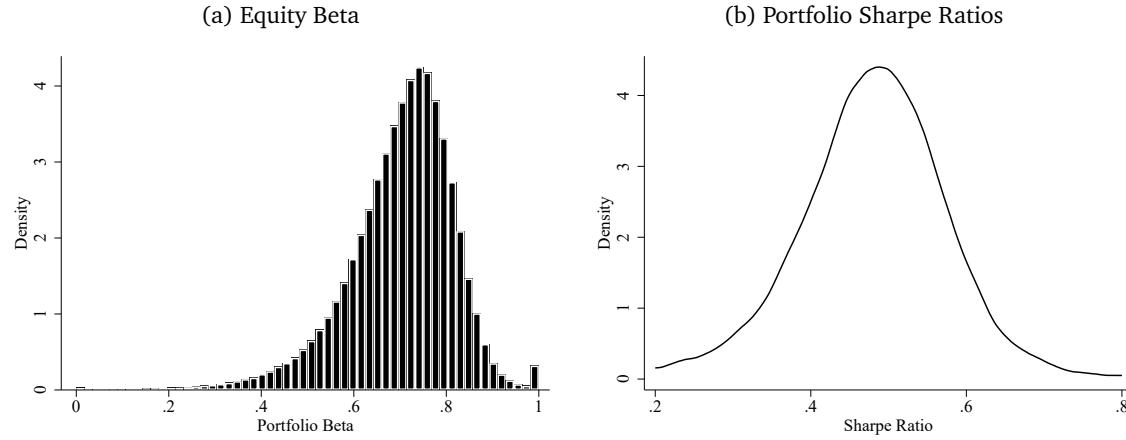
Next, we consider the counterfactual scenario in which all investors share the same beliefs about stock returns, $\bar{\mu}^s_t$, which are again linear combinations of stock factor loadings and estimated common factor beliefs. To estimate common factor beliefs, we regress the common belief $\bar{\mu}_t$ in our fund counterfactual in Section 5 on fund factor loadings. Using these common beliefs, we compute the optimal portfolio $\bar{\omega}^s_{mt}$, and evaluate changes in consumer surplus between ω^s_{mt} and $\bar{\omega}^s_{mt}$ following eq. (10). Similar to the counterfactual exercise in Section 5, we use plan specific risk aversion to compute both ω^s_{mt} and $\bar{\omega}^s_{mt}$. Due to the large number of individual stocks, this counterfactual is more computationally demanding, so we randomly draw 20% of plans in 2016 and restrict to plans with fewer than 1000 unique stocks.

Figure A5 illustrates much larger costs of belief distortion, when investors allocate across individual stocks instead of diversified mutual funds. In this stock counterfactual, the welfare cost of belief distortion is on average 7.5%, with a median of 4.7%. One reason is that individual stocks are further away from the efficient frontier.

H Appendix Tables and Figures

We include Figure A1 and A2 and Tables A11, A14, and A15, which are mentioned in the main text, as well as the figures and tables pertaining to the preceding Appendix sections.

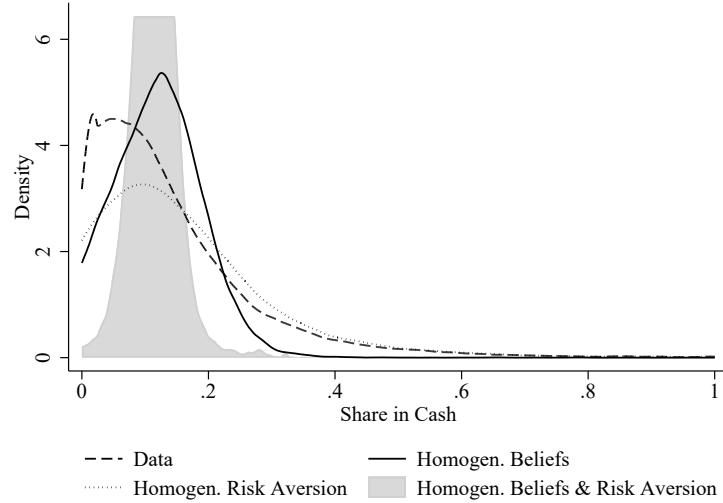
Figure A1: Market Exposure and Sharpe Ratios



Notes: Figure A1a displays the distribution of average equity beta across 401(k) portfolios. Observations are at the plan-by-year level. We compute the average equity beta for a 401(k) plan as the dollar weighted average equity beta across each fund available in the plan. For scaling purposes we truncate the distribution of equity betas at 0 and 1. Figure A1b displays the density of implied Sharpe ratios based on plan-level idiosyncratic expected returns and portfolio allocations in 2016.

Alt Text: Distributions of average equity beta and portfolio sharp ratios at the plan level.

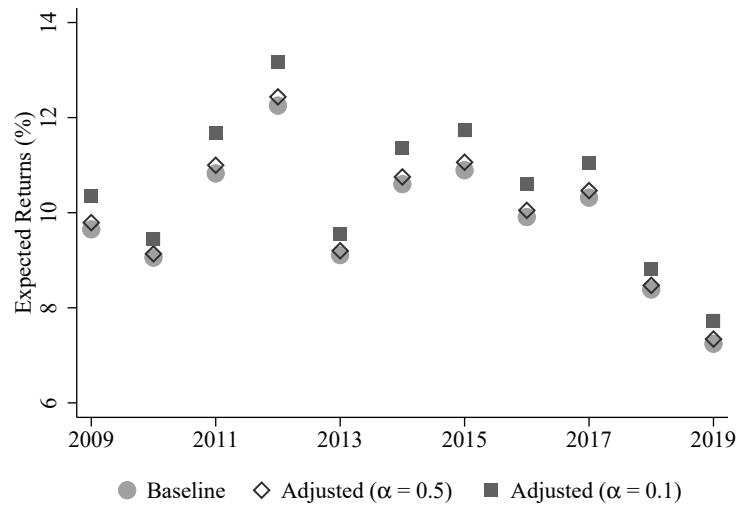
Figure A2: Counterfactual Allocations to Cash



Notes: Figure A2 displays actual and counterfactual densities of cash allocations by plan in 2016. The dashed line indicates the actual distribution of cash allocations across plans. The solid line indicates the counterfactual (optimal) allocations under the assumption that every investor has identical beliefs about each fund. The dotted line indicates the counterfactual allocations when investors have identical risk aversion parameters. The gray shaded area indicates allocations when investors share identical beliefs and risk aversion.

Alt Text: Counterfactual distributions of cash allocations without heterogeneity in beliefs or aversion.

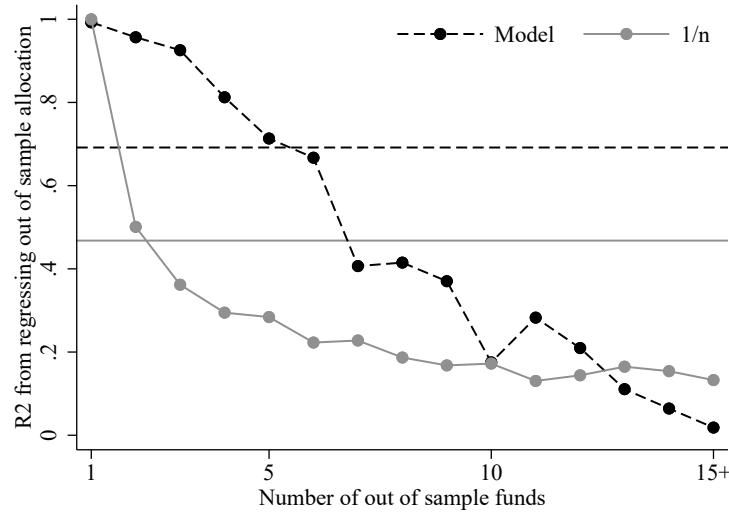
Figure A3: Expected Market Returns with Adjusted Beliefs



Notes: Figure A3 displays our mean expected market returns by year while adjusting for potential within-plan correlation in beliefs and individual risk preferences.

Alt Text: Average expected market returns with adjustments for within-plan correlation between beliefs and risk preferences.

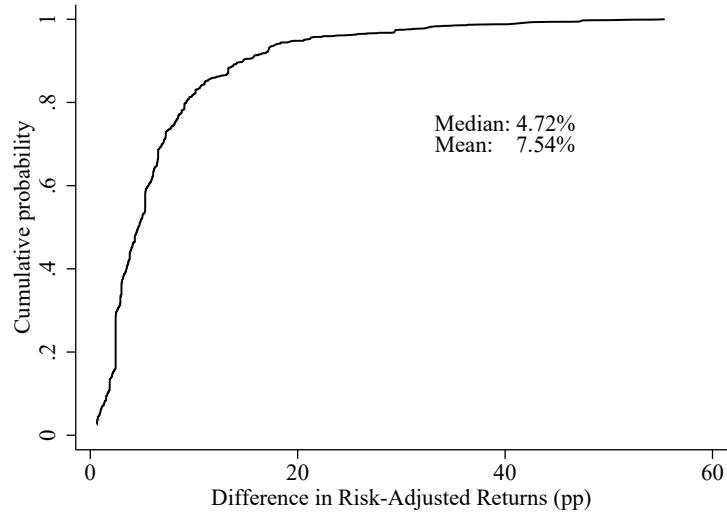
Figure A4: R^2 from Out of Sample Prediction with Regularization



Notes: Figure A4 displays R^2 of out of sample prediction with regularization based on the structure of mean-variance portfolio allocation. Black dashed line shows out of sample R^2 . Black dashed line with markers shows R^2 for plans with different number of out of sample funds. Gray lines show R^2 from $1/n$ allocation and fund category fixed effects.

Alt Text: R^2 of out of sample prediction with regularization based on the structure of mean-variance portfolio allocation.

Figure A5: Stock-Level Counterfactual: Welfare Costs of Belief Distortions



Notes: Figure A5 displays the distribution of $\Delta Consumer Surplus_{mt}$ at the plan level for the year 2016. We construct portfolios of individual stocks instead of mutual funds in 401(k) plans in Figure 12. Due to computation concerns, we randomly draw 20% of plans in 2016 and restrict to plans with fewer than 1000 unique stocks held by their mutual funds. The variable $\Delta Consumer Surplus_{mt}$ is calculated as the difference in consumer surplus if the investor held the optimal stock portfolio (given average beliefs in the economy) relative to the investors chosen portfolio (eq. 10). For reporting purposes, the distribution is censored at 5 percentage points.

Alt Text: Cumulative probability of counterfactual welfare changes in allocations across individual stocks.

Table A1: Estimation with Alternative Instruments

VARIABLES	(1) OLS	(2) Hausman Baseline	(3) Hausman Non 401k Funds	(4) Hausman Separate Accounts	(5) Cost Shifter Custodian Fee	(6) Waldfogel	(7) Provider Switch × Index Funds	(8) Provider Switch & Merger × Index Funds
Fee	-0.049*** (0.009)	-0.193*** (0.017)	-0.198*** (0.017)	-0.232* (0.136)	-0.275* (0.158)	-0.166* (0.095)	-0.407* (0.238)	-0.397* (0.232)
Observations	4,946,470	4,932,059	4,931,931	4,765,243	4,233,776	4,723,597	4,946,470	4,946,470
Plan-Year FE	X	X	X	X	X	X	X	X
Category-Year FE	X	X	X	X	X	X	X	X

Notes: Table A1 displays estimates corresponding to eq. (4). We report OLS estimate in column (1) and two-stage least squares estimates with different instruments in columns (2) to (8). Observations are at the investment option-by-plan-by-year level over the period 2009-2019. Standard errors are clustered at the plan level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Two-stage least squares estimates corresponding to eq. (4) with different instruments.

Table A2: Asset Allocation vs. Demographics

VARIABLES	(1) US Equities	(2) US Equities	(3) Bonds	(4) Bonds	(5) Cash	(6) Cash	(7) Intl. Equities	(8) Intl. Equities
Age	-0.667*** (0.119)	-0.170 (0.127)	0.180*** (0.049)	-0.098 (0.072)	0.802*** (0.102)	0.580*** (0.111)	-0.307*** (0.046)	-0.263*** (0.082)
ln(Income)	0.398*** (0.148)	-0.026 (0.198)	-0.649*** (0.087)	-0.028 (0.105)	0.540*** (0.158)	-0.090 (0.189)	-0.246*** (0.078)	0.108 (0.123)
ln(Home Value)	0.158 (0.115)	0.087 (0.232)	-0.400*** (0.067)	-0.344** (0.151)	0.305** (0.127)	0.345* (0.206)	-0.107 (0.066)	-0.095 (0.150)
College	0.772*** (0.112)	0.843*** (0.193)	0.346*** (0.068)	-0.214 (0.130)	-1.536*** (0.104)	-0.790*** (0.184)	0.275*** (0.064)	0.174 (0.129)
Employed	0.144* (0.080)	0.082 (0.081)	-0.133*** (0.051)	-0.080 (0.054)	0.028 (0.075)	-0.006 (0.078)	-0.007 (0.050)	0.016 (0.054)
Black	-0.174** (0.069)	-0.179 (0.151)	-0.131*** (0.048)	-0.073 (0.101)	0.781*** (0.072)	0.045 (0.146)	-0.412*** (0.044)	0.181* (0.094)
Hispanic	-0.620*** (0.090)	-0.504*** (0.166)	-0.085 (0.061)	-0.049 (0.113)	0.910*** (0.098)	0.584*** (0.160)	-0.271*** (0.062)	-0.017 (0.109)
Unionized	-0.401 (0.245)	-0.618** (0.245)	-0.679*** (0.179)	-0.443** (0.174)	3.705*** (0.275)	3.679*** (0.270)	-2.314*** (0.138)	-2.322*** (0.143)
Sector Equity Beta	0.311*** (0.074)	0.226** (0.096)	-0.233*** (0.050)	-0.081 (0.066)	-0.326*** (0.073)	-0.281*** (0.109)	0.169*** (0.043)	0.127** (0.061)
Share Retired	-0.469*** (0.058)	-0.401*** (0.055)	0.190*** (0.036)	0.135*** (0.034)	0.685*** (0.059)	0.645*** (0.057)	-0.344*** (0.033)	-0.318*** (0.033)
ln(Avg. Acct. Bal.)	1.028*** (0.070)	0.888*** (0.069)	-0.276*** (0.057)	-0.146** (0.057)	0.128** (0.064)	0.105* (0.062)	-0.743*** (0.043)	-0.725*** (0.043)
Observations	243,268	243,166	243,268	243,166	243,268	243,166	243,268	243,166
R-squared	0.110	0.171	0.031	0.099	0.078	0.149	0.032	0.096
Year FE	X	X	X	X	X	X	X	X
Naics FE		X		X		X		X
County FE		X		X		X		X

Notes: Table A2 displays the regression results corresponding to a linear regression model. The dependent variable is the portfolio weight of the corresponding asset class. The independent variables, other than the dummy variable *Union*, are all standardized such that they are in units of standard deviations. Observations are at the plan-by-year level over the period 2009-2019. Standard errors are clustered 2-digit NAICs by county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of portfolio weights across asset classes on demographic variables.

Table A3: 401(k) Menus vs. Demographics

VARIABLES	(1) US Equity Funds	(2) Bond Funds	(3) Cash Funds	(4) Intl. Equity Funds
Age	0.004*** (0.001)	-0.002*** (0.001)	-0.000 (0.000)	-0.001* (0.001)
ln(Income)	-0.001 (0.001)	0.002** (0.001)	-0.001* (0.000)	0.000 (0.001)
ln(Home Value)	0.001 (0.002)	-0.000 (0.001)	0.000 (0.001)	-0.001 (0.001)
College	0.001 (0.002)	-0.002 (0.001)	0.000 (0.001)	0.001 (0.001)
Employed	-0.000 (0.001)	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Black	-0.001 (0.001)	0.000 (0.001)	-0.001 (0.000)	0.001 (0.001)
Hispanic	0.000 (0.001)	-0.000 (0.001)	-0.000 (0.000)	0.001 (0.001)
Unionized	0.010*** (0.002)	-0.004*** (0.001)	0.009*** (0.001)	-0.012*** (0.001)
Sector Equity Beta	0.001* (0.001)	-0.000 (0.001)	-0.001*** (0.000)	0.000 (0.001)
Share Retired	-0.004*** (0.000)	0.001*** (0.000)	0.004*** (0.000)	-0.000 (0.000)
ln(Avg. Acct. Bal.)	0.003*** (0.001)	-0.002*** (0.000)	0.004*** (0.000)	-0.004*** (0.000)
Observations	243,166	243,166	243,166	243,166
R-squared	0.088	0.067	0.075	0.081
Year FE	X	X	X	X
Naics FE	X	X	X	X
County FE	X	X	X	X

Notes: Table A3 displays the regression results corresponding to a linear regression model. Observations are at the plan-by-year level over the period 2009-2019. The dependent variable is the number of funds available in the 401(k) menu in a given asset class (e.g., US equities) divided by the total number of funds available in the 401(k) menu. Standard errors are clustered 2-digit NAICs by county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of asset class menu compositions on demographic variables.

Table A4: Relative Number of Equity Investment Option and Asset Allocation

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Relative No. Equity Options	19.4*** (0.74)	26.0*** (0.74)	26.6*** (0.78)	23.3*** (0.78)	28.3*** (0.77)	29.1*** (0.79)
Offer Company Stock		5.63*** (0.16)	5.40*** (0.18)		5.77*** (0.16)	5.47*** (0.17)
ln(Total Plan Asset)			0.12*** (0.045)			0.20*** (0.048)
Observations	20,199	20,199	20,199	20,197	20,197	20,197
R-squared	0.033	0.090	0.091	0.122	0.176	0.176
NAICS 2 FE				X	X	X

Notes: Table A4 displays regression results of equity allocation on relative number of equity funds. Observations are at plan-by-year level over the period 2009-2019, weighted by total plan asset. We restrict plans whose start dates on 5500 Forms are on or after 2009. The dependent variable is equity allocation, which includes US equity, international equity and 50% of multi-asset funds. Relative No. of equity is computed following Benartzi and Thaler (2001), where each investment option is weighted by how long it has been in the plan and how well it has performed. To measure performance, we use S&P 500 Index as proxy for return on US equity, Barclays Agg Bond Index for bonds, S&P Global BMI for international equity, S&P US Treasury Bill 0-3 Month Index for cash/stable value. We assume return for multi-asset is 50% S&P 500 Index and 50% Barclays Agg Bond Index. For additional controls, we consider an indicator for whether the plan includes company stocks, log of total plan assets, and fixed effects for 2-digit NAICS code of sponsors of the plans. Standard errors are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of equity allocations on relative number of equity funds.

Table A5: Portfolio Allocation vs. Expenses

VARIABLES	(1)	(2)
Expense Ratio	-0.576*** (0.003)	-0.672*** (0.007)
Observations	5,063,093	5,048,630
Plan-Year FE	X	X
Category-Year-Index FE	X	X
IV		X

Notes: Table A5 displays the regression results corresponding to a linear regression model (eq. 15). Observations are at the investment option-by-plan-by-year level over the period 2009-2019 where we exclude target date funds. The dependent variable is the log share of plan assets held in the investment option. Expense ratios are measured in terms of percentage points. We estimate column (2) using two-stage least squares. We instrument for expenses using Hausman-type instruments where we instrument for the expenses for a fund using the average expenses of other funds managed by the same fund manager in different Lipper objective categories. Standard errors are clustered at the plan level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of fund allocations on expenses.

Table A6: Welfare Costs of Belief Distortions vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	-0.022 (0.034)											0.011 (0.029)
ln(Income)		-0.096*** (0.021)										0.003 (0.039)
ln(Home Value)			-0.068*** (0.022)									-0.031 (0.029)
College				-0.123*** (0.034)								-0.021 (0.047)
Employed					-0.079*** (0.020)							-0.026 (0.018)
Black						-0.009 (0.015)						-0.015 (0.021)
Hispanic							0.088*** (0.026)					0.070** (0.025)
Unionized								0.025 (0.060)				-0.025 (0.038)
Sector Equity Beta									0.114*** (0.026)			0.098*** (0.024)
Share Retired										0.196*** (0.031)		0.214*** (0.031)
ln(Avg. Acct. Bal.)											-0.120*** (0.024)	-0.140*** (0.023)
Observations	48,192	48,192	48,192	48,192	48,192	48,192	48,192	48,192	48,192	48,192	48,192	48,192
R-squared	0.000	0.005	0.002	0.008	0.003	0.000	0.004	0.000	0.006	0.019	0.008	0.041

Notes: Table A6 displays the results from linear regressions of log of welfare costs of belief distortions ($\Delta Consumer Surplus_{mt}$ from eq. 10) on standardized demographic variables and a dummy variable for Unionized. Observations are at the plan level for 2016. Standard errors are clustered 2-digit NAICs level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of welfare costs of belief distortions on demographic variables.

Table A7: Expected Returns vs. Past Fund Returns

VARIABLES	(1)	(2)	(3)	(4)
Lag Fund Ret.	0.005*** (0.000)	0.016*** (0.001)	0.007*** (0.002)	0.005*** (0.000)
Lag Fund Ret. x New Investment				-0.000 (0.000)
Observations	4,499,736	672,910	79,041	4,499,736
R-squared	0.937	0.941	0.940	0.937
FE	X	X	X	X
New Funds		X		
New Plans			X	

Notes: Table A7 displays the regression results corresponding to a linear regression model. Observations are at the investment option-by-plan-by-year level over the period 2009 through 2019. The dependent variable is the expected returns of the fund. Standard errors are in parenthesis and are clustered at the plan level. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Regressions of expected returns of individual funds on their past returns.

Table A8: Predictability of Forecast Errors

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Lag Forecast Error	0.035*** (0.001)	0.002 (0.001)				
Lag Fund Ret.			-0.030*** (0.001)	-0.029*** (0.001)		
Change in Beliefs					-0.511*** (0.007)	-0.795*** (0.012)
Observations	2,400,158	2,395,689	4,494,924	4,494,868	2,402,780	2,398,321
R-squared	0.627	0.662	0.616	0.648	0.628	0.664
Year FE	X		X		X	
Plan-Year FE		X		X		X

Notes: Table A8 displays the regression results corresponding to a linear regression model (eq. 18). Observations are at the investment option-by-plan-by-year level over the period 2009 through 2019. Standard errors are clustered at the plan level. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Predictability of forecast errors based on estimated expected returns.

Table A9: Alternative Model Specifications

(a) Risk Aversion and Expected Market Returns

	Obs	Mean	Std. Dev.	Median
Risk Aversion	243,268	3.553	0.472	3.506
Risk Aversion: No Time-Varying Intercept	243,268	3.599	0.286	3.604
Risk Aversion: Including Target Date Funds	243,268	3.554	0.482	3.515
Risk Aversion: New Plans Only	4,772	3.562	0.000	3.562
Risk Aversion: Simplified Risk Measure	243,268	7.630	1.685	7.160
Risk Aversion: 55 Factor Model	243,268	4.073	0.666	3.942
Expected Return	243,268	9.558	2.329	9.469
Expected Return: Time-Varying Intercept	243,268	9.696	2.150	9.766
Expected Return: Including Target Date Funds	243,268	9.729	2.219	9.615
Expected Return: New Plans Only	4,772	9.922	2.255	10.039
Expected Return: Simplified Risk Measure	243,268	14.951	4.068	14.445
Expected Return: 55 Factor Model	243,268	11.030	2.527	10.916

(b) Correlation: Baseline vs. Alternative Specifications

	Expected Return	Risk Aversion
Model: No Time-Varying Intercept	0.883***	0.576***
Model: Including Target Date Funds	0.948***	0.995***
Model: New Plans Only	0.844***	
Model: Simplified Risk Measure	0.700***	0.416***
Model: 55 Factor Model	0.817***	0.552***

Notes: Table A9 displays the results for our alternative model specifications. Panel (a) displays mean, standard deviation, and median of the estimates of risk aversion and beliefs across our model specifications. Column (1) of Panel (b) displays the correlation between the estimated expected returns from the baseline model specification with the estimated expected returns from the other model specifications. Column (2) of Panel (b) displays the correlation between the estimated risk aversion from the baseline model specification with the estimated risk aversion from the other model specifications. Observations in both panels are at the plan-by-year level. *** p<0.01, ** p<0.05, * p<0.10.

Alt Text: Comparison of estimated risk aversion and expected market returns across alternative model specifications.

Table A10: Model Estimates Accounting for Labor Income Risk (θ and ψ)

VARIABLES	(1) Sponsor Risk	(2) Sponsor Risk	(3) Industry Risk	(4) Industry Risk	(5) Income Risk	(6) Income Risk	(7) Income Risk	(8) Income Risk
θ	-0.162*** (0.050)		-0.104*** (0.011)		-0.176*** (0.017)	-0.191*** (0.018)	-0.191*** (0.017)	-0.159*** (0.017)
ψ Sponsor	0.184*** (0.014)	0.054*** (0.003)						
ψ Industry			0.329*** (0.022)	0.096 (0.078)				
ψ Covariance					0.071*** (0.007)			0.113*** (0.014)
ψ Countercyclical Variance						-0.011 (0.011)		-0.037*** (0.005)
ψ Cyclical Skewness							0.007*** (0.002)	-0.000 (0.002)
Observations	259,459	334,772	4,727,392	5,956,422	4,914,370	4,914,370	4,914,370	4,914,370
Plan-Year FE	X	X	X	X	X	X	X	X
Category-Year-Index FE	X		X		X	X	X	X
Fund-Year FE		X		X				

Notes: Table A10 displays the regression results corresponding to a linear regression model (eq. 21). Observations are at the investment option-by-plan-by-year level over the period 2009-2019. In columns (1), (3), and (5)-(8), we estimate using two-stage least squares, and we instrument for expenses using Hausman-type as described in the text. In columns (2) and (4), we run standard OLS. Standard errors are clustered at the plan level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. **Alt Text:** Two-stage least squares estimates of risk aversion accounting for labor income risk.

Table A11: 401(k) Participation vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	0.009*** (0.003)											0.008*** (0.002)
ln(Income)		0.017*** (0.003)										-0.002 (0.003)
ln(Home Value)			0.030*** (0.004)									0.011** (0.004)
College				0.032*** (0.004)								0.014*** (0.004)
Employed					0.007*** (0.002)							0.006*** (0.002)
Black						-0.006** (0.003)						0.005 (0.003)
Hispanic							-0.011*** (0.004)					0.001 (0.004)
Unionized								0.017*** (0.005)				0.018*** (0.005)
Sector Equity Beta									0.012*** (0.002)			0.006*** (0.002)
Share Retired										0.011*** (0.001)		-0.001 (0.001)
ln(Avg. Acct. Bal.)											0.083*** (0.001)	0.082*** (0.001)
Observations	242,378	242,378	242,378	242,378	242,378	242,378	242,378	242,378	242,378	242,378	242,378	242,378
R-squared	0.181	0.181	0.181	0.182	0.181	0.181	0.181	0.181	0.181	0.182	0.264	0.265
Year FE	X	X	X	X	X	X	X	X	X	X	X	X
Naics FE	X	X	X	X	X	X	X	X	X	X	X	X
County FE	X	X	X	X	X	X	X	X	X	X	X	X

Notes: Table A11 displays the regression results corresponding to a linear regression model. Observations are at the plan-by-year level over the period 2009-2019. The dependent variable is fraction of eligible employees that participate in 401(k) plans. Standard errors are clustered 2-digit NAICs by county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. **Alt Text:** Regressions of 401(k) plan participation on demographic variables.

Table A12: Expected Size Factor Returns vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	0.037*** (0.013)											0.067*** (0.008)
ln(Income)		0.045** (0.017)										-0.099*** (0.013)
ln(Home Value)			0.041*** (0.010)									0.041*** (0.008)
College				0.094*** (0.010)								0.137*** (0.015)
Employed					0.029* (0.015)							-0.014** (0.007)
Black						0.007 (0.012)						0.005 (0.004)
Hispanic							-0.061*** (0.011)					-0.038*** (0.008)
Unionized								-0.016 (0.032)				0.023 (0.018)
Sector Equity Beta									-0.004 (0.021)			0.024* (0.012)
Share Retired										-0.038*** (0.005)		-0.039*** (0.005)
ln(Avg. Acct. Bal.)											0.017** (0.008)	0.005 (0.007)
Observations	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268
R-squared	0.106	0.108	0.107	0.124	0.104	0.103	0.111	0.103	0.103	0.106	0.103	0.139
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Notes: Table A12 displays the results from linear regressions of expected size factor returns on standardized demographic variables and a dummy variable for *Unionized*. Observations are at the investment plan-by-year level over the period 2009 through 2019. Standard errors are clustered 2-digit NAICs level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. Alt Text: Regressions of expected size factor returns on standardized demographic variables.

Table A13: Expected Value Factor Returns vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	0.056*										0.101***	
	(0.027)										(0.012)	
ln(Income)		0.058***									-0.163***	
		(0.018)									(0.016)	
ln(Home Value)			0.043**								0.044***	
			(0.016)								(0.012)	
College				0.142***							0.224***	
				(0.011)							(0.014)	
Employed					0.043**						-0.016**	
					(0.019)						(0.007)	
Black						0.019					0.011*	
						(0.014)					(0.006)	
Hispanic							-0.100***				-0.059***	
							(0.014)				(0.005)	
Unionized								-0.065*			-0.012	
								(0.037)			(0.025)	
Sector Equity Beta									-0.005		0.041**	
									(0.034)		(0.015)	
Share Retired										-0.032**	-0.034***	
										(0.011)	(0.009)	
ln(Avg. Acct. Bal.)											0.030**	0.009
											(0.012)	(0.014)
Observations	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268	243,268
R-squared	0.328	0.328	0.327	0.340	0.327	0.326	0.332	0.326	0.326	0.326	0.326	0.350
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Notes: Table A13 displays the results from linear regressions of expected value factor returns on standardized demographic variables and a dummy variable for *Unionized*. Observations are at the investment plan-by-year level over the period 2009 through 2019. Standard errors are clustered 2-digit NAICs level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10. Alt Text: Regressions of expected value factor returns on standardized demographic variables.

Table A14: Stock Market Exposure vs. Local Economic Conditions

VARIABLES	(1)	(2)	(3)	(4)	(5)
Pop. Growth	-0.098** (0.047)				-0.129** (0.051)
Home Price Growth		0.018* (0.009)			0.014 (0.011)
Establishment Growth			0.021 (0.019)		0.016 (0.019)
GDP Growth				0.012* (0.007)	0.013** (0.007)
Observations	407,714	425,206	431,589	424,859	394,986
R-squared	0.767	0.763	0.762	0.762	0.767
Year FE	X	X	X	X	X
Plan FE	X	X	X	X	X

Notes: Table A14 displays the regression results corresponding to a linear regression model. Observations are at the investment plan-by-year level over the period 2009 through 2019. The dependent variable is the share in equities. Standard errors are in parenthesis and are clustered at the county-by-year level. *** p<0.01, ** p<0.05, * p<0.10.
Alt Text: Regressions of equity allocations on local economic conditions.

Table A15: Stock Market Exposure vs. Employer Returns

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Firm Return (1 years)	0.004*** (0.001)	0.009*** (0.002)							0.005*** (0.001)	0.009*** (0.002)
Firm Investment			-0.001 (0.012)	0.042** (0.019)					0.007 (0.014)	0.038* (0.021)
Sales Growth					-0.000 (0.002)	0.012*** (0.004)			-0.001 (0.003)	0.005 (0.004)
Employment Growth							0.003 (0.003)	0.011* (0.006)	0.002 (0.004)	0.004 (0.006)
Observations	20,155	20,315	17,931	18,067	19,755	19,891	19,711	19,852	17,403	17,521
R-squared	0.803	0.131	0.777	0.134	0.805	0.130	0.805	0.129	0.781	0.133
Year FE	X		X		X		X		X	
Plan FE	X		X		X		X		X	
NAICS-Year FE		X		X		X		X		X

Notes: Table A15 displays the regression results corresponding to a linear regression model. Observations are at the investment plan-by-year level over the period 2009 through 2019. The dependent variable is the share in equities. Standard errors are in parenthesis and are clustered at the plan level. *** p<0.01, ** p<0.05, * p<0.10.
Alt Text: Regressions of equity allocations on employer performance.