Plan Menus, Retirement Portfolios, and Investors' Welfare

Marco Loseto*

March 2025

Abstract

The recent wave of lawsuits involving 401(k) retirement plans has raised concerns about whether employers act in workers' best interests. While existing evidence suggests employers' plan choices deviate from those of an informed regulator, it remains unclear whether employer intermediation creates value relative to a scenario where workers choose their own retirement plan. In this paper, I evaluate employer intermediation by assessing whether employer-selected plan menus outperform those workers would choose by themselves. I find that workers do not necessarily select better menus because, despite being more fee-sensitive than employers, they also chase past investment performance and prefer active funds, which tend to underperform over long horizons. As a result, policies targeting investment fees—rather than fully realigning employer and worker preferences—prove more effective.

KEYWORDS: 401(k) plans, menu choice, portfolio choice, fund fees.

JEL Classification: G11, G23, G28, G51, J32, L51.

*Bocconi University and IGIER, marco.loseto@unibocconi.it. This paper is based on the second chapter of my PhD dissertation. I am deeply grateful to Lars Peter Hansen, Ali Hortaçsu, and Scott Nelson for their continuous guidance and support. I would also like to thank Michele Fioretti, Nicola Gennaioli, Julien Sauvagnat and Nicolas Serrano-Velarde for their thoughtful feedback. For their helpful comments I thank seminar participants at the SITE 2024, the Yiran Fan memorial conference, the UChicago Economic Dynamics & Financial Markets working group, the UChicago IO lunch and the Booth Finance Seminar. Financial support from the Bradley fellowship, Stevanovich fellowship, Neubauer PhD fellowship and the Bank of Italy's Bonaldo Stringher scholarship is gratefully acknowledged. This research was also funded in part by the John and Serena Liew Fellowship Fund at the Fama-Miller Center for Research in Finance, University of Chicago Booth School of Business. All errors are my own.

1 Introduction

Employer-sponsored 401(k) plans are one of the largest retirement savings vehicles in the United States, holding nearly $8 trillion in assets.[1] Through these plans, workers invest in a menu of mutual funds to build long-term retirement wealth. Employers are responsible for selecting the menu of funds and monitoring the performance of those funds on behalf of workers. Because of this role, they are held to fiduciary standards that require them to act in workers' best interests, with particular attention to the investment fees workers pay to the funds they select.

Nonetheless, investment fees paid by workers vary substantially across 401(k) plans.[2] The wide dispersion in fees has led to numerous lawsuits, with workers alleging that employers violated their fiduciary duty by offering high-fee mutual funds instead of more affordable alternatives.[3] This wave of litigation has, in turn, spurred a growing literature aimed at understanding why employers offer high-fee funds within their plan menus.

Two main explanations have emerged. First, frictions in the market for 401(k) administrative services incentivize employers to shift plan provision costs onto workers by offering high-fee funds that engage in revenue sharing (Pool, Sialm and Stefanescu, 2022). Second, employers themselves do not necessarily value value low-cost investment options as much as an informed regulator would, either due to limited financial expertise or resource constraints that prevent them from covering plan costs directly (Bhattacharya and Illannes, 2022).

While these findings highlight that employer choices deviate from those of an informed regulator, whether employer intermediation is valuable to workers remains an open question. The answer is not obvious because workers themselves make investment decisions that are often not as optimal as those a financially savvy regulator would make.[4]

In this paper, I address this question by assessing whether employer-selected plan menus outperform those workers would choose on their own. This allows me to evaluate whether employer intermediation in the provision of 401(k) plans creates value relative to a world without it. The analysis requires separately measuring employers' and workers' preferences for fund attributes, such as investment fees. To do so, I leverage comprehensive data on 401(k) menus and portfolio allocations, together with an empirical model of employer plan menu choice and an empirical model of worker retirement portfolio choice.

[1] According to the 2019 Survey of Consumer Finances (SCF), for a working-age household, retirement accounts are the second-most commonly held type of financial asset after transaction accounts (https://www.federalreserve.gov/publications/files/scf20.pdf). See also Badarinza et al. (2016).

[2] The average investment fee can differ by as much as 40 basis points between plans at the 75th and 25th percentiles. To put this in perspective: assuming an annual return of 6%, if an employee with an annual income of $70,000 contributes 10% to their 401(k) and shifts from a plan at the 75th percentile to one at the 25th percentile, they could save approximately $95,000 in investment fees.

[3] An overview of some recent lawsuits is provided in Appendix D based on Mellman and Sanzenbacher (2018).

[4] See Reuter (2024) for a recent survey of the literature on plan design and workers' investment behavior in defined contribution retirement plans.

Following Egan, MacKay and Yang (2023), I exploit within plan variation in portfolio allocations to estimate a model of retirement portfolio choice and recover worker preference parameters. Motivated by the growing popularity of default options,[5] the empirical framework I propose differs in that I allow workers to be either inactive or active investors. Inactive workers invest only in the available default fund whereas active workers invest across the non-default funds available in the plan menu. This distinction is important because it allows me to separate workers' preferences from workers' inertia following the identification argument in Abaluck and Adams-Prassl (2021).

The portfolio choice problem of active workers is modelled in two steps: first, they allocate across broad asset classes (i.e., Equity, Bond, and Cash) in a mean-variance fashion; then, within each asset class, they make a single discrete choice, selecting the fund that best matches their preferences. This structure aligns with evidence showing that active investors typically hold fewer than a handful of funds,[6] while also capturing their incentives to diversify across broad asset classes.

To measure employer preferences, I follow a similar approach. Rather than variation in portfolio allocations, I exploit variation in the likelihood of a fund being offered in a plan menu to estimate a model of plan menu choice and recover employer preferences. Because employers can offer more than a single fund, I model their menu choice as a multiple-discrete choice problem, which I show can be estimated using the techniques developed in Berry, Levinsohn and Pakes (1995).

I allow employers to potentially value the same set of fund attributes that workers care about. I specify their preferences for fund characteristics using a random utility framework: employers rank funds within the same investment category (e.g., Small-Cap Growth) based on observed characteristics, such as fees and past performance, and select those that best match their preferences. In principle employers can include as many funds as they want, but adding too many can increase monitoring costs.[7] The model captures this assuming that the likelihood of adding funds decreases as the menu expands.

Beyond fund attributes, employers also need to consider that the provision of 401(k) plans entails large fixed costs, driven primarily by legal and regulatory requirements.[8] Larger employers have a cost advantage because they can spread these costs over more participants. I capture these costs by letting the set of funds employers can choose from depend on the size of the plan, with larger employers having access to funds that are on

[5]This growth was triggered by the 2006 Pension Protection Act which classified Target Date Funds (TDFs) as a default investment option for retirement plans. Previously, Money Market Funds were the common default. TDFs became more appealing because: (1) they offer savers a balanced portfolio of equity and bond funds, which is likely to guarantee higher return at long retirement horizons, and (2) the allocation shifts from equity to bond automatically as workers age toward retirement, thus requiring minimum oversight. TDFs assets in defined contribution retirement plans grew from 48 billions in 2005 to 1.2 trillions in 2021 (ICI, 2022). Nowadays, most workers hold just a single TDF (Vanguard (2023)).

[6]See the recent industry report by Vanguard (2023).

[7]Employers fiduciary duty extends beyond menu design to monitoring included funds and keeping investors informed. A larger menu can make this oversight more challenging (https://www.shrm.org/in).

[8]See for example the analysis in Deloitte (2013).

average cheaper.

When estimating employers' and workers' preferences, I focus on the year a 401(k) plan was introduced for the first time and analyze employers and workers choices in this introduction year. This is done to avoid confounding preferences for fund characteristics with potential switching costs. Compared to employers, I find that workers are about three times more sensitive to fund fees, but they are also more likely to chase past fund performance and prefer actively managed funds.

This divergence in employers' and workers' valuations for fund attributes has important implications in determining the value of employer intermediation. To evaluate whether workers would design better menus themselves, I simulate counterfactual 401(k) menus and portfolio allocations under a scenario where employer preferences align with those of workers.[9] The results show that while this alignment reduces investment fees, it also decreases the availability of passive index funds, which tend to outperform active funds over long horizons and are generally considered higher-quality products.[10] Consistent with the estimated preference misalignment, this occurs because, while workers are more fee-sensitive than employers, they are also more prone to chasing past performance through active funds.

To assess whether the menus workers would choose outperform those selected by employers on their behalf, I compare plan portfolio returns under two scenarios: one where employers internalize only workers' fee-sensitivity and another where they also internalize workers' preference for past performance and active funds. The results indicate that workers' tendency to chase past performance significantly erodes the return gains from aligning employer preferences with theirs. While in the first scenario workers benefit from higher returns relative to the status quo, they are not necessarily better off when employers also internalize their performance-chasing behavior.

From a practical perspective, these findings suggests that policies directly targeting investment costs can improve workers' outcomes more effectively than those aimed at realigning employer and worker preferences. In light of this, I conclude the paper by examining two policies that mandate the inclusion of low-cost options: one requiring employers to offer the cheapest equity index fund available and another requiring them to offer the cheapest default fund. As before, I simulate counterfactual 401(k) menus and portfolio allocations for each policy and then compute the implied fee savings for workers.

The results indicate that regulating the cost efficiency of default options generates fee savings more than three times larger than those achieved by mandating the inclusion of a low-cost equity index fund. This difference arises despite the fact that the cheapest

[9] In predicting these counterfactual menus, I maintain the assumption that the set of funds available to employers depends on their size, ensuring that worker-selected menus remain feasible within employers' resource constraints.

[10] A long-standing literature in finance finds that active fund performance does not persist, fails to predict future performance, and does not justify their higher fees (Jensen (1968), Gruber (1996), Carhart (1997), Gil-Bazo and Ruiz-Verdu (2009), Fama and Kenneth (2010)). These findings are even more pronounced over long investment horizons (Bessembinder, Cooper and Zhang (2023)).

index funds are significantly less expensive than the cheapest default funds and reflects three key drivers: (1) the first policy affects active workers, whereas the second impacts inactive workers, who constitute the majority in this market; (2) not every active worker will switch to the new low-cost index fund because they still value past performance and other fund attributes; and (3) active workers that pick the new low-cost index reallocate only part of their portfolio, as they have a strong preference for maintaining a diversified portfolio across broad asset classes, regardless of which funds are offered within each class.

Overall, this exercise highlights that workers' investment behavior plays a key role in shaping the effectiveness of policies aimed at reducing investment costs for 401(k) investors and at improving their retirement outcomes.

Related literature. This paper contributes to three strands of the literature.

First, it contributes to the growing literature on 401(k) retirement plans. Much of this literature focuses either on how workers make retirement portfolio decisions (Madrian and Shea (2001), Beshears, Choi, Laibson and Madrian (2009), Carroll, Choi, Laibson, Madrian and Metrick (2009), Choi (2015), Benartzi and Thaler (2001), Benartzi and Thaler (2007), Tang et al. (2010), Chalmers and Reuter (2020), Egan, MacKay and Yang (2021), Duarte et al. (2022)) or on how employer sponsors design their 401(k) plans (Pool, Sialm and Stefanescu (2016), Pool, Sialm and Stefanescu (2022), Badoer, Costello and James (2020), Bhattacharya and Illannes (2022), Gropper (2023), Yang (2023)). In contrast, this paper jointly considers both aspects within an empirical model that predicts how employers' plan menus and workers' portfolio allocations would change under different counterfactual scenarios. This dual focus is key to evaluating the effectiveness of plan design policies aimed at reducing workers' investment costs, as the counterfactual analysis in Section 5 demonstrates.

Second, this paper contributes to the literature on demand estimation within the discrete-choice framework (McFadden (1974), Berry (1994), Berry, Levinsohn and Pakes (1995), Conlon and Gortmaker (2020), Berry and Haile (2021), Gandhi and Nevo (2021)) and is closely related to studies on multiple-discrete choice demand. Some of these studies model choices of multiple units of products that are perfect substitutes (Hendel (1999), Dubé (2004)), while others analyze the selection of multiple products that are imperfect substitutes (Beggs, Cardell and Hausman (1981), Berry, Levinsohn and Pakes (2004)).[11] This paper is more closely related to the latter but differs in that the menu choice framework I consider does not require information on how consumers rank products and can be estimated using workhorse demand estimation techniques.[12] More broadly, the frame-

[11] Models of this second type have been widely applied to study school choice and infer students' preferences from school rankings. See Agarwal and Somaini (2020) for a recent review.

[12] Alternatively, the choice of purchasing multiple products can be modeled as a single-discrete choice over the set of possible product bundles. This approach has been used to account for product complementarity in discrete choice models (see, for example, Gentzkow (2007) and Wang (2024)). However, a key challenge with this approach is that the set of possible bundles can quickly become too large for

work can be applied, beyond the 401(k) context, to recover the preferences of any type of intermediary in decision-making environments where they design choice sets on behalf of consumers or investors.

Lastly, my paper joins a growing literature that applies empirical industrial organization techniques to study financial markets. Clark, Houde and Kastl (2021) and Einav, Finkelstein and Mahoney (2021) provide comprehensive reviews of recent contributions. Related applications include research on pensions outside the 401(k) setting (Hastings, Hortaçsu and Syverson (2017), Illanes (2017), Luco (2019), Illanes and Padi (2021), Boehm (2024)) and studies on the market for financial advice (Egan (2019), Bhattacharya, Illanes and Padi (2020), Robles-Garcia (2020), Egan, Ge and Tang (2022), Guiso et al. (2022)).

The reminder of the paper proceeds as follows. Section 2 describes the institutional settings of the US market of 401(k) plans and the data used in the analysis. Section 3 provides descriptive evidence regarding the determinants of variation in investment fees and retirement portfolio allocations, which motivates several modeling assumptions. Section 4 introduces the empirical framework which I use to study counterfactual policies in Section 5. Section 6 concludes.

2 Institutional Settings and Data Sources

In this section I describe what is an employer-sponsored retirement plan and overview its administrative structure. After that I describe the data sources and provide some summary statistics about my sample of retirement plan menus and the investment options available therein.[13]

2.1 What is an employer-sponsored retirement plan?

Employer-sponsored retirement plans are savings vehicles designed to assist workers in accumulating wealth for their retirement years. Although there are various types of employer-sponsored retirement plans, the most common is the defined contribution (DC) plan. At its core, a defined contribution plan is a retirement plan in which an employee makes regular contributions. The final amount available upon retirement is not predetermined but instead depends on the contributions made and on the returns obtained from the investment options available in the plan.

Under a DC plan, workers' contributions represent a percentage of their salary which is subtracted from their paycheck before taxes, thereby reducing their current taxable income. The gains in the retirement account grow tax-deferred, implying taxes are not due until funds are withdrawn during retirement years. Withdrawals from a DC plan

computational feasibility. This issue is particularly pronounced in the context of 401(k) menu choice, where employers choose from thousands of funds and can design menus of virtually any size.

[13] Appendix C offers a detailed discussion on the data construction.

before a certain age (typically between 59 and 60) can result in penalties. After reaching retirement age, participants might withdraw distributions as lump sums, period payments or annuities. Importantly, all withdrawals at this point are subject to standard income tax.

Employers play a crucial role in the provision and design of these type of plans. First, many employers offer to match a portion of the worker's contributions. For example, an employer might contribute 50 cents for every dollar the employee contributes, up to a certain percentage of the worker's salary (typically around 5%). More importantly, employers are in charge of selecting and monitoring which investment options are to be included in the plan. Once the worker decides what percentage of their income to contribute, they typically have the autonomy to allocate their contributions, together with the part matched by their employer, across the options available within their plan. In many cases, to help less financially savvy employees or those who do not make an active investment choice, plan sponsors include options, such as TDFs, known as Qualified Default Investment Alternative (QDIA) to be used as default option when an employee contributes to the plan without specifying how their contribution should be invested.

Figure 1 illustrates the administrative structure of a 401(k) plan. Employer sponsors typically hire external plan administrators, known as recordkeepers, to assist in selecting the menu of funds and performing administrative tasks such as maintaining account balances and complying with regulatory standards. Recordkeepers are compensated in two ways: part of their compensation comes directly from workers in the form of administrative fees, usually a fixed dollar amount per plan participant. The second part of their compensation is indirect, coming from the investment fees workers pay to the funds, which rebate a portion of their fee revenues to the plan's recordkeeper. Many recordkeepers, such as Fidelity, also provide investment funds, and it is common for retirement plans to include funds from the recordkeeper's product line. For proprietary funds, all fee revenues accrue to the recordkeeper. Evidence from industry practitioners indicates that workers bear almost all plan costs, with investment fees accounting for more than 80% of total fees.[14]

Overall, a retirement plan consists in a menu of funds, spanning a broad range of investment categories, selected by employers and recordkeepers. Under the Employee Retirement Income Security Act (ERISA), employers are fiduciaries to their employees and are subject to litigation risk if their retirement plan is not designed in the employees' best interest. The inclusion of high-cost investment options or the lack of low-cost options are among the main triggers of some recent lawsuits, as I discuss in Appendix D.

[14]The analysis in Deloitte (2013) finds that employees pay, on average, 87% of total fees and that investment fees account for 82% of total fees, where total fees are defined as administrative fees plus investment fees.

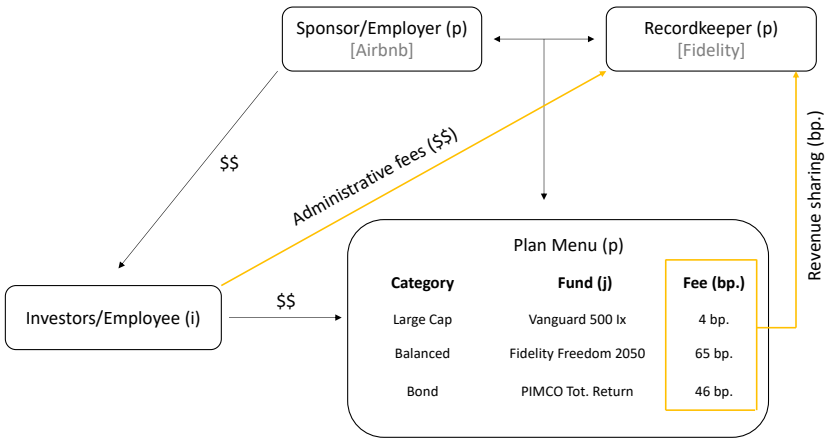


Figure 1: Administrative structure of an employer-sponsored 401(k) plan.

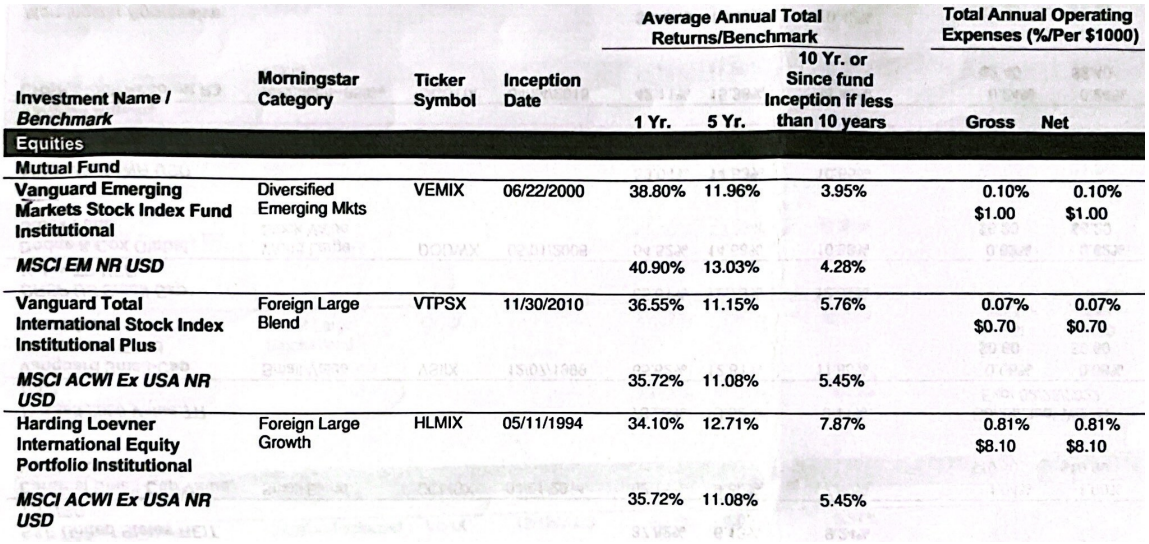
2.2 Data

The primary data source for this study is collected from Form 5500. This form is annually filed by employers with the Department of Labor (DOL) in adherence with the ERISA regulations. Within this form, Schedule H provides detailed information about retirement plan menus. Specifically, it contains information regarding the investment options offered within an employer's retirement plan and the plan-level dollar allocation to each of these options.[15]

Although Form 5500 filings are publicly available from the DOL website, information about plan menus comes in a non-standardized format, stored in pdf images, that would need to be digitized manually. I acquired the digitized version of these filings from BrightScope who collects and digitizes these data directly from DOL filings. I complement this data with additional information available from Form 5500, including the number of plan participants, employers' and workers' contributions, and the administrative fees paid directly to the recordkeeper.

Workers and employers have access to more information than just the list of available funds. Employers are required to provide workers with up-to-date information about the characteristics of the funds in the menu, including past investment performance and fees. Figure 2 illustrates a concrete example of the information provided to them. For each fund, workers know the investment category, the fund inception date, and the fund's past performance over the previous year as well as over longer horizons. To assess past performance, workers compare the fund's return to that of its benchmark, which is also part of the available information. Finally, workers have access to information about the fees (also known as expense ratios) that each fund charges its investors. I obtain data

[15]Figure A.1 shows Airbnb 2017's 401(k) plan menu, as reported in form F5500.



Investment Name / *Benchmark*	Morningstar Category	Ticker Symbol	Inception Date	Average Annual Total Returns/Benchmark 1 Yr.	5 Yr.	10 Yr. or Since fund Inception if less than 10 years	Total Annual Operating Expenses (%/Per $1000) Gross	Net
Equities								
Mutual Fund								
Vanguard Emerging Markets Stock Index Fund Institutional	Diversified Emerging Mkts	VEMIX	06/22/2000	38.80%	11.96%	3.95%	0.10% $1.00	0.10% $1.00
MSCI EM NR USD				40.90%	13.03%	4.28%		
Vanguard Total International Stock Index Institutional Plus	Foreign Large Blend	VTPSX	11/30/2010	36.55%	11.15%	5.76%	0.07% $0.70	0.07% $0.70
MSCI ACWI Ex USA NR USD				35.72%	11.08%	5.45%		
Harding Loevner International Equity Portfolio Institutional	Foreign Large Growth	HLMIX	05/11/1994	34.10%	12.71%	7.87%	0.81% $8.10	0.81% $8.10
MSCI ACWI Ex USA NR USD				35.72%	11.08%	5.45%		

Figure 2: Example of information on fund characteristics available to employers and workers. It includes details about fund investment category, inception date and expense ratio (last two columns). Information on fund annual returns is provided together with information on the returns of the benchmark the fund is compared to. For example, the Harding Loevner International Equity Portfolio fund underperformed its benchmark in the previous year by almost 2%.

on fund returns, inception dates, and expense ratios from the Center for Research in Security Prices (CRSP) and merge this information with the main dataset using fund tickers. Benchmark return data is obtained from Morningstar.

For my analysis, I focus on the introduction year of 401(k) plans first created between 2016 and 2020. This approach helps avoid confounding preferences for fund characteristics with employers' and workers' inertia or switching costs. Plan menus and portfolio allocations tend to be sticky over time, and the static menu choice I consider cannot distinguish whether employers select a fund more frequently because they prefer it over other alternatives or because it was already part of the menu in previous years. Similarly, I cannot determine whether workers invest more in a fund because they prefer it or because updating their previous portfolio allocations is costly.

There are approximately 3,700 new plans created between 2016 and 2019. I restrict the attention to those with a number of plan participants greater than 100 and with a number of investment options in between 5 and 50.[16] I also focus on plans in which at least 80% of the options are mutual funds and exclude plans offering brokerage accounts.[17] The resulting sample consists of approximately 1,900 plans.

Table 1 presents summary statistics at the plan level. The average plan in the sample holds nearly $13 million in assets and has around 600 participants. Both measures of

[16]Less than 1% of plans offers more than 50 investment options.

[17]In the sample of new plans, half of them have more than 95% of options classified as mutual funds. Separate accounts are the second most common investment option. For those, I do not have information about characteristics such as fees. Approximately 5% of plans offer brokerage accounts, I exclude these plans because I do not observe the funds workers can access through these accounts.

	N	Mean	Std	1%	25%	50%	75%	99%
Plan participants (#)	1956	613.73	2966.90	104.00	150.00	226.50	409.00	5877.80
Total assets (mln.)	1955	12.39	50.95	0.01	0.26	1.43	6.41	249.61
Account balance (thous.)	1956	20.13	37.59	0.03	1.03	5.43	21.77	184.87
Employer contribution (%)	1587	28.52	16.45	0.00	18.89	31.01	38.00	80.81
Num. of options	1956	24.91	6.89	9.00	21.00	25.00	29.00	44.00
Index fund (%)	1956	39.81	28.78	0.00	13.79	41.52	62.07	100.00
Has TDF (%)	1956	89.47	30.70	0.00	100.00	100.00	100.00	100.00
Auto enrollment (%)	1956	30.62	46.10	0.00	0.00	0.00	100.00	100.00
Integrated recordkeeper (%)	1478	60.49	48.90	0.00	0.00	100.00	100.00	100.00
Proprietary fund (%)	1478	19.82	30.59	0.00	0.00	3.45	28.09	100.00
Avg. expense ratio (pp.)	1955	0.43	0.27	0.08	0.24	0.34	0.56	1.29
Administrative fees (%)	1940	0.37	0.97	0.00	0.03	0.14	0.36	3.92

Table 1: Summary statistics for retirement plans established between 2016 and 2020. Statistics correspond to the plan's introduction year. Account balance is defined as total plan assets divided by the number of participants. Employer contribution is the employer's share of total contributions (in percentage). Administrative fees are expressed as a percentage of plan assets.

plan size are right-skewed due to the presence of large plans, with the median plan having approximately $1.5 million in assets and 226 participants. On average, employers contribute about 30% of total contributions.

The typical plan offers 25 investment options,[18] of which 40% are index funds. Nine out of ten plans include a Target Date Fund (TDF), and 30% of plans automatically enroll employees unless they opt out. More than half of the plans are administered by recordkeepers who are also vertically integrated into the provision of mutual funds, and, on average, about 20% of the funds offered are proprietary.

In terms of fees, the average plan expense ratio is about 43 basis points, while administrative fees paid directly to the recordkeeper amount to about 0.37% of plan assets. The distribution of administrative fees is right-skewed, with the median plan paying around 0.14%—half the average fee level.

3 Descriptive Evidence

I now provide a series of facts to motivate some of the modelling assumptions. I start by describing what explains variation in investment fees across sponsors, where investment fee is defined as the plan-level average expense ratio. After that, I present some facts on workers' retirement portfolios both at the plan-level and at the individual-level.

[18] TDFs from the same Target Date Series are counted as a different investment options.

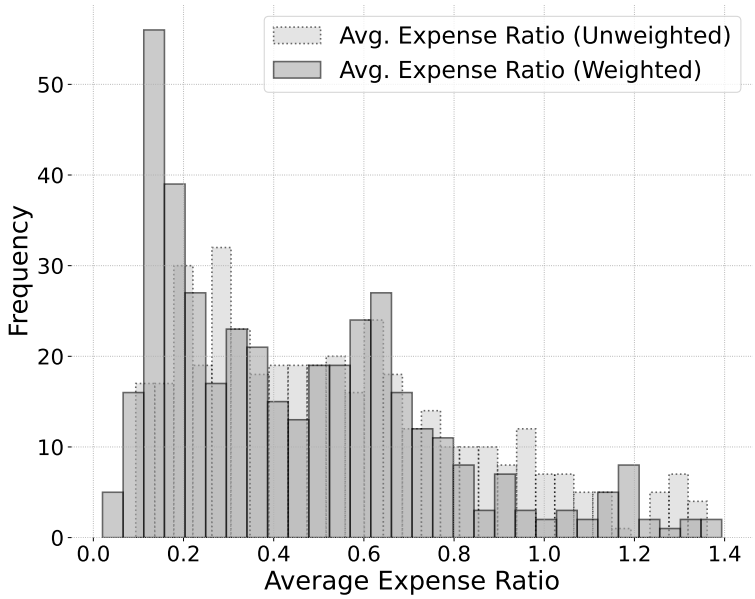


Figure 3: Distribution of the average expense ratio in newly established plans (2016 cross-section). The light gray distribution represents the unweighted average expense ratio, computed as the simple mean of fund expense ratios within each plan. The dark gray distribution represents the asset-weighted average expense ratio, where fund expense ratios are weighted by their share of plan assets.

3.1 What explains variation in plan expenses?

Investment fees that workers pay to funds vary significantly across plans. Figure 3 plots the average expense ratio for the 2016 cross-section of newly introduced plans. Both the unweighted and asset-weighted distributions display substantial dispersion, with the difference in investment fees between plans at the 90th and 10th percentiles being as large as 70 basis points. The asset-weighted distribution is slighted shifted to the left suggesting that workers tilt their portfolio allocations towards cheaper funds.

To better understand what explains this variation, I regress the plan-level average expense ratio on several plan characteristics, and report the results in Table 2. In the first specification, I regress the average fee on the share of employer contributions and the administrative fees paid directly to the plan recordkeeper. Both variables are not statistically significant, and remain so across different specifications, suggesting that workers facing higher investment fees are not compensated with more generous matching rates or lower administrative fees. In fact, the results indicate that employers offering cheaper funds may be also offering more generous contribution rates.

The second specification controls for two measures of plan size: the number of participants and the average account balance. Consistent with the presence of fixed costs in administering these types of plans, both measures are statistically significant and negatively correlated with investment fees, indicating that larger plans tend to offer less expensive funds.

	Coeff.	Std. Error	Coeff.	Std. Error	Coeff.	Std. Error
Employer contributions (%)	-0.073*	(0.039)	-0.059	(0.039)	-0.021	(0.026)
Administrative fees (pp.)	-0.332	(0.684)	-0.610	(0.637)	-0.072	(0.375)
log(# of particp.)			-1.788***	(0.652)	-0.685*	(0.416)
Avg. account balance (thous.)			-0.046***	(0.011)	-0.047***	(0.008)
Auto enrollment					-0.017	(0.861)
# of asset classes					-1.923***	(0.606)
Index fund (%)					-0.653***	(0.016)
# of funds					0.471***	(0.097)
Avg. performance (pp.)					-3.657***	(0.470)
R2	0.084		0.091		0.589	
N	1,573		1,573		1,573	

Table 2: Linear regression of plan average expense ratio on plan characteristics. Year fixed effects are included. Average expense ratio is measured in basis points. Robust standard errors are in parentheses. *** indicates 1% significance, ** indicates 5% significance, and * indicates 10% significance.

In the final specification, I control for plan characteristics, including information about the composition of funds in the menu. Among these characteristics, factors related to the composition of the menu, such as the share of index funds, have significant power in explaining variation in investment fees, with the regression R2 increasing more than six folds.

Overall, the variation in fees seems to be explained primarily by the characteristics of the funds in the menu, rather than by employers' contribution rates or administrative fees.[19] Based on this evidence, in modeling employers' and workers' preferences, I will focus on fund characteristics and abstract from how matching rates and administrative fees are determined. To capture the presence of fixed costs, I will let employers' choice set depend on the size of the plan measured in terms of number of plan participants, with larger employers having access to cheaper funds.

3.2 How do 401(k) plan participants invest?

The role of employer sponsors is limited to providing the menu of funds. After enrolling in the plan, workers are free to allocate their retirement contributions as they see fit. To understand how workers choose their portfolios, Table 3 presents the average plan-level portfolio share across different asset classes. Notably, the vast majority of plan assets are allocated to Target Date Funds (TDFs), with an average portfolio share of about 60%.

[19]Revenue sharing from high-fee funds may also cover unobserved services beyond standard investment and administrative functions. Examples could be advising or educational services for employees or employers. These unobserved services could contribute to variation in investment fees across plans. In Table A.1, I assess their potential impact by including recordkeeper and employer industry fixed effects to the last regression specification in Table 2. Recordkeeper fixed effects could capture variation in unobserved services across recordkeepers, while industry fixed effects account for demand-driven differences. Neither set of fixed effects significantly explains fee variation beyond fund characteristics, suggesting that the influence of unobserved services on investment fees may be limited.

Asset Class	Mean	SD
TDFs	0.61	0.25
US Equity	0.25	0.18
Cash	0.09	0.10
Non-US Equity	0.06	0.06
Balanced	0.08	0.14
Bond	0.03	0.05
Other	0.02	0.02

Table 3: Average portfolio share by asset class. BrightScope sample.

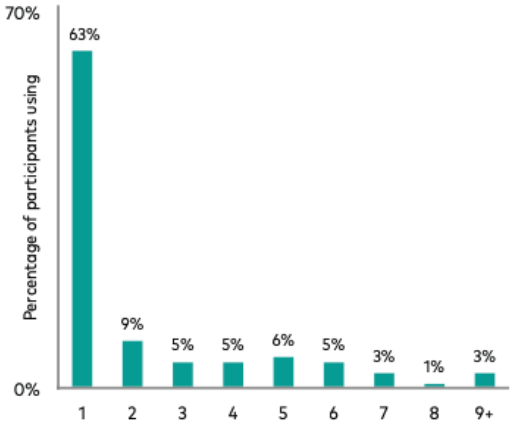


Figure 4: Distribution of number of funds held. Source: Vanguard (2023).

Evidence of the widespread use of TDFs also emerges from individual retirement account data. Figure 4 shows the distribution of the number of funds used by participants in a sample of plans administered by Vanguard. As of 2022, more than 60% of retirement accounts held just a single fund, and over 90% of these single-fund holders invested in TDFs.[20]

Why are TDFs so popular? Their surge in popularity is primarily due to their role as default investment options in retirement plans. These funds automatically adjust their asset allocation to become more conservative as the target date—typically the investor's expected retirement year—approaches. This automatic rebalancing makes TDFs an attractive choice for retirement savings, requiring minimal oversight from investors. Furthermore, their designation as the default option in many employer-sponsored retirement plans, following the 2006 Pension Protection Act, has significantly contributed to their growth. According to the Investment Company Institute (ICI, 2022), assets held in TDFs within defined contribution plans increased dramatically, from $48 billion in 2005 to approximately $1.2 trillion by 2021.

In Section 4.3, I build on this evidence to model workers' portfolio choice. I assume that a worker can be either inactive or active. Inactive workers do not consider the funds available in the plan menu and invest exclusively in the default TDF that best matches their expected retirement date. Conversely, active workers disregard the TDF and choose to form their optimal retirement portfolio from the other funds available. This assumption is motivated by the idea that TDFs are an all-in-one investment solution, which is supported by recent empirical evidence from individual retirement accounts. Specifically, Mitchell and Utkus (2021) analyze a sample of approximately 1.2 million retirement accounts from 880 plans administered by Vanguard, finding that only about 110,000 accounts invest in both TDFs and other funds.

[20] See Figure A.2.

4 Empirical Model

This section develops the empirical framework used to identify and estimate employers' and workers' preferences. First, I formalize employers' menu choice problem and estimate it from variation in the observed 401(k) plan menus. After that, I model workers' retirement portfolio problem and estimate their preferences from variation in the observed portfolio allocations.

4.1 Employer employer sponsors' menu choice problem

Employer p needs to choose a set of mutual funds, indexed by j, to be included in their 401(k) plan. Following the discrete choice literature, I specify employers' preferences for fund characteristics with a random utility framework. Employer p's utility from including fund $j \in J_p$ in the menu is

$$u_{pj} = x_j'\theta_p + \xi_j + \varepsilon_{pj},$$

where J_p is the set of funds available to p, x_j is a vector of observed fund characteristics, θ_p is a vector of preference parameters, ξ_j is a scalar capturing any characteristic unobserved to the researcher and ε_{pj} is an idiosyncratic T1EV taste shock.

This preference specification, where employers derive utility from fund characteristics, is convenient for two reasons. First, since workers naturally care about fund characteristics, it's reasonable to assume employers care as well, especially because their fiduciary duty requires them to choose a menu of funds in the workers' best interest. Second, this framework is empirically tractable and enables me to cast the estimation of employers' preferences into a demand estimation type of problem.

Unlike workers, employers, especially the smaller ones, need to consider that the provision of 401(k) plans is costly. I do not include these costs directly into employers' preferences, rather I capture them by letting the set of funds employers can choose from depend on the size of the plan. Specifically, I categorize employers into two groups—small (fewer than 500 plan participants) and large (500 or more plan participants)—and restrict them to choose among the funds offered by employers within their group. This captures the fact that smaller employers may offer expensive funds to cover for higher plan provision costs.[21]

I model employers' menu choice problem as a three stages process. First, employers choose which investment category, indexed by $g \in \{1, ..., G\}$, they want to offer. Second, they choose how many funds, n, to include within each investment category. Lastly, they include the n funds, $(j_1, ..., j_n)$ with the highest utility draws $u_{pj_1} > u_{pj_2} > ... > u_{pj_n}$.

[21] Table A.6 compares the choice sets of small and large employers. Large employers on average offer cheaper funds, more index funds and more TDFs. In practice, plan provision costs can also be covered directly through administrative fees, but such fees typically make up only a small portion of total plan costs, while fund fees represent the largest share (Deloitte (2013)).

Letting λ_{pg} the probability that employer p offers category g and $q_{pg}(z)$ the probability that sponsor p offers z funds within g, the model-implied probability that sponsor p includes fund j in its menu is:[22]

$$\underbrace{\phi_{jp}(x,\theta_p)}_{\text{fund } j \text{ inclusion probability}} = \underbrace{\lambda_{pg}(x,\theta_p)}_{\text{prob category } g \text{ included}} \sum_n \underbrace{(1-Q_{pg}(n))}_{\text{prob num of funds} \geq n} \underbrace{\phi^n_{pj}(x,\theta_p)}_{\text{prob } u_{pj} \text{ is } n^{th} \text{ highest}} \tag{1}$$

where $Q_{pg}(n) \equiv \sum_{z<n} q_{pg}(z)$ if $n > 1$ and 0 otherwise.

I assume that the probability of offering category g is a logit function of the investment category's inclusive value,[23] which I define as:

$$\lambda_{pg}(x,\theta_p) = \frac{\exp(I_g(x,\theta_p))}{1+\exp(I_g(x,\theta_p))} \quad \text{with} \quad I_g(x,\theta_p) \equiv \mathbb{E}_q\left[\frac{1}{n}\sum_{s=1}^{n}\mathbb{E}_\varepsilon[u_{pj_s}]\right].$$

This definition generalizes the concept of inclusive value from the single-discrete choice framework to the multiple-discrete choice setting. The key difference is that, in the multiple-discrete choice setting, the inclusive value is a function of the expected stream of utility across all the products chosen, which means it depends not only on the expected maximum but also on higher-order statistics of the utility distribution. If employers were forced to choose a single fund (i.e., $q_{pg}(1) = 1$), expression (1) would collapse to the familiar logit choice probability:[24]

$$\phi_{jp}(x,\theta_p) = \lambda_{pg}(x,\theta_p)\phi^1_{pj}(x,\theta_p) = \frac{\exp(x_j'\theta_p + \xi_j)}{\sum_{k\in g}\exp(x_k'\theta_p+\xi_k)}.$$

Once the decision of including category g is made, employers decide how many funds to offer. I assume that the number of options included is an integer drawn from a distribution $q_{pg}(n)$, which I calibrate to match the observed distribution of the number of options offered within each category, separately for small and large plans.[25] In the data, the probability of including more funds decays as the number of selected funds grows, consistent with employers facing some costs to adding more options to the menu.[26] These

[22] See Appendix B for a detailed derivation.

[23] This assumption implies that the decision of which investment category to offer is independent across investment categories, which allows me treat them as separate "markets" in estimation. Within each category, funds are substitutes to each other. In estimation, I use the investment category definition available in BrightScope, with the exception of TDFs, which I treat as a separate category. The categories are: 'Allocation', 'Commodities', 'EAFE', 'Emerging Markets', 'Foreign Bonds', 'Large Blend', 'Large-Growth U.S. Equities', 'Large-Value U.S. Equities', 'Mid Cap U.S. Equities', 'Mid-Growth U.S. Equities', 'Mid-Value U.S. Equities', 'Real Estate', 'Small Cap U.S. Equities', 'Small-Growth U.S. Equities', 'Small-Value U.S. Equities', 'TDF', 'Total Bond Market', 'World Equity'.

[24] The single-discrete choice problem can always be factored in two stages: first, choose between buying from the available products (i.e., the *inside goods*) or the outside option (i.e., the *outside goods*), then choose the best among the inside goods. In this case, the choice between inside and outside goods depends on the inclusive value of the inside goods. See Appendix B for more details.

[25] See Appendix B for more details on the calibration procedure.

[26] This decay is common across all investment categories and is particularly pronounced for $n \geq 3$. See, for example, Figure A.3.

costs arise partly because employer sponsors have the fiduciary duty to oversee the funds in the menu and ensure each of them is still a prudent investment option.[27]

In the last stage of the menu choice problem, employers select the n funds that best match their preferences. After taste shocks realize, employers rank each fund and select the n ones $(j_1, ..., j_n)$ with the highest match values $u_{j_1} > ... > u_{j_n}$. Under the assumption that taste shocks are T1EV, the probability that fund j is ranked in the n^{th} position has the following closed-form expression

$$\phi_{pj}^{n} = \sum_{(j_1,...,j_{n-1}) \in g/\{j\}} \prod_{z=1}^{n-1} \frac{\exp(v_{pj_z})}{\sum_{z'=z}^{J_{pg}} \exp(v_{pj_{z'}})} \cdot \frac{\exp(v_{jp})}{\sum_{z''=n}^{N_{gp}} \exp(v_{j_{z''}p})}$$

where $v_{pj} \equiv x_j'\theta_p + \xi_j$ denotes employer p's mean utility.[28]

Overall, equation (1) represents employer p demand for investment funds belonging to investment category g. Assuming employer preferences θ_p follow some parametric distribution F_θ, we can derive fund j's aggregate demand as

$$\phi_j(x, F_\theta) \equiv \int \lambda_{pg}(x, \theta_p) \sum_n (1 - Q_{pg}(n)) \phi_{pj}^{n}(x, \theta_p) dF_\theta(\theta_p). \tag{2}$$

The data counterpart to (2) corresponds to the share of plans that include fund j as part of their retirement menu. Thus, employer preferences can be estimated by matching the inclusion probabilities implied by (2) with the observed ones.

4.2 Identification and estimation of employers' preferences

The menu choice problem derived in the previous section implies a system of inclusion probabilities that, given employer preferences and fund characteristics, predicts the share of plan menus offering any given fund. Despite the multiple-discrete nature of the problem, the demand system in (2) can be inverted to recover employers' preferences. Specifically, as I show next, the contraction mapping proposed in Berry, Levinsohn and Pakes (1995) to estimate discrete choice demand systems for differentiated products generalizes to this multiple-discrete choice setting, allowing me to cast the problem of estimating employer preferences into a standard demand estimation problem.

Before defining the contraction mapping, it is useful to decompose employers' preference parameters into a common component and a component that varies across employers. Specifically, let the vector $\theta_p = \eta + \Pi\nu_p$, and redefine employers' mean utility as

$$v_{pj} = \underbrace{x_j'\eta + \xi_j}_{\equiv \delta_j} + x_j'\Pi\nu_p,$$

where η is a vector of parameters common across employers, ν_p is a vector of employer

[27] See discussion in a recent article https://www.shrm.org/in.

[28] See Appendix B for a detailed derivation.

characteristics whose distribution F_v is known and Π is a conformable matrix of preference parameters, capturing heterogenity in employer preferences.[29] The system of inclusion probabilities in (2) can be defined as function of $\delta = (\delta_1, ..., \delta_{J_g})$ and Π:[30]

$$\phi_j(\delta, \Pi) \equiv \int \lambda_{pg}(\nu_p, \delta, \Pi) \sum_n (1 - Q_{pg}(n)) \phi^n_{pj}(\nu_p, \delta, \Pi) dF_v(v_p). \tag{3}$$

The data counterpart to (3), is the share of plan menus that offer fund j, which I refer to as the observed inclusion probability and denote it by $\hat{\phi}_j$. For given Π, in Appendix B I show that the following operator $T : \mathbb{R}^{J_g} \to \mathbb{R}^{J_g}$

$$T(\delta) = \delta + \log(\hat{\phi}) - \log(\phi(\delta)) \tag{4}$$

is a contraction with modulus less than one. Following BLP (1995), this implies that, for any given vector of parameters (η, Π), we can compute the unique $\delta(\Pi)$ that matches observed and predicted inclusion probabilities and in turn compute the model residuals $\xi(\eta, \Pi) = \delta(\Pi) - X\eta$. Assuming we have a vector of observables Z_j such that $\mathbb{E}[Z_j \xi_j] = 0$, we can search iteratively for (η, Π) that minimizes the implied GMM objective.[31]

A key step in proving that (4) defines a contraction mapping is showing that the norm of the gradient of T is bounded by one. As I show in Appendix B, this holds when the norm of the gradient of λ_{pg} is itself bounded by one:

$$\sum_{k=1}^{J_g} \frac{\partial \lambda_{pg}(\delta, \nu_p)}{\partial \delta_k} = \frac{\partial \lambda_{pg}}{\partial I_g} \sum_{k=1}^{J_g} \frac{\partial I_g(\delta, \nu_p)}{\partial \delta_k} = \frac{\partial \lambda_{pg}}{\partial I_g} < 1. \tag{5}$$

Intuitively, how the probability of including category g changes with δ is the only thing that matters because, conditional on including the category, increasing all mean utility levels by the same amount does not change the inclusion probabilities. Moreover, when λ is a function of the category's inclusive value, $I_g(\delta, \nu_p)$, inequality (5) is satisfied as long as $\partial \lambda_{pg} / \partial I_g$ is lower than one. This is because, for any product k and positive integer s, one can show that the gradient of the expected s^{th} highest utility is the s^{th} ranking probability

$$\frac{\partial}{\partial \delta_k} \mathbb{E}_\epsilon \left[u_{pj_s} \right] = \phi^s_{pk}(\delta), \tag{6}$$

which in turn implies that increasing all mean utility levels by the same amount increases

[29] The vector ν_p could also capture taste shocks for particular fund characteristics.

[30] I am assuming that all employers face the same set of funds $\{1, ..., J_g\}$ in category g. In practice, I group employers into large and small plans, and assume the choice set is the same within these two groups. In estimation, I compute and match the observed inclusion probabilities separately for these two groups.

[31] Although the estimation algorithm closely follows BLP (1995), the multiple-discrete choice aspect is an important departure. Modeling the menu choice problem this way is essential to match observed inclusion probabilities, which, unlike product market shares, do not need to sum to one.

	Coefficient	Std. Error	WTP (pp.)	WTP Std. Error
log(Age)	0.097***	(0.011)	0.206***	(0.025)
Index fund	0.270***	(0.034)	0.572***	(0.090)
Integrated provider	0.220***	(0.017)	0.467***	(0.050)
Short performance	0.003**	(0.002)	0.007**	(0.003)
Long performance	0.016***	(0.002)	0.034***	(0.005)
Expense ratio	-0.472***	(0.028)		
N	8,290			
R2	0.669			
$Q_{pg}(n)$	calibrated			

Table 4: Estimates of employer utility parameters obtained from the inversion of fund inclusion probabilities in (3). The distribution of the number of options $Q_{pg}(n)$ is calibrated within each investment category, separately for small and large plans. Fixed effects at the plan size group level, year level, and category level are included. Fund performance and fund expenses are in percentage points (pp.). Robust standard errors are in parentheses. *** indicates 1% significance, ** indicates 5% significance and * indicates 10% significance.

the category's inclusive value by one,

$$\sum_{k=1}^{J_g} \frac{\partial I_g(\delta)}{\partial \delta_k} = \sum_{k=1}^{J_g} \mathbb{E}_q \left[\frac{1}{n} \sum_{s=1}^{n} \phi_{pk}^s(\delta) \right] = \mathbb{E}_q \left[\frac{1}{n} \sum_{s=1}^{n} \sum_{k=1}^{J_g} \phi_{pk}^s(\delta) \right] = \mathbb{E}_q \left[\frac{1}{n} n \right] = 1.$$

Equation (6) is a generalization of the Williams-Daly-Zachary (WDZ) theorem, which states that, in the single-discrete choice random utility model, the choice probability corresponds to the gradient of the expected utility of the highest-ranked product.[32] In Appendix B, I show that the WDZ theorem applies to higher-order statistics of the utility distribution, implying that the probability of a product being ranked s^{th} is the gradient of the expected utility from the s^{th} highest draw.

When taste shocks follow a T1EV distribution, the expectations of these higher-order statistics have an analytical expression, reducing the need to simulate the category's inclusive value at each iteration of the contraction mapping. To obtain these closed-forms, I leverage recursive properties of order statistics to express their expectations in terms of the maxima of sub-samples formed by progressively excluding the highest-ranked draws. Since these maxima have well-known analytical expressions, the expectations of the higher-order statistics do as well.[33]

Table 4 presents the estimates of employer preferences for fund characteristics. The vector of observed characteristics includes fund age, a dummy for whether a fund is an

[32] Intuitively, WDZ is the envelope theorem. The result can be verified analytically when taste shocks are distributed as T1EV, in which case, the inclusive value takes the familiar log-sum expression $\mathbb{E}_\epsilon[u_{pj_1}] = \log\left(\sum_{j=1}^{J_p} v_{pj}(\delta_j)\right)$, so that $\frac{\partial \mathbb{E}_\epsilon[u_{pj_1}]}{\partial \delta_k} = \phi_{pk}^1 = \frac{\exp(v_{pk})}{\sum_j \exp(v_{pj})}$ See Shi, Shum and Song (2018) for a general proof.

[33] I provide further details in Appendix B.

index fund, a dummy for whether the fund provider is vertically integrated into the provision of recordkeeping services, two measures of past investment performance, and the fund's expense ratio. I compute past performance as the difference between funds' net of fees returns and their benchmarks' returns, both over the previous year (short-term performance) and over the past ten years (long-term performance).[34] The estimated parameters are all significant and have the expected signs. Employers prefer older funds, index funds and funds that charge lower investment fees. They also prefer funds that outperformed their benchmark, especially if this performance has been sustained over time. The third column translates employers' preferences parameters in percentage points, scaling them by the coefficient on the fund expense ratio. Employers are willing to pay nearly 60 bp. in fees for an index fund and about 47 bp. for a fund whose provider is vertically integrated in the provision of recordkeeping services. This finding is consistent with the documented evidence that employer sponsors tend to favor the inclusion of funds affiliated with their plan recordkeeper (Pool, Sialm and Stefanescu (2016)).[35]

The estimates in Table 4 are based on the assumption that fund characteristics, including fees, are exogenous to employers' unobserved demand shocks. In the current setting, there are at least two reasons that make this assumption not unreasonable. First, mutual fund fees are uniform across all types of investors and are not set at the plan level, implying that employer demand is unlikely to be the primary driver of fund fees. Although demand from employers might correlate with that of non-401(k) investors at some more aggregate level—such as when a particular investment category becomes widely attractive—this specification controls for both year and investment category fixed effects. A similar argument applies to unobserved factors, such as brand effects, which may influence the demand of both 401(k) and non-401(k) investors. The estimates in Table A.2 show that controlling for fund brands reduces the magnitude of the fee coefficient. This is reassuring because fixed effects typically absorb much of the unobserved product quality that could potentially drive demand and, in the presence of price endogeneity, we would expect the magnitude of the fee coefficient to increase rather than decrease after brand fixed effects are included.

Second, for most funds, investments from 401(k) plans represent only a small portion of the fund's total AUM, again suggesting that employers demand is unlikely to be the primary driver of funds' pricing decisions. However, an important exception are TDFs. For the median TDF, 401(k) AUM accounts for more than 50% of total AUM (Table A.3), indicating that employer demand may play a significant role in their pricing decisions. To address this price endogeneity concern, Table A.2 presents two additional robustness checks. First, I re-estimate employer preference parameters after excluding TDFs from the estimation sample and show that the magnitude of the fee coefficient remains broadly

[34] If a fund is less than 10 years old then the long-term performance is measured since the fund inception.

[35] The practice of favoring the inclusion of proprietary funds can be seen as an extreme form of revenue sharing, as fee revenues of affiliated funds accrue entirely to the plan recordkeeper.

consistent with the one estimated in Table 4. Second, I use an Hausman-type of IV strategy where I instrument a fund's expense ratio with the average expense ratio of non-TDF funds from the same fund provider. In this case the estimated fee coefficient is only slightly larger, suggesting again that price endogeneity concerns may be limited, even for TDFs.

The baseline specification also assumes that employers have homogeneous preferences. However, the framework can accommodate heterogeneity in employer preference parameters in standard ways, such as by introducing random coefficients or interacting fund characteristics with employer characteristics. In this case, one must first compute the integral on the right-hand side of equation (2) and then numerically invert the demand system.[36] In Table A.4, I let the fee coefficient vary with plan assets-per-participant (a measure of plan size) and identify the interaction term between size and fees by matching the data correlation between assets-per-participant and average plan fees.[37] The estimated interaction term suggests that larger employers are more sensitive to fees than smaller employers, potentially because larger employers tend to be more sophisticated and have more bargaining power against the plan recordkeeper (Bhattacharya and Illannes (2022)). Another possible explanation is that larger employers may favor cheaper funds to mitigate liability risk, as the vast majority of 401(k) lawsuits involve large employers (Mellman and Sanzenbacher (2018)).[38]

4.3 Investors' portfolio choice problem

Worker i allocates its retirement contributions across the funds available in plan p. Motivated by the growing importance of TDFs and their prominent role as default options in 401(k) plans, I assume worker i can be either inactive or active. If inactive, worker i invests everything in the available default option (e.g., a TDF), denoted by d_i. This default option is worker-specific because it is often determined based on workers' characteristics, such as age. Younger workers are defaulted in TDFs with a target date that is further in the future relative to older worker. Conversely, an active worker, disregards the default option and invests across the non-default funds available in the plan. This assumption is motivated by the fact that TDFs are constructed to be a one-fits-all solution for retirement investing and by the empirical observation that workers rarely invests in both the default and non-default funds (Mitchell and Utkus (2021)). I denote the probability that worker i is active and considers the non-default funds by κ, which I parametrize as

[36] Observe that, unlike the single-discrete logit case, in this multiple-discrete choice setting numerical inversion is required even when preferences are homogeneous.

[37] This data moment is added to the list of moments setting the covariance between characteristics and model residuals to zero.

[38] Figure A.4 plots the distribution of the estimated employers' sensitivity to fees.

a logit function of the default characteristics,

$$\kappa(x_{d_i};\psi_i) = \frac{\exp(x'_{d_i}\psi_i)}{1+\exp(x'_{d_i}\psi_i)}.$$

This specification allows for the possibility that workers' wake up when the default option is not particularly attractive.

Workers who decide to invests in the non-default funds do so in two steps. First, they allocate their contributions across three broad asset classes Equity, Bond and Cash; and, second, within each of these asset classes, they make a single discrete choice decision, selecting only one fund. This two-step structure captures workers' diversification incentives and, at the same time, matches the evidence that workers typically invest in less than an handful of funds.

The portfolio allocation across asset classes is performed in a mean-variance fashion. Formally, letting $b \in \{\text{Equity}, \text{Bond}, \text{Cash}\}$, the portfolio allocation in asset class b, ω_{ib} solves

$$\max_{(\omega_{ib})_b} \sum_b \alpha_{ib}\omega_{ib} - \frac{1}{2}\sum_b \omega_{ib}^2 \quad \text{s.t} \quad \sum_b \omega_{ib} = 1. \tag{7}$$

where α_{ib} is given by

$$\alpha_{ib} = \underbrace{\mu_{ib}}_{\text{asset class } b \text{ fixed effect}} + \rho \underbrace{\log\left(\sum_{j\in b}\exp(V_{ij})\right)}_{\text{asset class } b \text{ inclusive value} \equiv I_{ib}}.$$

In standard portfolio choice theory, α_{ib} would be interpreted as worker i (risk-adjusted) subjective expected return from investing in asset class b.[39] I assume that this subjective expectation can be decomposed into an asset class fixed effect μ_{ib}, capturing worker i's expectations about the broad asset class, and into an inclusive value component, capturing worker i's expectations from investing in the funds that are actually available in the plan menu. The inclusive value corresponds to the utility worker i expects from choosing the best fund within category b in the second step of the portfolio choice problem. As before, following the discrete choice literature, I assume the realized utility from choosing to invest in fund j is $u_{ij} = V_{ij} + \varepsilon_{ij}$, with the mean-utility $V_{ij} = x'_j\theta_i + \xi_j$ specified as a linear function of the same set of fund characteristics entering employers' preferences. Under the assumption that taste shocks are T1EV distributed, the standard log-sum-exp expression for the inclusive value obtains.

How much the types of funds available in a plan matter for the allocation at the asset class level depends on the parameter ρ. If ρ is close to zero, then workers' allocation

[39] Assuming investing in the Cash asset class carries no risk, the mean-variance solution to worker i's portfolio problem is $\omega_{ig} = \alpha_{ig}$. In a mean-variance formulation that separates between expected returns, risk and risk aversion, α_{ig} =(expected return)/(risk aversion×return variance).

between Equity, Bond and Cash does not depend on the specific characteristics of the funds available. Conversely, larger values of ρ imply that workers portfolio allocations at the asset class level reflect the quality of the funds available in their plan.[40] Empirical evidence points toward a small ρ, as fund fixed effects only explain a small share of the observed variation in asset class portfolio shares across plans (Table A.5). This is consistent with workers prioritizing their diversification needs across broad asset classes, perhaps following well-known rule of thumb for the allocation across Equity, Bond and Cash, and regardless of the quality of the funds available in the plan.

Combining everything together, we can derive worker i portfolio allocation to fund j as function of fund characteristics x, parameters that influence the probability of being an active worker ψ_i and preference parameters θ_i:

$$\underbrace{\omega_{ij}(x;\psi_i,\theta_i,d_i)}_{\text{portfolio share}} = \begin{cases} 0 & \text{if } j \in D/\{d_i\} \\ 1 - \underbrace{\kappa(x_{d_i};\psi_i)}_{\text{active prob}} & \text{if } j = d_i \\ \kappa(x_{d_i};\psi_i) \quad \underbrace{\omega_{ib}(x;\theta_i)}_{\text{asset class allocation}} \quad \underbrace{\omega_{ij|b}(x;\theta_i)}_{\text{prob choosing } j \text{ in } b} & \text{if } j \notin D \end{cases} \tag{8}$$

where D is the set of default options available in the plan and $\omega_{ij|b}$ is the standard logit choice probability implied by the single-discrete choice problem.

Overall equation (8) represents the model-implied portfolio share of worker i. Because I only observe portfolio allocations at the plan level, I need to aggregate allocations of workers within a plan. Assuming ψ_i and θ_i follow some parametric distribution $F_{(\psi,\theta)}$, then the plan level portfolio share for fund j in plan p is given by

$$\omega_{pj}(x;F_{(\psi,\theta,d)}) = \sum_{d\in D} \int \omega_{ij}(x;\psi_i,\theta_i,d) dF_{\psi,\theta}(\psi_i,\theta_i|d_i=d)\Pr\{d_i=d\}. \tag{9}$$

where $\Pr\{d_i = d\}$ is the share of employees in plan p whose default option is d. The data counterpart to (9) corresponds to the portfolio share of fund j in plan p. Workers' preferences and the probability of workers' being active can be estimated matching the model-implied portfolio shares defined in (9) with the observed ones.

4.4 Identification and estimation of workers' preferences

The assumption that active workers disregard their default option allows me to identify and estimate the probability of workers being active (κ) from variation in the portfolio share of default options alone. This follows directly from equation (8). Letting d denote

[40] Equity and Cash asset classes are available in all plans in the sample.

an arbitrary default option in plan p, the portfolio share of d in p is

$$\omega_{pd} = \Pr\{d_i = d\} \left(1 - \int \kappa(x_d; \psi_i) dF_\psi(\psi_i | d_i = d)\right),$$

implying that ψ can be estimated from variation in ω_{pd}. The formal identification argument leverages the asymmetry of the Slusky matrix implied by the demand function defined in (8), which belongs to a broader class of demand models commonly used to analyze consumer choices with imperfect consideration.[41]

Because a plan can offer multiple default options, estimating ψ requires knowing the share of employees in plan p whose default option is d, denoted as $\Pr\{d_i = d\}$. Since this information is not directly available in my data, I calibrate $\Pr\{d_i = d\}$ using industry level data on workers' age distribution, available from the Current Population Survey (CPS). The rationale for this approach is that the most common default options are TDFs, and employees are defaulted into specific TDFs based on their age. Younger employees are typically assigned to TDFs with target retirement dates further in the future. Accordingly, I calibrate the share of employees in plan p assigned to TDF d as the share of workers in p's industry whose age corresponds to a retirement year closest to d's target date.[42]

Table 5 reports the estimates of the active probability parameters. The specification assumes that ψ is homogeneous, which implies the following linear specification:

$$\log\left(\frac{1 - \frac{\omega_{pd}}{\Pr\{d_i=d\}}}{\frac{\omega_{pd}}{\Pr\{d_i=d\}}}\right) = x'_{dp}\psi + \epsilon_{dp},$$

where $\Pr\{d_i = d\}$ on the left hand side is estimated from the CPS, as described above.

The estimates indicate that fund characteristics, such as fees and past performance, do not significantly influence the probability of workers being active investors. Instead, plan design features, such as the presence of automatic enrollment, appear to be more relevant. Consistent with the well-established literature on automatic enrollment (Madrian and Shea (2001)), employees in plans with automatic enrollment are more likely to be inactive. More interestingly, this pattern also holds for plans where the default option is a proprietary fund—i.e., a fund from the plan recordkeeper's product line. This finding could reflect workers' preferences for TDFs from well-known fund families or recordkeepers' incentives to offer their proprietary TDFs in plans where employees are more likely to be inactive investors.

Because 401(k) plans are among the largest sources of inflows for TDFs, fee endogeneity presents a concern. In the first specification in Table 5, fees are treated as exogenous.

[41] More specifically, the portfolio choice problem considered here is a Default Specific Model of the type considered in Abaluck and Adams-Prassl (2021), who formalize the identification argument for discrete-choice demand models with imperfect consideration. See also: Ho, Hogan and Scott Morton (2017), Hortaçsu, Madanizadeh and Puller (2017), and Heiss, McFadden, Winter, Wuppermann and Zhou (2021).

[42] I provide more details on the calibration procedure in Appendix C.

	Coefficient	Std. Error	Coefficient (2SLS)	Std. Error (2SLS)
Auto enrollment	-0.206***	(0.037)	-0.202***	(0.054)
Proprietary fund	-0.193***	(0.042)	-0.229***	(0.053)
Index fund	0.049	(0.047)	0.225	(0.146)
Short performance	0.001	(0.009)	0.006	(0.019)
Long performance	-0.009*	(0.005)	-0.007	(0.009)
Expense ratio	-0.078	(0.090)	0.479	(0.519)
N	14,332		14,332	
R2	0.197			
R2 (1st stage)			0.482	
R2 (exog.)			0.072	
F-stat (exog.)			27.881	

Table 5: Estimates of active probability parameters. Year, target date and sponsor's industry fixed effects are included. Standard errors are in parentheses. *** indicates 1% significance, ** indicates 5% significance and * indicates 10% significance.

The resulting estimate for the fee coefficient is negative, suggesting—contrary to economic intuition—that workers are more likely to consume default options when they are more expensive. To address this issue, the second specification uses the TDF series inception date as an instrument for fees.[43] The resulting estimate is positive, though still not statistically significant, and aligns with the intuition that workers prefer cheaper TDFs, ceteris paribus. The rationale for using the inception date as an instrument lies in its ability to capture fund families' cost advantages in launching TDFs. Since TDFs are designed for inactive and inattentive investors, fund families had strong incentives to enter the market as early as possible. More efficient fund families—particularly those also providing recordkeeping services—entered the market earlier and, due to the inertia of defaulting investors, were able to maintain larger margins.

Having estimated workers' probability of being active, I now consider the estimation of workers' preferences. This time I use variation in portfolio allocations of non-default funds, together with the estimating equations implied by the portfolio choice problem described in the previous section. In taking active workers portfolio choice problem to the data, I work with a slightly different version of the model. Specifically, I approximate the mean-variance portfolio allocation between asset classes defined in equation (7) with a mean-entropy type of portfolio problem:

$$\max_{(\omega_{ib})_b} \sum_b \alpha_{ib}\omega_{ib} - \sum_b \omega_{ib}\log(\omega_{ib}) \quad \text{s.t} \quad \sum_b \omega_{ib} = 1. \tag{10}$$

The empirical advantage of working with (10) is that it accounts for the absence

[43] A TDF Series, like the Fidelity Freedom series, is a set of TDFs with different target dates offered by the same fund provider. Typically, TDFs within the same series have the same fee. This implies that the instrument variation in Table 5 is at the fund provider level, which prevents me to control for fund brand fixed effects.

of short sales without imposing short-selling constraints directly.[44] This feature is realistic, as workers cannot short the funds available in their 401(k) plans. The optimal allocation derived from problem (10) follows the well-known logit formulation: $\omega_{ib} = \exp(\alpha_{ib})/\sum_{b'}\exp(\alpha_{ib'})$. This flexible empirical specification is the standard approach for modeling consumer demand in industrial organization and is increasingly used to estimate asset demand from holdings data (Koijen and Yogo (2019)). Conceptually, problem (10) preserves the risk-return trade-off embedded in the standard mean-variance formulation, with the only difference that the entropy of the portfolio is used as a measure of risk rather than its variance.[45] In addition, as I show in Appendix B, the mean-entropy portfolio problem can be micro-founded as a first-order approximation of the mean-variance portfolio problem.[46]

By combining the two steps of portfolio choice of active workers, in Appendix B I show that the following estimating equation obtains,

$$\log\left(\frac{\omega_{pj}}{\sum_{j'\notin D_p}}\right) - \log(\omega_{p0}) = \mu_b + \tilde{\theta}' x_j + (1-\rho)\log(\omega_{pj|b}) + \tilde{\xi}_{jp} \tag{11}$$

where D_p is the set of TDFs in plan p, ω_{p0} is the portfolio share of the Cash category in plan p, $\omega_{pj|b}$ is the portfolio share of fund j within asset class b, $\tilde{\theta} \equiv \rho\theta$ and $\tilde{\xi}_{jp} \equiv \rho\xi_{jp}$. The specification assumes that employees preference parameters are homogeneous and that the utility from the Cash category $b = 0$ is normalized to $\alpha_0 = 0$. This normalization is needed because the levels of utilities are not identified.[47] Although the estimating equation resembles the one obtained when estimating a nested-logit demand model, the model I consider here is not a nested-logit because workers solve a continuous portfolio choice problem across asset classes rather than choosing a single group of products.

Table 6 presents the estimated preference parameters for worker investors, as implied by the empirical specification in (11), alongside the employers' preference estimates from Table 4. Compared to employers, workers are more responsive to certain fund characteristics, such as long-term past performance and fees. The signs of the estimated parameters are consistent between the two, with the exception of the coefficient on the index fund dummy. Interestingly, the point estimate suggest that workers tend to prefer active funds over index funds, while employers exhibit a strong preference for the latter.

One interpretation of this result is that workers, conditional on fees, may prefer active funds because they aim to chase past performance in pursuit of higher potential returns.

[44] Not having to impose short-selling constraints is also particularly useful when simulating workers' portfolios under different counterfactual scenarios.

[45] The entropy measures the distance between two distributions. The objective in (10) is equivalent to measuring the distance of ω_{ib} with respect to an equally weighted portfolio.

[46] The approximation follows when assets are (or are perceived to be) uncorrelated. In the context of my application, 401(k) workers oftentimes follow naive diversification strategies (Benartzi and Thaler (2001)), making this assumption not unreasonable.

[47] To see this observe that replacing α_{ib} with $\alpha_{ib} + C$ for all b, does not affect the solution of the portfolio problem in (10).

	Employee		Employer	
	Coefficient	Std. Error	Coefficient	Std. Error
log(Age)	0.421**	(0.188)	0.097***	(0.011)
Index fund	-0.311	(0.360)	0.270***	(0.034)
Short performance	0.001	(0.020)	0.003**	(0.002)
Long performance	0.064**	(0.029)	0.016***	(0.002)
Expense ratio	-1.472***	(0.561)	-0.472***	(0.028)
ρ	0.033***	(0.003)		
R2	0.910			
N	18,097			

Table 6: Estimates of workers utility parameters. Fixed effects at the plan level and at the investment category level are included as well as a dummy for whether the fund provider offers administrative services to 401(k) plans. Standard errors are in parentheses. *** indicates 1% significance, ** indicates 5% significance and * indicates 10% significance.

In contrast, employers prioritize offering index funds, which are generally perceived as safer options. Compared to active funds, index funds are less likely to underperform, which could help mitigate employers' liability risk. Overall, this divergence may reflect a fundamental difference in incentives: while workers may extrapolate past performance and focus on maximizing returns, employers are more concerned with managing risk and providing stable investment options.

Lastly, the estimates in Table 6 suggest that the characteristics of the funds available within a plan have a limited influence on the portfolio allocations of active workers. This conclusion is supported by the parameter ρ, which measures the extent to which the types of funds available within a given asset class affect workers' allocations to that asset class. The parameter is small, approximately 0.03, indicating that fund-level characteristics play a minor role. More notably, this finding aligns with the observation that, after controlling for asset class fixed effects, fund fixed effects explain only about 4% of the variation in portfolio allocations at the asset class level (Table A.5). This finding further suggests that workers, regardless of the specific funds available to them, tend to maintain stable allocations across Equity, Bond and Cash asset classes.

5 Counterfactuals

In this section, I use the estimated demand models to conduct two sets of counterfactual exercises. The first set examines whether workers would choose better menus by themselves. To explore this, I simulate employers' plan menu choices under a counterfactual scenario where employer preferences align with those of workers, and then assess how the counterfactual menus differ from the observed ones.

The second set of counterfactuals evaluates the effects of plan design policies aimed

at reducing workers' investment costs. I consider two policies: one requiring employers to include the cheapest equity index fund available and another mandating the inclusion of the cheapest default option available. For each policy, I simulate employers' plan menu choices, workers' portfolio allocations, and calculate the resulting fee savings.

To account for differences in plan provision costs, in both counterfactuals I maintain the assumption that employers' choice sets depend on their size. As before, I categorize employers into small (fewer than 500 participants) and large (500 or more participants), and restrict each group to choose among funds offered by employers of the same size group. This aims at ensuring that the counterfactual menus remain feasible within the constraints faced by employers.[48]

5.1 Aligning employers' and workers' preferences

What would happen if employees could choose which funds to include in their 401(k) plan? To address this question, I use the menu choice framework developed in Section 4.1 to predict how plan menus would change if employers' marginal valuations for fund characteristics were aligned with those of workers.

Figure 5 summarizes the results of this counterfactual analysis. Aligning employers' and workers' sensitivity to fees reduces the average plan expense by nearly 25%. This cost reduction is primarily driven by a shift from active to passive options, as implied by the 40% increase in the share of index funds available. Additionally, the number of investment categories declines, likely because employers—now more fee-sensitive—prefer to exclude the most expensive categories.

The results change when employers' preferences are further aligned with workers' marginal valuation of fund past performance and fund investment strategy (i.e., active vs. passive). In this scenario, the reduction in plan expenses is smaller (about 15%), and the share of index funds in the counterfactual menus decreases slightly (by about 3%). This occurs because, conditional on fees, workers exhibit a stronger preference for active funds and are more inclined to chase past investment performance. Since active funds are, on average, more expensive than passive funds, accommodating workers' preference for actively managed funds leads to higher plan costs.

Ultimately, workers care about their retirement outcomes. To evaluate whether the menus they would choose outperform those selected by employers on their behalf, I use the portfolio choice model developed in Section 4.3 to predict workers' portfolio allocations under the counterfactual in which employers' preferences are aligned to those of workers. I then compare the returns of this counterfactual portfolio with those of the status quo portfolio, where employers' and workers' preferences remain misaligned.

Table 7 presents the results of this exercise. I focus on 401(k) plans introduced in 2016, the first year in my sample, and compute the cumulative return of plan portfolios

[48] Table A.6 shows differences in choice sets of small and large plans. The choice set for large sponsors features on average cheaper funds, more index funds and more TDF Series.

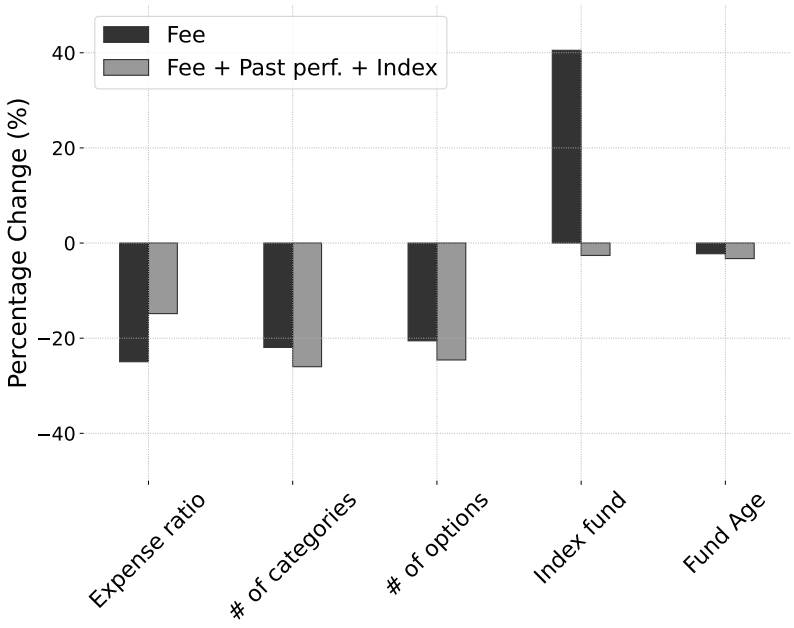


Figure 5: Percentage change in plan menu characteristics when employers' preferences are aligned to workers' preferences.

through 2023. This allows me to analyze the longest investment horizon for which fund returns are available in CRSP. When computing plan returns over time, I hold plan menus and portfolio allocations—chosen by employers and workers, respectively—constant.[49] To capture differences in outcomes between active and inactive workers, I compute returns separately for the portfolio of non-TDF options and for the TDFs available within a given plan.

Aligning employers' fee sensitivity with that of workers improves returns of the non-TDF portfolio by an average of 2.6 percentage points over eight years, or equivalently, by 32 basis points per year. When employers further internalize workers' preference for past performance and active funds, this gain decreases by about 30%, with returns rising by only 1.8 percentage points relative to the status quo. One reason workers benefit in both scenarios is that the counterfactual menus include fewer bond funds and more equity funds with a Blend tilt (Figure A.5). This shift occurs because employers—now more fee-sensitive—tilt their menus towards investment categories that contain cheaper funds.[50] Ultimately, workers fare better when employers internalize only their fee sensitivity because counterfactual menus in this scenario include more passive funds, which tend to outperform active funds on a net return basis.

The results for TDFs highlight similar patterns. Aligning employers' fee sensitivity with that of workers improves the returns of the TDF Series by an average of 65 basis points over an eight-year horizon. Compared to the non-TDF portfolio, these gains are

[49] In practice, employers and workers can adjust plan menus and portfolios over time, but such changes are rare.

[50] Typically, the cheapest market index funds are blended funds holding a mix growth and value stocks.

	Δ cumulative return (pp.)	Δ avg. annual return (pp.)
Non TDF		
Fee	2.61	0.32
Fee + past perf. + index	1.83	0.23
TDF		
Fee	0.65	0.08
Fee + past perf. + index	-0.11	-0.01

Table 7: Average counterfactual return relative to the status quo. Returns are net of fees and are computed over an 8-year investment horizon from 2016 to 2023. Magnitudes are reported in percentage points.

smaller, likely because returns across different TDF Series are less heterogeneous. More interestingly, when employers also internalize workers' preference for past performance and active funds, the returns of the TDF Series are, on average, 11 basis points lower than in the status quo. This suggests that, relative to other types of active funds, the underperformance of active TDFs is more pronounced—likely because TDFs tend to charge excessive management fees (Brown and Davies (2021)).

So far, the analysis reveals that while workers would select menus that are cheaper than those chosen by employers on their behalf, their tendency to chase past investment performance through active funds reduces the gains from having employers fully internalize their preferences. This effect is particularly evident when examining counterfactual returns for TDFs, which turn out to be lower in a world where employers' and workers' preferences are aligned. Conversely, when employers internalize only workers' fee sensitivity, return outcomes improve for both active and inactive workers, as counterfactual menus offer more passive options. Overall, this evidence suggests that policies aimed at reducing plan costs could enhance workers' retirement outcomes—an issue I explore in the next section.

Before proceeding, it is important to acknowledge that the counterfactual analysis in Table 7 is based on an eight-year investment horizon, which is significantly shorter than the typical 30- to 40-year horizon relevant for retirement investors. That said, I would expect the results in Table 7 to hold even more strongly over extended investment periods, as the ability of active funds to outperform market benchmarks tends to diminish over time (Bessembinder, Cooper and Zhang (2023)).

5.2 Mandating the inclusion of low-cost options

Many 401(k) plans include mutual funds that are significantly more expensive than comparable alternatives available in the market. The prevalence of high-fee investment options has been a major catalyst for the recent wave of 401(k) lawsuits (Mellman and Sanzenbacher (2018)). In this section, I use the empirical framework developed in Sec-

	Δ Fees (bp.)	Fee savings/year ($)	Fee savings ($)
Mandating low-cost index	-5.29	15.87	9,586.38
Mandating low-cost default	-18.12	54.37	33,566.15

Table 8: Fee savings under different counterfactual policies relative to the status quo. Fee savings per-year assumes a retirement account balance of $30,000. Fee savings are over 40 years, assuming an annual income of $70,000, a contribution rate of 9% and an annual return of 6%.

tion 4 to evaluate the effects of plan design policies aimed at reducing workers' investment costs. The policies considered mandate employers to include at least one low-cost option, either a market index fund or a low-cost default fund. To evaluate their effectiveness, I first simulate the counterfactual menus under each policy scenario, then simulate workers' portfolio allocations, and finally calculate the implied fee savings relative to the observed status quo.

Table 8 summarizes the results of these counterfactual analyses. Mandating the inclusion of a low-cost index fund reduces the asset-weighted average plan expense by approximately 5 basis points, which translates to an annual savings of $16 for a worker with a $30,000 retirement account balance. Over a 40-year investment horizon, these savings compound to nearly $10,000 for a worker earning $70,000 annually and contributing 9% of their income to their 401(k) each year.

Strikingly, fee savings are more than three times larger when the policy mandates the inclusion of a low-cost default option, such as a low-cost Target Date Fund (TDF). In this scenario, the average plan expense decreases by 18 basis points, resulting in annual fee savings of $54 and more than $30,000 over 40 years for the same worker.

The stark differences in policy outcomes stem from the following reasons. A policy requiring the inclusion of a low-cost equity index fund primarily benefits active workers who make their own portfolio decisions. In contrast, a policy mandating the inclusion of a low-cost TDF directly impacts inactive workers, who rely on the plan's default option and constitute the majority of participants in this market. Additionally, active workers will only reallocate part of their portfolio toward the newly added index fund, as they have a strong preference for maintaining a diversified portfolio across broad asset classes, regardless of which funds are offered within each class.[51] Conversely, inactive workers will have 100% of their assets automatically defaulted into the low-cost TDF. Overall, these results underscore the importance of accounting for workers' investment behavior when evaluating the effectiveness of policies aimed at reducing investment costs in 401(k) retirement plans.

[51] This incentive is captured by the preference parameter ρ.

6 Conclusions

This paper evaluates whether employer intermediation in the provision 401(k) retirement plans benefits workers by examining whether the plan menus employers select outperform those workers would choose on their own. While existing evidence emphasize that employers' choices can be far from what an informed regulator would choose, my findings suggest that workers' own investment behavior complicates the evaluation of employer intermediation. Workers are more sensitive to fund fees than employers, yet they also chase past performance and prefer active funds. As a result, workers' menus do not necessarily improve upon those provided by employers.

By simulating counterfactual plan menus, I find that aligning employer preferences with those of workers reduces investment fees but also decreases access to passive index funds, which tend to deliver higher long-term returns. Further analysis of workers' retirement outcomes reveals that employers' intermediation is beneficial in certain aspects: internalizing only workers' fee sensitivity improves outcomes, but fully aligning workers' preferences, including their performance-chasing behavior, erodes those a significant portion of those benefits.

From a policy perspective, my findings suggest that directly regulating investment costs is more effective than attempting to realign employer and worker preferences. Simulations of cost-targeting policies show that mandating the inclusion of low-cost default funds generates significantly larger fee savings than requiring the inclusion of low-cost index funds. This is primarily because default options impact inactive workers, who make up the majority of 401(k) participants. Additionally, the low-cost mandate index fund policy is less effective because active investors prefer to diversify across broad asset classes, meaning they only reallocate a portion of their portfolio to the newly added low-cost index fund rather than fully shifting their investments.

Overall, my results underscore the importance of accounting for workers' investment behavior when evaluating retirement plan design and policy interventions. While reducing investment fees is beneficial, policies must consider how investors respond to menu choices to ensure meaningful improvements in retirement outcomes.

References

Abaluck, J. and Adams-Prassl, A. (2021), 'What do consumers consider before they choose? identification from asymmetric demand responses', *The Quarterly Journal of Economics* 136(3), 1611–1663.

Agarwal, N. and Somaini, P. (2020), 'Revealed preference analysis of school choice models', *Annual Review of Economics* 12(1), 471–501.

Badarinza, C., Campbell, J. Y. and Ramadorai, T. (2016), 'International comparative household finance', *Annual Review of Economics* 8(1), 111–144.

Badoer, D., Costello, C. and James, C. (2020), 'I can see clearly now: The impact of disclosure requirements on 401(k) fees', *Journal of Financial Economics* 136, 471–489.

Beggs, S., Cardell, S. and Hausman, J. (1981), 'Assessing the potential demand for electric cars', *Journal of Econometrics* 16, 1–19.

Benartzi, s. and Thaler, R. (2001), 'Naive diversification strategies in defined contribution saving plans.', *American Economic Review* 91(1), 2997–3053.

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

Berry, S. (1994), 'Estimating discrete-choice models of product differentiation', *The RAND Journal of Economics* 25(2), 242–262.

Berry, S. and Haile, P. (2021), 'Foundations of demand estimation', *Handbook of Industrial Organization, Volume 4* p. 1–62.

Berry, S., Levinsohn, J. and Pakes, A. (1995), 'Automobile prices in market equilibrium', *Econometrica* 63(4), 841–890.

Berry, S., Levinsohn, J. and Pakes, A. (2004), 'Differentiated products demand systems from a combination of micro and macro data: The new car market', *Journal of Political Economy* 112(1), 68–105.

Beshears, J., Choi, J., Laibson, D. and Madrian, B. (2009), 'The importance of default options for retirement saving outcomes: Evidence from the united states', *The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States.* University of Chicago Press.

Bessembinder, H., Cooper, M. J. and Zhang, F. (2023), 'Mutual fund performance at long horizons', *Journal of Financial Economics* 147(1), 132–158.

Bhattacharya, V., Illanes, G. and Padi, M. (2020), 'Fiduciary duty and the market for financial advice', working paper.

Bhattacharya, V. and Illannes, G. (2022), 'The design of defined contributions plans', NBER working paper.

Boehm, E. (2024), 'Intermediation, choice frictions, and adverse selection: Evidence from the chilean pension market', *working paper* p. 310–345.

Brown, D. and Davies, S. (2021), 'Off target: On the underperformance of target-date funds', SSRN working paper.

Carhart, M. M. (1997), 'On persistence in mutual fund performance', *The Journal of Finance* 52(1), 57–82.

Carroll, G., Choi, J., Laibson, D., Madrian, B. and Metrick, A. (2009), 'Optimal defaults and active decisions', *The Quarterly Journal of Economics* 124(4), 1639–1674.

Chalmers, J. and Reuter, J. (2020), 'Is conflicted investment advice better than no advice?', *Journal of Financial Economics* 138(2), 366–387.

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

Clark, R., Houde, J.-F. and Kastl, J. (2021), *The industrial organization of financial markets*, Elsevier, p. 427–520.

Conlon, C. and Gortmaker, J. (2020), 'Best practices for differentiated products demand estimation with pyblp', *RAND Journal of Economics* 51(4), 1108–1161.

David, H. A. and Nagaraja, H. N. (2003), *Order Statistics*, Wiley.

Deloitte (2013), 'Inside the structure of defined contribution/401(k) plan fees', p. 1611–1663.

Duarte, V., Fonseca, J., Goodman, A. and Parker, J. (2022), 'Simple allocation rules and optimal portfolio choice over the lifecycle', NBER working paper.

Dubé, J.-P. (2004), 'Multiple discreteness and product differentiation: Demand for carbonated soft drinks', *Marketing Science* 23(1), 66–81.

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

Egan, M., Ge, S. and Tang, J. (2022), 'Conflicting interests and the effect of fiduciary duty: Evidence from variable annuities', *The Review of Financial Studies* 35(12), 5334–5386.

Egan, M., MacKay, A. and Yang, H. (2021), 'Recovering investor expectations from demand for index funds', *Review of Economic Studies* 89(5), 2559–2599.

Egan, M., MacKay, A. and Yang, H. (2023), 'What drives variation in investor portfolios? estimating the roles of beliefs and risk preferences.', NBER working paper.

Einav, L., Finkelstein, A. and Mahoney, N. (2021), *The IO of selection markets*, Elsevier, p. 389–426.

Fama, E. and Kenneth, F. (2010), 'Luck versus skill in the cross-section of mutual fund returns', *The Journal of Finance* 65(5), 1915–1947.

Gandhi, A. and Nevo, A. (2021), 'Empirical models of demand and supply in differentiated products industries', *Handbook of Industrial Organization, Volume 4* p. 1–62.

Gentzkow, M. (2007), 'Valuing new goods in a model with complementarity: Online newspapers', *American Economic Review* 97(3), 713–744.

Gil-Bazo, J. and Ruiz-Verdu, P. (2009), 'The relation between price and performance in the mutual fund industry', *The Journal of Finance* 64(5), 2153–2183.

Gropper, M. (2023), 'Lawyers setting the menu: The effects of litigation risk on employer-sponsored retirement plans', working paper.

Gruber, M. (1996), 'Another puzzle: The growth in actively managed mutual funds', *The Journal of Finance* 51(3), 783–810.

Guiso, L., Pozzi, A., Tsoy, A., Gambacorta, L. and Mistrulli, P. E. (2022), 'The cost of steering in financial markets: Evidence from the mortgage market', *Journal of Financial Economics* 143(3), 1209–1226.

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

Heiss, F., McFadden, D., Winter, J., Wuppermann, A. and Zhou, B. (2021), 'Inattention and switching costs as sources of inertia in medicare part d', *American Economic Review* 111(9), 2737–2781.

Hendel, I. (1999), 'Estimating multiple-discrete choice models: An application to computerization returns', *Review of Economic Studies* 66(2), 423–446.

Ho, K., Hogan, J. and Scott Morton, F. (2017), 'The impact of consumer inattention on insurer pricing in the medicare part d program', *The RAND Journal of Economics* 48(4), 877–905.

Hortaçsu, A., Madanizadeh, S. A. and Puller, S. L. (2017), 'Power to choose? an analysis of consumer inertia in the residential electricity market', *American Economic Journal: Economic Policy* 9(4), 192–226.

Illanes, G. (2017), 'Switching costs in pension plan choice', working paper.

Illanes, G. and Padi, M. (2021), 'Retirement policy and annuity market equilibria: Evidence from chile', NBER working paper.

Jensen, M. C. (1968), 'The performance of mutual funds in the period 1945–1964', *The Journal of Finance* 23(2), 389–416.

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

Luco, F. (2019), 'Switching costs and competition in retirement investment', *American Economic Journal: Microeconomics* 11(2), 26–54.

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

McFadden, D. (1974), 'Conditional logit analysis of qualitative choice behavior', *Frontiers in Econometrics* pp. 105–142.

Mellman, G. and Sanzenbacher, G. (2018), '401(k) lawsuits: What are the causes and consequences?', Center for Retirement Research, Boston College(18-8).

Mitchell, O. S. and Utkus, S. P. (2021), 'Target-date funds and portfolio choice in 401(k) plans', *Journal of Pension Economics and Finance* 21(4), 519–536.

Pool, V., Sialm, C. and Stefanescu, I. (2016), 'It pays to set the menu: Mutual fund investment options in 401(k) plans', *The Journal of Finance* 521(4), 1179–1812.

Pool, V., Sialm, C. and Stefanescu, I. (2022), 'Mutual fund revenue sharing in 401k plans', NBER working paper.

Reuter, J. (2024), 'Plan design and participant behavior in defined contribution retirement plans: Past, present, and future', *NBER working paper* .

Robles-Garcia, C. (2020), 'Competition and incentives in mortgage markets: The role of brokers', working paper.

Shi, X., Shum, M. and Song, W. (2018), 'Estimating semi-parametric panel multinomial choice models using cyclic monotonicity', *Econometrica* 86(2), 737–761.

Tang, N., Mitchell, O., Mottola, G. and Utkus, S. (2010), 'The efficiency of sponsor and participant portfolio choices in 401(k) plans', *Journal of Public Economics* 94(2), 1073–1085.

Vanguard (2023), 'How america saves 2023.'.

Wang, A. (2024), 'Demand with complementarities: A market share approach', SSRN working paper.

Yang, H. (2023), 'What determines 401(k) plan fees? a dynamic model of transaction costs and markups', working paper.

A Additional Figures and Tables

Plan Sponsor: Airbnb, Inc.
Employer Identification Number: 26-3051438
Plan Number: 001
Schedule H, Line 4(i)

(a)	(b) Identity of issue, borrower, lessor, or similar party	(c) Description of investment, including maturity date, rate of interest, collateral, par, or maturity value	(d) Cost	(e) Current value
*	Fidelity 500 Index Fund	Registered Investment Company	**	$ 12,419,697
	JPMorgan SmartRetirement 2045 Fund Class I	Registered Investment Company	**	12,134,364
	JPMorgan SmartRetirement 2055 Fund Class I	Registered Investment Company	**	10,659,751
	JPMorgan SmartRetirement 2050 Fund Class I	Registered Investment Company	**	6,899,616
	JPMorgan SmartRetirement 2040 Fund Class I	Registered Investment Company	**	4,131,471
	JPMorgan SmartRetirement 2035 Fund Class I	Registered Investment Company	**	2,892,780
*	Fidelity International Index Fund	Registered Investment Company	**	2,190,805
	T. Rowe Price Blue Chip Growth Fund	Registered Investment Company	**	2,183,653
*	Fidelity Mid-Cap Index Fund	Registered Investment Company	**	2,145,429
*	Fidelity Small Cap Index Fund	Registered Investment Company	**	2,036,228
*	Fidelity Government Money Market Fund	Registered Investment Company	**	1,253,430
	T. Rowe Price Mid-Cap Growth Fund	Registered Investment Company	**	1,236,211
	Metropolitan West Total Return Bond	Registered Investment Company	**	1,218,146
	Undiscovered Managers Behavioral Value	Registered Investment Company	**	818,923
	AB High Income Fund Advisor Class	Registered Investment Company	**	786,374
	JPMorgan SmartRetirement 2025 Fund Class I	Registered Investment Company	**	739,193
	JPMorgan SmartRetirement 2030 Fund Class I	Registered Investment Company	**	710,843
	JPMorgan Value Advantage Fund	Registered Investment Company	**	700,599
	JPMorgan SmartRetirement Income Fund Class I	Registered Investment Company	**	609,830
*	Fidelity Advisor Real Estate Fund	Registered Investment Company	**	556,645
	American Funds EuroPacific Growth Fund Class R-5E	Registered Investment Company	**	552,702
	Oppenheimer Developing Markets Fund Class Y	Registered Investment Company	**	492,797
	American Funds Capital World Growth and Income Fund Class R-5E	Registered Investment Company	**	471,587
	Franklin Small Cap Growth Advantage	Registered Investment Company	**	320,053
	Harbor Mid Cap Value Fund	Registered Investment Company	**	277,250
	JPMorgan SmartRetirement 2020 Fund Class I	Registered Investment Company	**	198,666
	JPMorgan SmartRetirement 2060 Fund Class I	Registered Investment Company	**	127,569
*	Participant loans	Interest rates range from 4.25% to 5.25%, maturing through December 2027		563,532
				$ 69,328,144

* Indicates party-in-interest
** Information is not required as investments are participant directed

Figure A.1: Airbnb 401(k) plan menu in 2017 as reported in form F5500.

		2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Percentage of participants holding a single fund		44%	48%	51%	55%	59%	60%	62%	60%	61%	63%
Percentage of single-fund participants using	Cash	11%	8%	7%	6%	5%	4%	4%	4%	4%	3%
	Bond funds	1%	1%	1%	1%	<0.5%	1%	<0.5%	<0.5%	<0.5%	<0.5%
	Traditional balanced funds	13%	5%	4%	6%	5%	5%	5%	2%	1%	1%
	Target-date funds	69%	81%	84%	83%	87%	87%	88%	91%	92%	93%
	Diversified equity funds	4%	4%	3%	3%	2%	2%	2%	2%	2%	2%
	Company stock	2%	1%	1%	1%	1%	1%	1%	1%	<0.5%	<0.5%

Source: Vanguard 2023.

Figure A.2: Distribution of single-fund holder by asset class and over time (Vanguard (2023)).

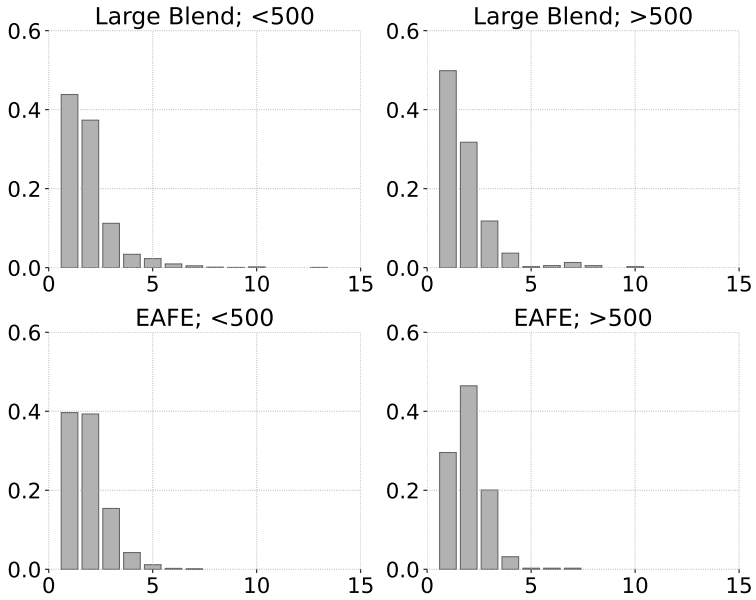


Figure A.3: Distribution of the number of options offered for Large Blend and EAFE categories in small and large plans respectively.

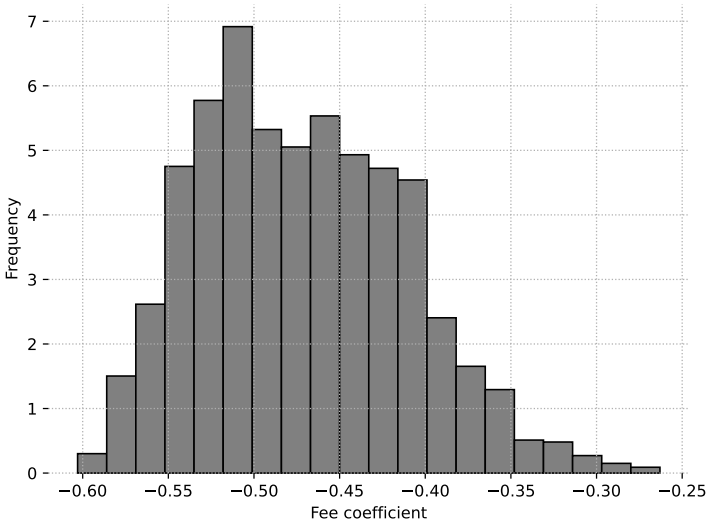


Figure A.4: Distribution of employers' sensitivity to fees, based on employer preference estimates reported in Table A.4.

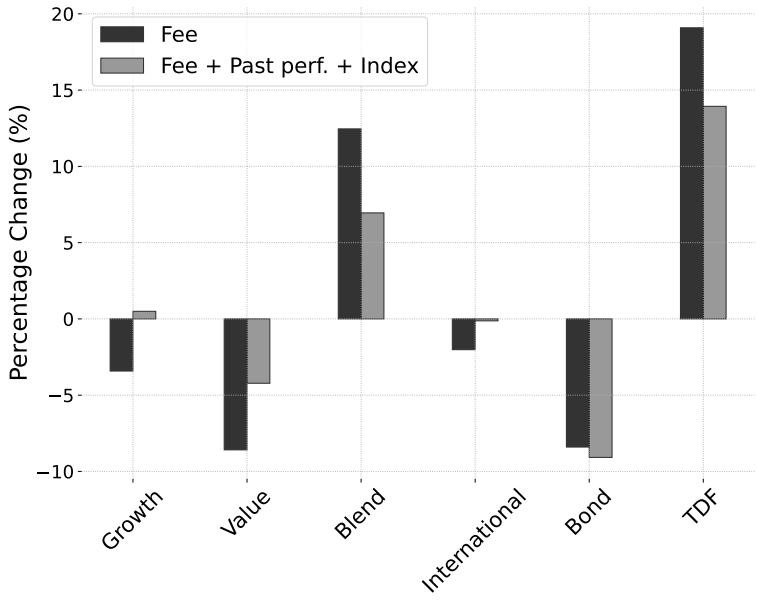


Figure A.5: Percentage change in the average share of investment options available from a given investment category when employers' preferences are aligned to workers' preferences.

	Coeff.	Std. Error	Coeff.	Std. Error	Coeff.	Std. Error
Employer contributions (%)	-0.021	(0.026)	-0.039	(0.026)	-0.032	(0.026)
Administrative fees (%)	-0.072	(0.375)	0.050	(0.386)	-0.033	(0.387)
log(# of particp.)	-0.685*	(0.416)	-0.949**	(0.426)	-1.041**	(0.435)
Avg. account balance (thou.)	-0.047***	(0.008)	-0.048***	(0.009)	-0.049***	(0.010)
Auto Enrollment	-0.017	(0.861)	-0.373	(0.863)	-0.320	(0.878)
# of asset classes	-1.923***	(0.606)	-1.827***	(0.610)	-1.814***	(0.612)
Index fund (%)	-0.653***	(0.016)	-0.656***	(0.018)	-0.656***	(0.018)
# of funds	0.471***	(0.097)	0.442***	(0.102)	0.452***	(0.104)
Avg. performance (%)	-3.657***	(0.470)	-3.705***	(0.479)	-3.666***	(0.483)
R2	0.589		0.601		0.600	
N	1,573		1,573		1,573	
Year FE	Y		Y		Y	
Reckeeper FE	N		Y		Y	
Industry FE	N		N		Y	

Table A.1: Linear regression of plan average expense ratio on plan characteristics. Average expense ratio is measured in basis points. Industry fixed effects are at the 2-digit NAICS. Robust standard errors are in parentheses. ** indicates 5% significance and * indicates 10% significance.

	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error
log(age)	0.089***	(0.011)	0.104***	(0.012)	0.070***	(0.012)
index fund	0.249***	(0.034)	0.201***	(0.041)	0.367***	(0.036)
integrated provider	0.214***	(0.018)	0.199***	(0.019)		
short performance	0.003*	(0.002)	0.003*	(0.002)	0.003**	(0.002)
long performance	0.016***	(0.002)	0.016***	(0.002)	0.017***	(0.002)
expense ratio	-0.484***	(0.028)	-0.599***	(0.057)	-0.352***	(0.031)
N	8,127		8,290		8,290	
R2	0.664		0.668		0.690	
$Q_{pg}(n)$	calibrated		calibrated		calibrated	
Fund brand FE	N		N		Y	
Exclude TDF Series	Y		N		Y	
Method	OLS		IV		OLS	

Table A.2: Estimates of employer utility parameters obtained from the inversion of fund inclusion probabilities in (3). The distribution of the number of options $Q_{pg}(n)$ is calibrated within each investment category, separately for small and large plans. Fixed effects at the plan size group level, year level, and category level are included in all specifications. The IV specification instruments a fund's expense ratio with the average expense ratio of non-TDF funds offered by the same fund provider. Fund performance and fund expenses are in percentage points (pp.). Robust standard errors are in parentheses.*** indicates 1% significance, ** indicates 5% significance and * indicates 10% significance.

	Mean	Median	Mean (weighted)	Median (weighted)
Non TDF	20.09	12.43	15.44	10.17
TDF	51.34	50.00	58.05	55.86

Table A.3: Percentage of fund AUM held by 401(k) retirement plans. Statistics are based on all funds available in the introduction year of retirement plans first established between 2016 and 2020. The share of a fund's AUM held by 401(k) plans is calculated as the ratio of the fund's total AUM in all 401(k) plans observed in BrightScope to its total AUM reported in CRSP. In the last two columns show the asset-weighted average and median share of fund AUM held by 401(k).

	Coefficient	Std. Error
Expense ratio	-0.223***	(0.028)
Expense ratio × log(Plan size)	-0.030***	(0.000)
log(Age)	0.097***	(0.011)
Index fund	0.269***	(0.034)
Integrated provider	0.220***	(0.017)
Short performance	0.003**	(0.002)
Long performance	0.016***	(0.002)
Expense ratio (Average)	-0.472	
N	8,290	

Table A.4: Estimates of employer utility parameters obtained from the inversion of fund inclusion probabilities in (3). The distribution of the number of options $Q_{pg}(n)$ is calibrated within each investment category, separately for small and large plans. Fixed effects at the plan size group level, year level, and category level are included. Fund performance and fund expenses are in percentage points (pp.). Plan size is measured in terms of assets-per-participant. Robust standard errors, computed with 2-step GMM, are in parentheses. *** indicates 1% significance, ** indicates 5% significance and * indicates 10% significance.

	Asset Class Allocation	Fund Allocation
Asset Class FE	0.58	0.07
Fund FE | Asset Class FE	0.04	0.22

Table A.5: Share of variation in asset class allocations and fund-level allocations explained by Asset Class and Fund fixed effects respectively. Asset classes are Equity, Balanced (including TDFs), Bond and Cash.

	Small	Large
# of tickers	2708	1434
# categories	18	18
Age (month)	152	156
Index fund	0.154	0.194
Expense ratio	0.746	0.675
TDF Series	0.016	0.028
Short performance	-0.114	0.073
Long performance	0.574	0.744
Integrated provider	0.473	0.519

Table A.6: Comparisons of fund choice sets for small and large plans. Small plans have less than 500 plan participants. Large plans have more than 500 plan participants. TDFs are grouped into serie. TDF Series is the share of options that are TDF Series.

B Derivation and Proofs

Derivation of inclusion probability. Letting λ_{gp} the probability that p chooses to offer category g, $q_{pg}(z)$ the probability that p includes z funds from category g and $\phi_j^{1:z}$ the probability that j is included conditional on z options, the (unconditional) probability that j is included in sponsor p menu can be rearranged as follows

$$\begin{aligned}
\phi_{jp} &= \lambda_{gp}\left(\sum_z q_{pg}(z)\phi_{jp}^{1:z}\right) \\
&= \lambda_{gp}\left(\sum_z q_{pg}(z)\sum_{n=1}^{z}\phi_{jp}^{n}\right) \\
&= \lambda_{gp}\sum_n\sum_{z\geq n} q_{pg}(z)\phi_{jp}^{n} \\
&= \lambda_{gp}\sum_n \phi_{jp}^{n}(1-Q_{pg}(n))\sum_{n\geq z}\frac{q_{pg}(n)}{(1-Q_{pg}(z))} \\
&= \lambda_{gp}\sum_n (1-Q_{pg}(n))\phi_{jp}^{n}
\end{aligned}$$

where the second equality uses the fact that the probability of j being included conditional on z options is equivalent to the probability that fund j is one of the first z highest utility draws i.e., $\phi_{jp}^{1:z} = \sum_{n=1}^{z}\phi_{pj}^{z}$; and the fourth equality uses the definition of $Q_{pg}(n)$. The final expression coincides with expression (1) provided in the main text.

The previous expression boils down to the familiar single discrete choice probability if sponsors choose only one fund. To see this, let $q_{pg}(1) = 1$ and note that fund j inclusion probability reduces to

$$\begin{aligned}
\phi_{jp} &= \lambda_{gp}\phi_{jp}^{1} \\
&= \frac{\exp(I_g)}{1+\exp(I_g)}\frac{\exp(v_{pj})}{\sum_{k\in g}\exp(v_{pk})} \\
&= \frac{\sum_k \exp(v_{pk})}{1+\sum_{k\in g}\exp(v_{pk})}\frac{\exp(v_{pj})}{\sum_{k\in g}\exp(v_{pk})} = \frac{\exp(\delta_j)}{1+\sum_{k\in g}\exp(\delta_k)}
\end{aligned}$$

where the second equality assumes that λ_{gp} is a logit function of the inclusive value and that taste shocks are T1EV; and the third equality uses the fact that with T1EV shocks $I_g = \log(\sum_{k\in g}\exp(v_{kp}))$.

Derivation of ranking probabilities. Consider a simple example where we have four options $\{j,k,l,m\}$ and we want to compute the probability that option j is ranked 2nd. For simplicity, I drop sponsor subscript p. The probability that j is ranked 2nd equals to the sum of all possible utility rankings in which u_j is the second highest:

$$\begin{aligned}
\phi_j^2 = &\Pr\{u_k > u_j > u_l > u_m\} + \Pr\{u_k > u_j > u_m > u_l\} \\
&+ \Pr\{u_l > u_j > u_k > u_m\} + \Pr\{u_l > u_j > u_m > u_k\}
\end{aligned}$$

$$+ \Pr\{u_m > u_j > u_k > u_l\} + \Pr\{u_m > u_j > u_l > u_k\}. \tag{12}$$

Next, consider any of the six rankings above, for example the first one, and note that

$$\Pr\{u_k > u_j > u_l > u_m\} = \frac{\exp(V_k)}{\sum_{s\in\{j,k,l,m\}}\exp(V_s)} \cdot \frac{\exp(V_j)}{\sum_{s'\in\{j,l,m\}}\exp(V_{s'})} \cdot \frac{\exp(V_l)}{\sum_{s''\in\{l,m\}}\exp(V_{s''})} \tag{13}$$

Expression (13) can be derived analytically by integrating over the T1EV extreme value shocks and is known as ranked-ordered-logit (ROL).[52] Expression (13) applies analogously to all six terms in (12) and implies that ϕ_j^2 only depends on how the first two choices are ranked but not on the order of the 3rd and 4th choices. To see this consider the sum of the first two terms on the RHS of (12) and note that from (13) we can factor out the first two factors of each addend and that the sum of the last factors equals one. Overall we obtain

$$\begin{aligned}&\Pr\{u_k > u_j > u_l > u_m\} + \Pr\{u_k > u_j > u_m > u_l\} = \\ &= \frac{\exp(V_k)}{\sum_{s\in\{j,k,l,m\}}\exp(V_s)} \cdot \frac{\exp(V_j)}{\sum_{s'\in\{j,l,m\}}\exp(V_{s'})}.\end{aligned}$$

Applying the same steps for all three lines in (12), we obtain the expression in the main text for the case of $n = 2$. The general case is the natural generalization of the above steps.

Calibration of the distribution of the number of options included within category. I calibrate the distribution of the number of options included separately for each investment category and for small and large plans. Small plans have less than 500 plan participants whereas large plans have more than 500.

Letting $\hat{q}_{sg}(n)$ the observed share of plans in group $s \in \{\text{small}, \text{large}\}$ offering n investment options in category g, the calibrated distribution $q_{sg}(n)$ is found by solving:

$$\begin{aligned}\min_{(q_{sg}(n))_{n=1}^{\bar{n}}} &\sum_{n=1}^{\bar{n}} (q_{sq}(n) - \hat{q}_{sq}(n))^2 \\ \text{s.t} &\sum_{n=1}^{\bar{n}} n q_{sq}(n) = \bar{q} \\ &\sum_{n=1}^{\bar{n}} q_{sq}(n) = 1\end{aligned}$$

where $\bar{q}$ is the average number of options offered in category g by plans in group s. In the calibration, I also set $\bar{n} = 3$ because in 98% of cases sponsors offer less than 3 funds within an investment category.

[52] For more details see Beggs, Cardell and Hausman (1981).

Conditions for the contraction mapping. Recall the operator $T : \mathbb{R}^J \to \mathbb{R}^J$ defined in the main text is

$$T_j(\delta) = \delta_j + \log(\hat{\phi}_j) - \log(\phi_j(\delta))$$

where $\phi_j(\delta)$ is fund j inclusion probability defined in equation (3) as

$$\phi_j(\delta) = \sum_n \int \lambda(\delta, \nu)(1 - Q(n))\phi^n(\delta, \nu)dF(v).$$

To apply the BLP contraction argument, we need the partial derivatives of T_j to sum to a positive number that is less than one for any j. Letting $\phi_j^{1:n} \equiv \sum_{s=1}^n \phi_j^s$ and recalling that $\sum_n (1 - Q(n))\phi_j^n = \mathbb{E}_q[\phi_j^{1:n}]$ we obtain

$$\begin{aligned}
\sum_{k=1}^J \frac{\partial T_j(\delta)}{\partial \delta_k} &= 1 - \frac{1}{\phi_j(\delta)} \sum_{k=1}^J \int \frac{\partial}{\partial \delta_k} \lambda(\delta, \nu)\mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]dF(\nu) \\
&= 1 - \frac{1}{\phi_j(\delta)} \int \left(\sum_{k=1}^J \frac{\partial \lambda(\delta, \nu)}{\partial \delta_k} \right) \mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)] + \lambda(\delta, \nu) \underbrace{\left(\sum_{k=1}^J \frac{\partial \mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]}{\partial \delta_k} \right)}_{=0} dF(\nu) \\
&= 1 - \frac{1}{\phi_j(\delta)} \int \left(\sum_{k=1}^J \frac{\partial \lambda(\delta, \nu)}{\partial \delta_k} \right) \mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]dF(\nu) \qquad (14)
\end{aligned}$$

where the third equality follows because increasing all mean-utility levels does not alter the ranking of products, conditional on the category's inclusion.

From (14), we see that to ensure that the sum of partials is positive and less than one it suffices to choose λ to be increasing in δ_k (e.g., the outside and inside goods are substitutes) and such that $\sum_{k=1}^J \frac{\partial \lambda(\delta, \nu)}{\partial \delta_k} < 1$. When λ is a logit function of the category inclusive value $I(\delta, \nu)$ these properties are satisfied:

$$\begin{aligned}
\sum_{k=1}^J \frac{\partial T_j(\delta)}{\partial \delta_k} &= 1 - \frac{1}{\phi_j(\delta)} \int \lambda(\delta, \nu)(1 - \lambda(\delta, \nu)) \left(\sum_{k=1}^J \frac{\partial I(\delta, \nu)}{\partial \delta_k} \right) \mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]dF(\nu) \\
&= 1 - \frac{1}{\phi_j(\delta)} \int \lambda(\delta, \nu)(1 - \lambda(\delta, \nu)) \underbrace{\left(\sum_{k=1}^J \mathbb{E}_q[n^{-1}\phi_k^{1:n}(\delta, \nu)] \right)}_{=1} \mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]dF(\nu) \\
&= 1 - \frac{1}{\phi_j(\delta)} \int \lambda(\delta, \nu)(1 - \lambda(\delta, \nu))\mathbb{E}_q[\phi_j^{1:n}(\delta, \nu)]dF(\nu) < 1.
\end{aligned}$$

where the second equality follows from applying the WDZ theorem for higher-order statistics (proved below). Specifically, observe that

$$\sum_{k=1}^J \frac{\partial I(\delta, \nu)}{\partial \delta_k} = \sum_{k=1}^J \mathbb{E}_q \left[n^{-1} \sum_{s=1}^n \frac{\partial}{\partial \delta_k} \mathbb{E}[u_{j_s}] \right]$$

$$
\begin{aligned}
&= \sum_{k=1}^{J} \mathbb{E}_q \left[n^{-1} \sum_{s=1}^{n} \phi_k^s \right] \\
&= \sum_{k=1}^{J} \mathbb{E}_q \left[n^{-1} \phi_k^{1:n} \right] \\
&= \mathbb{E}_q \left[n^{-1} \sum_{k=1}^{J} \phi_k^{1:n} \right] \\
&= \mathbb{E}_q \left[n^{-1} n \right] = 1.
\end{aligned}
$$

The theorem also requires own and cross-partials to be positive, which ensures that bounding the sum of the partials is equivalent to bounding the norm of the jacobian of the operator T. It is easy to check that these conditions hold here: increasing the mean utility of product j, increases the probability that j is included, whereas increasing the mean utility level $k \neq j$ reduces j's inclusion probability because products are substitutes. This proves that T is a contraction in R^J.

Williams-Daly-Zachary (WDZ) theorem for higher-order statistics. Following the notation in the main text, let $u_j = v_j + \varepsilon_j$ denote product j's random utility and assume taste shocks $\varepsilon = (\epsilon_1, ..., \epsilon_J)$ are drawn from a distribution that is absolutely continuous with respect to the Lebesgue measure. Further, denote the r^{th} order statistic of $(u_1, ..., u_J)$ as $U_{(r)}(v, \varepsilon)$, with $U_{(J)}(v, \varepsilon) < ... < U_{(1)}(v, \varepsilon)$ and define the expected utility from the r^{th}-ranked product as

$$
I_r(v) \equiv \mathbb{E}_\varepsilon \left[U_{(r)}(v, \varepsilon) \right].
$$

We want to show the following:

$$
\lim_{h \to 0} \frac{1}{h} \mathbb{E}_\varepsilon \left[I_{(r)}(v_k + h, v_{-k}) - I_{(r)}(v) \right] = \phi_k^r(v) \equiv \Pr\{u_k \text{ is the } r^{th}\text{-highest}\}.
$$

The logic of the proof follows Lemma 2.1 in Shi, Shum and Song (2018) and leverages the dominated convergence theorem. For any realization of ε we have two cases two consider. First, if $U_{(r)}(v_{-k}, \varepsilon_{-k}) < v_k + \varepsilon_k < U_{(r-1)}(v_{-k}, \varepsilon_{-k})$, then for h sufficiently small

$$
\frac{1}{h} \left(U_{(r)}(v_k + h, v_{-k}, \varepsilon) - U_{(r)}(v, \varepsilon) \right) = \frac{v_k + h + \varepsilon_k - (v_k + \varepsilon_k)}{h} = 1.
$$

Second, if $v_k + \varepsilon_k < U_{(r)}(v_{-k}, \varepsilon_{-k})$ or $v_k + \varepsilon_k > U_{(r-1)}(v_{-k}, \varepsilon_{-k})$, then for h sufficiently small

$$
\frac{1}{h} \left(U_{(r)}(v_k + h, v_{-k}, \varepsilon) - U_{(r)}(v, \varepsilon) \right) = 0.
$$

Because the distribution of ε is absolutely continuous we also have that $\Pr\{v_k + \varepsilon_k = U_{(r)}(v_{-k}, \varepsilon_{-k})\} = 0$ and $\Pr\{v_k + \varepsilon_k = U_{(r-1)}(v_{-k}, \varepsilon_{-k})\} = 0$. Lastly, observe that, for any

h and almost all realization of ε,

$$\left|U_{(r)}(v_k+h,v_{-k},\varepsilon)-U_{(r)}(v,\varepsilon)\right|\leq|v_k+h+\varepsilon_k-(v_k+\varepsilon_k)|=|h|.$$

Thus, the dominated convergence theorem implies that

$$\begin{aligned}
&\lim_{h\to 0}\mathbb{E}_\varepsilon\left[\frac{1}{h}\left(U_{(r)}(v_k+h,v_{-k},\varepsilon)-U_{(r)}(v,\varepsilon)\right)\right]=\\
&\quad=\mathbb{E}_\varepsilon\left[\lim_{h\to 0}\frac{1}{h}\left(U_{(r)}(v_k+h,v_{-k},\varepsilon)-U_{(r)}(v,\varepsilon)\right)\right]\\
&\quad=\mathbb{E}_\varepsilon\left[\mathbf{1}\{U_{(r)}(v_{-k},\varepsilon_{-k})<v_k+\varepsilon_k<U_{(r-1)}(v_{-k},\varepsilon_{-k})\}\right]\\
&\quad=\Pr\{u_k \text{ is the } r^{th}\text{-highest}\}.
\end{aligned}$$

Representation of order statistics in terms of maxima. I describe how to express expectations of order statistics in terms of maxima of sub-samples formed by progressively removing the highest ranked draws. To do so, I use results on recursive relations between order statistics for non-iid random variables, as detailed in David and Nagaraja (2003) Chapter 5.

Let $(U_1,...,U_J)$ be a random vector with continuous c.d.f F, and denote by $U_{(r),J}$ the r^{th} order statistic of a J dimensional sample. From David and Nagaraja (2003) Chapter 5, we know that for any x, the cdf of the r^{th} order statistic, can be expressed as

$$F_{(r),J}(x)=\sum_{s=r}^{J}(-1)^{s-r}\binom{s-1}{r-1}\sum_{1\leq i_1<...<i_{J-s}\leq J}F_{(s),s}^{(i_1,...,i_{J-s})}(x),\tag{15}$$

where $F_{(s),s}^{(i_1,...,i_{J-s})}$ is the c.d.f of the maximum after $J-s$ draws, indexed by $(i_1,...,i_{J-s})$, have been dropped. The last term in (15) represents the probability that exactly s draws are lower or equal than x. This probability is obtained by summing over all possible permutations $(i_1,...,i_{J-s})$ of $\{1,...,J\}$, such that $i_1<...<i_{J-s}$.[53]

The intuition behind equation (15) is easily derived when considering the expectation of the second highest draw i.e., $r=J-1$, for the case of $J=3$ and $r=2$. In this case

$$\begin{aligned}
F_{2,3}(x)&=\Pr\{U_1,U_2,U_3\leq x\}+\Pr\{U_1,U_2\leq x\wedge U_3>x\}\\
&\quad+\Pr\{U_1,U_3\leq x\wedge U_2>x\}+\Pr\{U_2,U_3\leq x\wedge U_1>x\}\\
&=\Pr\{U_1,U_2,U_3\leq x\}+(\Pr\{U_1,U_2\leq x\}-\Pr\{U_1,U_2,U_3\leq x\})\\
&\quad+(\Pr\{U_1,U_3\leq x\}-\Pr\{U_1,U_2,U_3\leq x\})+(\Pr\{U_2,U_3\leq x\}-\Pr\{U_1,U_2,U_3\leq x\})\\
&=(\Pr\{U_1,U_2\leq x\}+\Pr\{U_1,U_3\leq x\}+\Pr\{U_3,U_2\leq x\})-2\Pr\{U_1,U_2,U_3\leq x\})\\
&=\left(F_{(2),2}^{(1)}(x)+F_{(2),2}^{(2)}(x)+F_{(2),2}^{(3)}(x)\right)-2F_{(3),3}(y)(x).
\end{aligned}$$

[53] There are $\binom{J}{s}$ of such permutations.

Because, equation (15) holds point-wise, by differentiating and integrating over x, it is satisfied also in terms of moments,

$$\mathbb{E}[U_{r:J}] = \sum_{s=r}^{J}(-1)^{s-r}\binom{s-1}{r-1}\sum_{1\leq j_1<\ldots<j_{J-s}\leq J}\mathbb{E}^{(j_1,\ldots,j_{J-s})}[U_{s:s}]. \tag{16}$$

When the random vector $(U_1, ..., U_J)$ is distributed as a T1EV extreme value distribution with location parameters $(\delta_1, ..., \delta_J)$, the expectation of any order statistic is analytical because expected maxima have well-known closed-form expressions,

$$\mathbb{E}^{(j_1,\ldots,j_{J-s})}[U_{s:s}] = \log\left(\sum_{k\neq(j_1,\ldots,j_{J-s})}\exp(\delta_k)\right).$$

For example, under the T1EV distribution, the analytical expression for the expectation of the second highest draw is,

$$\mathbb{E}\left[U_{(J-1):J}\right] = \sum_{i_1=1}^{J}\log\left(\sum_{k\neq i_1}\exp(\delta_k)\right) - (J-1)\log\left(\sum_{k}\exp(\delta_k)\right).$$

Mean-Entropy approximation. Let B the number of asset classes available to worker i and consider a first-order approximation of $\log(\omega_{ib})$ around the equally weighted portfolio:

$$\log(\omega_{ib}) \approx \log\left(\frac{1}{B}\right) + B\left(\omega_{ib} - \frac{1}{B}\right)$$

which can be rearranged as

$$\omega_{ib} \approx \frac{1}{B} - \frac{1}{B}\log\left(\frac{1}{B}\right) + \frac{1}{B}\log\left(\omega_{ib}\right).$$

Next, multiply the above on both sides by ω_{ib} and sum across b to obtain:

$$\sum_b \omega_{ib}^2 \approx \frac{1}{B} + \frac{1}{B}\sum_b \omega_{ib}\log(\omega_{ib}) - \frac{1}{B}\log(B).$$

The mean-variance problem in (7) can be then approximated as

$$\max_{(\omega_{ib})_b}\sum_b \alpha_{ib}\omega_{ib} - \frac{1}{2B}\sum_b \omega_{ib}\log(\omega_{ib}) + \lambda\left(1 - \sum_b \omega_{ib}\right),$$

where λ is the Lagrange multiplier on the constraint $\sum_b \omega_{ib} = 1$. Taking the first order conditions we have that

$$\alpha_{ib} - \frac{1}{2B} - \lambda = \frac{1}{2B}\log(\omega_{ib}) \quad \Rightarrow \quad \omega_{ib} = \frac{\exp\left(2B\alpha_{ib}\right)}{\exp\left(1 + 2B\lambda\right)}$$

Summing up over asset classes b, we obtain

$$\exp(1+2B\lambda)=\sum_b \exp\left(2B\alpha_{ib}\right) \Rightarrow \omega_{ib}=\frac{\exp\left(2B\alpha_{ib}\right)}{\sum_{b'}\exp\left(2B\alpha_{ib'}\right)}.$$

The mean-entropy portfolio allocation and the mean-variance portfolio allocation ($\omega_{ib}=(\alpha_{ib}-B^{-1}\sum_{b'}\alpha_{ib'})+\frac{1}{B}$) carry the same economic content. Both allocations tilt toward assets with a larger α_{ib}, and both converge to the equally weighted portfolio if α_{ib} are equal across b. The main difference is that the mean-variance solution allows for short positions.

Estimating equation implied by active workers' portfolio choice. The portfolio share of a non-default fund j in plan p is given by:

$$\omega_{pj}(x;F_{\psi,\theta})=\sum_d \int \kappa(x_d;\psi_i)\omega_{ib}(x;\theta_i)\omega_{ij|b}(x;\theta_i)dF_{\psi,\theta}(\psi_i,\theta_i|d_i=d)\Pr\{d_i=d\},$$

which under the assumption of homogeneous preferences boils down to

$$\omega_{pj}(x)=\left(\sum_d \kappa(x_d;\psi)\Pr\{d_i=d\}\right)\omega_b(x;\theta)\omega_{pj|b}(x;\theta). \tag{17}$$

Next, observe that the portfolio share of TDF d in plan p is

$$\omega_{pd}(x_d)=\Pr\{d_i=d\}\left(1-\kappa(x_d;\psi)\right)$$

which implies that,

$$\sum_{j\notin D_p}\omega_{pj}(x)=\sum_{d\in D_p}\Pr\{d_i=d\}\kappa(x_d;\psi). \tag{18}$$

Combining (17) and (18) and rearranging, one obtains:

$$\begin{aligned}\frac{\omega_{pj}}{\sum_{j'\notin D}\omega_{pj'}} &= \omega_b(x;\theta)\omega_{pj|b}(x;\theta)\\ &= \frac{\exp(\alpha_{pb})}{\sum_{b'}\exp(\alpha_{pb'})}\frac{\exp(V_{jp}(x;\theta))}{\exp(I_{pb}(x;\theta))}\\ &= \frac{1}{\sum_{b'}\exp(\alpha_{pb'})}\frac{\exp(\mu_b+V_{jp}(x;\theta))}{\exp((1-\rho)I_{pg}(x;\theta))}\end{aligned}$$

where we substituted for $V_{jp}(x;\theta)=x_j'\theta+\xi_{pj}$; $I_b(x;\theta)=\log\left(\sum\exp(V_{jp})\right)$ and $\alpha_{pb}=\mu_b+\rho I_{pb}(x;\theta)$. Letting $b=0$ denote the Cash asset class, after normalizing $\alpha_{p0}=0$, we can divide both sides by $\omega_{p0}=1/\sum_{b'}\exp(\alpha_{pb'})$ and take logs to obtain

$$\log\left(\frac{\omega_{pj}}{\sum_{j'\in D}\omega_{pj'}}\right)-\log(\omega_{p0})=\mu_b+x_j'\theta-(1-\rho)I_{pb}(x;\theta)+\xi_{pj}.$$

To obtain the expression in the main text, we need to substitute for I_{pb} into the previous equation from:

$$\log\left(\frac{\omega_{pb}}{\omega_{p0}}\right) = \mu_b + \rho I_{pb}.$$

After substituting for $\rho I_{pb}(x;\theta)$ and rearranging terms one obtains

$$\begin{aligned}\log\left(\frac{\omega_{pj}}{\sum_{j'\in D}\omega_{pj'}}\right) - \log(\omega_{p0}) &= \mu_b + \rho(x_j'\theta) + (1-\rho)\left(\log\left(\frac{\omega_{pj}}{\sum_{j'\in D}\omega_{pj'}}\right) - \log(\omega_{pb})\right) + \rho\xi_{pj} \\ &= \mu_b + \tilde{\theta}' x_j + (1-\rho)\log(\omega_{pj|b}) + \tilde{\xi}_{jp}.\end{aligned}$$

C Data construction

This section provides details about the data sources and the construction of key variables. I construct a unified dataset using multiple sources: BrightScope, which digitizes the Department of Labor's (DOL) Form 5500 filings; the public Form 5500 files available from the DOL website; mutual fund characteristics and returns from CRSP; and benchmark returns from Morningstar. These data sources allow me to link plan-level characteristics to fund-level metrics in a consistent panel.

BrightScope and DOL's public files. I begin with the BrightScope plan-level data, which contain basic information about employers and their retirement plans. These data include employer identifiers, industry classifications, plan names, DOL-filing identifiers, and indicators for whether a plan offers a brokerage account, has automatic enrollment, or files the short-form F5500-SF. I exclude plans that file the F5500-SF, as they do not report information on available investment options.

Each plan typically has a unique DOL-filing identifier per year. For the small number of cases where duplicate filings exist, I retain only the most recent submission. I use this identifier to merge additional plan-level information available from the DOL website. For each year, I download plan attributes—such as total participants, net assets, and plan inception date—from the main Form 5500 files and extract financial details (including contributions and administrative expenses) from relevant schedules attached to the main filing.

BrightScope also provides information on third-party service providers, which I use to identify each plan's recordkeeper. I isolate recordkeepers by looking at providers' service codes, as defined in the form F5500 instructions. The data also includes information on providers' direct and indirect compensation, which I use to measure administrative fees at the plan level. If multiple administrators are listed for a plan, I assume the recordkeeper with the largest market share is the primary provider.[54] If a plan's recordkeeper is also a mutual fund provider, I classify the plan as vertically integrated and flag all funds within the plan that belong to the recordkeeper's product line as proprietary.

Finally, I merge BrightScope's investment menu data with the plan-level dataset. The investment menu data, structured at the plan-by-fund level, include fund tickers, investment categories, an index fund indicator, the type of investment option (e.g., mutual fund, separate account), and the share of plan assets held in each fund. I classify target-date funds (TDFs) based on fund names and use the plan inception date to identify the year a plan was first introduced.

CRSP and Morningstar data. I incorporate fund characteristics and return data from CRSP, merging them into the BrightScope dataset using fund tickers. Since BrightScope

[54] Almost all plans are matched to at most two recordkeepers, and market share is defined as the share of plans administered in a given year.

filings are annual, I convert CRSP monthly returns into annualized returns and compute average returns over the past year and the past ten years (or since inception if the fund is younger than ten years).

I retrieve fund benchmarks and benchmark returns from Morningstar and merge them into the CRSP data using fund tickers. For funds with multiple benchmarks, I compute benchmark returns as a weighted average of the prospectus benchmarks, with weights reflecting the fund's portfolio allocation across them. If a fund's benchmark is missing, I assign the most common benchmark used by funds in the same Morningstar investment category.

The resulting dataset links plan-level attributes, provider information, and investment menus with fund-level characteristics and benchmark adjusted performance measures.

Calibration of workers' age distribution. In order to estimate the probability of workers being active (κ) I need to measure the share of workers assigned to each of the available default options. To see this formally, recall that the model-implied portfolio share of TDF d in plan p is

$$\omega_{pd} = \Pr\{d_i = d\} \left(1 - \int \kappa(x_d; \psi_i) dF_\psi(\psi_i | d_i = d)\right), \tag{19}$$

which implies that F_ψ can be estimated from variation in ω_{pd}. The challenge here is that I do not observe the share of plan p employees whose default is d, denoted by $\Pr\{d_i = d\}$.

To overcome this challenge, I calibrate $\Pr\{d_i = d\}$ using industry level data on workers' age distribution available from the CPS. Because the BrightScope data contains information on sponsors' industries, the CPS provides me with an estimate of the distribution of workers' age for that particular employer. Having a measure of workers' age distribution is helpful because employees are assigned to their default TDF depending on their age, with younger employees defaulted in TDFs with a target retirement date further in the future.

I define industries at the 2 digit NAICS Code. For each industry l, I compute the mean and standard deviation of the share of workers within the CPS age groups: 20-24, 25-34, 35-44, 45-54, 55-64 and over 65.[55] Letting μ_l^a and σ_l^a the mean and standard deviation of age group a in industry l, and letting $d(a)$ be the TDF in plan p whose target date is the closest to group a retirement date, I assume $\Pr\{d_i = d(a)\} \sim N(\mu_l^a, \sigma_l^a)$. I then estimate it as

$$\mathbb{E}\left[\Pr\{d_i = d(a)\} \middle| \Pr\{d_i = d(a)\} > \omega_{pd}\right] = \mu_l^a + \sigma_l^a \frac{\phi_z(\omega_{pd})}{1 - \phi_z(\omega_{pd})},$$

where the conditioning set follows from the fact that the observed portfolio share bounds $\Pr\{d_i = d(a)\}$ from below, as implied by equation (19).

[55] The industries used in CPS are more granular than the NAICS 2 digits, so that, within a given 2 digits NAICS, there are multiple sub-industries.

To match TDFs and age groups I proceed as follows. First, I extract the target retirement date from each TDF's name. When missing, I use use the modal target date of TDFs in the same Morningstar investment category and/or the same Morningstar benchmark and/or the same BrightScope investment category. The set of available target dates ranges from 2005 to 2065 in steps of 5 years. I assume that the over 65 age group has retirement dates 2005, 2010 and 2015; the age group 55-64 has retirement dates 2020 and 2025; the age group 45-54 has retirement dates 2030 and 2035; the age group 35-44 has retirement dates 2040, 2045 and 2050; the age group 25-34 has retirement dates 2055 and 2060; the age group 20-24 has retirement date 2065.

D 401(k) Lawsuits

This appendix provides details about some recent 401(k) lawsuits. I start by providing some background on 401(k) regulations building on (Mellman and Sanzenbacher (2018)) and then offer few specific examples of recent lawsuits.

The design of 401(k) retirement plans is governed by the Employee Retirement Income Security Act of 1974 (ERISA), with the Department of Labor (DOL) in charge of updating and enforcing such regulation. The law specifies that plan sponsors (i.e., employers) have a fiduciary duty to their plan investors requiring them to design and administer the plan in the 'sole benefit' of plan participants.

While the regulation is clear about the role of plan sponsors as fiduciaries, it provides almost no guidance on how to fulfill such duty in practice. For example, not much is said about how plan fiduciaries should select the type and number of investment options or determine a reasonable level of fees. Instead of laying out specific regulations or guidance, the DOL's general approach to overseeing 401(k)s has been through its own enforcement actions or through privately initiated litigation. Overall, plan fiduciaries are often left to guess what practices comply with ERISA and may only become aware of an alleged violation from a DOL investigation or lawsuit.

Typically there are two reasons that trigger 401(k) lawsuits. First, the inclusion of inappropriate investment options and second, the inclusion of options charging excessive fees. The former was the most common cause after the Great Recession mainly as a consequence of the inclusion of poor-performing employers' own stock. However, this kind of lawsuit has become less common since a 2014 Supreme Court ruling in the case of *Dudenhoeffer v. Fifth Third Bancorp* indicating that plan fiduciaries will not be held liable for failure to predict the future performance of the employer's stock. Since then, most lawsuits involved allegations of excessive investment and administrative fees. In what follows, I describe a few recent examples.

Allen v. M&T Bank Corp (2016). The Plaintiff (Allen) alleges that the Defendant (M&T) breached their fiduciary duties by retaining their proprietary funds within the plan despite the availability of similar lower cost and better performing investment options. According to the plan's Form 5500 filed for 2010, of the 22 mutual fund investment options in the Plan, 8 were from proprietary M&T mutual funds, representing over 30% of all mutual fund investments. However, these proprietary mutual funds charged significantly higher fees than average for performance that most often trailed both the Fund benchmarks and the mutual fund averages.

The Plaintiff provides some specific examples. For instance, the Wilmington Large Cap Value Institutional lagged the performance of a more reasonably priced alternative, Vanguard Equity Income Fund Admiral Shares. The Wilmington fund charged an expense ratio of 1.17%, higher than the Large Cap Value average of 0.83% and the 0.21% fee

charged by the Vanguard Equity Income Fund Admiral Shares. A similar observation is made for the Wilmington Funds Small Cap Growth Institutional Fund charging an expense ratio of 1.39% against the 0.40% fee charged by Vanguard Stragic Small Cap Equity Fund.

Creamer v. Starwood Hotels & Resorts Worldwide Inc (2016). The Plaintiff (Creamer) alleges that the Defendant (Starwood Hotels & Resorts Worldwide, Inc.) serially breached its fiduciary duties in the management, operation and administration of its employees' 401(k) plan. It failed to ensure that fees charged to participants were reasonable. It caused plan participants who invested in index funds to pay seven times more than a reasonable fee. Indeed, the Starwood Plan received from the BrightScope rating service a score of only 61. The top BrightScope rating for peer plans was 90. The Plaintiff highlights that this difference would require sixteen years of additional by Starwood employees to reach the same level of savings as peer plan participants. Starwood participants lost savings of $110,871 per participant.

The Plaintiff also provides some specific examples. For instance, the BlackRock LifePath Index funds (the plan TDF) just hold other BlackRock index funds. BlackRock Life Path 2050 Index Fund institutional shares have net operating expenses of 0.20%. The 2050 Index Fund is a fund that invests all of its assets in other BlackRock funds. 52% of the Life Path Index Fund was invested in the BlackRock Russell 1000 Index Fund. The Russell 1000 Index fund had net operating expenses of 0.08%. Thus, the fee paid by plan participants is 0.20% plus 0.08% for a total of 0.28%. In contrast, the Vanguard Institutional Index Fund Institutional Shares had a total expense ratio of only 0.04% so the plan has chosen funds with fees that are 700% more than the comparable Vanguard fund - a difference of 24 basis points

McCorvey v. Nordstrom, Inc. (2017). The Plaintiff (McCorvey) alleges that the Defendant (Nordstrom) failed to adequately and prudently manage the plan. It allowed unreasonable fees to be incurred by participants and failed to use lower cost investment vehicles. The annual operating fees charged for many of the plan's investment options were substantially higher than reasonable management and operating fees of comparable funds, both index and actively managed funds. These fees were up to 16 times higher than comparable index funds, and up to 2.7 times higher than comparable actively managed funds.

The Plaintiff highlights that the high fee funds in the Nordstrom plan could have been easily replaced by lower cost index funds, TDF, or actively managed funds. For example, the PIMCO Total Return charging 46 basis points in fees could have been replaced by the Vanguard High Dividend Yield Index Fund charging 15 basis points or the Vanguard Growth and Income Fund Admiral Shares (an active fund) charging 23 basis points. Similarly, the average expense ratio for the set of TDFs available in Nordstrom (42 basis

points) could have been five times lower by replacing it with Vanguard Institutional Target Date Funds whose average expense was around 9 basis points.

Pledger v. Reliance Trust (2015). The Plaintiff (Pledger) alleges that the Defendant (Reliance Trust) breached its fiduciary by providing to the plan investment options that contained unreasonable management fees when cheaper versions of the same investments were available to the plan, as were other high-quality, low-cost institutional alternatives. The Plaintiff also alleges that the plan recordkeeper (Insperity) and the Defendant engaged in self-dealing by offering higher-cost investments to the plan's participants, because Reliance selected those investments in order to pay a larger amount of revenue-sharing to the recordkeeper.

Schapker v. Waddell & Reed Fin. Inc (2017). The Plaintiff (Schapker) alleges that the Defendants (WR Financial) selected the investment opportunities made available to the plan participants. During the Class Period, more than 97% of the investment opportunities made available to the plan participants were established and managed by WR Financial or its affiliates. Only one unaffiliated investment option—out of dozens of funds offered each year—was ever offered to plan participants. Because nearly all the investment options the Defendants made available to plan participants were established and managed by WR Financial or its affiliates, the Defendants caused the plan to pay its own Sponsor, WR Financial.

Further the Plaintiff points out that the fees charged to plan participants for their investments were in excess of the fees typically charged by unaffiliated companies for comparable mutual funds and products, and the performance levels of the investment options within the plan were worse than the performance achieved by unaffiliated companies for comparable mutual funds and investment products. Defendants could have selected comparable investment products from unaffiliated companies that cost less and performed better than the proprietary branded investment products to which the Defendant limited the plan participants.