NBER WORKING PAPER SERIES

WHAT DRIVES VARIATION IN INVESTOR PORTFOLIOS? EVIDENCE FROM RETIREMENT PLANS

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

We thank John Campbell, Xavier Gabaix, Sam Hanson, Andrei Shleifer, Adi Sunderam and the seminar participants at Harvard Business School, Indiana University, Oklahoma State University, and the University of Toulouse. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

What Drives Variation in Investor Portfolios? Evidence from Retirement Plans
Mark L. Egan, Alexander MacKay, and Hanbin Yang
NBER Working Paper No. 29604
December 2021
JEL No. G0,G11,G12,G40,G5,G51,J32

ABSTRACT

We study empirical patterns in investment behavior using a comprehensive data set of defined contribution plans. Using plan-level portfolio allocation data for the near universe of 401(k) plans over the period 2009-2019, we document substantial differences in investment behavior across plans. Plans with wealthier and more educated participants tend to have higher equity exposure while plans with more retirees and minorities tend to have lower equity exposure. These patterns cannot be explained by differences in 401(k) menus or participation costs. To help interpret these facts, we use a revealed preference approach to estimate investors' expectations of stock market returns and risk aversion, where we allow investors to have heterogeneous risk aversion and subjective and potentially biased beliefs. We find that there is substantial variation in both beliefs and risk aversion across investors and over time, and that both sources of variation help explain investors' portfolio decisions. We also provide new evidence to understand how investors form beliefs. We find that investors extrapolate beliefs from past fund returns even when they initially allocate portfolios in new plans. We also find that investors extrapolate beliefs about the market from the past performance of their employer, which suggests that investor experience helps shape beliefs.

Mark L. Egan
Harvard Business School
Baker Library 365
Soldiers Field
Boston, MA 02163
and NBER
megan@hbs.edu

Hanbin Yang
Harvard University
hanbin.v.yang@gmail.com

Alexander MacKay
Morgan 217
Harvard Business School
Soldiers Field
Boston, MA 02163
amackay@hbs.edu

1 Introduction

Defined contribution plans account for the bulk of equity participation in the US and roughly one third of retirement assets.[1] 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). Defined contribution plans typically provide participants with a fixed menu of investment options, often mutual funds, each with its own fee. This unique structure, where investors choose from a limited menu of investment options with different costs, provides insight into how investors form portfolios in terms of their risk preferences, beliefs, and biases.

We study empirical patterns in investment behavior using a comprehensive data set of retirement plans offered by firms to their employees. Our data includes plan-level portfolio allocations for the near universe of 401(k) plans from 2009 through 2019. We document high 401(k) participation rates—74 percent on average—and substantial differences in investment behavior across plans. For instance, consider the share of retirement assets allocated to US equity funds. The allocation is 44 percent on average but ranges widely across plans. The 10th percentile of plans has 17 percent of assets in US equities, compared to 64 percent for the 90th percentile.

We examine the determinants of heterogeneity in equity allocation. We show that systematic heterogeneity in equity allocation can be explained by demographic factors such as income, age, race, and housing wealth. Investment patterns differ greatly for employees in different sectors, but there is also substantial within-sector heterogeneity. While much of the previous literature on 401(k)s has focused on plan design and has shown how plan design can have a substantial impact on 401(k) participation, the differences in allocations across 401(k) plans we observe are quite large and are not explained by differences in menus or naive investment strategies. Instead, the patterns suggest that investors make conscious choices based on risk and fund fees. To understand heterogeneity in investment choices, we develop a model of portfolio allocation to recover investors' expected returns and risk aversion that rationalize the investment patterns we observe, conditional on demographic factors and the menu offered in each plan. We study how beliefs and risk aversion depend systematically on observable characteristics, how beliefs evolve over time, and the potential biases in beliefs.

Our data on defined contribution 401(k) plan allocations comes from BrightScope Beacon. As of 2021, Americans held roughly $7 trillion in 401(k) assets.[2] BrightScope Beacon provides detailed annual plan and fund level information for ERISA defined contribution plans, covering 97 percent of plan filings and 98 percent of plan assets. Specifically, BrightScope Beacon provides details on the menu and plan-level aggregate fund allocations for 70,000 different 401(k) plans over the period 2009-2019. The typical 401(k) plan offers 26 different investment

[1] https://www.ici.org/system/files/2021-06/21_rpt_recsurveyq1.pdf
[2] https://www.ici.org/faqs/faq/401k/faqs_401k

options. Thus, our data set has 450,000 plan-by-year observations and 11 million fund-by-plan-by-year observations. We also observe details regarding the plan sponsor (e.g., employer) and each investment option/fund available in the 401(k) plans. We supplement the investment option data using data from CRSP mutual fund that provides detailed historical return and fee data for the investment options available in 401(k) plans. We also match our 401(k) data with American Community Survey (ACS) data at the industry-by-county-by-year level to proxy for participant demographics.

We document substantial heterogeneity in allocations across plans and find that plan allocations are highly correlated with participant demographics. 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. These differences in portfolio allocations we document cannot be explained by either the composition of the menu—plan menus are largely uncorrelated with participant demographics—or participation costs, which have been emphasized in prior research. We examine allocation decisions conditional on participation.

Building on the previous literature, we examine how fees and the composition of the 401(k) investment menu impact portfolio allocations. Investors disproportionately choose funds with lower fees. We estimate that a 10 basis point (bp) increase in fund expense ratios is associated with a 5.9% decrease in demand, suggesting that fees play an important role in allocation decisions. Previous work indicates that the choices of investors are driven by the menu of funds (Benartzi and Thaler, 2001, 2007). We find that, while the menu of funds is slightly correlated with investor choices, the relationship is substantially weaker than previously documented, which is consistent with the evidence in Huberman and Jiang (2006). One potential reason for this finding is that 401(k) sponsors have substantially increased the number of investment options available 401(k) menus over the past 30 years, which has made the menus themselves less important. We also find variation in 401(k) holdings over time: adjusting for returns, the one-year autocorrelation in fund holdings is 0.69-0.88, which indicates that some investors actively rebalance their portfolios over time.

Overall, the empirical evidence suggests that there is substantial variation in portfolio holdings across plans, that this variation is correlated with participant demographics, and that little of this variation is explained by differences in menu composition. Consequently, we instead focus on how differences in risk aversion and beliefs explain variation in holdings.

To interpret the decisions of investors, we model an investor's portfolio decision as a mean-variance optimization problem. Investors with heterogeneous risk aversion and subjective and potentially biased beliefs select portfolios to maximize returns net of risk. When forming portfolios, investors trade off their subjective expectations with the corresponding additional risk, where the risk of the asset is weighted by the investor's risk aversion. To separately identify expectations and preferences for risk, we exploit variation in the expenses associated with funds available in an investor's 401(k) menu. Variation in expenses allows us to quantify how investors trade off changes in fees, which shift expected returns, with risk. The key identifying assumption is that the variation in fund expenses is orthogonal to investors' expected returns of that fund. We use Hausman-type instruments to help ensure the variation in expenses we exploit is orthogonal to investor beliefs (Hausman, 1996). The implied beliefs and risk aversion allow us to study investment behavior independent of the funds and fees offered by different plans.

We use the model to estimate the time-varying distributions of risk aversion and beliefs across investors using our plan-level data from BrightScope. In the data, we observe the average portfolio allocations for each 401(k) plan, which allows us to recover the average expected returns across plan participants for each investment option available in the 401(k) menu. For example, in 2019, the average participant in the IBM sponsored IBM 401(k) Plus Plan held 4% of their portfolio in the Vanguard Russell 1000 Value Index and 2% in the Vanguard European Stock Index. Using our framework, we can then separately recover the average IBM 401(k) participant's expectations about the return of both the Vanguard Russell 1000 Value Index and the Vanguard European Stock Index as of 2019. Thus, we recover average investor expectations within a 401(k) plan at the investment option-by-401(k) plan-by-year level.

We recover reasonable time-varying distributions of both risk aversion and beliefs that are consistent with previous research and realized returns. In our baseline specification, we estimate an average constant relative risk aversion parameter of 5.6, which is comparable to what other researchers have found in the literature.[3] The average investor in our sample behaved as if she expected the market return to be 11.5% over the period 2009-2019. To put this in perspective, the compound annual growth rate of the S&P 500 was 11.2% over the same period. We find that accounting for heterogeneity in both risk aversion and beliefs is important for fitting the investment patterns we find in the data. Our simple two parameter model with risk aversion and beliefs explains more than 50% of the variation in equity holdings across plans and indicates that both risk aversion and beliefs play important roles in explaining an investor's overall equity exposure. As expected, investors with more optimistic beliefs about the market have higher equity (and lower cash) exposure, while more risk averse investors have lower equity (and higher cash) exposure. We find very similar estimates of risk aversion and expected returns if we estimate the model using data from newly introduced plans in the year of inception and exclude default options, such that the allocation decisions reflect the active choices of plan participants.

Next we try to understand the drivers of the cross-sectional variation in beliefs and risk aversion. We find that more educated investors tend to have more optimistic market expectations. Conversely, older and minority investors tend to have more pessimistic market expectations.

[3] For 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, and Meeuwis (2019) estimate relative risk aversion of 5.4.

We also find that an investor's beliefs are correlated with their work experience. For example, investors working in the real estate sector are 34% (3.4 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.

We also find that risk aversion varies with observable investor characteristics. 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 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. There is a long literature documenting that investors extrapolate their beliefs across a number of settings.[4] We first find that investors extrapolate their beliefs from fund past returns. One feature of our data is that there is a fair amount of turnover in 401(k) plans and 401(k) menus. We take advantage of such turnover to show that investors also extrapolate from past returns when they initially form portfolios when a 401(k) plan was introduced, and so the extrapolation cannot be explained by inattention or inertia in rebalancing. Investors potentially form their beliefs based on the past returns reported in 401(k) plan brochures.

Next, we show how investors' personal experience influences their beliefs. For the subset of publicly traded employers, we examine how investors' expectations of the market return change in response to the performance of their employer. We find that investors' expectations, averaged at the plan level, are positively correlated with the past performance of their employer, as measured by returns, investment, employment growth, and sales growth, suggesting that investors become more optimistic about the market when their employer is doing well. We find that this result holds when comparing an investor's expectations relative to other investors employed in the same sector at the same time (i.e., industry-by-year fixed effects). This suggests that an investor's expectations reflect their personal experience as has been shown in other settings (e.g., Malmendier and Nagel, 2011, 2015). Our findings also suggest that investors potentially form broader conclusions about the market based on the recent performance of their firm.

Given these factors driving the formation of beliefs, we next explore the rationality of investor expectations. We find that investor forecast errors are predictable and consequently violate full information rational expectations, similar to the evidence documented in the literature.[5] Consistent with the evidence presented in Coibion and Gorodnichenko (2015) and

[4]For 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).

[5]For example, see Bacchetta et al. (2009); Coibion and Gorodnichenko (2012, 2015); Amromin and Sharpe

Bordalo et al. (2018), we find evidence that investor forecast revisions, measured as changes in beliefs, are also correlated with forecast errors. The relationship between changes in forecasts and forecast errors is negative, which suggests that investors overreact to news.

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

Related Literature

Our paper relates closely to the literature on household finance. There exists a long theoretical literature dating back to Markowitz (1952) and Merton (1969) on how households should invest.[6] There exists a parallel literature examining how investment allocations vary with investor characteristics. Consistent with the previous literature, both in terms of empirical findings and theoretical predictions, we document 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); and is negatively correlated with age (Cocco et al., 2005) and minority status (Campbell, 2006). We build on the literature by documenting how portfolio allocations vary across 401(k) plans. 401(k) plans provide a unique setting to study household portfolio choices because we observe the suite of investments available to participants and we can use variation in plan design/expenses to separately identify investor beliefs and risk aversion. Our empirical model provides insight into how differences in risk aversion and beliefs contribute to the patterns we see in the data.

To this end, our paper also relates on the literature using a revealed-preference approach to estimating beliefs and risk aversion across investors. Our methodology relates most closely to Egan et al. (forthcoming) who build and estimate a structural model to recover investor expectations and risk aversion using data from leveraged exchanged traded funds. We employ a similar identification strategy, building on the work of Barseghyan et al. (2013), where we use exogenous variation in expenses to separately identify beliefs and risk aversion. Our framework also builds on the work in Shumway et al. (2009) who use a revealed-preference approach to understand fund manager beliefs. A key distinction between our work and that of Shumway et al. (2009) is that we focus on separately identifying risk aversion and beliefs of retail investors. In contrast, Shumway et al. (2009) do not separately identify risk aversion and beliefs and instead recover beliefs of institutional investors up to an affine transformation that

(2014); Greenwood and Shleifer (2014); Gennaioli et al. (2016); Bordalo et al. (2019) among others.

[6] See Campbell et al. (2002), Campbell (2006), Gomes et al. (2020), and Cochrane (2021) for an a discussion of the literature.

is scaled by risk aversion and translated by an investors' shadow value of borrowing. This type of revealed preference approach to asset pricing, by focusing on quantities rather than prices or returns, is also conceptually related to Berk and van Binsbergen (2016), Koijen and Yogo (2019a), and Heipertz et al. (2019). In particular, Koijen and Yogo (2019a) develop a demand system asset pricing model where investors have heterogeneous preferences, which has been extended to study exchange rates (Koijen and Yogo, 2019b), cryptocurrencies (Benetton and Compiani, 2021), bonds (Bretscher et al., 2020), and global equities (Koijen et al., 2019).

The two main methods used to study investor beliefs use either survey data or asset prices to measure beliefs. One strand of literature uses survey data to measure investor beliefs (Vissing-Jorgensen, 2003; Ben-David et al., 2013; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014 Nagel and Xu, 2019). While some researchers have been skeptical of surveys, recent evidence suggests they produce valuable information (Greenwood and Shleifer, 2014; Giglio et al., 2019). Our paper is motivated by Giglio et al. (2019), who document substantial and persistent heterogeneity in beliefs across retail investors. Using novel survey and account level data from Vanguard, Giglio et al. (2019) find that beliefs are reflected in the portfolios of investors, especially when investors are attentive, are actively trading, and hold tax-advantaged accounts. When looking at tax-advantaged accounts, Giglio et al. (2019) 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 baseline estimates imply a 1pp increase in beliefs about the stock market returns is correlated with a 3.25pp increase in equity share.[7] Despite using completely different samples and methodologies, we find estimates that are in line with Giglio et al. (2019). Our approach provides insight into both investor risk aversion and beliefs when such survey data is unavailable.

The other strand of literature uses data on asset prices to recover investor beliefs (Ross, 2015).[8] While that literature uses data on asset prices to recover the distribution of beliefs of a single representative investor, we use data on investment flows to recover the distribution of beliefs as well as 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.

Lastly, our paper relates to the literature on retirement savings (see Benartzi and Thaler, 2007 and Choi, 2015 for a discussion of the literature). A 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); and

[7]The results in column (3) of Table 11 indicate that a 1 standard deviation increase in beliefs (2.51pp) is associated with a 8.167pp increase in equity share.

[8]Other recent examples include Jensen et al. (2019), Martin and Ross (2019), and d'Arienzo (2020). There is also a related strand of literature that focuses on robust identification of investor beliefs (Chen et al., 2020; Ghosh and Roussellet, 2020; Ghosh et al., 2020).

Carroll et al., 2009) 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 (83% at the median plan) in our sample. Another strand of literature focuses on menu design and fees (Pool et al., 2016; Pool et al., 2020; Bhattacharya and Illanes, 2021). By contrast, we focus on the asset allocation decisions conditional on both participation and the 401(k) menu. Bekaert et al. (2017) document how both menu design and investor characteristics are related with international equity exposure in 401(k) plans. Previous work (Benartzi and Thaler, 2001, 2007) emphasizes the importance of behavioral frictions in 401(k) asset allocation decisions. In a similar theme, we find that the beliefs that rationalize investor choices violate full information rational expectations and that investors extrapolate their beliefs.[9] A growing body of research documents that such adaptive expectations could have significant implications for the macroeconomy and financial markets (Bordalo et al., 2018; Gennaioli and Shleifer, 2018; Bordalo et al., 2018; Malmendier et al., 2020).

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 97% of plan filings and 98% 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 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

[9] Previous research such as Bacchetta et al. (2009); Coibion and Gorodnichenko (2012, 2015); Amromin and Sharpe (2014); Greenwood and Shleifer (2014); Gennaioli et al. (2016); Bordalo et al. (2019) among others have found that beliefs violate full information rational expectations. See Malmendier (2021) for an overview of the literature on experience effects in finance.

contributions, and the share of participants that are retired.

Lastly, 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

We start by documenting substantial heterogeneity in 401(k) holdings across plans and over time. We group investment options into six major asset classes: US equities, bonds, cash, target date funds, alternatives, and international equities. One minor complication in computing equity exposure is that some funds invest across asset classes (i.e., allocation funds such as Bridgewater's All Weather fund). When calculating US equity and bond shares, we assume that non-target-date allocation funds hold sixty percent of their assets in US equities and forty percent in bonds.

Figure 1 displays the portfolio weights for six major asset classes across plan-by-year observations. The average plan holds 44% of the 401(k) assets in US equities,[10] 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 (Figure 1a). 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%.

Figure 2 displays the average portfolio weights for each of the major asset classes over time. In panels (a) and (b) we compute portfolio weights both excluding and including target date funds because target date funds tend to be the default option in most plans following the Pension Protection Act of 2006. One can see that there is substantial variation in the average holdings over time. Around the time of the financial crisis, investors increased the weight held in cash and bonds at the expense of US equities and international assets. Another key trend in the industry has been the rise of allocation assets, which are primarily comprised of target date funds, consistent with the evidence documented in Parker et al. (2020).[11]

Table 1 displays the summary statistics corresponding to our baseline database from BrightScope. Panel 1a displays plan level summary statistics. 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 64% 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.

[10] Excluding multi-asset/allocation funds, the average is 39%.

[11] Appendix Figure A1 shows a version where we do not attribute non target date allocation funds to US equity and bond assets. The sharp rise of allocation assets is mostly driven by target date funds, and the trends of equity and bond are similar when we do not consider non-target-date allocations.

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, which is consistent with estimates from the Survey of Consumer Finances.[12] The high participation rates are a relatively new phenomenon in the US. For example, in 1988 only 57% of eligible employees participated in 401(k) plans (Choi, 2015). Participation rates remained relatively high and constant over our sample period of 2009-2019.[13] The high participation rates may be a direct result of the earlier research, such as Madrian and Shea (2001), Choi et al. (2002), Choi et al. (2007), and Beshears et al. (2009), which emphasize how automatic enrollment increases 401(k) participation. While there has been concern about the lack of retirement savings in the US, these summary statistics suggest that the low retirement savings rates are driven by 401(k) plan eligibility rather than 401(k) plan participation.

We also observe detailed information on each investment option. BrightScope Beacon classifies each investment option into eight different types of investment vehicles. The vast majority of investment options are structured either as mutual funds (61%) or separate accounts (19%).

Table 1b displays investment option level summary statistics. BrightScope Beacon provides the latest expense ratios for each investment option as of 2019, and we are able to obtain historical expense ratio data for those investment options structured as mutual funds using data from CRSP. As of 2019 the equal weighted average expense ratio was 57 bps.

Overall, the summary statistics presented in Figure 1 and Table 1 indicate that there is substantial variation in plan characteristics and holdings. One important caveat is that, because our data is aggregated at the plan level rather than at the individual level and because each plan has 1,261 members on average, our summary statistics may understate the true heterogeneity in individual holdings. In the remainder of the paper we explore the drivers of portfolio heterogeneity

3 How Do Asset Allocations Vary Across Investors?

The simple summary statistics presented in Section 2 indicate that there is substantial heterogeneity in asset allocations across plans. In this section we explore what drives the heterogeneity in asset allocation across plans. We find that asset allocations are highly correlated with participant demographics and employment. We also explore how much of the heterogeneity is driven by participant allocation decisions versus heterogeneity in 401(k) menus. We find that, while features of investment menus, such as the number of equity funds available, are correlated with investment decisions, differences in investment menus across plans do not explain the facts we document about the investment allocation decisions of investors.

[12] https://crr.bc.edu/wp-content/uploads/2020/10/IB_20-14.pdf

[13] Appendix Figure A2 displays participation and employer contribution rates over time. In Appendix Table A1, 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.

We also examine other elements of the investment allocation process. We find evidence that investors appear to make at least partially informed decisions when selecting investment options. Investment decisions are sensitive to the expense ratios and investors appear to rebalance their 401(k) portfolios over time. We also investigate whether investors use naive "1/N" diversification strategies by equally distributing their portfolio across all investment options. This type of behavior was first documented by Benartzi and Thaler (2001). Consistent with Benartzi and Thaler (2001), we find some evidence that investment allocations are correlated with the composition of the menu; however, a naive strategy only explains a small fraction of the variation in holdings and fails to explain the main patterns we observe in the data.

3.1 Asset Allocation and Investor Characteristics

There exists a long literature in household finance examining asset allocation. Previous research has focused on the effects of age, income, wealth, education, and employment on asset allocation decisions, among others. In this section we examine how 401(k) portfolio allocation decisions vary across participant characteristics previously emphasized in both the theoretical and empirical literature. One unique feature of our setting is that all plan participants have access to all investment options available in the 401(k) menu. Thus, the patterns in allocation decisions we document are not driven by participation costs, which has been emphasized in other settings.[14] In the second half of the paper, we build and estimate a structural model of portfolio choice to understand how systematic differences in risk aversion and beliefs explain the facts we document in this section.

We examine how investment allocations vary across investor demographics in the following regression:

$$Share\,in\,US\,Equities_{kt} = X_{kt}\beta + \mu_t + \epsilon_{kt}. \tag{1}$$

Observations are at the 401(k) plan-by-year level. The dependent variable $Share\,in\,US\,Equities_{kt}$ reflects the share of assets held in equities in plan k 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.[15]

We consider demographics, industry, and plan variables in X_{kt}. First, we consider data on employment and demographics. Our employment and demographic data comes from the American Community Survey (ACS) and is measured at the county-by-industry-by-year level. We merge the ACS data with our 401(k) data based on the year, sponsor/employers industry (i.e., 2 digit NAICS), and county headquarters. Note that because we do not perfectly observe participant demographics, this may introduce measurement error in our demographic covari-

[14]See Campbell (2006) and Gomes et al. (2020) for a discussion of the literature.

[15]As mentioned in Section 2.2, we assume that non-target-date allocation funds hold sixty percent of their assets in US equities. Our main findings are robust under other assumptions such as if we include target date funds, exclude all allocation funds, etc.

ates and could attenuate some of our results. Following the literature we focus on age, income, housing wealth, and race.

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 Table 2. We include time fixed effects in each specification. Columns (1)-(11) display univariate regressions, and the specification reported in column (12) includes the full set of controls. For ease of interpretation the independent variables are in units of standard deviation. For example, the results in column (12) indicate that a one standard deviation increase in the $\ln(Average\ Account\ Balance)$ is correlated with a 1.04 pp increase in allocation to US equity.

Income and Wealth: Equity exposure is positively correlated with income, home value, and financial wealth (measured by average account balance). The results in column (2) indicate that a one standard deviation increase in log income is associated with a 1.34 pp increase in allocation to US equity, although the effect become smaller and insignificant once we include other controls due to multicollinearity in column (12). The existing theoretical predictions regarding income and equity allocation is mixed. For example, Cocco et al. (2005) shows how income is analogous to a safe asset, and hence is positively correlated investment in risky equity; however, other theoretical works highlight how income risk can also crowd out equity allocation (Lynch and Tan, 2011; Storesletten et al., 2004).

We also find that plans with larger average account balances tend to tilt their portfolios towards equities. The results in column (11) indicate that a one standard deviation increase in the average account balance is correlated with a 1.29 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. While our results are consistent with those findings, our results provide some additional evidence on the mechanism. Because we are looking at 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; instead, we observe variation in the portion allocated to equity funds by those that participate in 401(k) plans.

We find a similar positive relationship between home wealth and equity exposure, although the effect becomes insignificant in column (12). Similar to income, the existing theoretical predictions regarding home value and equity exposure are mixed.[16]

[16]For example, 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

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 0.69 and 0.49 pp decline in US equity holdings, respectively (column 12). 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; thus, younger investors allocate more to risky equity. 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.[17]

Other Demographics: We also find that more educated households have higher equity allocation, and that education explains more variation in equity allocation than any of the other demographic factors. The results in column (4) indicate that a one standard deviation increase in the share of college educated individuals is correlated with a 1.55 pp increase in equity allocation, consistent with the findings in Campbell (2006) and Black et al. (2018). This relationship 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 0.86 and 0.39 percentage point decreases in equity exposure. Campbell (2006) and Chiteji and Stafford (2000) also find that minorities have lower equity shares. In Section 5 we explore whether these differences in equity exposure across minorities and non-minorities are potentially driven by differences in beliefs or risk aversion.

Employment: The connection between labor income and equity allocation is more nuanced once we consider income risk. If income is highly correlated with equity, income risk will crowd out financial asset risk leading to lower exposure to risky assets. Measuring income risk by the equity beta corresponding to the sector of the plan sponsor, we find a modest positive but insignificant effect on allocation to US equity. One explanation is that the actual correlation between income and equity is low for the average household. (Campbell et al., 2009; Davis and Willen, 2013)

Table 3 displays the distribution of equity exposure by the 2-digit NAICs of the employer. The results indicate that average equity exposure varies dramatically across sectors, ranging from 51.8 percent in Educational Services to 64.2 percent in Information. Such variation could potentially be consistent with background risk. Households with higher undiversified labor

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.

[17] Empirical estimates tend to be mixed due to the identification challenge of collinearity among cohort, time and age effect. Using Norwegian administrative data, Fagereng et al. (2017) find that risky asset share of stock market participants is a decreasing function of age. However, Samwick and Poterba (1997) and Ameriks and Zeldes (2004) find evidence of hump-shaped patterns based on US data.

risks may effectively be more risk averse and should invest more cautiously (Heaton and Lucas, 2000; Viceira, 2001). In addition, in sectors with more flexible labor conditions, households can adjust labor supply in response to investment returns, and thus increase willingness to take financial risk (Bodie et al., 1992, Farhi and Panageas, 2007). 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 Educational Services sector would be substantially riskier than employment in the Information sector. Instead, some of the differences across sectors may be explained by differences in risk aversion and beliefs, in addition to underlying risk. Our results in Table 3 are probably best explained by a mixture of these factors. We delve into these sector differences in equity exposure further in Section 5 to understand if they can be explained by differences in beliefs and/or risk aversion.

Other Asset Classes: The differences in holdings across plans extends to other asset classes as well. Table 4 displays the regression results where we replicate eq. (1) for the other main asset classes. The dependent variable in column (1) is the portfolio share in US equities, in column (2) is the share in bonds, in column (3) is the share in cash, and in column (4) is the share in international equities. A couple of interesting patterns emerge in Table 4. 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.[18]

3.2 Understanding How Investors Form Portfolios

There exists a long theoretical literature illustrating how rational investors should be investing their portfolios dating back to Merton (1969) and a corresponding empirical literature documenting that portfolio theory often fails to match how households invest in practice (Benartzi and Thaler, 2007; Cochrane, 2021). Here, we explore which factors appear to drive investor portfolio decisions. Based on the previous literature, we focus on how investors form portfolios based on expenses (Hortaçsu and Syverson, 2004) and the composition of the menu (Benartzi

[18] In Appendix Table A3, we replicate Table 4 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. We also show in Appendix Table A2 that the menus themselves are largely uncorrelated with participant demographics.

and Thaler, 2001). We also explore the rebalancing behavior of investors.

Our results suggest that investor decisions appear at least partially informed and attentive: investors are sensitive to fees and appear to rebalance their portfolios over time. While we find evidence suggesting that the composition of the 401(k) menu is correlated with investment decisions, the evidence is weaker than what has been documented previously in the literature.[19] These facts are important for motivating the empirical model we build and estimate in Sections 4 and 5. In particular, an investor's sensitivity to fees is a key moment for separately identifying risk aversion from beliefs in our quantitative model.

3.2.1 Responding to Fees

In any portfolio choice model, investors trade off risk with expected returns. Measuring how investors respond to exogenous changes in fund expense ratios provides insight into this trade-off as expense ratios directly impact the expected returns of the fund. This also provides some insight into the optimality of an investor's investment decisions.

We start with a simple cut of the data by looking at the equal-weighted distribution of fund expenses relative to the asset weighted distribution of fund expenses in Figure 3. Panels (a) and (b) show a stark contrast between the equal weighted and asset weighted distributions of expenses. The asset weighted distribution is shifted dramatically to the left relative to the equal weighted distribution. The average fund appearing on an investor's 401(k) menu charges an expense ratio of 57 bps; however, the average expense ratio paid by investors is 26 bps. This is driven in part because investors tend to tilt their portfolio allocations towards inexpensive funds and provides prima facie evidence that investor demand is sensitive to fees.

We examine the relationship more formally in the following demand specification:

$$\ln Share_{kmt} = \alpha p_{kt} + Fixed\,Effects + \xi_{kmt}. \tag{2}$$

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 k at time t. Fund expense p_{kt} is the independent variable of interest. We include plan-by-year, investment category-by-year, and index fund-by-year fixed effects. 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.[20] While we present eq. (2) as a simple linear specification, by including plan–by-year fixed

[19]Prior work suggests that some investors may follow naive diversification strategies by allocating their portfolio equally across the funds in their 401(k) menu, which suggests that the composition of the 401(k) menu has a large impact on investment allocations.

[20]For example, when designing 401(k) menus, larger plans may have access to funds with lower expense ratios or may be able to negotiate lower expense ratios.

effects, eq. (2) directly corresponds to the workhorse discrete-choice demand model developed in Berry (1994) that is commonly used in the industrial organization literature.[21]

One concern with estimating demand is that fund expenses are potentially endogenous. For example, if investors are particularly optimistic about a fund (e.g., high ξ_{kmt}) the fund provider may find it optimal to increase the fund expense ratio. This type of endogeneity would typically bias our estimate of α upwards (i.e., α is less negative) such that investors appear less sensitive than they actually are. To account for the potential endogeneity of fees, we instrument for fees using Hausman-type instruments. Specifically, we use the average fee charged by the mutual fund provider in other Lipper objective investment categories in the same year. We report our demand estimates in Table 5. 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 increase in fees is associated with a 5.9% 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 -0.35.[22] Collectively, the results suggest that investors respond to expenses, which is perhaps not surprising given that expenses are often a salient feature when investors make 401(k) decisions.

3.2.2 Comparing Allocations to Naive Diversification

Previous results in the literature have suggested that some investors follow naive diversification strategies in which they simply split their allocation evenly across all of the options in their retirement plans (Benartzi and Thaler, 2001). Investor behavior along these lines would enable employers to adjust the share of equity allocation by simply increasing the relative number of equity options in the retirement plan.

Using our data, we consider an implication of these naive diversification strategies for aggregate investment trends. Holding fixed the balance in each plan in each year, we simulate the

[21]Following 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. 2 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; Egan, Lewellen, and Sunderam, 2017; Wang, Whited, Wu, and Xiao, 2018), bonds (Egan, 2019), credit default swaps (Du et al., 2019), insurance (Koijen and Yogo, 2016, 2018), mortgages (Benetton, 2018) and investments more generally (Koijen and Yogo, 2019a,b; Koijen et al., 2019).

[22]We compute the demand elasticity assuming a market share of 1/26 and fee of 0.61%.

counterfactual holdings if all investors simply allocated their funds evenly across all funds in the menu and, alternatively, evenly across the eight categories of funds in our data (e.g., Bond Funds, Cash / Stable Value, International Stock, US Large Cap Stock, etc.).

Figure 4(a) shows the allocation to US equities in our data and the counterfactual "One Over N" naive investment strategies. From 2009 through 2019, excluding target date funds, there has been a steady increase in the share of retirement assets allocated to US equities, as shown by the solid line. Conversely, if investors were simply allocating funds evenly across funds (dashed line) or fund categories (dash-dotted line), US equities would have declined as a share of assets. Over time, US equity funds are making up a smaller share of fund choices, even as they constitute a greater share of retirement assets. In Figure 4(b) and (c), we compare allocations to US equities with the predicted allocations based on the 1/N strategy. Panel (b) displays the observed distribution of equity allocations with those predicted by the 1/N strategy. The results suggest that the observed dispersion in equity allocations is much more disperse than what would be implied if all investors used a 1/N strategy. Furthermore, panel (c) indicates that the 1/N strategy explains very little of the variation in equity allocations.

We also replicate the baseline analysis presented in Benartzi and Thaler (2001). Appendix Table A4 shows that in our setting the menu composition is correlated with investment decisions, though we find a weaker relationship. Using the same specification, we obtain R^2 ranging from 0.03 to 0.18 whereas R^2 ranges from 0.25 to 0.62 in Benartzi and Thaler (2001). This weaker relationship with naive investment strategies is consistent with the findings in Huberman and Jiang (2006), where "the available fund mix and number of funds offered hardly explains participants' choices of funds." 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 its possible that evenly splitting allocations across all investment options is more challenging when the investor faces a broader menu.

3.2.3 Rebalancing Behavior

Lastly, we examine investor rebalancing behavior. Figure 2 displays average holdings over time. Investors' equity exposure is slightly decreasing over this period. Given that the S&P 500 Index increased almost 200% during our sample period, the fact that the investors' shares in US equities did not increase dramatically over the same period suggests that investors must be rebalancing their portfolios over this period. Survey evidence shows that 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.[23]

[23]See 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.

To understand investor's rebalancing behaviors more systematically, we calculate the autocorrelation in plan holdings in Table 6. Specifically, we calculate the variable $Expected\ Portfolio\ Weight_{ijt}$ which assumes that the portfolio weight of fund i grows by the return of fund i relative to the total return of the 401(k) portfolio over the same period. The results in column (1) indicate that the correlation between $Expected\,Portfolio\,Weight_{ijt}$ and $Portfolio\,Weight_{jt}$ is 0.69. Part of the reason investors rebalance is because 401(k) menus turnover quite frequently. If we restrict our attention to those plans that have been outstanding for at least one year, roughly 30% of 401(k) investment options were not available in the previous year. If the investment option is not available in the previous year, $Expected\,Portfolio\,Weight_{ijt}$ is zero by construction. Turnover among 401(k) menus and providers induces participants to rebalance their portfolios. In columns (2) we replicate our analysis where we restrict our attention to only those investment options that were available in the previous year. Not surprisingly, the autocorrelation in holdings is higher if we exclude changes in the investment menu. Overall, the results suggest that, while there is persistence in portfolios, there is also variation in investor portfolios over time. Since fund fees are relatively persistent in the data and investors' beliefs might also be persistent, one might naturally expect there to be a large amount of persistence in portfolios over time even if investors are actively rebalancing their portfolios.

Overall, our results suggest that investors' portfolio allocations appear to be at least partially informed, in the sense that investors respond to exogenous changes in fees and rebalance their portfolios over time. This suggests that investors put some thought into their 401(k) allocation decisions and that these decisions reflect some information about investor expectations and risk aversion. In the next section, we develop and estimate a model of portfolio choice that allows us to interpret the facts documented in Sections 3.1 and 3.2 in terms of differences in investor beliefs and risk aversion. An important point is that we do not impose rationality of beliefs in the model. In fact, as documented in Section 6, we find evidence that investor expectations violate full information rational expectations, that investors over-react, and that investors are subject to misatributtion bias.

4 Model

Motivated by the above findings, we model each investor's 401(k) portfolio allocation as a mean-variance decision problem where each investor trades off her subjective and potentially biased expectation of the return of investing an additional dollar in one of the investment options available in her 401(k) plan 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. As described further below, exogenous variation in the fees investors pay for different investment options in their 401(k) menus allows us to separately identify beliefs from risk aversion. Intuitively, the variation in fees allows us to measure how investors trade off expected returns versus risk, which directly translates to risk aversion.

We also use our estimates of beliefs and risk aversion to better understand the portfolio allocations of investors. Analyzing portfolio allocations, without our structural framework, 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 identical sets of investors faced different plans menus, they may have different portfolio allocations. With sufficient variation in funds within a menu, our methodology would recover the same set of beliefs despite the different observed allocations.

Our quantitative model also provides insight into our previous results from Section 3.1 where we document how equity exposure varies with age, race, wealth, and across industries. For example, we find that educated individuals tilt their portfolios towards stocks because they have more optimistic beliefs about stock returns. Conversely, differences in risk aversion, rather than beliefs, helps explain why wealthier individuals hold more equities.

4.1 Investor's Problem

Each investor i must form portfolios from the set of securities $k = 1, ...K_i$ and a risk-free asset. We assume investors have mean-variance preferences with risk aversion λ. Investors choose the $K \times 1$ vector of weights $\boldsymbol{\omega}_i$ to maximize

$$\max_{\omega} \boldsymbol{\omega}_i'(\boldsymbol{\mu}_i - \boldsymbol{p}) + (1 - \boldsymbol{\omega}_i'\mathbf{1})R_F - \frac{\lambda}{2}\boldsymbol{\omega}_i'\Sigma\boldsymbol{\omega}_i,$$

where $\boldsymbol{\mu}_i$ is a vector of investor i's expectations of fund returns, $\boldsymbol{p}$ is a vector of fund expenses, R_F is the risk-free return, Σ is the $K \times K$ covariance matrix of expected fund returns, and λ is risk aversion. The corresponding first order condition is

$$\boldsymbol{\mu}_i - \boldsymbol{p} - \mathbf{1}R_F = \lambda\Sigma\boldsymbol{\omega}_i.$$

We then have $\kappa_i \in \{1, ..., K_i\}$ first order conditions for every investor.

4.2 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

$$\Sigma_t = \boldsymbol{b}_t I_L \boldsymbol{b}_t' + \boldsymbol{\sigma_t} I_K \boldsymbol{\sigma_t}'.$$

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 security k as

$$\mu_{ikt} - p_{kt} - R_F = \lambda \left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K} b_{jlt}\omega_{ijt} \right) + \sigma_k^2 \omega_{ikt} \right). \tag{3}$$

The term on the left hand side reflects the expected return net of fees of 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 security k.

In the data, we do not observe each investor i's portfolio but instead observe the aggregated portfolio for all investors participating in the same defined contribution retirement plan m. 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. 3) across all individuals participating in defined contribution plan m as

$$\left(\frac{1}{\sum_{i \in \mathcal{I}_m} A_i} \right) \sum_{i \in \mathcal{I}_m} A_i \left(\mu_{ikt} - p_{kt} - R_F \right) = \lambda \left(\frac{1}{\sum_{i \in \mathcal{I}_m} A_i} \right) \sum_{i \in \mathcal{I}_m} A_i \left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K} b_{jlt}\omega_{ijt} \right) + \sigma_k^2 \omega_{ikt} \right).$$

This simplifies to

$$\bar{\mu}_{kt}^{(m)} - p_{kt} - R_F = \lambda \left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K} b_{jlt}\bar{\omega}_{jt}^{(m)} \right) + \sigma_k^2 \bar{\omega}_{kt}^{(m)} \right),$$

where $\bar{\mu}_{kt}^{(m)}$ 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 $\bar{\omega}_{kt}^{(m)}$ is the average (dollar-weighted) portfolio weight. We assume all individuals in the same plan have the same risk aversion λ.

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

$$\left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K} b_{jlt} \bar{\omega}_{jt}^{(m)} \right) + \sigma_k^2 \bar{\omega}_{kt}^{(m)} \right) = \theta p_{kt} + \epsilon_{kt} \tag{4}$$

where the parameter θ is the negative inverse of risk aversion (i.e., $\theta = \frac{-1}{\lambda}$) and ϵ_{kt} is equal to average investor beliefs divided by risk aversion (i.e., $\epsilon_{kt} = (\bar{\mu}_{kt}^{(m)} - R_F)/\lambda$). Eq. (4) 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_{kt}). Given exogenous variation in fees, we are able to recover the parameter θ and consequently risk aversion λ. And given risk aversion, we can recover average beliefs as $\lambda \epsilon_{kt} = \bar{\mu}_{kt}^{(m)} + R_F$.

4.3 Implementation

4.3.1 Risk

To estimate risk aversion and recover investor beliefs we need to estimate the factor structure of fund returns ($\boldsymbol{b}_t, \boldsymbol{\sigma}_t$). We estimate the factor structure using a 55-factor model following Shumway et al. (2009). Our factors include the Fama-French 5 factors, momentum, and 49 industry portfolios. For empirical convenience, we orthognalize and standardize the factors such that each factor has unit variance. To aid interpretation, the first factor is the market factor and the other factors are orthogonalized relative to the first factor.

In our data, we observe daily returns for mutual funds and stocks, which comprise roughly two-thirds of the investment options. For these investment options we can directly estimate the factor loadings using returns data. For the non-mutual fund and stock investment options, such as separate accounts, we do not observe high-frequency data. However, we do observe their investment category classification as per Morningstar and BrightScope. We use these classifications to determine the risk associated with these investment options.

Specifically, we use the following methodology to calculate risk for each of the investment options. We first estimate the factor loadings for each mutual fund and equity in CRSP using weekly return 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. For those investment options where we do not observe a ticker, we calculate the risk associated with the investment option based on the average risk of all other funds that belong to the same Morningstar category in the same year.[24]

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 1a, with the idea that investors think of risk in terms of broad asset classes

[24]For a handful of options we do not observe the Morningstar category. 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.

(e.g., bonds, international stocks, cash, etc.). We find estimates of beliefs and risk aversion using our alternative methodology that are highly correlated with our baseline estimates. We provide comparison statistics in Table A8. In Appendix A.2, we also explore the case where investors account for labor income risk and find that investors behave as if they are not averse to labor income related risks.

4.3.2 Expenses

We determine fund expenses/fees 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 the parameter θ in eq. (4).

To help address this concern, we include plan-by-year fixed effects and fund classification-by-year[25] fixed effects in our main empirical specification. Thus, we allow for the fact that fees may rise endogenously in response to 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. For example, suppose that (i) Fidelity anticipates that participants in IBM's 401(k) plan have relatively optimistic beliefs about Fidelity's Large Cap Growth Index Fund relative to the other investment options in IBM's 401(k) plan (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, 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 for all 401(k) plans helps alleviate these endogeneity concerns.[26]

Nevertheless, to account for the potential endogeneity of fees, we instrument for fees using Hausman-type instruments as used in Section 3.2.1. Specifically, we use the average fee charged by the same mutual fund provider in other Lipper investment objective categories. 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

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

[26]While mutual fund expenses are set uniformly, the expenses for other types of investment options, such as separate accounts, could vary across 401(k) plans. We focus on mutual funds in our empirical analysis because we only observe historical expense ratios for investment options structured as mutual funds.

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.

4.3.3 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 A8, we find qualitatively similar estimates if we include target date funds in our analysis.

4.4 Estimation

We estimate the empirical analog of the investor's first order conditions for choosing an optimal portfolio (eq. 4) in the following regression specification:

$$\varsigma^2_{mkt} = \theta p_{mkt} + \phi_{mt} + \phi_{j(k)t} + \epsilon_{mkt}, \tag{5}$$

where

$$\varsigma^2_{mkt} = \left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K_i} b_{jlt} \bar{\omega}^{(m)}_{jt} \right) + \sigma^2_k \bar{\omega}^{(m)}_{kt} \right)$$

and ϕ_{mt} and $\phi_{j(k)t}$ are plan-by-year and fund type-by-year fixed effects. Here, we introduce subscript m to denote 401(k) plans, and $j(k)$ to denote fund type based on the fund's classification in both Morningstar and BrightScope as well as whether the fund is an index fund. 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. (5). Our estimates allow us to recover risk aversion, as $\widehat{\lambda} = -\frac{1}{\hat{\theta}}$.

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, our estimated fixed effects, and the residual from eq. (5):

$$\widehat{\bar{\mu}_{kt}^{(m)} - R_F} = -\frac{1}{\hat{\theta}}\left(\hat{\phi}_{mt} + \hat{\phi}_{j(k)t} + \hat{\epsilon}_{mkt}\right). \tag{6}$$

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:

$$\widehat{\bar{\mu}_{kt}^{(m)} - R_F} = \delta_{mt} b_{1kt} + \eta_{mkt}, \tag{7}$$

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. (7). We estimate eq. (7) to recover each investor's belief, averaged at the plan level, about the stock market.

4.5 Identification and Interpretation

We estimate risk aversion by measuring how investors trade off risk and expected returns in eq. (5). Specifically, we estimate risk aversion by examining how investors' portfolios change, in terms of marginal changes in risk, in response to plausibly exogenous changes in expense ratios, which shift expected returns. Given risk aversion and the marginal risk of an investment, we calculate investors' expected returns such that the marginal risk, scaled by risk aversion, is equal to the expected return. Thus, we assume that (i) investors only trade off risk and expected returns when making portfolio decisions, and (ii) we, as the econometrician, can correctly measure investor's beliefs about risk. It is important to emphasize that we do not impose rational beliefs in our analysis and our framework allows for behavioral biases and mistakes in investor beliefs. As we show later in Section 6, investor beliefs violate full information rational expectations, are extrapolative, and suffer from misattribution bias. While we believe our model provides an intuitive and straightforward way to estimate risk aversion and beliefs, there are a couple of elements of our model that merit further discussion. Furthermore, it is useful to understand how the interpretation of our estimates would change if assumptions (i) and (ii) were violated.

First, we assume that investors understand and agree and the risk of their portfolio, and that we, as the econometrician, assess risk in the same way. If investors have heterogeneous beliefs about risk or use a different model for assessing risk this could introduce measurement error into the dependent variable ς^2_{mkt}. Suppose we observe a noisy measure of risk $\tilde{\varsigma}^2_{mkt} = \varsigma^2_{mkt} + \varepsilon_{mkt}$, where ε_{mkt} is the measurement error. Provided that the measurement error ε_{mkt} is orthogonal to our Hausman instrument, our instrumental variable estimate of risk aversion will still be consistent even if the measurement error is not idiosyncratic (e.g., $Cov(\varsigma^2_{mkt}, \varepsilon_{mkt}) \neq 0$). While this does not impact our measurement of risk aversion, it will impact the beliefs we recover in the data. Rather than recovering beliefs $\bar{\mu}^{(m)}_{kt} - R_F$, we will recover beliefs plus the measurement error in risk, $\bar{\mu}^{(m)}_{kt} - R_F + \varepsilon_{mkt}$.

Since we average investor portfolios at the plan level, any measurement error that is mean zero and is uncorrelated across investors within a plan will not affect our results. However, if the measurement error is either not mean zero or correlated across investors within a plan, the estimates of beliefs we recover will potentially be biased. How such bias would affect our interpretation of beliefs depends on the specific exercise we perform in the following sections. For example, when using eq. (7) to recover beliefs about the stock market, we would require ε_{mkt} to be independent from loading on market factor b_{1kt}. If investors perceive investment options with high market beta to be riskier compared to our dependent variable ς^2_{mkt}, then our recovered stock market beliefs would include a term corresponding to $Cov(b_{1kt}, \varepsilon_{mkt}) > 0$ and overestimate the true stock market beliefs.

Second, we assume that investors actively trade-off and equate marginal risk and return when making investment decisions. 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 $\boldsymbol{\zeta_i}$:

$$\lambda \Sigma \boldsymbol{\omega}_i = \boldsymbol{\mu}_i - \boldsymbol{p} - \mathbf{1} R_F + \boldsymbol{\zeta_i}. \tag{8}$$

One could interpret this optimization error as either a true error term or it could be capturing 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 if we were to observe a noisy measure of risk. Our estimate of risk aversion would still be consistent, but our estimates of beliefs would reflect this preference (and potentially be biased) if this optimization error was either not mean zero and/or correlated across investors within a plan.

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 portfolio every period (Gabaix, 2019). Generally speaking this would result in our estimate of risk aversion being biased upwards because investors will appear as if they are insensitive to expected returns/fees. In other words, it appears as if investors are unwilling to take on additional risk after an increase in fund expected returns.[27] If we were to systematically over-estimate risk aversion, this would result in us also over-estimating investor optimism

[27]Consider the a simple example where all investors in plan have the same beliefs at any given moment time, but only a fraction π of investors are attentive and update their portfolios. Also for convenience, assume that the factor loadings of the funds do not change over time such that $\Sigma = \Sigma_t \forall t$. Let $\boldsymbol{\omega}_t$ denote the optimal portfolio weights given that an investor updates her portfolio at time t:

regarding fund returns because investors equate expected returns scaled by risk aversion with risk.

To help address the potential concern regarding inattentive investors and its potential impact on risk aversion, 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 by definition reflects an active choice of the participant. We discuss this robustness check in Section 5 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 have a huge impact on our estimate of risk aversion. We also find similar estimates of beliefs when we restrict our sample to the first year each 401(k) plan is introduced.

5 Estimates of Risk Aversion and Beliefs

Here we present our baseline estimates of risk aversion and beliefs and examine how they vary across investor demographics and characteristics. As documented in Section 3.1, we find substantial heterogeneity in investment portfolios across investors and that this heterogeneity is highly correlated with investor demographics. We use our model estimates to further understand why portfolios differ across investors and how much is driven by differences in investors' beliefs versus risk aversion.

$$\lambda \Sigma \boldsymbol{\omega}_t = \boldsymbol{\mu}_t - \boldsymbol{p_t} - \mathbf{1} R_{F,t}.$$

What we observe in the data not the optimal portfolio weights at time t but rather some weighted function of the current and past optimal weights $\bar{\boldsymbol{\omega}}_t = \pi \boldsymbol{\omega}_t + \sum_{l=1}^{\infty} \pi(1-\pi)^l \boldsymbol{\omega}_{t-l}$. We can rewrite our estimation equation as:

$$\Sigma \bar{\boldsymbol{\omega}}_t = \frac{1}{\lambda}\left[\pi\left(\boldsymbol{\mu}_t - \boldsymbol{p}_t - \mathbf{1}R_{F,t}\right) + \sum_{l=1}^{\infty} \pi(1-\pi)^l \left(\boldsymbol{\mu}_{t-l} - \boldsymbol{p}_{t-l} - \mathbf{1}R_{F,t-l}\right)\right].$$

When we regress the term $\Sigma\bar{\omega}_t$ on the fund expense ratios $\boldsymbol{p}_t$ using two-stage least squares with our instrument $\boldsymbol{Z}_t$, our estimate $\hat{\frac{1}{\lambda}}$ will converge to

$$plim\, \hat{\frac{1}{\lambda}} = \frac{1}{\lambda}\left[\sum_{l=0}^{\infty} \pi(1-\pi)^l \frac{Cov(Z_t, p_{t-1})}{Cov(Z_t, p_t)}\right],$$

where we have assumed that our instrument Z_t is orthogonal to past changes in beliefs. If we further assume that $p_t = \phi p_{t-1} + \epsilon_t$ where $Cov(\epsilon_t, Z_t) = 0$ and $-1 < \phi < 1$, we can show that

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

5.1 Risk Aversion

We report our baseline model estimates corresponding to eq. (5) in Table 7. In the model reported in column (1), we keep the parameter θ and consequently risk aversion fixed across 401(k) plans. In columns (2) and (3) we allow θ and risk aversion to vary across plans based on plan characteristics/demographics. We estimate an average risk aversion parameter of 5.3-5.6 across the three specifications, which is in line with what other researchers have found in the literature as discussed in the Introduction.

We find that accounting for heterogeneity in risk aversion, as discussed further below, is important for explaining investment decisions. The interaction terms in Table 7 indicate how demographics are correlated with the parameter θ. Recall that the parameter θ corresponds to the inverse of risk aversion ($\theta = -\frac{1}{\lambda}$). For ease of interpretation, we report the corresponding marginal effects of demographics on risk aversion in Table 8. We find evidence that older plan participants behave as if they are more risk averse. The results in column (2) of Table 8 indicate that a one standard deviation increase in age is associated with a 0.47 (9%) 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 0.58 (11%) increase in risk aversion. Wealthier investors 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 (9%) decrease in risk aversion, although insignificant. Lastly, we find mixed evidence that risk aversion is correlated with race. We find some evidence that plans with a greater share of black participants appear slightly more risk averse while plans with a greater share of Hispanic investors appear slightly less risk averse.

Lastly, in columns (3) of Tables 7 and 8 we allow risk aversion to vary in the year the 401(k) plan was first introduced. As discussed in Section 4.5, if investors are inattentive then they may appear more risk averse in the data than they actually are. Consistent with this, we find that investors behave as if their risk aversion is 0.97 (18%) lower in the year of inception. Consequently, when constructing our estimates of risk aversion and beliefs in the remainder of our analysis, we set the variable $New\, 401(k)\, Plan$ equal to one to account for inattention with the idea that this reflects the true preferences of investors. In Appendix Table A8, 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.

Figure 5 displays the 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 distribution of risk aversion looks fairly constant over time and the time series variation in risk aversion comes from time series variation in plan characteristics. In Appendix Table A8, we also allow the average level of risk aversion to vary year-to-year (i.e., interact fees with time dummy variables in eq. (5)); however, the results suggests that the average level of risk aversion does not vary significantly year-to year. Consistent with the estimates reported in Tables 7 and 8, Figure 5 illustrates that there is substantial 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 5.3.

5.2 Investor Beliefs

Figure 6 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 Column 3 in Tables 7 and 8). The bright red solid line displays the average belief across plans over time. The results suggest that investors became more optimistic coming out of the financial crisis as the mean expected return increased from roughly 10% to 12%. Investors remained optimistic through 2018, and then the average expected return fell below to below 10% in 2019. The average expected return over our sample is 11.50%, which is remarkably close to the realized return of the S&P 500 over this period. The compound annual growth rate (CAGR) of the S&P 500 over the period 2009-2019 was 11.22%.

There is also substantial heterogeneity in beliefs across plans. In Figure 6, 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 2012, the 10th percentile expected return is 9 pp and the 90th percentile is 14 pp. The standard deviation in expected market returns across plans within a year is 2.9 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 7 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 2.6 pp, which is close to the plan-year standard deviation of 2.9. Thus, our estimates imply that relative pessimism and relative optimism about market returns are persistent features of retirement plans.

Note that the our analysis examines the cross-sectional dispersion in the average plan beliefs, where each plan is a collection of individuals. Given that median plan has more than 200 participants, and to the extent that there is variation in investor beliefs within plans, the dispersion shown in Figure 6 and Figure 7 could understate the individual-level dispersion in beliefs by an order of magnitude.

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 (corr=0.45) 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. In other words, we examine how the covariates

X_{mt} explain variation in beliefs that is orthogonal to risk aversion in the following regression:

$$\delta^*_{mt} = X'_{mt}\Gamma + \nu_{mt}. \tag{9}$$

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.[28] We control for the same set of industry and plan characteristics as in our previous analysis in Section 3.1.

Table 9 displays the estimates corresponding to eq. (9). We include year fixed effects in each specification. Columns (1)-(11) display univariate regressions and column (12) includes the full set of control 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 ln(Income) is associated with a 0.19 pp increase in expected market returns. Similarly, a one standard deviation increase in the fraction of college educated individuals is associated with a 0.27 pp increase in expected market returns.

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 0.09pp (0.09 pp) decrease in expected returns. These differences in market expectations could be driven by differences in 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 Table 10. The results suggest that there is substantial heterogeneity across sectors. Investors from the most optimistic sector, Real Estate, expect the market return to be 40% higher than investors from the least optimistic sector, Accommodation and Food Services. It is interesting to examine how both beliefs and equity allocations vary across sectors by comparing Tables 3 and 10. For example, investors in the Information Sector have the highest equity allocations but are only in the 60th percentile in terms of investor expectations of market returns. This suggests that the investor risk aversion, in addition to beliefs, plays an important role why investors in the Information sector have high equity exposure. We also find evidence that there is substantial heterogeneity in beliefs within a sector. The average interquartile range of beliefs within a sector is 3.2 pp. In other words, within a sector those investors in the 75th percentile of the beliefs distribution expect the market return to be 33% higher than investors in the 25th percentile of the beliefs distribution.

[28]We calculate δ^*_{mt} as the residual from the regression of δ_{mt} on the parameter θ_{mt},which corresponds to the inverse of risk aversion.

5.3 What Explains Holdings? Beliefs vs. Risk Aversion

Our results in the previous section indicate that there is substantial heterogeneity in beliefs and risk aversion across investors. We examine how dispersion in beliefs and risk aversion explain variation in equity exposure across plans in the following regression:

$$Equity\,Share_{mt} = \beta\lambda_{mt} + \gamma\delta_{mt} + \epsilon_{mt}. \tag{10}$$

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 11 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 pp (19% 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 -5.7 pp (8% relative to the mean allocation) decrease in an investor's equity allocation. Conversely, an investor's expectations of the market return are negatively correlated with her cash holdings, and an investor's risk aversion is positively correlated with her cash holdings. 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 52% of the variation in equity exposure.

5.4 Interpreting Heterogeneity in Portfolio Allocations in Terms of Risk Aversion and Beliefs

We find that there is substantial heterogeneity in both risk aversion and beliefs across investors and that heterogeneity in both risk aversion and beliefs explain a substantial amount of variation in portfolio allocations. These findings provide a useful lens for understanding why portfolio allocations vary across investors as documented in Section 3.1. Our framework allows us to understand if the differences in equity exposure across investors are driven by differences in beliefs, risk aversion, or both. Variation in portfolio allocations could also be explained by differences in 401(k) menus. However, as shown in Appendix Table A2, we find little relationship between participant demographics and the composition of 401(k) menus.

For example, we find a positive relationship between the share of college educated individuals and equity exposure. This appears to be driven primarily by differences in beliefs about market returns rather than differences in risk aversion or 401(k) menus. In fact, we find a positive relationship between the share of college educated individuals and risk aversion and a negative relationship between the share of college educated individuals and the number of equity 401(k) options. Conversely, both differences in risk aversion and beliefs help explain why older investors tend to have lower equity exposure. Older investors appear both more risk averse in the data and pessimistic about future returns. We also find that differences in subjective beliefs, rather than differences in 401(k) menus or risk aversion, help explain why equity exposure varies across ethnicities, which could be a function of differences in trust in financial markets (Guiso et al., 2008; Gennaioli et al., 2015).

5.5 Alternative Specifications and Robustness

We consider several alternative specifications to assess the robustness of the estimated parameters. First, we allow for a time-varying intercept in our estimates of risk aversion to account for the fact that risk aversion may vary systematically over time in ways that are not correlated with our observed plan and investor characteristics. Second, we re-estimate the model including target date funds, which are excluded from our baseline analysis. Third, to account for potential inertia in investor behavior, we estimate the model using only new plans.[29] For all three of these specifications, we find very similar estimates of risk aversion and expected returns. Results are reported in Table A8. The mean risk aversion ranges from 4.4 to 4.7, close to our baseline estimate of 4.5. The mean expected return ranges from 11.5 to 12.5, similar to our baseline estimate of 11.5. 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 an alternative measure of risk based on a simplified factor structure, which we describe in Section 4.3.[30] As in the above specifications, we find that the estimates of risk aversion and beliefs are highly correlated with our baseline estimates.

6 Evidence on the Formation of Beliefs

Investor beliefs play a critical role in determining investor portfolios and vary substantially across investors. 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 how investors extrapolate

[29]Because this greatly restricts our sample, we only estimate a single risk aversion parameter for this specification.

[30]We construct the factors for our alternative measure of risk by forming equal weighted portfolios based on the broad BrightScope categories reported in Table 1a. While the baseline and alternative measures of risk are highly correlated (0.94), the standard deviation of our alternative measure of risk is roughly 40% smaller than the standard deviation of our baseline measure of risk. This helps explain why we estimate slightly higher average risk aversion (8.666) with our alternative measure of risk relative to our baseline estimates.

their beliefs on two dimensions. First, we find that investors extrapolate their beliefs from past fund returns. Given that past returns are often a salient feature of 401(k) brochures/documents, investors are likely to extrapolate by looking at last year's returns. Using variation in 401(k) menus over time, we also show that investors extrapolate from past returns for both existing funds and funds newly added to their menus. Thus, the extrapolation we document cannot be explained by a lack of rebalancing.

To understand the role of experience in shaping extrapolation, we examine how investors extrapolate their beliefs based on the past performance of their employer. Following strong performance of their employer, as measured by investment, hiring, sales growth, and stock returns, investors become more optimistic about the market. This result holds when we compare investors working in the same sector and at the same time but for different employers. This suggests that investors potentially form broader beliefs about the future performance of the market based on their local experience with the success of their employer.

Lastly, we 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, appear to depend on past market returns, and also appear to depend on recent performance by their employer. 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.

6.1 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:

$$\bar{\mu}_{kt}^{(m)} = \rho Ret_{kt-1} + \upsilon_{kt}. \tag{11}$$

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 ($\bar{\mu}_{kt}^{(m)}$). The independent variable Ret_{kt+1} measures the past monthly return of investment option k averaged over year $t-1$ to t and is annualized. Table 12 displays the estimates corresponding to eq. (11). We examine extrapolation across three different subsets of the data: (i) the full data set in columns (1), (2), and (5); (ii) fund-by-plan observations in the first year the fund was added to the plan in column (3)[31]; and (iii) fund-by-plan observations corresponding to the first year a 401(k) plan was introduced in column (4). Samples (ii) and (iii) allow us to examine how investors

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

extrapolate their beliefs about funds they have not previously held their 401(k).

We find evidence that investors extrapolate their beliefs from past returns both in terms of the level of beliefs (columns 1 and 3-5) and changes in beliefs (column 2). The results in columns (3)-(5) indicate that investors extrapolate their beliefs from past returns for funds they did not hold in the past. The results in column (3) indicate that a ten percentage point increase in last year's return is correlated with an 0.46 pp increase in expected returns. In column (5) 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 negative coefficient which indicates that investors only extrapolate their beliefs 20% less from the returns of new funds versus returns of existing funds. The results in columns (3)-(5) show that the extrapolation we observe is not simply a function of investor inattention or inertia in portfolio rebalancing. While we believe it is useful to examine how past returns impact investor beliefs, in Appendix Table A5, we show that we find similar results if examine portfolio weights rather than beliefs.

6.2 Extrapolation from Employer Economic Conditions

An advantage in our setting is that we observe details on the investor's employer, the fund sponsor. This allows us to explore how investors' beliefs depend on their employment. Using the sponsor's EIN number, 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:

$$\delta_{mt} = \varphi Performance_{mt} + \eta_{mt}. \tag{12}$$

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 participating 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 13. 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 industry-by-year fixed effects, which allow us to effectively compare the beliefs of two investors working in the same industry at the same time but for different firms. The results in column (3) indicate that investors become 0.41 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 0.16 pp increase in expected return of the market. Thus, investors extrapolate not only from recent fund returns, but also based on the performance of their employer. This suggests that investors may misattribute the performance of their employer to the performance of the economy more generally.

6.3 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:

$$\varepsilon_{mkt+1} = Ret._{kt+1} - \bar{\mu}_{kt}^{(m)} \tag{13}$$

where $Ret._{kt+1}$ measures the monthly return of investment option k averaged over year t to $t+1$ and is annualized. The term $\bar{\mu}_{kt}^{(m)}$ is the average participant in plan m's expected return of fund k.[32] We test the predictability of forecast errors in the following regression model:

$$\varepsilon_{mkt+1} = \alpha_0 + \alpha_1 X_{mkt} + \eta_{mkt+1}. \tag{14}$$

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

Table 14 displays the estimation results corresponding to eq. (14). In short, we find overwhelming evidence that forecast errors are predictable. The results in columns (1) and (2) indicate that forecast errors are persistent. We also find that investors tend to over predict fund returns following past positive fund returns (columns 3 and 4). This is consistent with our finding, discussed in the proceeding section, 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.

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 US and

[32] In Appendix A.1, we also explore forecast errors over longer horizons (e.g., five years) and find similar patterns.

international equities and less in cash. In contrast, the investment options available to plan participants do not vary systematically with participant characteristics.

To understand the patterns we document, we propose 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 heterogeneous expense ratios, offers an ideal setting for our framework.

We find that there is substantial heterogeneity in both risk aversion and beliefs across investors and that both explain meaningful variation in investors' portfolios. 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.

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 full information rational expectations. Investors appear to overreact to recent news and extrapolate their beliefs. Investors extrapolate their beliefs from fund returns even when they have no prior experience with the funds (e.g., new funds added to the 401(k) plan). We also document that investors' employment experience also influences their beliefs. Investors become more optimistic about the stock market following strong financial performance from their employer.

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 potentially have important implications for asset prices. Our framework can also be easily applied in other settings to provide insight about investor beliefs and risk aversion, which could be particularly valuable when survey data is unavailable.

References

Ameriks, J. and S. P. Zeldes (2004). How do household portfolio shares vary with age. Technical report, working paper, Columbia University.

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.

Bacchetta, P., E. Mertens, and E. Van Wincoop (2009). Predictability in financial markets: What do survey expectations tell us? *Journal of International Money and Finance 28*(3), 406–426.

Bach, L., Laurent E. Calvet, and Paolo Sodini (2020). Rich Pickings? Risk, Return, and Skill in Household Wealth. *THE 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. (2018). Leverage regulation and market structure: A structural model of the UK mortgage market. Working paper.

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

Berk, J. B. and J. H. van Binsbergen (2016). Assessing asset pricing models using revealed preference. *Journal of Financial Economics 119*, 1–23.

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.

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.

Bodie, Z., R. C. Merton, and W. F. Samuelson (1992). Labor supply flexibility and portfolio choice in a life cycle model. *Journal of economic dynamics and control 16*(3-4), 427–449.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Over-reaction in macroeconomic expectations. Working paper, National Bureau of Economic Research.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance 74*(6), 2839–2874.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *The Journal of Finance 73*(1), 199–227.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma (2020). Swiss finance institute working paper. *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., J. F. Cocco, F. J. Gomes, and P. J. Maenhout (2009). *Investing Retirement Wealth: A Life-Cycle Model*. University of Chicago Press.

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

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.

Chen, X., L. P. Hansen, and P. G. Hansen (2020). Robust identification of investor beliefs. *National Bureau of Economic Research*.

Chiteji, N. S. and F. P. Stafford (2000). Asset ownership across generations. Technical report, 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.

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.

Cochrane, J. H. (2021). Portfolios for long-term investors. Technical report, National Bureau of Economic Research.

Coibion, O. and Y. Gorodnichenko (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy 120*(1), 116–159.

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.

d'Arienzo, D. (2020). Maturity increasing overreaction and bond market puzzles. *Available at SSRN 3733056*.

Davis, S. J. and P. S. Willen (2013). Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice. Technical report, Working Papers.

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, SSRN 2845567.

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 (2017). The cross section of bank value. Working paper, National Bureau of Economic Research.

Egan, M. L., A. MacKay, and H. Yang (Forthcoming). Recovering investor expectations from demand for index funds. *Review of Economic Studies*.

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. pp. 87.

Farhi, E. and S. Panageas (2007). Saving and investing for early retirement: A theoretical analysis. *Journal of Financial Economics 83*(1), 87–121.

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. and A. Shleifer (2018). *A crisis of beliefs: Investor psychology and financial fragility*. Princeton University Press.

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

Ghosh, A., A. G. Korteweg, and Q. Xu (2020). Recovering heterogeneous beliefs and preferences from asset prices. *Working Paper*.

Ghosh, A. and G. Roussellet (2020). Identifying beliefs from asset prices. *Working Paper*.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2019). Five facts about beliefs and portfolios. Working paper, National Bureau of Economic Research.

Gomes, F., M. Haliassos, and T. Ramadorai (2020). Household Finance. *Journal of Economic Literature, forthcoming*.

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.

Hausman, J. A. (1996). Volume title: The economics of new goods volume author/editor: Timothy f. bresnahan and robert j. gordon, editors volume publisher: University of chicago press volume isbn: 0-226-07415-3 volume url: http://www. nber. org/books/bres96-1.

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

Heipertz, J., A. Ouazad, and R. Rancière (2019). The transmission of shocks in endogenous financial networks: A structural approach. Working paper, National Bureau of Economic Research.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *The Quarterly Journal of Economics 119*(2), 403–456.

Huberman, G. and W. Jiang (2006). Offering versus choice in 401 (k) plans: Equity exposure and number of funds. *The Journal of Finance 61*(2), 763–801.

Jensen, C. S., D. Lando, and L. H. Pedersen (2019). Generalized recovery. *Journal of Financial Economics 133*(1), 154–174.

Koijen, R. S., R. J. Richmond, and M. Yogo (2019). Which investors matter for global equity valuations and expected returns? Working paper, SSRN 3378340.

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

Koijen, R. S. and M. Yogo (2018). The fragility of market risk insurance. Working paper, National Bureau of Economic Research.

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, SSRN 3383677.

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. (2021). Experience effects in finance: Foundations, applications, and future directions. *Review of Finance 25*(5), 1339–1363.

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.

Malmendier, U., D. Pouzo, and V. Vanasco (2020). Investor experiences and financial market dynamics. *Journal of Financial Economics 136*(3), 597–622.

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

Martin, I. and S. Ross (2019). Notes on the yield curve. *Journal of Financial Economics 134*(3), 689–702.

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

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

Nagel, S. and Z. Xu (2019). Asset pricing with fading memory. Working paper, National Bureau of Economic Research.

Parker, J. A., A. Schoar, and Y. Sun (2020). Retail financial innovation and stock market dynamics: The case of target date funds. Technical report, National Bureau of Economic Research.

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. *Available at SSRN 3752296*.

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

Samwick, A. A. and J. M. Poterba (1997). *Household portfolio allocation over the life cycle*. National Bureau of Economic Research.

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

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.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance 56*(2), 433–470.

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.

Wang, Y., T. M. Whited, Y. Wu, and K. Xiao (2018). Bank market power and monetary policy transmission: Evidence from a structural estimation. 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.

Tables and Figures

Figure 1: Distribution of Holdings

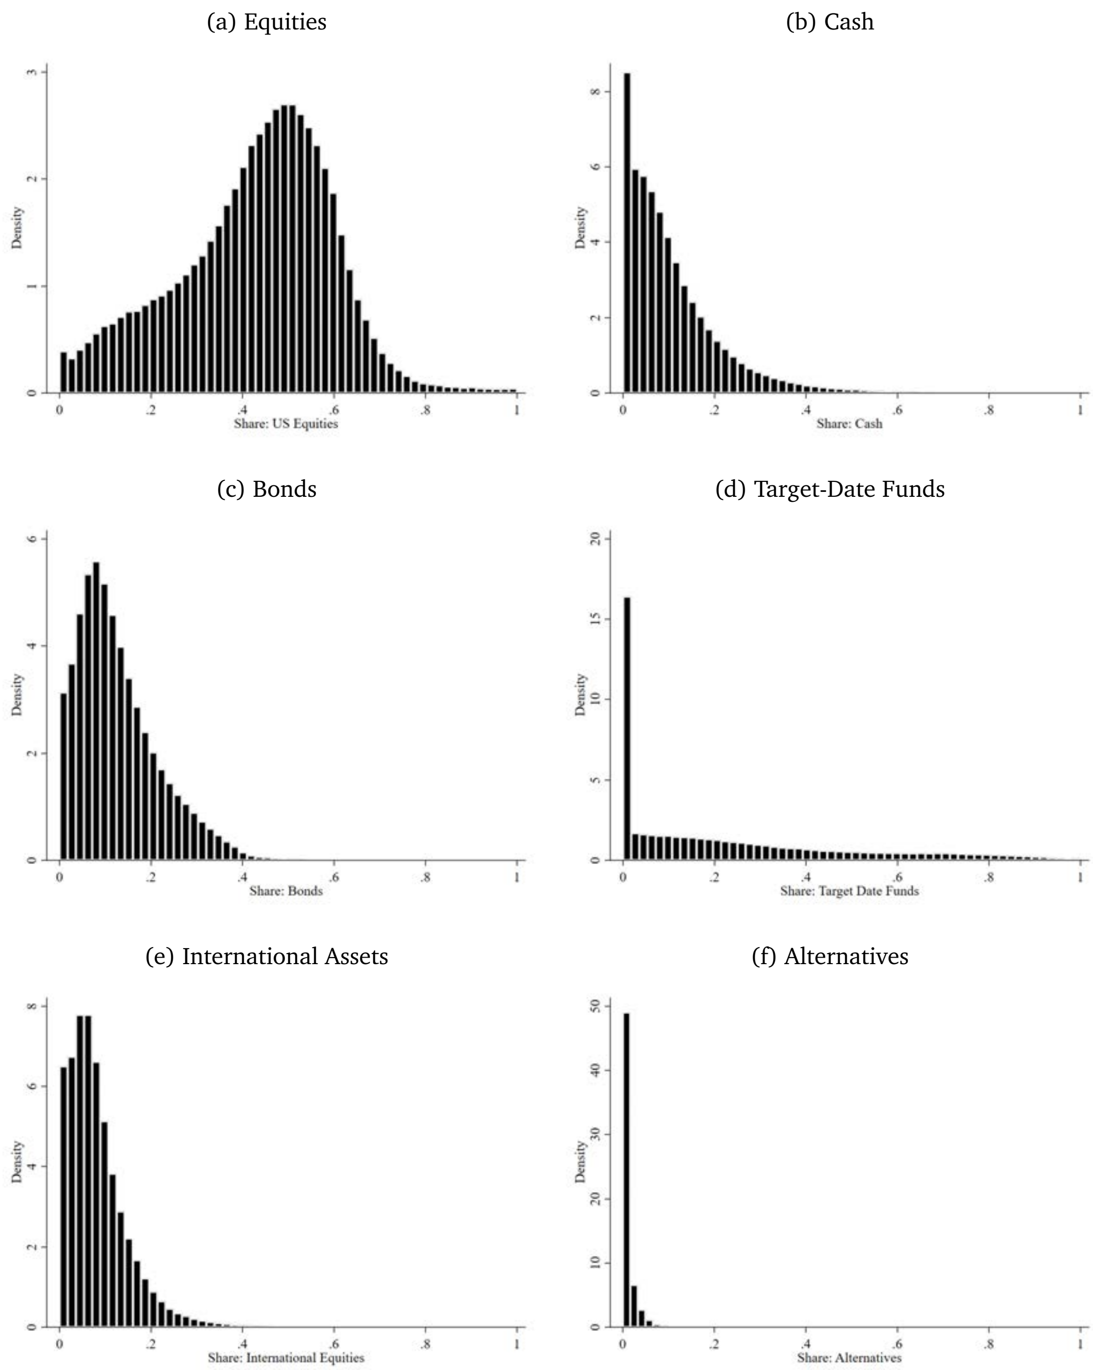

Figure 1 displays the distribution of holdings across 401(k) plans. Observations are at the plan-by-year level over the period 2009-2019 for those plans with at least five investment options.

Figure 2: Holdings Over Time

(a) Holdings Over Time, Excluding Target Date Funds

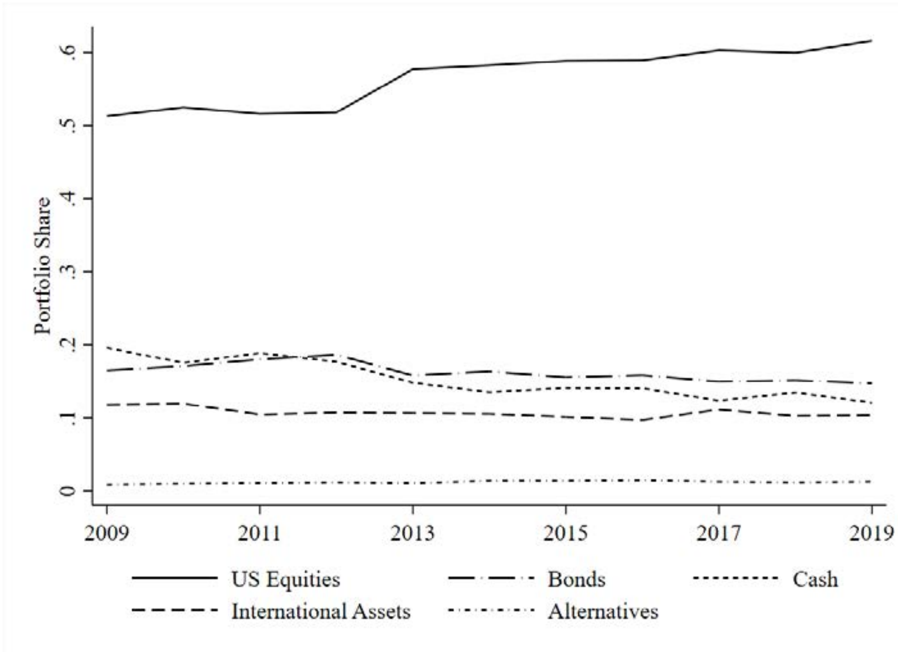

(b) Holdings Over Time, Including Target Date Funds

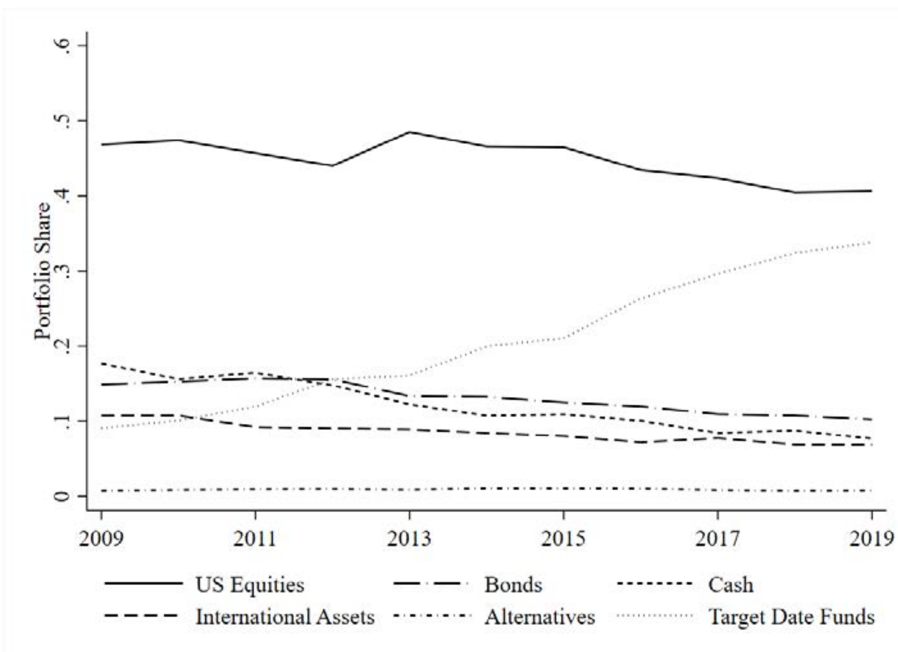

Figure 2 displays the equal-weighted average holdings across plans over the period 2009-2019. In panel (a) we calculate portfolio shares excluding target date funds. In panel (b) we calculate portfolio shares including target date funds.

Figure 3: Fund Expenses

(a) Fund Expenses (Equal Weighted)

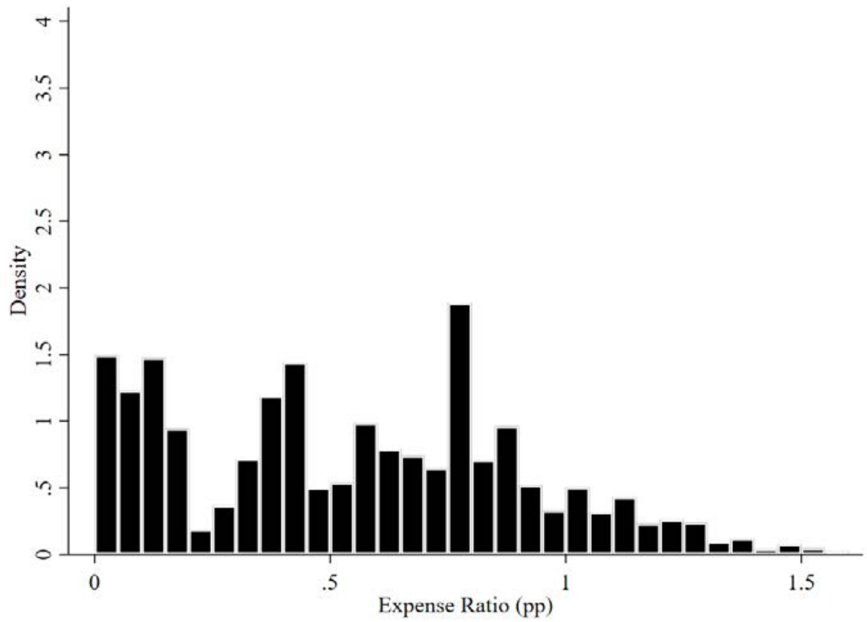

(b) Fund Expenses (AUM Weighted)

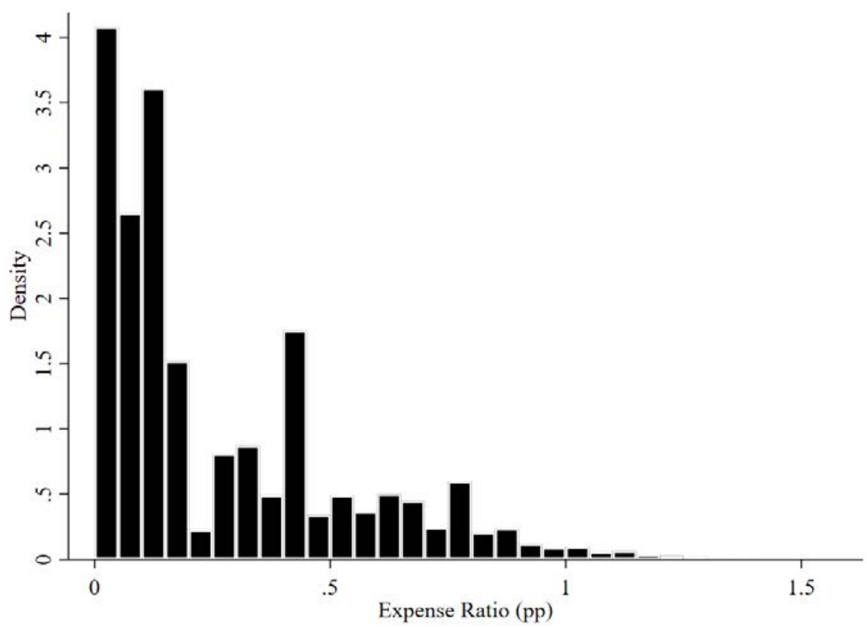

Figure 3 displays the distribution of fund expenses. Observations are at the fund-by-plan level as of 2019 as reported by BrightScope. Panel (a) displays the equal weighted distribution of fund expenses. Panel (b) displays the asset weighted distribution of fund expenses.

Figure 4: Equity Allocations and Naive Diversification

(a) Observed Allocations Compared to Naive Diversification

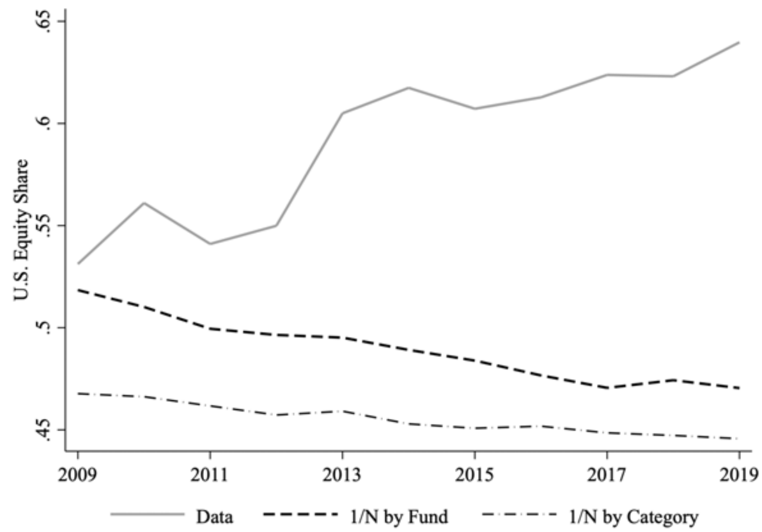

(b) Observed vs Predicted Allocations

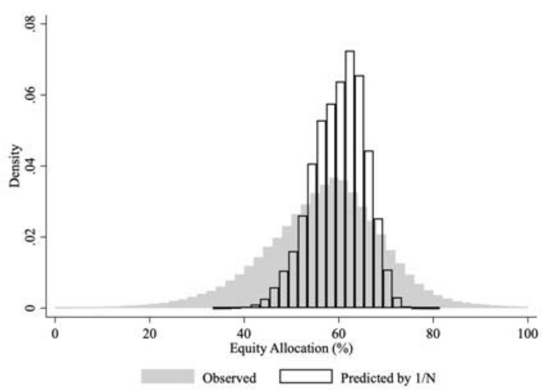

(c) Observed vs Residualized Allocations

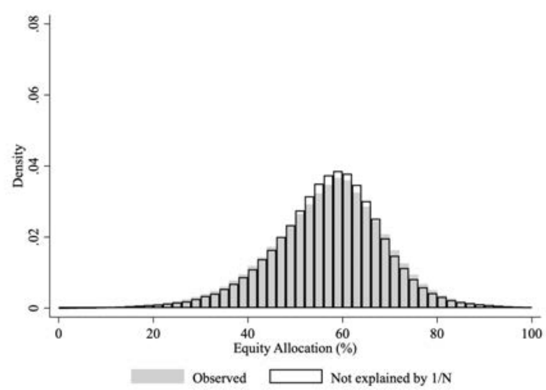

Figure 4 panel (a) displays the share of assets held in US equities over the period 2009-2019 and the expected share of assets held in US equities if all investors used either a naive 1/N strategy by fund or 1/N strategy by investment category. When computing the share held in US equities, we drop all target date fund assets and assume that remaining non-target date allocation funds hold 60% in US equities. Panels (b) and (c) compare allocations to US equities with the predicted allocations based on the 1/N strategy and allocations not explained by the 1/N strategy, respectively. To predict allocation, we regress observed equity shares on equity shares implied by 1/N strategy by fund, along with year and 2 digit NAICs fixed effects.

Figure 5: Risk Aversion Over Time

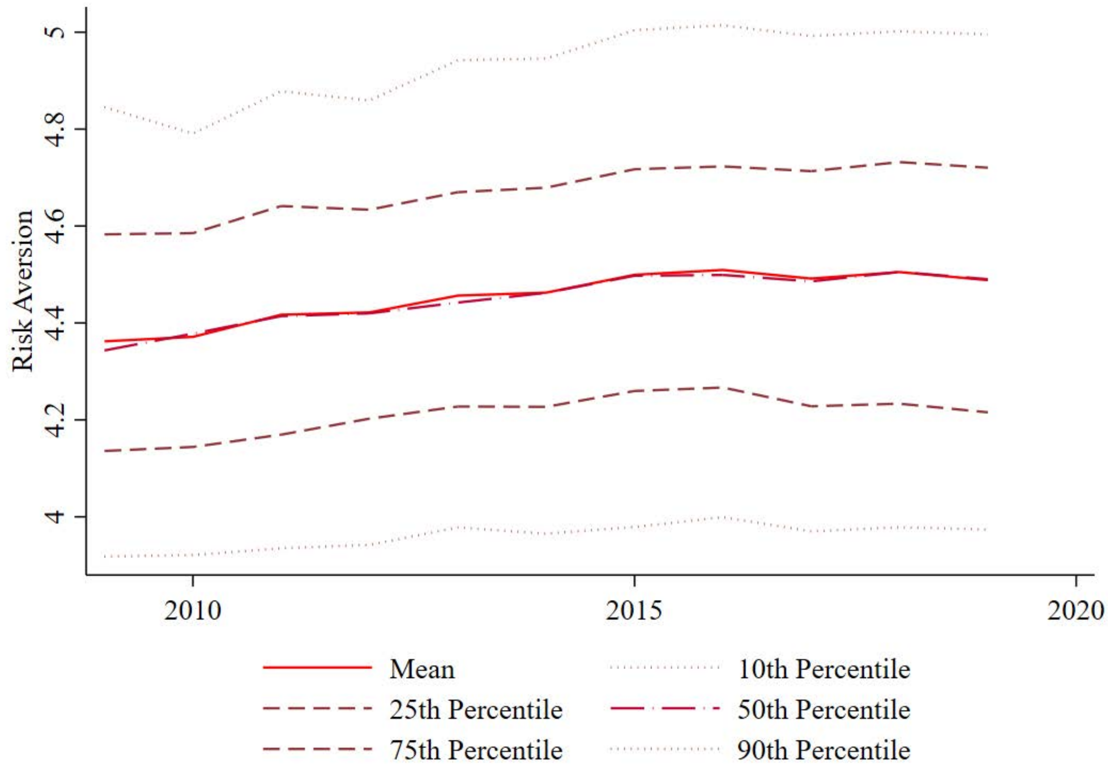

Figure 5 displays estimated risk aversion over time. Risk aversion corresponds to our model estimates reported in column (3) of Table 7. When computing risk aversion, we set the dummy variable *New 401(k) Plan* equal to one for each plan.

Figure 6: Distribution of Investor Beliefs About the Stock Market Over Time

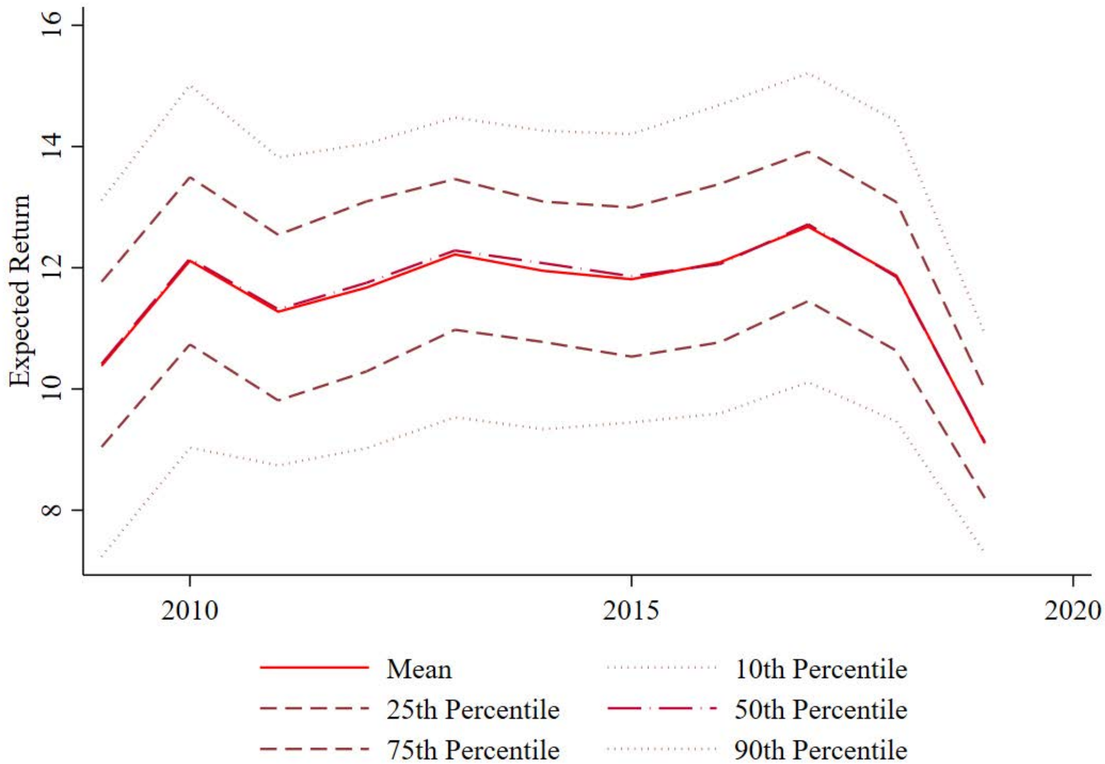

Figure 6 displays the estimated distribution of investor expectations of market returns. The estimates correspond to the specification reported in column (3) of Table 7. When computing risk aversion and beliefs, we set the dummy variable *New 401(k) Plan* equal to one for each plan.

Figure 7: Cross-Section of Investor Beliefs About Stock Market Returns

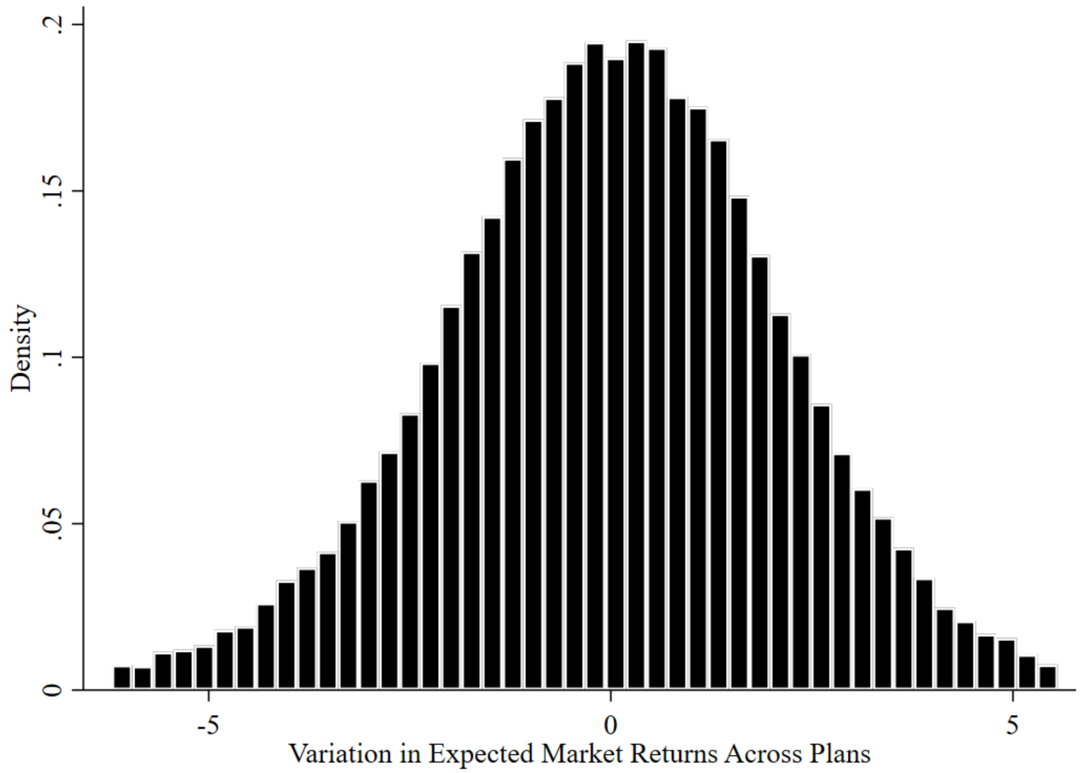

Figure 7 displays the estimated cross-sectional distribution of investor expectations of market returns. The estimates correspond to the specification reported in column (3) of Table 7. When computing risk aversion and beliefs, we set the dummy variable *New 401(k) Plan* equal to one for each plan. Expectations 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. Observations are at the plan level. To account for outliers we truncate the distribution at the 1% and 99% percentile.

Table 1: Summary Statistics

(a) Plan Summary Statistics

	Count	Mean	SD	Median
Total Assets (millions)	442,631	84.749	689.657	10.722
Number of Plan Participants	425,075	1,261	92,360	223
Number of Investment Options	442,631	26.297	13.835	26.000
Average Account Balance	424,136	66,082	5.33e+05	45,324
Plan Participation Rate	405,832	0.738	0.922	0.833
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

	Count	Mean	SD	Median
Volatility	1.06e+07	0.138	0.042	0.149
Expense Ratio (bp; BrightScope)	1.86e+06	0.569	0.383	0.590
Expense Ratio (bp; CRSP)	6.60e+06	0.606	0.432	0.610

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. (5) and is annualized.

Table 2: Equity Allocation vs. Demographics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	-0.300											-0.686***
	(0.230)											(0.220)
ln(Income)		1.343***										0.381
		(0.294)										(0.374)
ln(Home Value)			1.011***									0.161
			(0.303)									(0.192)
College				1.554***								0.771**
				(0.308)								(0.296)
Employed					0.967***							0.138
					(0.311)							(0.154)
Black						-0.390						-0.195
						(0.226)						(0.125)
Hispanic							-0.855***					-0.609***
							(0.246)					(0.150)
Unionized								-1.057				-0.377
								(0.708)				(0.454)
Sector Equity Beta									0.620			0.523
									(0.725)			(0.411)
Share Retired										-0.372***		-0.487***
										(0.112)		(0.083)
ln(Avg. Acct. Bal.)											1.292***	1.038***
											(0.135)	(0.237)
Observations	243,284	243,284	243,284	243,284	243,284	243,284	243,284	243,284	240,166	243,284	243,284	240,166
R-squared	0.084	0.095	0.090	0.099	0.089	0.084	0.087	0.083	0.084	0.084	0.092	0.110
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Table 2 displays the regression results corresponding to a linear regression model. The dependent variable is the share of the 401(k) held in US equities. 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. 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 level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Table 3: Equity Allocations Across and Within Sectors

Sector	Mean	p25	Median	p75
Information	64.2	55.2	62.6	69.2
Utilities	63.2	52.6	61.1	67.9
Management of Companies and Enterprises	62.9	51.8	59.5	66.9
Professional, Scientific, and Tech. Services	61.7	54.2	61.1	67.7
Real Estate and Rental and Leasing	61.1	51.1	59	66.2
Manufacturing	61	49.1	57.3	64.7
Retail Trade	60.2	48.8	57.4	64.9
Finance and Insurance	60.1	51	58.9	65.4
Wholesale Trade	59.5	49.7	57.8	65.1
Agriculture, Forestry, Fishing and Hunting	59.5	46.7	55.2	63.6
Admin and Support and Waste Services	58.7	49.3	58.2	65.7
Public Administration	58.6	46.5	53.1	60.2
Transportation and Warehousing	58	47.9	56.4	63.9
Other Services (except Public Administration)	57.6	48.9	57.1	64.4
Accommodation and Food Services	57.6	48.2	57.6	65.6
Arts, Entertainment, and Recreation	57.2	51	59.5	66.5
Health Care and Social Assistance	56.9	47.8	55.7	62.5
Mining, Quarrying, and Oil and Gas Extraction	56.3	47.2	56.3	64.2
Construction	56.2	47.4	56.8	64.8
Educational Services	51.8	47.3	55.5	63.9

Table 3 displays the distribution of US equity allocations (i.e., share of plan assets held in US equities) across sectors (2-digit NAICS). 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.

Table 4: Asset Allocation vs. Demographics

VARIABLES	(1) US Equities	(2) Bonds	(3) Cash	(4) International Equities
Age	-0.686***	0.201*	0.820***	-0.323***
	(0.220)	(0.108)	(0.163)	(0.097)
ln(Income)	0.381	-0.684**	0.596**	-0.253**
	(0.374)	(0.311)	(0.279)	(0.118)
ln(Home Value)	0.161	-0.411**	0.305	-0.101
	(0.192)	(0.162)	(0.301)	(0.119)
College	0.771**	0.406	-1.575***	0.262*
	(0.296)	(0.294)	(0.357)	(0.149)
Employed	0.138	-0.131	0.015	0.010
	(0.154)	(0.099)	(0.159)	(0.068)
Black	-0.195	-0.134	0.821***	-0.424***
	(0.125)	(0.086)	(0.145)	(0.083)
Hispanic	-0.609***	-0.082	0.911***	-0.282***
	(0.150)	(0.091)	(0.201)	(0.096)
Unionized	-0.377	-0.656	3.639***	-2.298***
	(0.454)	(0.614)	(0.451)	(0.251)
Sector Equity Beta	0.523	-0.432	-0.350	0.154
	(0.411)	(0.504)	(0.366)	(0.099)
Share Retired	-0.487***	0.199***	0.684***	-0.335***
	(0.083)	(0.057)	(0.127)	(0.072)
ln(Avg. Acct. Bal.)	1.038***	-0.279	0.110	-0.735***
	(0.237)	(0.295)	(0.393)	(0.076)
Observations	240,166	240,166	240,166	240,166
R-squared	0.110	0.031	0.078	0.031
Year FE	X	X	X	X

Table 4 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 level and at the county level and are in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 5: Portfolio Allocation vs. Expenses

	(1)	(2)
Expense Ratio	-0.518***	-0.590***
	(0.003)	(0.006)
Observations	5,063,120	5,049,239
R-squared	0.435	0.435
PlanxYear FE	X	X
CategoryxYear FE	X	X
Index FundxYear FE	X	X
IV		X

Table 5 displays the regression results corresponding to a linear regression model (eq. 2). 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 2-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. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 6: Autocorrelation in Portfolio Weights

	(1)	(2)
Expected Portfolio Weight	0.694***	0.879***
	(0.005)	(0.008)
Observations	4,542,432	3,117,521
R-squared	0.482	0.772
Excluding Newly Added Funds		X

Table 6 displays the one year autocorrelation in portfolio weights. Observations are at the investment option-by-plan-by-year level over the period 2009-2019 where we restrict our attention to those 401(k)s that were available for at least a year. For ease of interpretation, all dependent and independent variables are standardized such that coefficient estimates are equivalent to correlation coefficients. We compute *Expected Portfolio Weight* under the assumption that the portfolio weight of a fund grows by the return of fund relative to the total return of the 401(k) portfolio over the same period (assuming no rebalancing). Column (1) includes investment options that were not available in the fund menu in the prior year, and hence *Expected Portfolio Weight* is equal to zero for these options. In column (2) we restrict the sample to those investment options that were available in the prior year. Standard errors are clustered at the plan level and are in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$..

Table 7: Baseline Model Estimates (θ)

	(1)	(2)	(3)
Fee	-0.179***	-0.190***	-0.190***
	(0.017)	(0.017)	(0.017)
X Age		0.017**	0.017**
		(0.007)	(0.007)
X Frac Black		0.009*	0.009*
		(0.005)	(0.005)
X Frac Hispanic		-0.005	-0.004
		(0.010)	(0.010)
X Frac College		0.021**	0.021**
		(0.009)	(0.009)
X log(Median Family Income)		-0.017	-0.017
		(0.011)	(0.011)
X log(Median House Value)		0.006	0.006
		(0.010)	(0.010)
X Frac Employed		-0.000	-0.000
		(0.007)	(0.007)
X Unionized		0.008	0.008
		(0.015)	(0.015)
X Share Retired		-0.002	-0.002
		(0.006)	(0.006)
X log(Avg. 401(k) Balance)		-0.006	-0.006
		(0.006)	(0.006)
X New 401(k) Plan			-0.035*
			(0.019)
Observations	4,932,664	4,528,711	4,528,711
R-squared	0.928	0.928	0.928
Fixed Effects	X	X	X
Avg. Risk Aversion	5.6	5.3	5.3

Table 7 displays the regression results corresponding to a linear regression model (eq. 5). Observations are at the investment option-by-plan-by-year level over the period 2009-2019. The dependent variable is the additional risk of investing a dollar in a given investment option, given the other portfolio holdings in the plan. The independent variables, other than the dummy variables *Union* and *New 401(k) Plan*, are all standardized such that they are in units of standard deviations. We estimate each specification using 2-stage least squares. We instrument for expenses and the corresponding interaction terms using Hausman-type instruments as described in the text. Because each observation reflects the average behavior of plan participants, we weight each observation by the total assets of the 401(k) plan. All specifications include plan-by-year, Morningstar investment category-by-BrightScope investment category-by-year fixed effects, and index-fund-by-year fixed effects. In the bottom panel we calculate average risk aversion where we set all of the interaction terms equal to zero. Standard errors are clustered at the plan level and are in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 8: Interpreting (θ) in terms of Risk Aversion

	(1)	(2)	(3)
Avg. Risk Aversion	5.590***	5.264***	5.271***
	(0.543)	(0.474)	(0.475)
X Age		0.469*	0.471*
		(0.243)	(0.243)
X Frac Black		0.236*	0.239*
		(0.137)	(0.138)
X Frac Hispanic		-0.125	-0.124
		(0.284)	(0.285)
X Frac College		0.583**	0.583**
		(0.272)	(0.273)
X log(Median Family Income)		-0.466	-0.464
		(0.339)	(0.340)
X log(Median House Value)		0.157	0.156
		(0.279)	(0.280)
X Frac Employed		-0.000	-0.000
		(0.194)	(0.194)
X Unionized		0.226	0.233
		(0.440)	(0.442)
X Share Retired		-0.049	-0.051
		(0.170)	(0.170)
X log(Avg. 401(k) Balance)		-0.153	-0.154
		(0.171)	(0.172)
X New 401(k) Plan			-0.969*
			(0.555)

Table 8 presents the estimates of (θ) from Table 7 in terms of risk aversion and the marginal effects of each independent variable on risk aversion. The independent variables, other than the dummy variables *Union* and *New 401(k) Plan*, are all standardized such that they are in units of standard deviations. Standard errors are clustered at the plan level and are in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 9: Residualized Variation in Expected Market Returns vs. Demographics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	-0.158***											-0.203***
	(0.041)											(0.050)
ln(Income)		0.194**										0.050
		(0.073)										(0.071)
ln(Home Value)			0.177***									0.001
			(0.052)									(0.048)
College				0.270***								0.189**
				(0.081)								(0.077)
Employed					0.130*							0.010
					(0.070)							(0.036)
Black						-0.105**						-0.089***
						(0.040)						(0.028)
Hispanic							-0.118*					-0.090**
							(0.058)					(0.034)
Unionized								-0.641***				-0.475***
								(0.146)				(0.122)
Sector Equity Beta									0.227			0.215**
									(0.135)			(0.096)
Share Retired										-0.146***		-0.135***
										(0.023)		(0.022)
ln(Avg. Acct. Bal.)											0.094*	0.060
											(0.045)	(0.063)
Observations	243,284	243,284	243,284	243,284	243,284	243,284	243,284	243,284	240,166	243,284	243,284	240,166
R-squared	0.006	0.010	0.008	0.019	0.004	0.003	0.003	0.006	0.003	0.005	0.002	0.042
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Table 9 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 residual variation in expected market returns, residualized on risk aversion. Expected market returns are calculated based on the specification reported in column (3) of Table 7 where we set the dummy variable *New 401(k) Plan* equal to one for each plan. 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.

Table 10: Expected Market Returns Across and Within Sectors

Sector	Mean	p25	Median	p75
Real Estate and Rental and Leasing	13.3	11.5	13.4	15.1
Management of Companies and Enterprises	13	10.4	12.1	13.9
Public Administration	12.3	10.6	12.3	13.9
Professional, Scientific, and Tech. Services	12.3	11.3	12.8	14.2
Transportation and Warehousing	12	10	11.6	13.2
Finance and Insurance	11.9	10.2	11.9	13.4
Health Care and Social Assistance	11.8	10.3	12.1	13.7
Information	11.8	10.9	12.6	14
Wholesale Trade	11.3	9.7	11.3	12.8
Other Services (except Public Administration)	11.2	9.8	11.5	13.1
Utilities	11.1	10	11.4	13.1
Arts, Entertainment, and Recreation	11.1	9.5	11.1	12.6
Manufacturing	11	9.4	11	12.5
Admin and Support and Waste Services	10.8	9.8	11.5	13.1
Educational Services	10.8	10.6	12.5	14.4
Mining, Quarrying, and Oil and Gas Extraction	10.5	8.8	10.4	11.9
Retail Trade	10.2	9	10.5	11.9
Agriculture, Forestry, Fishing and Hunting	10.1	8.4	10.1	11.8
Construction	9.9	8.5	10.1	11.5
Accommodation and Food Services	9.4	8.3	9.8	11.2

Table 10 displays the distribution of expected market returns across sectors (2-digit NAICS). Expected market returns are calculated based on the specification reported in column (3) of Table 7 where we set the dummy variable *New 401(k) Plan* equal to one for each plan. Observations are at the plan-by-year level over the period 2009-2019.

Table 11: 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.)	-3.877***	-5.676***	-3.186***	-4.919***	2.593***	3.866***
	(0.272)	(0.139)	(0.203)	(0.120)	(0.240)	(0.208)
Expected Returns (Std.)	9.585***	13.032***	8.167***	11.347***	-6.903***	-9.332***
	(0.353)	(0.254)	(0.272)	(0.213)	(0.409)	(0.376)
Observations	243,284	243,284	243,284	243,284	243,284	243,284
R-squared	0.520	0.803	0.351	0.596	0.301	0.455
Year FE		X		X		X

Table 11 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 share of the plan portfolio in equities in columns (1) and (2); in US equities in columns (3) and (4); and cash in columns (5) and (6). When computing the shares, we drop all target date fund assets and assume that remaining non-target date allocation funds hold 60% in US equities. Expected market returns and risk aversion are calculated based on the specification reported in column (3) of Table 7 where we set the dummy variable *New 401(k) Plan* equal to one for each plan. 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.

Table 12: Expected Returns vs. Past Fund Returns

VARIABLES	(1) Beliefs	(2) Δ Beliefs	(3) Beliefs	(4) Beliefs	(5) Beliefs
Lag Fund Ret.	0.014***	0.001**	0.042***	0.022**	0.015***
	(0.001)	(0.001)	(0.002)	(0.009)	(0.001)
Lag Fund Ret. x New Investment					-0.004***
					(0.001)
Observations	4,500,266	2,397,686	973,937	79,087	4,500,266
R-squared	0.938	0.894	0.943	0.942	0.938
FE	X	X	X	X	X
New Funds			X		
New Plans				X	

Table 12 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. Expected returns are calculated based on the specification reported in column (3) of Table 7 where we set the dummy variable *New 401(k) Plan* equal to one for each plan. Each specification include plan-by-year, investment category (MorningstarXBrightScope)-by-year, and index fund-by-year fixed effects. Standard errors are in parenthesis and are clustered at the plan level. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 13: Expected Market Returns vs. Employer and Industry Returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Firm Return (1 years)	0.007***	0.005***							0.006***	0.004***
	(0.002)	(0.002)							(0.002)	(0.002)
Firm Investment			0.041***	0.025***					0.036***	0.021***
			(0.004)	(0.004)					(0.004)	(0.004)
Sales Growth					0.006***	0.005***			0.002***	0.002**
					(0.001)	(0.001)			(0.001)	(0.001)
Employment Growth							0.009***	0.006***	0.003**	0.003**
							(0.001)	(0.001)	(0.001)	(0.001)
Observations	11,743	11,738	10,479	10,474	11,458	11,452	11,448	11,441	10,088	10,081
R-squared	0.222	0.365	0.247	0.380	0.225	0.365	0.225	0.365	0.244	0.377
Year FE	X	X	X	X	X	X	X	X	X	X
NAICSxYear FE		X		X		X		X		X

Table 13 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 expected return of the market. The expected return of the market is calculated based on the specification reported in column (3) of Table 7 where we set the dummy variable *New 401(k) Plan* equal to one for each plan. Standard errors are in parenthesis and are clustered at the plan level. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 14: Predictability of Forecast Errors

	(1)	(2)	(3)	(4)	(5)	(6)
Lag Forecast Error	0.085***	0.042***				
	(0.001)	(0.001)				
Lag Fund Ret.			-0.198***	-0.195***		
			(0.003)	(0.003)		
Change in Beliefs					-0.602***	-0.997***
					(0.008)	(0.012)
Observations	2,400,338	2,395,869	4,495,452	4,495,396	2,402,960	2,398,501
R-squared	0.617	0.656	0.605	0.643	0.617	0.659
Year FE	X		X		X	
PlanxYear FE		X		X		X

Table 14 displays the regression results corresponding to a linear regression model (eq. 14). Observations are at the investment option-by-plan-by-year level over the period 2009 through 2019. The dependent variable is investor's forecast errors as measured per eq. (13). *Lag Forecast Error* measures investors forecast error in the previous period. *Lag Fund Ret.* measures the annual fund return in the previous year. *Change in Beliefs* measures the change in investor's beliefs about the expected returns of the fund over the previous year. Standard errors are clustered at the plan level. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

A Additional Analysis and Robustness

A.1 Relation to Future Returns

We examine the relationship between investor expectations and return in the following regression:

$$Ret_{kt'} = \bar{\mu}_{kt} + \eta_{kt.} \tag{15}$$

Observations are at the fund-by-year level. The dependent variable is the return of fund k over the period t to t', where we examine the forecastability of returns over a one, three, and five year horizon. We control for the mean expected return of fund k at time t across plans $(\bar{\mu}_{kt})$ and the interquartile range of expected returns of fund k at time t across plans.

Table A6 displays the corresponding estimates. We examine the predictability of returns over a one year horizon in columns (1) and (4); over a three year horizon in columns (2) and (5); and over a five year horizon in columns (3) and (6) . In columns (1)-(3) we do not control for fund risk, while in columns (4)-(6) we control for fund risk as measured by the fund's factor loadings.[33] Controlling for fund risk is important because otherwise investor expectations could just be capturing differences in fund risk. Consistent with this intuition, we find a positive and significant relationship between investor expectations and future returns in columns (1)-(3) when we omit risk controls. However, once we start controlling for differences in risk in columns (4)-(6), the relationship between investor expectations and future returns disappears. Thus, investor expectations do not forecast future returns once we account for known differences in risk.

A.2 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:

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

Rearranging the terms yields:

$$\left(\sum_{l=1}^{L} b_{klt} \left(\sum_{j=1}^{K} b_{jlt}\omega_{ijt} \right) + \sigma_k^2 \omega_{ikt} \right) = \theta p_{kt} + \psi \left(\sum_{l=1}^{L} b_{klt} b_{wlt} \right) + \epsilon_{kt}, \tag{16}$$

[33] Specifically, we control for the time-varying factor loadings the 55 factors used to calculate portfolio risk. We also allow coefficients on the factor loadings to vary over time.

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., $\epsilon_{kt} = (\bar{\mu}_{kt}^{(m)} - R_F)/\lambda$), and ψ is equal to $-\varpi$.

We estimate the empirical equivalent of eq. (16) as

$$\varsigma_{mkt}^2 = \theta p_{mkt} + \psi \xi_{mkt}^2 + \phi_{mt} + \phi_{j(k)t} + \epsilon_{mkt}, \tag{17}$$

where:

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

and

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

The term ξ_{mkt}^2 captures the additional risk of investing in asset k due to labor income risk. We proxy for the factor loadings for labor income risk using the equity factor loadings corresponding to the industry of the plan sponsor m.

Table A7 displays the corresponding estimates. We estimate a similar inverse risk aversion coefficient θ as in our baseline specification in column (1). In column (2) we include fund-by-year fixed effects, which absorbs the term θ. The object of interest is the parameter $\psi = -\varpi$. Note that we estimate $\psi > 0$ which implies a negative weight ϖ such that investors behave as if they are risk seeking with respect to their labor income risk. One caveat is that the additional risk due to labor income $\xi_{mkt}^2 = \left(\sum_{l=1}^{L} b_{klt} b_{wlt} \bar{\omega}_{jt}^{(m)} \right)$ could be correlated with investor beliefs μ, which would make it endogenous in eq. (17). 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.

B Additional Tables and Figures

Figure A1: Holdings Over Time

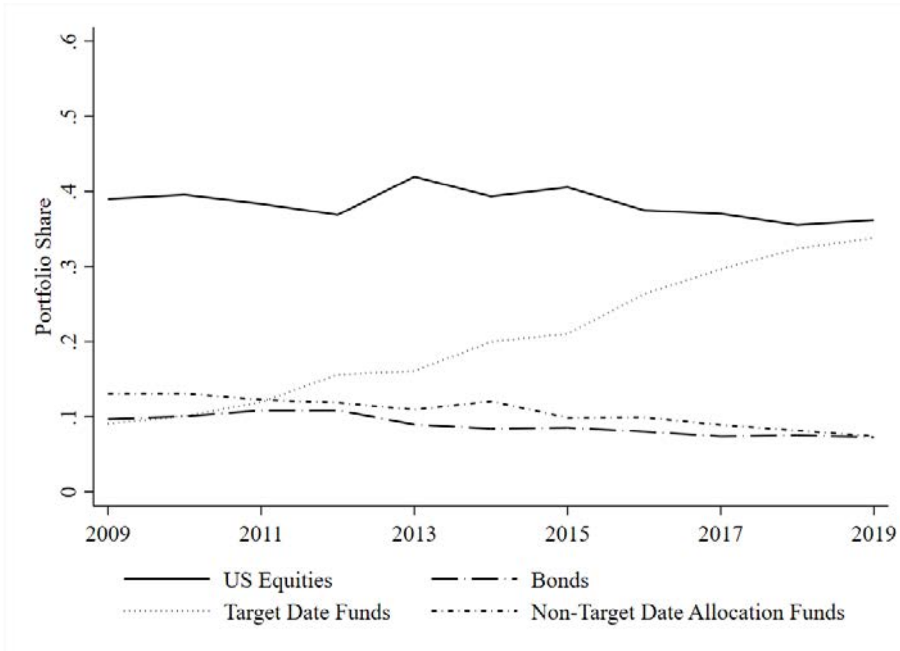

Figure A1 displays the equal-weighted average holdings of target date and non target date allocation funds, as well as US equity and bond assets without considering allocation funds across plans over the period 2009-2019.

Figure A2: Participation and Employer Contributions Over Time

(a) Participation Over Time

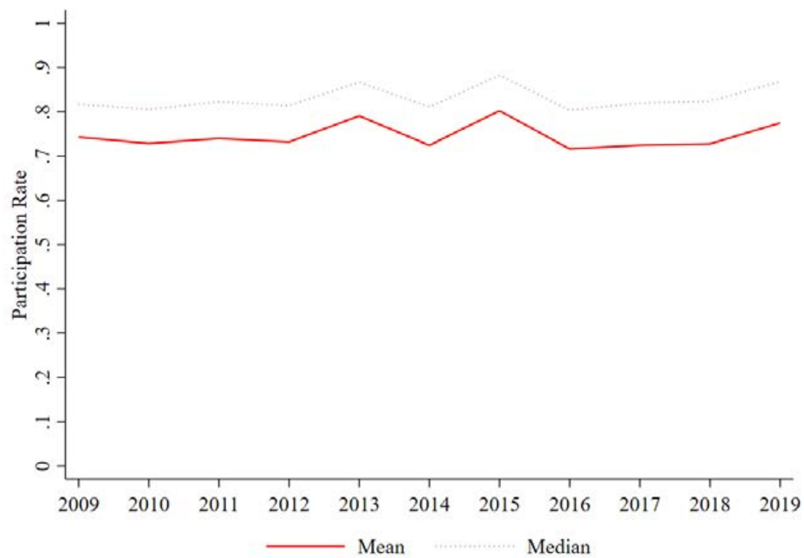

(b) Employer Contributions (Share of Total Contributions)

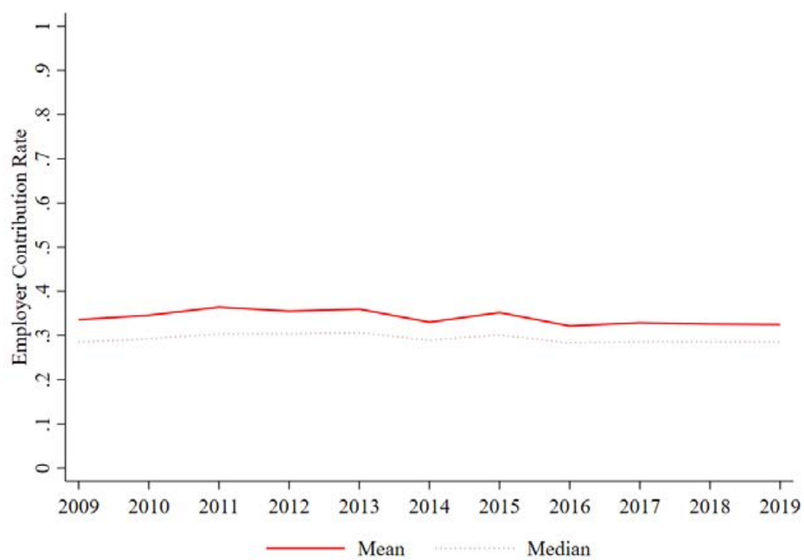

Figure A2 panel (a) displays the average and median 401(k) participation rate across 401(k) plans over the period 2009-2019. We measure the participation rate as the share of individuals who participate in the plan relative to the number of individuals who are eligible to participate. Panel (b) displays the average and median employer contribution rate across 401(k) plans. The employer contribution rate is measured as the employer's 401(k) contribution relative to the total 401(k) contribution (i.e., employer contribution plus employee contribution).

Table A1: 401(k) Participation vs. Demographics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Age	0.042**											0.026**
	(0.019)											(0.012)
ln(Income)		0.044***										-0.003
		(0.014)										(0.010)
ln(Home Value)			0.008									-0.015*
			(0.007)									(0.008)
College				0.037**								0.017
				(0.016)								(0.011)
Employed					0.048***							0.024***
					(0.013)							(0.005)
Black						-0.016**						-0.006*
						(0.006)						(0.003)
Hispanic							-0.029***					-0.009*
							(0.008)					(0.004)
Unionized								0.024				0.025*
								(0.017)				(0.014)
Sector Equity Beta									0.021			0.017*
									(0.020)			(0.010)
Share Retired										0.014***		-0.002
										(0.002)		(0.001)
ln(Avg. Acct. Bal.)											0.099***	0.089***
											(0.005)	(0.005)
Observations	242,737	242,737	242,737	242,737	242,737	242,737	242,737	242,737	239,620	242,737	242,737	239,620
R-squared	0.039	0.043	0.011	0.033	0.046	0.014	0.023	0.010	0.012	0.012	0.152	0.188
Year FE	X	X	X	X	X	X	X	X	X	X	X	X

Table A1 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 level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Table A2: 401(k) Menus vs. Demographics

VARIABLES	(1) US Equity Funds	(2) Bond Funds	(3) Cash Funds	(4) International Funds
Age	0.002	-0.000	0.000	-0.001**
	(0.001)	(0.001)	(0.000)	(0.001)
ln(Income)	0.004**	-0.001	-0.001	-0.002*
	(0.002)	(0.001)	(0.001)	(0.001)
ln(Home Value)	0.002*	-0.000	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
College	-0.005**	0.002	-0.001*	0.002**
	(0.002)	(0.001)	(0.001)	(0.001)
Employed	0.001	-0.000	0.000	-0.000
	(0.001)	(0.001)	(0.000)	(0.001)
Black	0.004***	-0.000	0.001**	-0.003***
	(0.001)	(0.000)	(0.000)	(0.001)
Hispanic	0.001	-0.000	-0.001*	-0.001
	(0.001)	(0.001)	(0.000)	(0.001)
Unionized	0.012***	-0.006**	0.010***	-0.012***
	(0.003)	(0.002)	(0.002)	(0.002)
Sector Equity Beta	-0.002	0.001	-0.000	0.000
	(0.003)	(0.002)	(0.001)	(0.001)
Share Retired	-0.004***	0.002***	0.004***	-0.001
	(0.001)	(0.000)	(0.000)	(0.000)
ln(Avg. Acct. Bal.)	0.003**	-0.003***	0.004***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	240,166	240,166	240,166	240,166
R-squared	0.030	0.012	0.034	0.016
Year FE	X	X	X	X

Table A2 displays the regression results corresponding to a linear regression model (eq. 10). 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 level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

Table A3: Asset Allocation vs. Demographics

VARIABLES	(1) US Equities	(2) Bonds	(3) Cash	(4) International Equities
Age	-0.754***	0.214**	0.810***	-0.232***
	(0.173)	(0.087)	(0.157)	(0.075)
ln(Income)	0.191	-0.655**	0.643**	-0.137
	(0.322)	(0.274)	(0.249)	(0.101)
ln(Home Value)	0.075	-0.396**	0.349	-0.034
	(0.183)	(0.161)	(0.289)	(0.092)
College	0.976***	0.327	-1.502***	0.143
	(0.259)	(0.255)	(0.331)	(0.117)
Employed	0.103	-0.109	-0.012	0.032
	(0.140)	(0.090)	(0.150)	(0.052)
Black	-0.349***	-0.122	0.792***	-0.241***
	(0.106)	(0.080)	(0.138)	(0.058)
Hispanic	-0.642***	-0.068	0.941***	-0.217***
	(0.146)	(0.078)	(0.191)	(0.071)
Unionized	-0.867*	-0.408	3.112***	-1.555***
	(0.415)	(0.526)	(0.448)	(0.135)
Sector Equity Beta	0.588*	-0.473	-0.330	0.150
	(0.293)	(0.461)	(0.341)	(0.118)
Share Retired	-0.317***	0.127**	0.487***	-0.303***
	(0.082)	(0.047)	(0.113)	(0.050)
ln(Avg. Acct. Bal.)	0.898***	-0.165	-0.114	-0.468***
	(0.228)	(0.267)	(0.363)	(0.058)
US Equity Funds	4.919***			
	(0.123)			
Bond Funds		3.396***		
		(0.118)		
Cash Funds			5.454***	
			(0.334)	
International Funds				4.653***
				(0.099)
Observations	240,166	240,166	240,166	240,166
R-squared	0.212	0.143	0.108	0.319
Year FE	X	X	X	X

Table A3 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 *US Equity Funds*, *Cash Funds*, *Bond Funds*, and *International Funds* are 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. 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 level and at the county level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

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

	(1)	(2)	(3)	(4)	(5)	(6)
Relative No. Equity Options	19.4***	26.0***	26.6***	23.3***	28.3***	29.1***
	(0.74)	(0.74)	(0.78)	(0.78)	(0.77)	(0.79)
Offer Company Stock		5.63***	5.40***		5.77***	5.47***
		(0.16)	(0.18)		(0.16)	(0.17)
ln(Total Plan Asset)			0.12***			0.20***
			(0.045)			(0.048)
Constant	57.6***	50.8***	48.0***	55.1***	49.3***	44.7***
	(0.48)	(0.51)	(1.14)	(0.51)	(0.52)	(1.18)
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

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. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table A5: Portfolio Weights vs. Past Fund Returns

VARIABLES	(1) Weight	(2) Δ Weight	(3) Weight	(4) Weight	(5) Weight
Lag Fund Ret.	0.001***	0.000***	0.001***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Lag Fund Ret. x New Investment					-0.001***
					(0.000)
Observations	6,825,623	3,800,916	1,478,618	131,198	6,825,623
R-squared	0.435	0.156	0.462	0.436	0.436
FE	X	X	X	X	X
New Funds			X		
New Plans				X	

Table A5 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 weight of the fund in the investor's portfolio. Each specification include plan-by-year, investment category (MorningstarXBrightScope)-by-year, and index fund-by-year fixed effects. Standard errors are in parenthesis and are clustered at the plan level. *** p<0.01, ** p<0.05, * p<0.10.

Table A6: Return Predictability

	(1)	(2)	(3)	(4)	(5)	(6)
	Future Fund Returns by Investment Horizon					
	1yr	3yr	5yr	1yr	3yr	5yr
Mean Exp.	0.186***	0.381***	0.451***	0.013	-0.028	0.071***
	(0.003)	(0.003)	(0.003)	(0.022)	(0.021)	(0.025)
Observations	79,242	68,913	49,775	79,242	68,913	49,775
R-squared	0.454	0.427	0.439	0.567	0.574	0.546
Year FE	X	X	X	X	X	X
Risk Controls				X	X	X

Table A6 displays the regression results corresponding to a linear regression model (eq. 15). Observations are at the fund-by-year level over the period 2009 through 2019. The dependent variable is the future return measured over a 1 year, 3 year, and 5 year horizon and is annualized. Robust standard errors in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table A7: Model Estimates Accounting for Labor Income Risk(θ and ψ)

	(1)	(2)
θ	-0.106***	
	(0.010)	
ψ	0.282***	0.188***
	(0.019)	(0.033)
Observations	4,720,031	5,946,885
R-squared	0.949	0.956
PlanxYear FE	X	X
CategoryxYear FE	X	
Index FundxYear FE	X	
FundxYear FE		X

Table A7 displays the regression results corresponding to a linear regression model (eq. 17). Observations are at the investment option-by-plan-by-year level over the period 2009-2019. The dependent variable is the additional risk of investing a dollar in a given investment option, given the other portfolio holdings in the plan. We estimate each specification using 2-stage least squares. We instrument for expenses and the corresponding interaction terms using Hausman-type as described in the text. Because each observation reflects the average behavior of plan participants, we weight each observation by the total assets of the 401(k) plan. All specifications include plan-by-year, Morningstar investment category-by-BrightScope investment category-by-year fixed effects, and index-fund-by-year fixed effects. Standard errors are clustered at the plan level and are in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table A8: Alternative Model Specifications

(a) Risk Aversion and Expected Market Returns

VARIABLES	(1) Obs	(2) Mean	(3) Std. Dev.	(4) Median
Risk Aversion: Baseline	243,284	4.512	0.454	4.503
Risk Aversion: Time-Varying Intercept	243,284	4.690	0.897	4.491
Risk Aversion: Simplified Risk Measure	243,284	8.661	0.909	8.541
Risk Aversion: Including Target Date Funds	243,284	4.439	0.450	4.433
Risk Aversion: New Plans Only	4,773	4.615	0.000	4.615
Expected Return: Baseline	243,284	11.496	2.511	11.533
Expected Return: Time-Varying Intercept	243,284	11.824	2.713	11.632
Expected Return: Simplified Risk Measure	243,284	15.623	3.130	15.651
Expected Return: Including Target Date Funds	243,284	11.523	2.278	11.577
Expected Return: New Plans Only	4,773	12.513	2.036	12.873

(b) Correlation: Baseline vs. Alternative Specifications

	(1) Expected Return	(2) Risk Aversion
Model: Time-Varying Intercept	0.797***	0.569***
Model: Simplified Risk Measure	0.815***	0.780***
Model: Including Target Date Funds	0.942***	0.995***
Model: New Plans Only	0.776***	0.000
Observations	243671	243671

Table A8 displays the results for our alternative model specifications. We estimate four alternative specifications. First, in the *Time-Varying Intercept Model* we allow mean risk aversion to vary year-by-year by interacting the variable *Expense Ratio* with year dummy variables in eq. (5). Second, in the *Simplified Risk Measure Model* we calculate the covariance of fund returns using a simplified factor model where we construct the factors by forming equal weighted portfolios based on the broad BrightScope categories reported in Table 1a. Third, in the *Including Target Date Funds Model* we include target date funds when computing portfolios, risk, and expected market returns. Fourth, in the *New Plans Only Model* we estimate the model using data from 401(k)'s in the year the plan is introduced. We focus on the year of inception because it captures the active decisions of investors. Because we have fewer observations in this sample, we keep risk aversion constant across investors/plans in the *New Plans Only Model*. 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. *** $p<0.01$, ** $p<0.05$, * $p<0.10$.