Plan Menus, Retirement Portfolios, and Investors' Welfare

Marco Loseto*

2 April, 2024

Please click here for the latest draft

Abstract

Employer-sponsored retirement plans are a crucial component of the US savings system. Many of these plans include funds substantially more expensive than comparable alternatives available in the marketplace. To understand why these high-cost investment options are provided in equilibrium and to quantify the effects of alternative plan design policies on investors' welfare, this paper introduces a structural model of plan menu choice and fee competition between funds. The model features a two-layer demand system: plan sponsors construct retirement menus, and plan investors form portfolios from the available options. Consistent with the presence of agency frictions, model estimates imply that sponsors are only half as responsive to funds' fees as investors and favor the inclusion of funds affiliated with the plan recordkeeper. In response, funds charge sizable margins to investors. This is especially evident for Target-Date funds (TDFs), whose estimated margins are nearly twice as large as the median of all funds. Because model estimates suggest that a sizable share of investors is inactive, counterfactual policies mandating the inclusion of low-cost default options can generate considerable welfare gains for plan investors.

KEYWORDS: 401(k), retirement plan design, portfolio choice, funds fees

JEL Classification: G11, G23, G28, G51, J32, L13, L51

*Kenneth C. Griffin Department of Economics, University of Chicago, mloseto@uchicago.edu. I am deeply grateful to Lars Hansen, Ali Hortaçsu, and Scott Nelson for their continuous guidance and support. I am also grateful to Doug Diamond, Lorenzo Garlappi, Niels Gormsen, John Heaton, Adam Jørring, Ralph Koijen, Federico Mainardi, Eric Richert and Francesco Ruggieri for their detailed and thoughtful feedback. For their helpful comments I thank seminar participants at the Bank of Italy, Bocconi, Boston College, Duke, the Federal Reserve Board, IESE, Rice, Maryland, USI Lugano, 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 defined contribution (DC) retirement plans are a crucial component of the US savings system, holding nearly $11 trillion in assets as of 2021.[1] These plans allow employees to allocate a portion of their pre-tax income towards retirement savings through a range of investment options, typically mutual funds, to build long-term wealth. For many workers, the assets held in DC plans are among the most important components of their balance sheets and are a significant determinant of their future retirement security.[2]

Despite their importance, many plans do not provide their investors (i.e., employees) with cost-efficient investment options. For instance, in 2019, over half of the plans failed to offer low-cost S&P 500 index trackers. Perhaps even more strikingly, one out of every five plans did not offer an equity fund with an investment fee below 10 basis points (Figure 1).[3]

A closer look at plan expenses reveals substantial dispersion across sponsors, with the difference in the average expense between plans at the 75th and 25th percentiles of about 40 basis points (Figure 2). 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.[4]

Beyond dispersion, plan expenses are also surprisingly high, with the asset-weighted average expense ratio for the median plan in the 2019 cross-section close to 40 basis points.[5] For context, in that same year, had a retail investor constructed a portfolio of Vanguard index funds to obtain exposure to all asset classes available in a typical retirement plan, the expense ratio would have been more than four times lower.[6] Additionally, more expensive plans do not produce better investment performance for their investors (Figure 3).[7] All things considered, it is unsurprising that employees have increasingly

[1] Among those, 401(k) are the largest totalling $7.7 trillions. As a share of the US retirement market assets, employer-sponsored DC plans account for 30%. If one includes individual retirement accounts (IRA), DC plans account for 63% of the US retirement assets. https://www.icifactbook.org/.

[2] According to the 2019 Survey of Consumer Finances (SCF), for a working age household, the average account balance in a DC plan (including IRAs) was nearly $270,000. More generally, 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, Campbell and Ramadorai (2016).

[3] On average, around 50% of plans do not offer the cheapest share class of a fund even though they meet its minimum investment requirement (Figure A2).

[4] The calculation assumes a working period of 40 years and that the annual return is the same. In practice, plans with higher expenses tend to have gross of fees returns that are even lower (Table A1).

[5] These patterns are not limited to the 2019 cross-section. Plan expenses were even higher before 2019. At the same time, the dispersion in plan expenses has been roughly stable over time (Figures A3) and remains even when comparing plans of similar size (Figures A4).

[6] In 2019, the expense ratio for a Vanguard equally-weighted portfolio of retail index funds, including its International equity index funds (VEMAX, VEUSX, VPADX), US Equity funds (VGSLX, VFIAX, VIMAX, VSMAX) and Bond Fund (VBTLX) is below 10 basis points. By retail, I mean that the minimum investment required is none or limited. Figure A5 compares the expense for the median plan against this portfolio of Vanguard index funds over time.

[7] Table A1 and Figure 3 show that the plan-level (gross of fees) performance tends to be lower for more expensive plans. Similar patterns have been found in the context of active investing (Gil-Bazo and Ruiz-Verdu (2009)) contrasting what frictionless models with rational investors predict (Berk and Green (2004)).

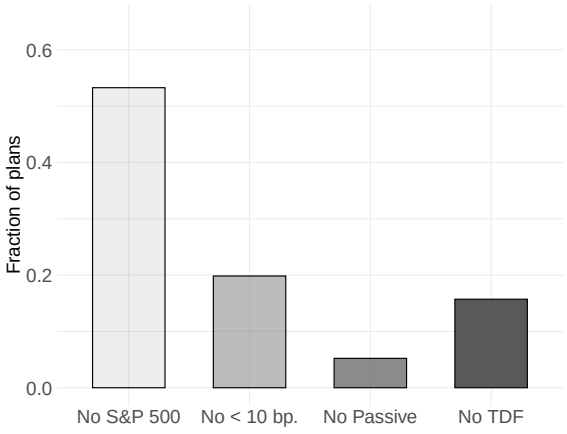

Figure 1: Plan quality for the year 2019. 'No S&P 500' is the share of plans without an S&P 500 tracker with a fee below 10bp.

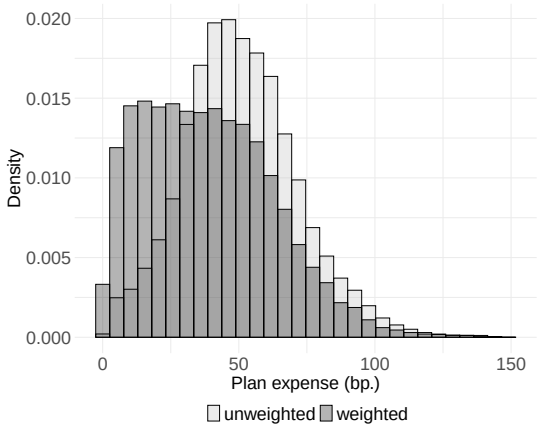

Figure 2: Distribution of asset-weighted and unweighted average plan expense ratio across plans for the year 2019.

sought to hold plan sponsors (i.e., employers) accountable for allegedly violating their fiduciary duties, with high investment costs emerging as the common theme in many recent 401(k) lawsuits (Mellman and Sanzenbacher (2018)).[8]

Why do many plans feature investment options that are less cost-efficient than comparable alternatives in the marketplace? And can policy regulating the design of retirement plans effectively help reduce these costs and improve investors' outcomes?

To address the first question we need to understand what drives sponsors' decisions to include high-cost funds in their retirement plan. Although sponsors have a fiduciary duty to design their plan in the investors' best interest, agency frictions may reduce sponsors' sensitivity to funds' fees and make them value attributes other than fees when constructing their plan menu. For example, sponsors may favor the inclusion of funds affiliated with their plan provider (a.k.a recordkeeper, Pool, Sialm and Stefanescu (2016)) or may favor the inclusion of costlier funds to reduce direct fees paid to the recordkeeper (Badoer, Costello and James (2020), Bhattacharya and Illannes (2022), Pool, Sialm and Stefanescu (2022)).[9]

Addressing the second question requires understanding how investors allocate their contributions across the options available in their plan. Their investment behavior is crucial to evaluate policies regulating the design of retirement plans. For instance, if many investors are inactive because of inertia (Madrian and Shea (2001), Beshears, Choi, Laibson and Madrian (2009), Choi (2015)), policies mandating the inclusion of low-cost options, like an S&P 500 tracker, might be ineffective. This is because inactive investors are unlikely to reallocate their contributions to the new low-cost option.

From a supply-side perspective, funds' competition incentives are key to rationalize

[8] I provide more details of some recent 401(k) lawsuits in Appendix E.

[9] Plan sponsors outsource administrative tasks such as maintaining employees' account balances to plan providers (a.k.a recordkeeper), which are often vertically integrated into fund provision. Employers can compensate plan providers either directly or indirectly by offering expensive funds that pay higher kickbacks. Part of this compensation can also cover for services, like financial advising, that I do not observe. Differences in unobserved services can explain the observed dispersion in plan expenses. Figure A6 shows that the dispersion in plan expenses remains substantial even when comparing employers within the same industry, size and plan provider. This suggests that variation in unobserved services across providers or employers cannot be the major source of differences in plan expenses.

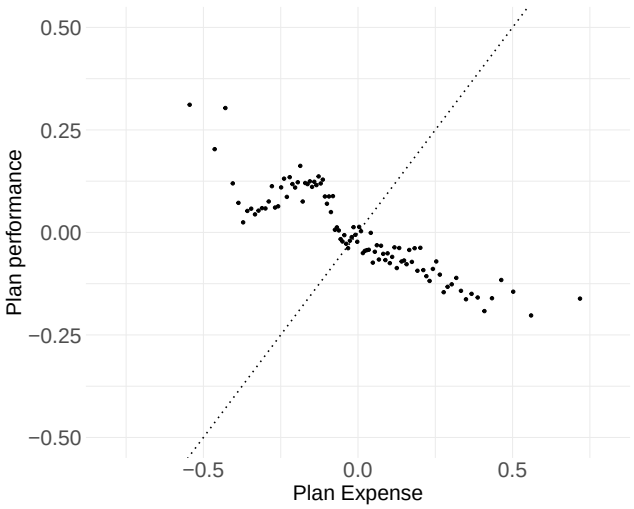

Figure 3: Binned scatter of plan (gross) performance and average plan expense corresponding to the specification in the second column of Table A1. Expenses and performance are yearly demeaned and measured in percentage points.

excessive fees and to predict how these fees will respond to policy changes. Understanding competition in this space requires a model that accounts for how funds set fees in response to sponsors' incentives to include them in their plan and in response to investors' incentives to substitute toward competitors' funds available in the plan.

In this paper, I combine tools from finance and industrial organization to develop a structural model of plan menu choice, retirement portfolio choice and fee competition between funds providers. The model features a two-layer demand system where, in the first layer, sponsors construct their retirement plan and, in the second layer, plan investors form their retirement portfolio from the options available in their plan menu.

Plan sponsors compose their menu by selecting investment options from the pool available in their recordkeeper's network of funds.[10] In modelling sponsors' preferences, I accommodate for two types of agency frictions. First, sponsors have a taste for funds' affiliation. Everything else equal, a fund affiliated with the plan recordkeeper will be more likely to be included in plan menus. Second, sponsors' fee sensitivity can be arbitrarily misaligned to that of investors. This captures how the compensation structure between sponsors and recordkeepers might incentivize the former to favor the inclusion of expensive funds. I model sponsors' menu choice as a two-stage multiple discrete choice problem. First, sponsors decide whether to include or not a particular investment category. Once this decision is made, they proceed to choose the options within that category that offer the highest indirect utilities. Sponsors face some cost to adding more than one option within each category, which the model captures by assuming that the number of options included is drawn randomly with decaying probability.[11]

[10] When forming their plan menus sponsors might not be aware of all available options. Part of this is because different recordkeepers have different pre-existing relationships with specific mutual fund providers. In fact, plan-level data indicates that recordkeepers' networks of funds are far from perfectly overlapping (Figure A7).

[11] Adding options may be costly for several reasons. First, sponsors fiduciary duty is not limited to the design of the retirement menu but also requires them to monitor the included options and provide plan investors with up-to-date information about their performance. Second, sponsors' asset base (i.e., the total contributions) is limited, and with too many options, it could be challenging to meet the minimum

In the second layer, plan investors allocate their contributions among the options available in their menu. Unlike in much of the prior IO literature, investors' preferences are tied to a standard portfolio problem, with risk aversion over the variance of returns. These preferences drive substitution patterns between funds, and unlike in most structural IO demand models, this framework generally allows funds to be either complements or substitutes to each other. In modelling how this type of demand system determines funds' fee-setting incentives, I rely on tools I develop in a related paper (Loseto (2023)), which exploits the close relationship between standard demand systems and the network games literature (Ballester, Calvó-Armenagol and Zenou (2006)). To capture investors' inertia, I allow for the possibility that some investors do not make an active investment decision. Instead, they default their asset allocation into the plan default option, typically a Target-Date fund (TDF).

I estimate sponsors' and investors' preferences using comprehensive plan-level data. This information is collected from 401(k) plan menus and asset allocations as reported by plan sponsors on form F5500 to the Department of Labor (DOL).[12] Specifically, I compute the probability of a given investment fund being included in a plan from the observed plan menus. The observed variation in these inclusion probabilities enables me to identify and estimate plan sponsors' preferences. Similarly, the observed variation in plan-level asset allocations allows me to identify plan investors' preferences. To account for the possible endogeneity of investment fees, I exploit the granularity of the data to control for unobserved demand shocks along several dimensions, including investment category fixed effects, sponsors fixed effects, funds' brand fixed effects and a fixed effect for passive funds. To further ensure that the residual variation in fees is uncorrelated with unobserved demand shocks, I use two types of cost-shifter instruments: funds' turnover ratios and a Hausman-type instrument. Funds' turnover measures trading-related transaction costs passed on to investors through higher fees (Pástor, Stambaugh and Taylor (2020)).[13] As for the Hausman strategy, it uses the average fee charged by funds within the same fund provider but from other investment categories as an instrument for funds' fees (Hausman (1996)).

Model estimates indicate that plan sponsors are less sensitive to fees than plan investors. This is particularly evident when comparing them to investors actively forming their retirement portfolios, who are over twice more elastic to fees than sponsors.[14] This misalignment in sponsors' and investors' elasticity to fees suggests that sponsors may not adequately internalize their employees' preferences when constructing retirement plan menus. The estimates also indicate that sponsors have a strong preference for including funds affiliated with their recordkeeper. Quantitatively, being affiliated with the plan recordkeeper increases the inclusion probability by 0.36 percentage points, a magnitude four times higher than what a 10 basis points reduction in funds' fees would lead to.

I model supply as a differentiated Bertrand oligopoly where funds providers set fees

investment requirements required by investment funds. In estimation, I calibrate this probability to match the observed distribution of the number of funds included within each category (Figure A1). In Appendix B, I offer a simple microfoundation for the optimal choice of the number of options to be included within each category building on the Stigler (1961) simultaneous search model.

[12] See Section 3 for more details on the data.

[13] Pástor, Stambaugh and Taylor (2020) show how funds' profit-maximizing behaviour implies that funds' optimal fee must depend on funds' trading costs which are a function of funds' turnover.

[14] In Appendix D I model sponsors' preferences as a weighted average of their true preferences and investors' preference and estimate that sponsors weigh their own preferences nearly three times more than investors' preferences.

simultaneously before sponsors make their plan menu decisions and plan investors form their portfolios. Funds compete along two margins. First, they compete to be included in sponsors' plan menus and internalize how agency frictions reduce sponsors' fee elasticity. Second, conditional on inclusion, they compete for plan investors' assets. When setting fees, they internalize investors' inertia, and that the likelihood of facing close competitors is low because sponsors tend to include no more than one fund in each investment category (Figure A1).[15]

After estimating sponsors' and investors' demands, I use the Nash-Bertrand equilibrium conditions to recover funds' marginal costs and price-cost margins. The median fund, spanning all fund types, charges a margin of 14 basis points. Given the median expense ratio, this translates to a markup of approximately 20%. Looking at passive funds and TDFs specifically, their estimated marginal costs indicate that these funds are more efficient than others. However, they do not fully transfer these efficiency gains to investors. This is especially evident for TDFs, where the estimated median markup stands at about 39%— nearly twice the median markup observed across all funds.

TDFs' pricing power arises from two forces. First, they are often funds affiliated with the plan recordkeeper. Sponsors value this affiliation, making these funds more likely to be included in plan menus. Second, following the 2006 Pension Protection Act, they were designated as a qualified default option (QDIA) for employer-sponsored retirement plans, allowing them to attract demand from inactive investors who are inelastic to fees. A recent study by Vanguard finds that the share of plan investors holding a single TDF increased from 20% in 2010 to 54% in 2019, suggesting that the fraction of inactive investors might be far from negligible.[16] Model estimates match this evidence closely, indicating that two out of five investors did not make an active investment decision and that the fraction of inactive investors more than doubled over time from around 25% in 2010 to nearly 60% in 2019.

In the last part of the paper, with the estimated demand and supply parameters, I explore a series of counterfactuals to quantify the effects of plan design policies on plan investors' welfare. In the first counterfactual, I eliminate agency frictions whereby plan sponsors favor the inclusion of funds affiliated with their recordkeeper. This policy turns out to be ineffective: sponsors simply substitute from affiliated funds to similarly expensive unaffiliated funds and overall costs for investors do not meaningfully decrease. In other words, eliminating sponsors' preference for affiliated funds does not make them more responsive to fees.

In the second set of counterfactuals, I consider policies that mandate the inclusion of a low-cost equity index fund tracking the S&P 500 and the inclusion of a low-cost TDF. Mandating the inclusion of low-cost equity S&P 500 funds leads to an increase in investors' surplus of about 2%. The increase is modest because investors value lower fees, but they also want to diversify across the available funds, dampening the incentive to substitute toward the low-cost index fund. Moreover, this policy only benefits active

[15]While most sponsors aim to provide investors with a broad range of investment categories (e.g., Large Cap Value, etc.) they typically offer only one fund per category (Figure A1). Although this plan structure helps manage risk, it could weaken competition between funds, as price competition is more intense when products are more alike e.g., when funds belong to the same category. Consistent with this Bertrand-type of intuition, plan-level data suggests that fees are inversely correlated with the number of options offered within a specific category (Table A2).

[16]The study is based on an examination of retirement plan data from 5 million defined contribution plan participants across Vanguard's recordkeeping business. See *'How America Saves'*, Vanguard (2022).

investors, leaving inactive investors' surplus unchanged. Mandating the inclusion of low-cost TDFs increases investors' surplus by about 11%, a magnitude more than five times larger than the previous policy. This policy is more effective because it also benefits inactive investors who reallocate their entire portfolio to the low-cost TDF.

Lastly, I consider a policy that caps funds' expense ratios at 50 basis points. Under this policy, sponsors can only offer funds with expenses below this cap.[17] Investors' outcomes further improve, with an increase in surplus of approximately 14%. This policy is effective because it affects the entire menu of options by limiting the inclusion of the most inefficient funds. It benefits inactive investors because it replaces the most inefficient TDFs, and it benefits active investors by reducing the costs of all options available in the plan menu.

The rest of the paper proceeds as follows. Section 2 describes how this paper contributes to the literature. Section 3 describes the data. Section 4 sets up the demand side of the model. Section 5 focuses on the supply side and characterizes the equilibrium fees. Section 6 estimates the demand side of model. Section 7 turns to the supply side and recovers funds' price-cost margins. Section 8 presents the results of policy counterfactuals and Section 9 concludes.

2 Contributions to the Literature

This paper contributes to the literature that studies retirement investing and the design of retirement plans. A large part of this literature has focused on the demand side by studying 401(k) enrollment decisions with a particular focus on the role of automatic enrollment into default options (Madrian and Shea (2001), Beshears, Choi, Laibson and Madrian (2009), Carroll, Choi, Laibson, Madrian and Metrick (2009), Choi (2015)), behavioural biases in retirement investing (Benartzi and Thaler (2001), Huberman and Jiang (2006), Benartzi and Thaler (2007), Tang, Mitchell, Mottola and Utkus (2010)), the demand for financial advice (Chalmers and Reuter (2020), Reuter and Richardson (2022)), and the implications of automatic enrollment on saving behaviour over the life-cycle (Choukhmane (2021), Duarte, Fonseca, Goodman and Parker (2022)).

Another part of the literature has looked at the supply side and has mainly focused on empirically examining the role of agency frictions between plan providers (i.e., record-keepers) and plan sponsors (i.e., employers). Pool, Sialm and Stefanescu (2016) show that plan providers vertically integrated into fund provision tend to favor affiliated funds, Badoer, Costello and James (2020) provide evidence of how plan providers trade-off direct fees from the sponsor with indirect fees paid by funds via revenue-sharing agreements and Pool, Sialm and Stefanescu (2022) show that funds paying revenue-sharing fees are more likely to be included in a plan and less likely to be deleted. More recently, Gropper (2023) shows how employers distort plan menus to reduce litigation risk. They do so by reducing the number of options offered, which reduces employees' retirement wealth. Building on this evidence, Bhattacharya and Illannes (2022) take a more structural perspective and develop a model where revenue-sharing fees are the outcome of a bargaining game between sponsors and recordkeepers. Yang (2023) instead models employers' dynamic

[17] Under this policy, most of active funds will not find profitable to operate in the retirement market. Nevertheless, investors' surplus improves because most of these funds do not generate enough alpha to justify their expenses (Jensen (1968), Gruber (1996), Carhart (1997), Fama and Kenneth (2010)).

decision to switch plan provider and shows how switching costs can rationalize a sizable share of the observed dispersion in plan expenses.

This paper lies in between these two strands of work. Its primary contribution is to develop and estimate a structural model of plan menu choice, retirement portfolio choice and fee competition between differentiated fund providers. I model sponsors' menu choice as a multiple discrete choice problem building on the workhorse discrete choice frameworks developed in Berry (1994) and Berry, Levinsohn and Pakes (1995). I accommodate for agency frictions by allowing sponsors' preferences to depend on funds' characteristics that relate to the identity of the plan recordkeeper, without modelling the latter as separate agents. For example, I allow sponsors to have a taste for whether or not a fund is affiliated with its recordkeeper. Moreover, by modelling and identifying sponsors' and investors' preferences separately, I can allow their sensitivity to expenses to be arbitrarily misaligned and capture agency frictions whereby sponsors favor expensive funds to reduce the direct fees paid to the recordkeeper.

A key goal of the paper is to quantify the effects of alternative plan design policies on investors' welfare. To this end, I incorporate investors' portfolio decisions into the model. After sponsors' menus have been chosen, I assume plan investors with quadratic preferences (Markowitz (1952)) form their retirement portfolio from the options available in their menu. In a recent contribution, Egan, MacKay and Yang (2023) also model retirement portfolio choice as a mean-variance problem and show how variation in expense ratios can identify investors' beliefs and risk aversion separately. I complement their work in two ways. First, I extend their methodology by also allowing investors to be inactive with some probability, in which case, investors' contributions are defaulted into some of the available TDFs. Second, I develop the supply side and characterize the Nash-Bertrand equilibrium fees when investors have mean-variance demand rather than discrete choice demand.

I assume funds compete in a differentiated oligopoly by setting fees simultaneously a la Nash-Bertrand and use the implied first-order conditions to recover funds' marginal costs and price cost margins. This links my paper to the empirical industrial organization literature on imperfect competition in the mutual fund industry. Most of this literature models investment decisions on the demand side as a standard discrete choice problem. This is done either because the focus is on competition between financially homogeneous products such as S&P index funds (Hortaçsu and Syverson (2004), Brown, Egan, Jeon, Jin and Wu (2023)), ESG and non-ESG funds that track the same underlying index (Baker, Egan and Sarkar (2022)) and variable annuities (Koijen and Yogo (2022)), or because investors are assumed to be risk neutral (Massa (2003), Roussanov, Ruan and Wei (2021)).

No paper in this literature has modeled fee competition between differentiated funds when investors have mean-variance preferences and characterized the resulting equilibrium fees. To the best of my knowledge this is the first paper to do so. Specifically, I show how the quadratic structure of investor preferences implies that equilibrium fees can be decomposed into three components. One of these components, which I refer to as 'Hotelling markdown', is equivalent to the vector of Bonacich centralities of a network in which funds are the network nodes and funds' substitution patterns are (inversely) proportional to the network edges. This markdown captures how much a monopolist should give up when it faces competitors that are closer in the space of product characteristics. A more central fund faces more similar competitors and must charge lower fees (e.g., has

a higher Hotelling markdown).

In a related paper (Loseto (2023)), I develop part of this decomposition in the context of price competition between multi-product firms selling differentiated products to consumers with quadratic preferences. Building on a growing literature on linear-quadratic network games (Ballester, Calvó-Armenagol and Zenou (2006), Ushchev and Zenou (2018), Pellegrino (2023)), I show how a firm's Bonacich network centrality fully summarizes its pricing power: firms that are more central charge lower prices because their products are not sufficiently differentiated from their competitors products.[18] The characterization of the equilibrium fee I provide in this paper is more general because the network is ex-ante unknown to players (e.g., funds do not know which plan will include them), and players' actions influence the resulting network structure (e.g., when setting fees, funds influence their plan inclusion probability).[19]

Lastly, my paper joins a growing literature at the intersection between industrial organization and financial economics that uses structural models to study market structure and competition. This literature has looked at mortgages (Allen, Clark and Houde (2014), Benetton (2021), Robles-Garcia (2021), Agarwal, Grigsby, Hortaçsu, Matvos, Seru and Yao (2021)), credit cards (Nelson (2020)), loans (Cuesta and Sepúlveda (2020), Benetton, Buchak and Robles-Garcia (2022)), banking (Egan, Hortaçsu and Matvos (2017), Buchak, Matvos, Piskorski and Seru (2018), Buchack, Matvos, Piskorski and Seru (2022)), municipal bonds (Brancaccio, Li and Schüroff (2020)), auctions (Hortaçsu and McAdams (2010), Kastl (2011), Richert (2021)), variable annuities (Koijen and Yogo (2022)) and the market of financial advice (Egan (2019), Bhattacharya, Illanes and Padi (2020)).

3 Institutional Setting 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.

3.1 What is an employer-sponsored retirement plan?

Employer-sponsored retirement plans are savings vehicles designed to assist employees in accumulating wealth for their retirement years. Although there are various types of employer-sponsored retirement plans, the most common is the defined contribution 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 pre-determined but instead depends on the contributions made and on the returns obtained from the investment options available in the plan.

[18] In Appendix F I describe and simulate a simple differentiated Bertrand-Network game.

[19] Grice and Guecioueur (2023) also study fee competition between investment providers from a network perspective. Building on Grice (2023), they propose a model of competition between fund families where the competitive network is micro-founded from investors' imperfect consideration. My model instead features competition at the fund level and belongs to the class of network games studied in Loseto (2023) where the underlying network is determined by assets' characteristics, thereby allowing for arbitrary substitution patterns between products. Moreover, in the model I am considering, the network structure is unknown to players because investment providers set fees before plan sponsors design their retirement menu.

Under a DC plan, employees 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 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 employee's contributions. For example, an employer might contribute 50 cents for every dollar the employee contributes, up to a certain percentage of the employee's salary (typically around 6%). More importantly, employers are in charge of selecting and monitoring which investment options, typically mutual funds, are to be included in the plan. Once the employee 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 4 summarizes graphically the structure of a DC plan. Plan sponsors typically hire recordkeepers to assist in the design of their retirement menu and to perform administrative tasks such as maintaining account balances. Most recordkeepers, like for example Fidelity, are also providers of investment funds and it is not uncommon for retirement plans to include funds from the recordkeeper product line. Overall, a retirement plan consists in a set of assets, spanning a broad range of investment categories, selected by sponsors and recordkeepers. Under the Employee Retirement Income Security Act (ERISA), plan sponsors 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 describe in Appendix E.

3.2 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 data regarding the investment options offered within an employer's retirement plan and the plan-level dollar allocation to each of these options.

Although Form 5500 filings are available for download 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 for the years 2010 to 2019 from BrightScope Beacon who collects and digitizes these data directly from the publicly available DOL filings. Overall, the data covers more than 90% of total plan assets in each year from 2010 to 2019, digitizing an average of 55,000 plans

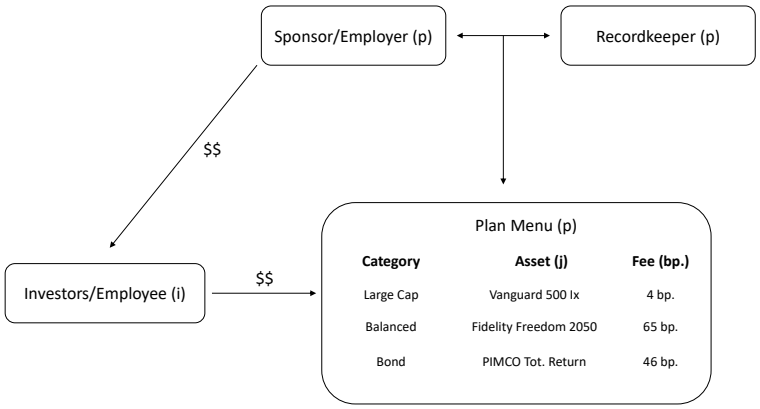

Figure 4: Administrative structure of a defined-contribution employer-sponsored retirement plan.

per-year.[20]

I complement this data with additional information from the DOL Form 5500, including the number of plan participants and the identity of each plan recordkeeper. Also, I obtain data on funds' historical expense ratios from the Center for Research in Security Prices (CRSP) and merge those in the main dataset using funds' tickers. BrightScope also provides data on funds' expense ratios, but the historical expenses are only available starting from 2016.[21]

Table 1 provides some summary statistics for the plans in my sample. Following Bhattacharya and Illannes (2022), I focus on plans whose number of participants is between 100 and 5000, representing roughly 95% of the whole sample.[22] The average plan has close to 30 millions dollars in assets and around 475 participants. Both measures of plan size are right-skewed due to the presence of large plans, with the median plan having assets ranging around 10 millions and 250 plan participants.

Looking at the characteristics of the investment menu, the average plan offers 25 investment options across 15 different investment categories. For the average plan, one out of four options is affiliated with the plan recordkeeper. Moreover, most of the plans offer at least a TDF fund and a passive fund. Turning to plan expenses, the average expense ratio charged by the average plan is of about 63 basis points. This is more than 10 basis point higher than its asset-weighted average, suggesting that plan investors tilt their allocations toward cheaper funds.

Table A4 reports some summary statistics for the funds offered in the sample of plan menus I observe. Excluding cash accounts and common stocks, there are 5600 distinct funds for which at least a ticker identifier is available.[23] In my sample of plans, the average

[20] See Table A3 for more details on the data coverage.

[21] Before 2016, BrightScope reports the most recent funds' expense ratio which I replace with the one obtained from CRSP.

[22] Vary large plan sponsors are more likely to engage in private negotiations with recordkeepers and mutual fund providers to obtain fees that are lower than fees available in the mutual fund market. These privately negotiated fees are not available in the BrightScope data.

[23] The same fund may have multiple tickers, one for each different share class.

	N	Mean	SD	p5	p25	p50	p75	p95
Tot. Assets (mln.)	60798	27.85	59.15	1.78	5.55	11.15	25.11	107.64
N. of participants	60798	475.19	597.08	118.00	160.85	250.10	494.83	1723.87
N. of options	60798	24.72	14.46	12.96	17.88	21.55	26.91	47.86
N. of categories	60798	14.96	3.27	9.55	13.00	15.00	17.01	20.00
Affiliated (%)	60798	25.13	27.56	0.00	0.64	14.33	43.24	82.33
Target Date (%)	60798	78.36	37.63	0.00	73.17	100.00	100.00	100.00
Passive (%)	60798	92.17	23.67	17.35	100.00	100.00	100.00	100.00
Avg. expense (pp.)	60798	0.63	0.21	0.26	0.49	0.63	0.76	0.98
Avg. expense (w.)	60798	0.51	0.24	0.11	0.35	0.51	0.66	0.93
Assets per participants ($ thous.)	60742	50.50	60.80	4.84	17.50	35.05	64.67	143.73
Equity	60798	0.60	0.09	0.46	0.56	0.61	0.66	0.73
Bond	60798	0.18	0.06	0.09	0.14	0.18	0.21	0.28
Balanced	60798	0.17	0.09	0.02	0.13	0.17	0.22	0.31
N. of years	60798	5.33	2.90	1.00	3.00	5.00	8.00	10.00
N. of recordkeepers	60798	1.20	0.47	1.00	1.00	1.00	1.00	2.00

Table 1: Plan level summary statistics for the years 2010 to 2019. Each variable is first averaged within plan across years and then tabulated across plans. The variable 'N' is the number of plans and the variable 'N. of years' is the number of years a plan is observed in the sample. Sample is for sponsors with number of participants between 100 and 5000.

fund manages 191 million dollars of retirement assets and has an average portfolio share of about 3%. The distribution of assets is right-skewed, with the median fund having 7 millions in assets. Interestingly, the share of plan assets allocated into a given fund varies significantly across plans, with the average standard deviation of the portfolio share across funds being as large as 3%.

The data also suggests that sponsors review their menu of investment options often. In Table A4, the penultimate row reports the average fraction of years a fund is included within a plan menu. On average, a fund is offered in only half of the years that I observe the plan menu, implying that sponsors regularly modify their menu offerings. The same does not appear to be true when looking at plans' recordkeepers, with more than 75% of plans having the same recordkeeper over the sample period (Table 1).

4 Demand

This section describes the two-layer demand side of the model. I start from the first layer where I describe sponsors' preferences and the menu choice problem they face. After that, I derive funds' plan inclusion probabilities implied by sponsors' demand. Next, I turn to the second layer of demand where I first describe investors' preferences and then derive individual and aggregate asset demand systems.

4.1 Sponsors' menu choice problem

Throughout the paper I will index plan sponsors by p and mutual funds by j. All vectors are in bold. The goal of sponsor p is to choose a set of mutual funds to include into its retirement plan. Typically sponsors hire recordkeepers to help in designing and managing their plan menu and, empirical evidence suggests that the set of funds sponsors consider to being with is strongly influenced by the recordkeeper identity. From a modelling point

of view, I capture this by allowing for heterogeneity in sponsors' consideration sets.[24] Formally, I assume that sponsor p considers with positive probability a subset $N_p \subset N$ of all mutual funds available. Empirically, I assume that fund j belongs to N_p if I observe at least one plan that includes j and shares the same recordkeeper as p. In other words, sponsors with the same recordkeeper have the same consideration set.

Sponsor p's random utility from including fund j is given by

$$u_{jp} = V_{jp}(\boldsymbol{\theta}_p) + \varepsilon_{jp} \tag{1}$$

with the non-random utility part, V_{jp}, defined as

$$V_{jp}(\boldsymbol{\theta}_p) = \boldsymbol{w}'_{jp}\boldsymbol{\theta}_p + \zeta_j \tag{2}$$

where $\boldsymbol{w}_j$ is the vector of fund j's observed characteristics including its expense ratio, past returns and an indicator for whether j is a fund affiliated with p's recordkeeper, ζ_j captures characteristics, possibly correlated with fees, unobserved to the econometrician and, ε_{jp} is a random preference shock distributed as T1EV. When modeling the preferences of sponsors I place no restrictions on their parameters, $\boldsymbol{\theta}_p$, thereby allowing for such preferences to be arbitrarily misaligned from those of plan investors. In Appendix D I show how plan investors' preferences can be nested into sponsors' preferences and how to interpret the parameters $\boldsymbol{\theta_p}$ as a weighted average between sponsors' true preference parameters and investors' preference parameters.[25]

The preference specification in (1) captures two types of agency frictions that have been documented in the literature. First, the indicator for funds' affiliation allows for the possibility that sponsors prefer to include funds affiliated with the plan recordkeeper (Pool, Sialm and Stefanescu (2016)).[26] If the preference coefficient on the affiliation indicator is positive, an affiliated fund will be more likely to be included than an otherwise identical unaffiliated fund. Second, sponsors might be willing to include expensive funds to reduce their direct payments to the recordkeeper (Badoer, Costello and James (2020), Bhattacharya, Illanes and Padi (2020)). If this incentive is strong enough, it will affect the estimated preference coefficient on funds' fees and will likely reduce sponsors' elasticity to funds' fees.

I assume that funds are classified into investment categories indexed by $g \in \{1, ..., G\}$ and model the menu choice as a two-stage decision problem. In the first stage, sponsors choose which category to offer in their plan and this decision is made independently across categories. For each selected investment category, in the second stage, sponsors evaluate the options available making within category comparisons and selecting the options providing the highest indirect utility.

To complete the decision problem, I need to specify how the number of options to be included within each selected category is chosen. I assume that this number is drawn randomly from a geometric distribution with parameter q, rather than modelling this

[24] In principle sponsors can choose/change their recordkeeper. Yet, the data suggests that sponsors tend to stick with the same recordkeeper over time (Table 1). This motivates why I to abstract from modelling recordkeepers and sponsors as separate agents.

[25] Identifying and estimating how much sponsors weigh their investors preferences requires additional assumptions. In Appendix D I provide more details and find that sponsors weigh their own preferences three times more than their investors preferences.

[26] The data suggests that affiliated fund are nearly twice more likely to be included in a plan, even when comparing employers within the same industry and with similar size (Figure A11).

choice as the outcome of a rational decision problem.[27] This implies that the probability of n options being included in a given category is given by:

$$q(n) \equiv q(1-q)^{n-1} \quad \text{for } n = 1, 2, ... \tag{3}$$

This modelling assumption is guided by the empirical observation that sponsors tend not to include more than one option per investment category. Figure A1 plots the empirical distribution of the number of options offered within investment category and shows how in nearly 70% of instances sponsors only include one option per category.[28] The probability then decays as the number of options included increases, consistent with the presence of some cost that sponsors incur when adding more than one option within the same category. In estimation, I calibrate q to match the observed empirical distribution allowing for heterogeneity at the year-recordkeeper-category level.

4.2 Funds' plan inclusion probabilities

In this section I derive funds' inclusion probabilities implied by sponsors' menu choice problem described in the previous section. To this end, I will analyze sponsors' decision problem backward.

Consider sponsor p who has chosen to offer category g and needs to select n investment funds within g. At this stage, p ranks all the options according to (1) and selects the n options $\{j_1, ..., j_n\}$ such that

$$u_{j_1} > u_{j_2} > ... > u_{j_n}. \tag{4}$$

Fund j will be included in p's plan if and only if u_j is ranked among first nth highest utilities. Letting j_z be the option with the zth highest utility, the probability that j is included in plan p is given by

$$\phi_{jp}^{1:n} \equiv \sum_{z=1}^{n} \phi_{jp}^{z}$$

where ϕ_{jp}^{z} is the probability that j is ranked in the zth position i.e.,

$$\phi_{jp}^{z} \equiv \Pr\{j = j_z\}.$$

Under the assumption that the random utility shocks are distributed as T1EV, an analytical expression for each ϕ_{jp}^{z} can be derived. For $z = 1$, ϕ_{jp}^{1} corresponds to the standard logit choice probability

$$\phi_{jp}^{1} = \frac{\exp(V_j(\boldsymbol{\theta}_p))}{\sum_{k \in g} \exp(V_k(\boldsymbol{\theta}_p))}. \tag{5}$$

For $z = 2$, ϕ_{jp}^{2} is the probability that j provides the 2nd highest utility which equals the sum of probabilities of all utility rankings where u_j is the 2nd largest utility. In a world

[27] As I explain in Appendix B, where I offer a simple microfoundation for the optimal choice of the number of options included within an investment category, incorporating such decision in the full model would considerably complicate its estimation.

[28] For each year-plan-category pair I count the number of funds offered and plot the resulting distribution in Figure A1.

in which there are only 4 options, say $\{j, k, l, m\}$,

$$\begin{aligned}\phi_{jp}^2 = &\Pr\{u_{kp} > u_{jp} > u_{lp} > u_{mp}\} + \Pr\{u_{kp} > u_{jp} > u_{mp} > u_{lp}\} \\ &+ \Pr\{u_{lp} > u_{jp} > u_{kp} > u_{mp}\} + \Pr\{u_{lp} > u_{jp} > u_{mp} > u_{kp}\} \\ &+ \Pr\{u_{mp} > u_{jp} > u_{kp} > u_{lp}\} + \Pr\{u_{mp} > u_{jp} > u_{lp} > u_{kp}\}.\end{aligned}$$

In Appendix B I show that the independence of irrelevant alternatives (IIA) property of the logit model implies ϕ_{jp}^2 can be written as follows for an arbitrary number of options:

$$\phi_{jp}^2 = \sum_{j_1 \neq j} \frac{\exp(V_{j_1 p})}{\sum_k \exp(V_{kp})} \frac{\exp(V_{jp})}{\sum_{k \neq j_1} \exp(V_{kp})}.$$

The above expression can be further generalized that to the case in which we want to compute the probability of j being ranked in an arbitrary position zth

$$\phi_{jp}^z = \sum_{(j_1,\ldots,j_{z-1}) \in g/\{j\}} \prod_{z'=1}^{z-1} \frac{\exp(V_{j_{z'} p})}{\sum_{z''=z'}^{N_{gp}} \exp(V_{j_{z''} p})} \cdot \frac{\exp(V_{jp})}{\sum_{z'''=z}^{N_{gp}} \exp(V_{j_{z'''} p})}$$

where $N_{gp} \subset N_p$ is the set of funds that p considers in category g.

Moving backward in sponsor p's menu choice problem, the probability that n options are chosen from investment category g is given by $q(1-q)^{n-1}$ so that, conditional on g being offered, the probability of j being included in p's plan is just

$$\sum_{n=1}^{\infty} q(1-q)^{n-1} \phi_j^{1:n}.$$

The choice of whether or not to include category g is assumed to depend on sponsors' expected utility from the highest ranked option, which under our T1EV assumption, is given by

$$\mathbb{E}[u_{j_1 p}] = \log\left(\sum_{k \in g} \exp(V_{kp}(\boldsymbol{\theta}_p))\right).$$

The probability that p decides to offer investment category g as part of its retirement plan equals

$$\lambda_{gp} = \frac{\exp(\mathbb{E}[u_{j_1 p}])}{1 + \exp(\mathbb{E}[u_{j_1 p}])}$$

which can be interpreted as the probabilistic outcome of a binary choice logit problem.

Combining all pieces together, it can be shown that the unconditional probability of fund j being included in sponsor p plan can be written as

$$\phi_{jp}(\boldsymbol{\theta}_p) = \lambda_{gp}(\boldsymbol{\theta}_p) \cdot \sum_{n=1}^{\infty} (1-q)^{n-1} \phi_{jp}^n(\boldsymbol{\theta}_p) \tag{6}$$

where I make explicit its dependence on the vector of preference parameters $\boldsymbol{\theta}_p$, $(1-q)^{n-1}$ is the probability that p includes a number of options greater or equal than n and ϕ_j^n is

the probability that j is ranked in the nth position.[29]

Equation (6) extends the expression of the logit choice probabilities to the case in which decision makers can select more than a single option and where the number of option chosen is determined by the parameter q. Expression (6) collapses to the standard discrete choice logit probability if we assume sponsors can only include one fund within each investment category. This corresponds to setting $q = 1$, which implies that

$$\phi_{jp}(\boldsymbol{\theta}_p) = \lambda_{gp}(\boldsymbol{\theta}_p) \cdot \phi^1_{jp}(\boldsymbol{\theta}_p) = \frac{\exp(V_{jp}(\boldsymbol{\theta}_p))}{1 + \sum_{k \in g} \exp(V_{kp}(\boldsymbol{\theta}_p))}.$$

In this more general context, the decision of not including investment category g plays the role of the outside option in standard discrete choice models, with mean indirect utility normalized to 0.

Overall, equation (6) represents sponsor p individual demand for investment funds belonging to investment category g. Assuming sponsors preferences $\boldsymbol{\theta}_p$ follow some distribution F_θ, we can derive fund j's aggregate demand as

$$\phi_j = \int \phi_{jp}(\boldsymbol{\theta}_p) dF_\theta(\boldsymbol{\theta}_p). \tag{7}$$

The data counterpart to (7) corresponds to the share of plans that include fund j as part of their retirement menu. In Section 6 I estimate the distribution of sponsors preference parameters by matching (7) to these observed inclusion probabilities.

4.3 Investors' retirement portfolio problem

Consider investor i who allocates its dollar contribution A across the investment funds available in plan p, indexed by $j \in \{1, ..., J_p\}$, and a cash account $j = 0$. In practice, not all plan investors make an active investment decision and many of them are automatically defaulted into one of the options available which often corresponds to a TDF fund or a balanced fund. I denote this default option $j = d$ and assume that with probability δ investor i does not make an active investment decision. In this case, i's contribution will be allocated entirely to fund d. Conversely, with probability $(1 - \delta)$ investor i makes an active investment decision and allocates A across all options available including d.

Conditional on making an active investment decision, investor i forms its retirement portfolio by choosing the vector of portfolio weights $\boldsymbol{a} \equiv (a_1, ..., a_{J_p})$ to maximize the following linear-quadratic utility:[30]

$$U_p(\boldsymbol{a}) \equiv \underbrace{\sum_{j=1}^{J_p} a_j(\boldsymbol{w}_j'\boldsymbol{\beta} - f_j + \xi_j)}_{\text{linear utility component}} - \underbrace{\frac{\gamma}{2}\sum_{j=1}^{J_p} a_j^2}_{\text{naive diversification}} - \underbrace{\frac{\gamma}{2}\sum_{j,k} g_{jk}a_j a_k}_{\text{funds substitutability}}. \tag{8}$$

The preferences defined in (8) capture the idea that investors value funds along three margins. The first is a 'perfect substitute' margin that pushes them to allocate their entire contribution to the fund with the highest linear utility component. This component depends on observed funds' characteristics $\boldsymbol{w}_j$, fees f_j, and on unobserved characteristics

[29] I provide details on the derivation in Appendix B

[30] In equation (8) I have already substituted for the portfolio share on the cash account a_0 using the constraint that $a_0 + \boldsymbol{a}'\mathbf{1} = 1$ and assuming that returns on the cash account are normalized to 0.

ξ_j. If this were the only margin, investors' portfolio problem would collapse to a standard discrete choice problem with investors' indirect utility for fund j given by $\boldsymbol{w}_j'\beta - f_j + \xi_j$.

In the context of investment choices, standard portfolio theory predicts that investors should diversify across assets to reduce risk (Markowitz (1952)). The second margin captures this incentive. Specifically, it captures investors' incentives to naively diversify across the available options which, in the context of retirement investing, has been shown to be a key determinant of individual portfolio allocations (Benartzi and Thaler (2001), Huberman and Jiang (2006)).

The third margin captures how investors perceive substitutability between funds. I assume that investors perceive fund j and fund k as substitutable if they belong to the same investment category, in which case the term $g_{jk} = 1$. If funds belong to different investment categories $g_{jk} = 0$. This implies that investors will have an incentive to diversify more across funds from different categories rather than across funds within the same category. As I explain in more detail later on, this assumption is motivated by the empirical evidence that investment categories fixed effects explain a substantial fraction of the observed plan-level portfolio allocation, suggesting that investors' allocate across styles rather than among individual funds (Barberis and Shleifer (2003)). Conversely, funds' loadings on standard risk factors (Fama and French (1992), Carhart (1997)) have negligible explanatory power, especially after controlling for investment categories fixed effects.

4.4 Plan asset demand system

Letting G_p be the $J_p \times J_p$ matrix whose (j,k) element is g_{jk}, under the previous assumptions, investor i's optimal portfolio allocation is given by

$$\boldsymbol{a}_i(\boldsymbol{f}_p) = \begin{cases} \boldsymbol{e}_d & \text{if } i \text{ inactive} \\ \frac{1}{\gamma}(I + G_p)^{-1}(W_p\boldsymbol{\beta} + \boldsymbol{\xi}_p - \boldsymbol{f}_p) & \text{if } i \text{ active} \end{cases} \tag{9}$$

where $\boldsymbol{e}_d$ is a unit vector that takes value of 1 in its d element corresponding to the default option and W_p is the matrix of observed characteristics for the funds available in plan p.

In the data I do not observe asset allocations at the individual level. Therefore, I need to aggregate individual investors demands to obtain the plan level demand system. Letting $\boldsymbol{s}_p$ be the J_p vector of plan p portfolio shares, A_p plan p total wealth and defining $\boldsymbol{\eta} \equiv (\boldsymbol{\beta}, \gamma, \delta)$ as the vector of demand parameters, we can sum demands across all investors in plan p to obtain:

$$\boldsymbol{s}_p(\boldsymbol{f}; \boldsymbol{\eta}) \equiv \sum_{i \in I_p} \frac{A}{A_p} \boldsymbol{a}_i(\boldsymbol{f}) = \delta \boldsymbol{e}_d + \frac{1-\delta}{\gamma}(I + G_p)^{-1}(W_p\boldsymbol{\beta} + \boldsymbol{\xi}_p - \boldsymbol{f}_p). \tag{10}$$

Empirically, I can use variation in the observed plan level allocations to estimate the demand system in (10). To this end, it is useful to multiply both sides of (10) by $(I + G_p)$ to obtain a demand system where only own fees and own demand shocks enter each equation:

$$\tilde{s}_p(\boldsymbol{f}; \boldsymbol{\eta}) = \delta \tilde{\boldsymbol{e}}_d + W_p\tilde{\boldsymbol{\beta}} - \tilde{\gamma}\boldsymbol{f}_p + \tilde{\boldsymbol{\xi}}_p. \tag{11}$$

where $\tilde{\boldsymbol{s}}_p \equiv (I + G_p)\boldsymbol{s}_p$, $\tilde{\boldsymbol{e}}_d \equiv (I + G_p)\boldsymbol{e}_d$, $\tilde{\boldsymbol{\beta}} \equiv \boldsymbol{\beta}(1-\delta)/\gamma$ and $\tilde{\boldsymbol{\xi}} \equiv \boldsymbol{\xi}(1-\delta)/\gamma$. Because

$\tilde{\boldsymbol{s}}_p$ is observed, I can estimate (11) via linear methods.

Before turning to the supply side, a couple of remarks are in order. First, so far I have assumed that investors' preferences are homogeneous. In this way, the parameters of the plan-level demand system are the same as the ones for the individual demand system. From an empirical perspective, the nature of the data and the linear structure of the demand system prevent me from learning about unobserved heterogeneity in preference parameters.[31] Nevertheless, in Appendix B, I show how to interpret the parameters in (10) as weighted averages of the heterogeneous individual parameters. Second, I am assuming that the default option is the same for each individual. This is done only for expositional purposes and in the estimation and in Appendix B I allow for the presence of multiple default funds.

Lastly, in setting the supply side profit maximization problem, I will be working with a version of equation (10) that I rearrange slightly as follows:

$$\boldsymbol{s}_p(\boldsymbol{f};\boldsymbol{\eta}) = \delta \boldsymbol{e}_d + \frac{(1-\delta)}{\gamma}(I - \mathcal{K}_p)(\boldsymbol{\mu}_p - \boldsymbol{f}) \tag{12}$$

where $\boldsymbol{\mu}_p \equiv W_p\boldsymbol{\beta} + \boldsymbol{\xi}_p$ and the matrix $\mathcal{K}_p$ is defined as

$$\mathcal{K}_p \equiv \tilde{G}_p(I + \tilde{G}_p'\tilde{G}_p)^{-1}\tilde{G}_p',$$

with $\tilde{G}_p$ a matrix with J_p rows, one for each fund in plan p, and a number of columns equal to the number of investment categories. The jth row of $\tilde{G}$, denoted $\tilde{\boldsymbol{g}}_j$, equals 1 in correspondence of fund j's category and 0 otherwise. The matrix G_p that appears in (10) is the outer product of $\tilde{G}_p$, e.g., $G_p = \tilde{G}_p\tilde{G}_p'$ with $g_{jk} = \boldsymbol{g}_j'\boldsymbol{g}_k$.

Rewriting aggregate asset demand as in (12) helps in visualizing own and cross substitution patterns across different assets. In particular, the fee elasticity between asset j and asset l is proportional to

$$\frac{\partial s_j}{\partial f_l} \propto \begin{cases} -(1-\kappa_{jj}) = -\left(1 - \tilde{\boldsymbol{g}}_j'(I + \tilde{G}_p\tilde{G}_p')^{-1}\tilde{\boldsymbol{g}}_j\right) & \text{if } j = l \\ \kappa_{jk} = \tilde{\boldsymbol{g}}_j'(I + \tilde{G}_p\tilde{G}_p')^{-1}\tilde{\boldsymbol{g}}_l & \text{if } j \neq l \end{cases} \tag{13}$$

which is always between $(-1, 1)$ if $j \neq k$ and between $(-1, 0)$ if $j = l$. If one defines $\tilde{G}$ as the matrix of funds' factor loadings, $\frac{\partial s_j}{\partial f_l}$ measures how close/correlated asset j and l are in terms of their risk exposures $\tilde{\boldsymbol{g}}_j$ and $\tilde{\boldsymbol{g}}_l$ respectively.[32] When $\tilde{G}$ is the matrix of funds categories fixed effects the same interpretation applies but the substitution patterns are by construction sparse.

5 Supply

I model supply as a differentiated Bertrand oligopoly where investment advisors set fees simultaneously before sponsors make their plan menu decisions and plan investors form

[31] Recall that I do not observe portfolio allocations at the individual level. One way to introduce heterogeneity would be adding interaction terms between funds' characteristics, such as fees, and observable plan characteristics. Egan, MacKay and Yang (2023) take this approach to uncover heterogeneity in investors' risk aversion.

[32] I provide more details in Appendix *B*

their portfolios. I assume that the same fund charges the same fee across different plans because, although funds can price discriminate by offering different share classes, almost all the observed variation in fees is across funds and not across share classes within the same fund.[33]

Let P be the number of plan sponsors, A_p the dollar wealth of plan p and $\mathcal{S}_{jp} \subseteq 2^{N_p}$ be the set of all possible menus where p includes fund j. Fund j chooses its fee f_j to maximize the following expected dollar profit

$$\max_{f_j} \quad P \cdot (f_j - c_j) \cdot \int_p A_p \left(\sum_{S \in \mathcal{S}_{jp}} \phi_p(S, \boldsymbol{f}; \boldsymbol{\theta}_p) s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p | S) \right) dF(A_p, \boldsymbol{\theta}_p, \boldsymbol{\eta}_p) \tag{14}$$

where $\phi_p(S, \boldsymbol{f}; \boldsymbol{\theta}_p)$ is the probability that sponsor p chooses plan menu S and, conditional on menu S, $s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p | S)$ is fund j's portfolio weight within plan p

$$s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p | S) = \begin{cases} \frac{1-\delta_p}{\gamma_p} \left((1 - \kappa^S_{jj})(\mu_{jp} - f_j) - \sum_{l \neq j, l \in S} \kappa^S_{jl} (\mu_{lp} - f_l) \right) & \text{if } j \neq d \\ \delta_p + \frac{1-\delta_p}{\gamma_p} \left((1 - \kappa^S_{dd})(\mu_{jp} - f_j) - \sum_{l \neq j, l \in S} \kappa^S_{dl} (\mu_{lp} - f_l) \right) & \text{if } j = d \end{cases}$$

where I made explicit the dependence of the elements of $\mathcal{K}$ on the realized menu S and allow investor parameters to depend on p.

Problem (14) is particularly complex to solve because it requires investment providers to internalize how a marginal increase in fees affects (i) investors portfolio allocation s_j conditional on a given menu S, (ii) the probability that plan menu S is chosen and (iii) trade-off these changes across all possible plan menus S and plan sponsors p. From a computational perspective, given the large number of funds available in the market, the number of possible plan menus in $\mathcal{S}_{jp}$ would be too large to make the computation of (14) and its derivatives feasible.[34] To overcome these difficulties, I will simplify funds' pricing problem in a way that allows me make the problem computationally tractable while at the same time preserving the two dimensions along which funds' compete in the retirement market, namely, competition for plan inclusion and competition for plan asset allocations.

To simplify problem (14) I will assume that funds only consider the effect of a marginal change in fees on their aggregate inclusion probability and not on the probability of any of the possible menus being selected by the sponsor. Equivalently, I assume funds do not take into account that changing fees influences the probability that a particular menu is chosen but only consider how it affects their total likelihood of being included. With this

[33] Additionally, in the context I am considering, 401(k) sponsors are almost always treated as institutional investors and many investment providers offer a specific share class for the retirement plans market.

[34] Goeree (2008) faces a similar problem when estimating a discrete choice demand model with imperfect consideration. She overcomes the computational burden by simulating consumers consideration sets. My case is more complex because investors' consideration sets, or equivalently plan menus, are the outcome of sponsors' menu choice problem. Moreover, consideration probabilities in my case are not independent for products that belong to the same investment category. Lastly, in my case prices affect consideration probabilities which means that derivatives of consideration probabilities will enter firms' first order conditions. This will be true for both own consideration probabilities but also competitors consideration probabilities.

assumption, fund j fee setting problem can be rewritten more compactly as

$$\max_{f_j} P \cdot (f_j - c_j) \cdot \int A_p \phi_{jp}(\boldsymbol{f}; \boldsymbol{\theta}_p) s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p) dF(A_p, \boldsymbol{\theta}_p, \boldsymbol{\eta}_p) \tag{15}$$

where $\phi_{jp}(\boldsymbol{f}; \theta_p)$ is the probability that p includes j as defined in (6) and $s_j(\boldsymbol{f}; \boldsymbol{\eta}_p)$ is the expected portfolio allocation of fund j within plan p and is given by

$$s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p) = \begin{cases} \frac{1-\delta_p}{\gamma_p} \left((1 - \bar{\kappa}^p_{jj})(\mu_{jp} - f_j) - \sum_{l \neq j} \bar{\kappa}^p_{jl}(\mu_{lp} - f_l) \right) & \text{if } j \neq d \\ \delta_p + \frac{1-\delta_p}{\gamma_p} \left((1 - \bar{\kappa}^p_{dd})(\mu_{jp} - f_j) - \sum_{l \neq j} \bar{\kappa}^p_{dl}(\mu_{lp} - f_l) \right) & \text{if } j = d \end{cases} \tag{16}$$

with,

$$\bar{\kappa}^p_{jl} \equiv \begin{cases} \sum_{S \in \mathcal{S}_{jp} \cap \mathcal{S}_{lp}} \frac{\phi_p(S, \boldsymbol{f}; \theta_p)}{\phi_{jp}(\boldsymbol{f}; \theta_p)} \kappa^S_{jl} = \phi_{lp}(\boldsymbol{f}, \theta_p) \cdot \mathbb{E}[\kappa^S_{jl} | j, l \in S] & \text{if } j \neq l \\ \sum_{S \in \mathcal{S}_{jp}} \frac{\phi_p(S, \boldsymbol{f}; \theta_p)}{\phi_{jp}(\boldsymbol{f}; \theta_p)} \kappa^S_{jj} = \mathbb{E}[\kappa^S_{jj} | j \in S] & \text{if } j = l \end{cases} \tag{17}$$

Problem (15) can be obtained from (14) after dividing and multiplying the term in the round brackets by ϕ_{jp} and exploiting the linear structure of s_{jp} to rewrite the expectation over S more compactly. The resulting term collapses to $s_{jp}(\boldsymbol{f}; \boldsymbol{\eta}_p)$ as defined in (16), which represents the asset demand from plan p investors that fund j expects before sponsor p chooses its plan menu. Asset characteristics affect this expected demand through the matrix $\bar{\mathcal{K}}^p$, whose (j, l) element, defined in (17), captures how much competitive pressure j expects from competitor l. The latter depends on how likely is fund l to be included in plan p and, conditional on that, on how close substitute asset j and asset l are.[35]

My restriction on funds' fee-setting behavior assumes that funds do not internalize how fees affect the elements $\bar{\kappa}_{jl}$. A constructive way to impose this restriction could be assuming that funds believe sponsors will include at most one fund per investment category, thereby assigning positive probability only to plan menus containing funds from different categories. Formally, this would require funds to have a biased belief $\hat{q} = 1$ about the parameter q governing the distribution of the number of funds included within each category. In practice, we know that most plans do not include more than one option per category (Figure A1), making this assumption perhaps not so unreasonable. Under this assumption, I show in Appendix B that $\bar{\kappa}_{jl}$ would not depend on f_j because plan inclusion decision are assumed to be independent across investment categories. Moreover, this assumption is one of the sufficient conditions that allows me to prove existence of a Bertrand-Nash equilibrium.[36]

5.1 Nash equilibrium fees

In this section I derive the Nash equilibrium fees implied by problem (15) assuming that funds take $\bar{\mathcal{K}}$ as given. Denoting by $s_j(\boldsymbol{f})$ fund j expected dollar asset allocation, the

[35] This measure of substitutatibility is given by the second term in (17) $\mathbb{E}[\kappa^S_{jl} | j, l \in S; \boldsymbol{f}]$. This is an expectation because κ^S_{jl} depends on the whole menu S and not only on fund j and fund l characteristics. Formally, this can be seen from the definition of κ_{jl} in (13) where κ_{jl} depends on the characteristics of all competitors through the weighting matrix $(I + \tilde{G}'_p \tilde{G}_p)^{-1}$.

[36] In Appendix B I show that when $\hat{q} = 1$, sponsors preferences are homogeneous and a particular dominance diagonal condition on the jacobian of the demand system is satisfied there exists a Bertrand-Nash equilibrium.

first order conditions with respect to f_j is given by the usual Bertrand pricing equation

$$s_j(\boldsymbol{f}) + (f_j - c_j) \cdot \frac{\partial s_j(\boldsymbol{f})}{\partial f_j} = 0. \tag{18}$$

The difference between the current setting and standard oligopolistic problems is that funds are competing along two dimensions, namely, they compete for being included in a plan and, conditional on plan inclusion, they compete for plan investors' allocations. These two layers of competition are enclosed in $\partial s_j(\boldsymbol{f})/\partial f_j$ which is made of the following two terms

$$\begin{aligned}\frac{\partial s_j(\boldsymbol{f})}{\partial f_j} = & \underbrace{\int A_p \frac{\partial \phi_{jp}}{\partial f_j}(\boldsymbol{f};\boldsymbol{\theta}_p) s_{jp}(\boldsymbol{f};\boldsymbol{\eta}_p) dF(A_p,\boldsymbol{\theta}_p,\boldsymbol{\eta}_p)}_{\text{demand loss from marginal sponsor}} \\ & \underbrace{- \int A_p (1-\delta_p)\gamma_p^{-1}\phi_{jp}(\boldsymbol{f};\boldsymbol{\theta}_p)(1-\bar{\kappa}^p_{jj}) dF(A_p,\boldsymbol{\theta}_p,\boldsymbol{\eta}_p)}_{\text{demand loss from marginal investor}} < 0. \end{aligned} \tag{19}$$

Expression (19) captures the classic expected revenue loss from marginal consumers not willing to buy at higher price. In this context, the reduction in demand comes from two forces (i) the marginal sponsor not willing to include j in its plan and (ii) the marginal plan investors reducing its allocation to fund j. Equation (18) then tells us that, for given competitors fees, fund j will choose f_j that equalizes the profit reduction from losing the marginal sponsors and investors to the profit gain from charging the inframarginal ones an higher fee.

The linear structure of investors asset demand allows me to go beyond this standard intuition and to offer a novel characterization of the Nash equilibrium fees that sheds light on the forces driving price competition in this market. To this end, I will define the following variables,

$$\begin{aligned} &\bar{\phi}_j \equiv \int A_p(1-\delta_p)\gamma_p^{-1}\phi_{jp} dF_p; \qquad \tilde{\kappa}_{jl} \equiv \bar{\phi}_j^{-1} \int A_p(1-\delta_p)\gamma_p^{-1}\phi_{jp}\bar{\kappa}^p_{jl} dF_p \\ &\bar{\mu}_j \equiv \bar{\phi}_j^{-1} \int A_p(1-\delta_p)\gamma_p^{-1}\phi_{jp}[I-\bar{\mathcal{K}}^p]'_j \boldsymbol{\mu}_p dF_p; \qquad \tilde{\boldsymbol{\mu}} \equiv (I-\tilde{\mathcal{K}})^{-1}\bar{\boldsymbol{\mu}} \\ &\bar{\delta} \equiv \bar{\phi}_j^{-1} \int A_p \phi_{jp}\delta_p dF_p; \qquad \iota_j \equiv -\bar{\phi}_j^{-1} \int A_p \frac{\partial \phi_{jp}}{\partial f_j} s_{jp}(f_j - c_j) dF_p \end{aligned}$$

where I suppressed all functions' arguments and $[I-\bar{\mathcal{K}}^p]_j$ denotes the jth row of the matrix $I-\bar{\mathcal{K}}^p$. Next, I rewrite (18) in vector form for all funds in terms of these variables to obtain

$$\bar{\delta}\boldsymbol{e}_d + (I-\tilde{\mathcal{K}})(\tilde{\boldsymbol{\mu}} - \boldsymbol{f}) - \boldsymbol{\iota} - (I - diag(\tilde{\mathcal{K}}))(\boldsymbol{f} - \boldsymbol{c}) = 0$$

By rearranging this system of Bertrand FOCs, I reach one of the paper's key findings, which decomposes equilibrium fees into three components:

$$\boldsymbol{f}^* = \underbrace{\frac{\tilde{\boldsymbol{\mu}} + \boldsymbol{c}}{2}}_{\text{monopolist fee}} - \underbrace{\boldsymbol{h}\left(\tilde{\mathcal{K}}, \frac{\tilde{\boldsymbol{\mu}} - \boldsymbol{c}}{2}\right)}_{\text{Hotelling markdown}} - \underbrace{\left(I - \frac{diag(\tilde{\mathcal{K}})}{2} - \frac{\tilde{\mathcal{K}}}{2}\right)^{-1} \frac{\boldsymbol{\iota}}{2}}_{\text{plan inclusion markdown}} \tag{20}$$

where for simplicity I assumed that there is no default fund.[37]

Equation (20) decomposes the fees charged by funds in an interior Nash equilibrium into three terms. The first term represents the vector of fees funds would charge as monopolist.[38] From these fees there are two types of markdowns that need to be subtracted to account for the two dimension of competition driving pricing incentives in this market. The plan inclusion markdown in (20) captures the optimal reduction in fees required to increase the probability of being included in a plan. If funds knew that they will be included with certainty i.e., $\phi_{jp} \equiv 1$, the plan inclusion markdown would disappear because $\iota_j = 0$. The Hotelling markdown $\boldsymbol{h}(\cdot)$ instead captures the optimal reduction in fees required to compete against similar funds. The simple Hotelling (1929) model predicts that when two firms located on a line are closer to each other, they will charge lower margins. The same intuition carries over in this more general setting. Funds that are closer to their competitors have an higher $\boldsymbol{h}$ an must lower their fee.

The natural question at this point is, what does being closer to competitors mean in this context and how does that relate to $\boldsymbol{h}$? Loosely speaking a fund is closer to its competitors when its characteristics are less differentiated from competitors' characteristics. In practice, this measure of proximity/differentiation is embedded in the asset demand cross-substitution patterns through the matrix $\mathcal{K}$ defined in (13). Fund j and fund l are closer substitutes if their characteristics are closer as measured by κ_{jl}. The proximity of each fund to all other competitors is summarized by the vector $\boldsymbol{h}(\cdot)$ which is defined as

$$\boldsymbol{h}\left(\tilde{\mathcal{K}}, \frac{\tilde{\boldsymbol{\mu}} - \boldsymbol{c}}{2}\right) \equiv \left(I - \frac{(\tilde{\mathcal{K}} - \kappa_0 I)}{2(1-\kappa_0)}\right)^{-1} \frac{(\tilde{\mathcal{K}} - \kappa_0 I)}{2(1-\kappa_0)} \frac{\tilde{\boldsymbol{\mu}} - \tilde{\boldsymbol{c}}}{2} \tag{21}$$

where for expositional purposes I assumed that $\tilde{\kappa}_{jj} \equiv \kappa_0$. Expression (21) is a measure of proximity because it is equivalent to the Bonacich network centrality measure studied in Bonacich (1987).[39] This network centrality measure appears in (20) because the Bertrand fee-setting game belongs to a broader class of network games first studied in Ballester et al. (2006).[40] The main insight from this literature is that Nash equilibrium actions will generally depend on a player's network centrality. In this case, funds' equilibrium fees depend on how central a fund is in the competitive network. A more central fund faces more similar competitors and charges lower margins.

So far, I have assumed that a Nash equilibrium exists. Given the complexity of funds' pricing problem deriving general results about existence and uniqueness is not an easy task. Nonetheless, in Appendix B I provide sufficient conditions for existence and uniqueness of an interior Bertrand-Nash equilibrium for some particular cases. Specifically, I start by showing that when funds know with certainty which plan menus will include them (e.g., $\phi_{jp} = 1$ if p includes j) then the following dominance diagonal condition

$$(1 - \tilde{\kappa}_{jj})(\tilde{\mu}_j - c_j) > \sum_{l \neq j} |\tilde{\kappa}_{jl}|(\tilde{\mu}_j - c_l) \quad \text{all } j,$$

ensures that the Bertrand-Nash equilibrium is interior, with $f_j^* \in (c_j, \tilde{\mu}_j)$ for all j, and

[37] All derivations are presented in Appendix B

[38] The price charged by a monopolist facing a linear demand $q(p) = a - p$ with marginal cost c, is $(a + c)/2$.

[39] For any zero-diagonal adjacency matrix A, positive scalar $\delta > 0$ and non-zero vector $\boldsymbol{u}$, the vector of Bonacich centralities is defined as $\boldsymbol{b}(A, \boldsymbol{u}) \equiv (I - \delta A)^{-1} \delta A \boldsymbol{u}$.

[40] More details are provided in Appendix F.

unique. Moreover, I am also able to show that when sponsors' preference are homogeneous and funds believe that sponsors will include at most one fund per category (e.g., $\hat{q} = 1$) there exists a Nash-equilibrium even when funds' do not know with certainty which plan menus will include them.

6 Demand Identification and Estimation

In this section I describe how to identify and estimate the model. I estimate sponsors' preferences from variation in the observed plan inclusion probabilities and investors' preferences from variation in the observed plan-level portfolio allocations. After that, I turn to the supply side and recover funds' marginal costs and markups using the demand estimates together with the Nash-Bertrand equilibrium conditions.

6.1 Identification and estimation of sponsors' preferences

The menu choice model developed in Section 4 provides us with an analytical expression for the expected probability that fund j is included in a retirement plan for a given distribution of sponsors' preference parameters $\boldsymbol{\theta}_p \sim F(\boldsymbol{\theta}_p; \bar{\boldsymbol{\theta}})$ which I assume to be parametrized by the vector $\bar{\boldsymbol{\theta}}$:

$$\phi_j(\bar{\boldsymbol{\theta}}) = \int \lambda_{gp}(\boldsymbol{\theta}_p) \cdot \sum_{n=1}^{\infty} (1-q)^{n-1} \phi_{jp}^n(\boldsymbol{\theta}_p) dF(\boldsymbol{\theta}_p; \bar{\boldsymbol{\theta}}), \tag{22}$$

where $\lambda_{gp}(\boldsymbol{\theta}_p)$ is the probability that plan p offers category g and $\phi_{jp}^n(\boldsymbol{\theta}_p)$ is the probability that j is the fund with the nth highest utility.

The data counterpart to equation (22) is the share of retirement plans that include fund j. The estimation strategy is then to find the vector of parameters $\bar{\boldsymbol{\theta}}$ that makes the model implied inclusion probabilities in (22) as close as possible to the observed ones. As $q \to 1$, expression (22) collapses to the standard random-coefficient logit formula for product market shares

$$\phi_j(\bar{\boldsymbol{\theta}}) = \int \frac{\exp(V_{jp}(\boldsymbol{\theta}_p))}{1 + \sum_{k \in g} \exp(V_{kp}(\boldsymbol{\theta}_p))} dF(\boldsymbol{\theta}_p; \bar{\boldsymbol{\theta}}),$$

considered in the workhorse demand models of Berry (1994) and Berry, Levinsohn and Pakes (1995).

In estimation, I compute the observed inclusion probability for each fund at the year-recordkeeper-category level, where I define an investment category as the interaction between the standard investment categories and an indicator for passive funds. In this way, an index fund and active fund from say Large-Cap-Growth would be classified into two different categories Large-Cap-Growth-Active and Large-Cap-Growth-Passive. I refer to a particular year-recordkeeper-category combination as a market, indexed by t, and denote the share of retirement plans in market t that offer fund j as $\hat{\phi}_{jt}$.

I allow for the possibility that funds' fees are correlated with sponsors' demand shocks ζ_{jt}. These shocks enter sponsors' mean utilities $V_{jpt}(\boldsymbol{\theta}_p) = \boldsymbol{w}'_{jpt}\boldsymbol{\theta}_p + \zeta_{jt}$ and are observed by market participants, including investment funds, but unobserved to the econometrician. If funds set fees after observing ζ_{jt} or have better information about these demand shocks, demand and supply simultaneity would bias preference parameters estimates. I account

for this type of price endogeneity in two ways. First, I exploit the granularity of the data to absorb unobserved heterogeneity in demand along three dimensions: (i) time by including year fixed effects, (ii) product quality by including funds' brand fixed effects, and (iii) financial characteristics by including investment category fixed effects. Second, I instrument funds' fees with funds' turnover ratios which capture trading-related costs that funds incur when selling and buying securities. The instrument is relevant as long as profit maximizing funds optimally pass these costs to investors through higher fees (Pástor, Stambaugh and Taylor (2020)). The identifying assumption is that the variation in funds' turnover ratio not explained by time, brand, investment category and passive fixed effects enters sponsors' demand only through fees.

A large literature in finance has studied the relationship between funds' turnover and funds' investment performance. The results are mixed: Carhart (1997) finds a negative cross-sectional relationship, Wermers (2000) and Kacperczyk, Sialm and Zheng (2005) find no relationship, and Pástor, Stambaugh and Taylor (2017) find a positive time-series relationship. Regardless of the sign of such relationship, if investors chase performance (Chevalier and Ellison (1997)) and my set of fixed effects does not control for that appropriately, the exclusion restriction might not hold because turnover may correlate with unobserved demand shocks. In Appendix C I provide supporting evidence for the validity of the instrument (e.g., the residual turnover) by showing that it does not correlate with standard measure of current and future investment performance for the funds in my sample.

To estimate sponsors preferences I use a nested-fixed point algorithm similar to the one developed in Berry et al. (1995).[41] To start with, I assume that the distribution of sponsors preference parameters is normal with mean $\boldsymbol{\mu}_{\bar{\theta}}$ and variance $\Sigma_{\bar{\theta}}$ (e.g., $\bar{\boldsymbol{\theta}} = (\boldsymbol{\mu}_{\bar{\theta}}, \Sigma_{\bar{\theta}})$) and write sponsor p mean utility as

$$V_{jt}(\bar{\boldsymbol{\theta}}) = \bar{v}_{jt} + \boldsymbol{w}'_{jt}\Gamma_{\bar{\theta}}\boldsymbol{\nu}_p$$

where $\bar{v}_{jt} \equiv \boldsymbol{w}'_{jt}\boldsymbol{\mu}_{\bar{\theta}} + \zeta_{jt}$ is the homogeneous component of preferences, $\boldsymbol{\nu}_p \sim N(\mathbf{0}, I)$ are random tastes for funds' characteristics and $\Gamma_\theta\Gamma'_\theta = \Sigma_\theta$. The estimation algorithm starts with a guess of $\bar{\boldsymbol{\theta}}$ and then for each market t finds the vector $\bar{\boldsymbol{v}}_t(\bar{\boldsymbol{\theta}})$ that matches observed and model-implied inclusion probabilities:

$$\hat{\boldsymbol{\phi}}_t = \boldsymbol{\phi}_t(\bar{\boldsymbol{v}}_t(\bar{\boldsymbol{\theta}})).$$

After that, the demand residuals $\boldsymbol{\zeta}_t(\bar{\boldsymbol{\theta}}) = \bar{\boldsymbol{v}}_t(\bar{\boldsymbol{\theta}}) - W_t\boldsymbol{\mu}_{\bar{\theta}}$ are computed for each market. The last step exploits the orthogonality condition between ζ_{jt} and an appropriate vector of instruments $\boldsymbol{Z}_{jt}$, $\mathbb{E}[\zeta_{jt}|\boldsymbol{Z}_{jt}] = 0$ to form the GMM norm

$$\boldsymbol{\zeta}(\bar{\boldsymbol{\theta}})'\boldsymbol{Z}\Omega(\bar{\boldsymbol{\theta}})\boldsymbol{Z}\boldsymbol{\zeta}(\bar{\boldsymbol{\theta}}). \tag{23}$$

The algorithm keeps searching over $\bar{\boldsymbol{\theta}}$ until (23) is minimized.

6.2 Estimates of sponsors' preferences

Table 2 presents the estimates of sponsors' preference parameters. The first column reports the estimates of the means of the preference distribution and the second column

[41] See Appendix B for more details.

reports the corresponding standard deviations. These estimates minimize the GMM objective (23) following the estimation algorithm I discussed in the previous section.

In the estimation of sponsors' preferences I include five characteristics, two of them are continuous and three of them are binary. The two continuous characteristics are funds' expense ratios, measured in basis points (bp.), and funds' returns in the previous year gross of expenses, measured in percentage points (pp.). Because I absorb investment category fixed effects, returns are relative to the average return of funds within the same category.

The three binary characteristics are indicators for whether a fund is affiliated with the sponsor's recordkeeper, for whether a fund is a target date and their interaction. Including an indicator for affiliated funds allows me to accommodate for the presence of agency frictions whereby sponsors favor the inclusion of funds belonging to their recordkeeper product line. The inclusion of an indicator for TDFs instead captures the possibility that sponsors have a preference for funds that rebalance plan investors allocation automatically as they age. These type of funds have been created specifically for retirement investing and, after the Pension Protection Act of 2006, qualify as default option for plan participants who do not make an active investment decision. Since then, TDFs' market share in the retirement market has been growing substantially and it is reasonable to think that sponsors may have a preference for such funds even if just to comply with current regulations and reduce liability risk.

The parameter estimates reported in the first column of Table 2 suggest that sponsors value more whether a fund is affiliated, and especially if it is an affiliated TDF, rather than how cost-efficient such fund is or how it performed relative to its investment category. The preference coefficient for funds' affiliation is large and significant. The preference coefficient on funds' expense ratios is negative and significant but its magnitude is small if compared with the coefficient on funds' affiliation; on average sponsors are willing to pay 44 bp (=0.88/0.02) more in fees for an affiliated fund. On the other hand, plan sponsors do not seem to value funds' returns gross of fees, as the estimated coefficient is close to zero. I also allow for heterogeneity around the mean of sponsors' sensitivity to fees and report the estimated standard deviation in the second column of Table 2. Although modest in magnitude, the estimated heterogeneity is significant at conventional significance levels.

To get a better understanding of the estimated magnitudes, I report the median marginal effect of each characteristic on the inclusion probabilities in the third column of Table 2. The marginal effect is the unit change in the expected probability of being included in a plan for a unit increase in the corresponding characteristic. For instance, the first number in the third column tells us that, on average, a ten basis points increase in expenses reduces the plan inclusion probability by almost 0.1 percentage points. On the other hand, the marginal effect of being an affiliated fund is almost four times larger. These magnitudes are far from being negligible given that the median inclusion probability in the estimation sample is roughly 0.52%. The marginal effect of a one percent increase in funds' gross returns is instead substantially smaller, which is not surprising given how small its corresponding preference coefficient is.

The bottom part of Table 2 presents some additional information including information about sponsors' elasticity to fees. With the model estimates, I compute the elasticity of inclusion probabilities to fees for each fund-market combination and report the median of this distribution in the bottom part of Table 2. The latter is around around 1.77,

Employers preference parameters			
	Mean	S.D.	Marginal Effect (pp.)
Expense Ratio (bp.)	-0.020	0.002	-0.008
	(0.002)	(0.000)	
Affiliated (dummy)	0.879	-	0.363
	(0.046)		
Target (dummy)	-0.371	-	-0.123
	(0.088)		
Target × Affiliated	0.194	-	0.081
	(0.096)		
Gross returns (pp.)	0.004	-	0.002
	(0.001)		
Median fee elasticity		-1.77	
q (Calibrated)		0.70	
GMM objective (df)		6.74 (1)	

Table 2: Two-step GMM estimates of plan sponsor preferences. Robust standard errors are reported in parentheses. Year, category, passive and fund brand fixed effects are included. For the marginal effects, inclusion probabilities are in percentage points.

suggesting that sponsors' demand is not too elastic to funds' fees.

As mentioned before, I allow for the possibility that funds' fees are endogenous but treat other characteristics as predetermined and independent of demand shocks.[42] In estimation I instrument for fees using a third order polynomial of funds' turnover ratios. Overall I have seven moments, four included characteristics and three instruments function of funds' turnover, to estimate six model parameters, five means and one standard deviation. The resulting GMM objective is of 6.74 and rejects the overidentifying restriction at the 1% level. This may be due to the fact that there does not seem to be too much heterogeneity in sponsors' preferences even though the model allows for it. The estimates for the homogeneous model are indeed similar (Table A5) and come with a GMM objective of 4.57 which, although not perfect, it is not rejected at conventional significance levels.

The nature of the data allows me to assess the heterogeneity of the estimates more directly. I do so along several dimensions. First, I split the sample based on plan size as measured by the number of plan participants and find that smaller plans are less responsive to fees and tend to have a stronger preference for affiliated funds than large plans (Table A6). This is consistent with smaller sponsors having less bargaining power and being less willing to pay or search for cheaper investment options (Bhattacharya and Illannes (2022)).[43] Second, I split the sample before and after 2014 and find that sponsors have become more elastic to fees over time (Table A7). This is consistent with sponsors, as well as plan investors, becoming more attentive to fees in response to regulatory interventions mandating the disclosure of funds' fees and performance (Kronlund, Pool,

[42]This assumption is often used in the empirical industrial organization literature where product characteristics are assumed to be determined before demand shocks are realized.

[43]The fact that smaller sponsors are more likely to include expensive funds could also reflect the presence of fixed costs in plan provision. For instance smaller sponsors might not have the asset base to access the cheapest share classes of a fund. The model accounts for this because sponsors' consideration sets depend on the size group in Table A6.

Sialm and Stefanescu (2021)).

In Table A8 I account for sponsors' inertia by including the lagged plan inclusion probability as an additional characteristic that affects sponsors' demand for a given fund. Intuitively, if sponsors' menus are sticky because of switching costs, we would expect funds' inclusion probabilities to persist over time. If inertia is an important driver of menu choices we would also expect past inclusion probabilities to explain a substantial fraction of the cross-sectional variation in the current inclusion probabilities, possibly reducing the importance of other characteristics such as fees and funds' affiliation. The estimates in Table A8 confirm these intuitions. The coefficient on past inclusion probabilities is large, positive and significant, suggesting that sponsors' inertia is an important driver of plan menu choices. Moreover, when accounting for inertia, the estimated sponsors' sensitivity to fees is lower, with an estimated elasticity to fees of about 1.3. The estimated parameters suggest that, on average, sponsors are willing to pay 17 basis point in fees for a 1% increase in the past inclusion probability.

As an additional robustness check, in the second part of Table A5, I estimate sponsors' preferences allowing for heterogeneity in the parameter q governing the distribution of the number of options included within each investment category. Specifically, I estimate q at the year-recordkeeper-category level and find substantially similar results. The reason is that the empirical distribution of the number of options included within each category is essentially the same along several dimensions of heterogeneity one might consider. For example, in Figures A8 and A9 I plot such distribution for small and large plans, measured by the number of plan participants, and before and after the year 2014. In both cases the distribution of the number of options included within an investment category is virtually unchanged. Similarly, in Figure A10 I plot the distribution of the number of options included within each category by broad asset classes and find that, except for bond funds, the distribution is almost identical to the one for the full sample.

6.3 Identification and estimation of investors' preferences

The identification of investors' preferences follows closely the logic for the identification of sponsors' preferences. The main difference is that the identifying variation comes from variation in the observed portfolio allocations rather than variation in plan inclusion probabilities.

The estimation of plan investors' preferences is different and simpler than the estimation of sponsors' preferences because investors' demand is linear in preference parameters or some known function of those. To see this recall that plan p portfolio shares are given by

$$\boldsymbol{s}_p(\boldsymbol{f}_p;\boldsymbol{\eta}_p) = \delta \boldsymbol{e}_p + \frac{(1-\delta)}{\gamma}(I+G_p)^{-1}(W_p\boldsymbol{\beta} + \boldsymbol{\xi}_p - \boldsymbol{f}_p) \tag{24}$$

which, after multiplying both sides by $I+G_p$, becomes a system of estimating equations whose RHS only depends on own demand shocks and whose LHS is an observed linear transformation of the observed plan-level portfolio allocations

$$\tilde{\boldsymbol{s}}_p(\boldsymbol{f}_p;\boldsymbol{\eta}_p) = \delta \tilde{\boldsymbol{e}}_p + W_p\tilde{\boldsymbol{\beta}} - \tilde{\gamma}\boldsymbol{f}_p + \boldsymbol{\xi}_p \tag{25}$$

with $\tilde{\boldsymbol{s}}_p \equiv (I+G_p)\boldsymbol{s}_p$, $\tilde{\boldsymbol{e}}_d \equiv (I+G_p)\boldsymbol{e}_d$, $\tilde{\boldsymbol{\beta}} \equiv \boldsymbol{\beta}(1-\delta)/\gamma$ and $\tilde{\boldsymbol{\xi}} \equiv \boldsymbol{\xi}(1-\delta)/\gamma$. Equation

(25) can be estimated via linear regression methods.

As before, I allow for the possibility that funds' fees are correlated with plan investors' demand shocks $\boldsymbol{\xi}_p$. If funds make their price-setting decision after observing $\boldsymbol{\xi}_p$, demand and supply simultaneity would bias preference parameters estimates. To account for this type of price endogeneity I follow the same approach I used for the identification of sponsors' preferences. First, I exploit the granularity of the data to absorb unobserved heterogeneity in demand by including time, funds' brand, passive and investment category fixed effects. In this case, because the estimation is at the fund-plan level I am also able to include plan/sponsors' fixed effects to absorb plan-level preference shocks. Second, I instrument funds' fees with funds' turnover ratio which captures trading costs that are pass on to investors through higher fees. Again, the identifying assumption is that variation in funds' turnover ratio not explained by time, brand, category, passive and plan fixed effects enters investors' demand only through fees.[44]

To assess the robustness of my estimates, I also implement an Hausman-type of identification strategy to account for the endogeneity of fees. Specifically, following Egan, MacKay and Yang (2023), I instrument the fee charged by any given fund with the average expense ratio charged by the same fund provider in other investment categories. This instrument will be relevant when a provider's cost of operating a mutual fund is correlated with its costs of operating its other mutual funds and when these costs are pass on to investors through fees. The instrument will be excluded if investors' residual demand shocks for any given fund (after controlling for the above series of fixed effects) are uncorrelated with the fees charged by the same fund provider on its funds in other investment categories. The resulting estimates are very similar to the ones I obtain when using funds' residual turnover as instrument for fees.

I estimate investors' preferences applying linear IV to equation (25) under the assumption that funds' turnover ratios Z_j are mean-independent of investors' demand shocks, formally, I require that $\mathbb{E}[\xi_{jp}|Z_j] = 0$. To compute the LHS of (25) I need to specify which asset characteristics form the matrix $G_p = \tilde{G}_p\tilde{G}_p'$. In classic portfolio theory $\tilde{G}_p$ would include characteristics capturing the correlation structure between assets. Perhaps the most natural way to proceed would be to construct $\tilde{G}_p$ after estimating funds' loadings onto some underlying risk factors from the time-series of funds' returns and then, for each plan, compute the variance-covariance matrix of the funds available, G_p.

Although common in practice, I do not follow this approach and instead use funds' classification into investment categories to construct $\tilde{G}_p$. The reason I do so is that funds' loadings on standard factors do not seem to explain the retirement portfolio allocations observed in the data, whereas investment category fixed effects do. Table A9 presents the R-squared from regressing observed portfolio allocations on categories and funds' factors loadings. Factors alone explain close to 4% of the observed variation in portfolio shares whereas investment categories fixed effects alone explain more than three times that. More importantly, after absorbing categories fixed effects, factors' R2 drops substantially to 0.1% suggesting that factors explanatory power was just proxing for investment category classifications. In a world in which investors' allocation decisions depend on assets' factor structure we would expect factor loadings to have some power in explaining the observed portfolio shares.

Based on this evidence, I use investment categories to model how plan investors interpret assets substitution patterns. I do so by creating a three level nesting structure

[44] In Appendix C I provide more details on using funds' turnover as instrument for funds' fees.

of asset categories. The first level consists into three broad asset classes Equity, Allocation and Bond. Then I create a second level for each of these classes. For instance, funds belonging to the Equity class are further classified into Equity-Large, Equity-Mid, Equity-Small and Equity-International. In the third level, Equity-Large fund are further classified into Equity-Large-Blend, Equity-Large-Growth and Equity-Large-Value and similarly for other second level categories. I consider each of these levels as a separate asset characteristic corresponding to a column of the matrix $\tilde{G}_p$; for instance if j is an Equity-Large-Value fund then the jth row of $\tilde{G}_p$ is a vector $\tilde{\boldsymbol{g}}_j$ that takes value 1 for the Equity, Equity-Large and Equity-Large-Value columns and 0 everywhere else. The outer product of this matrix of category indicators $G_p = \tilde{G}_p \tilde{G}_p'$ is then a three-block diagonal matrix whose element (j, l) element $\boldsymbol{g}_{jl} = \tilde{\boldsymbol{g}}_j' \tilde{\boldsymbol{g}}_l$ equals 3 if fund j and l belong to the same 3rd level category (e.g., both are Equity-Large-Value funds), equals 2 if they belong to the same 2nd level but to a different 3rd level category, equals 1 if they only belong to the same 1st level and equals 0 otherwise.[45]

6.4 Estimates of investors' preferences

Table 3 presents the estimates of investors' preference parameters based on the linear specification in equation (25). The first three columns present some OLS estimates whereas the fourth column reports the IV estimates from instrumenting funds' fees with funds' turnover ratios. The estimates reported correspond to the coefficients on the RHS of equation (25). Besides funds' expenses, I assume that past returns gross of fees and funds' affiliation enter the set of asset characteristics W_p determining the linear component of plan investors' preferences $W_p \boldsymbol{\beta} - \boldsymbol{f}_p + \boldsymbol{\xi}_p$. Investment categories also enter this linear component of investors' utility because I absorb category fixed effects in all specifications. If one interprets this linear component as investors' subjective expectations about assets' returns, the implicit assumption is that investors' subjective beliefs depend on funds' past returns relative to their corresponding investment category.

The OLS estimates broadly suggest that plan investors dislike fees, like returns and have a preference for funds' that are affiliated with their sponsor recordkeeper. A closer look at the magnitudes further reveals that plan investors care more about funds' returns than their sponsors because, in this case, the preference coefficient on returns and its corresponding marginal effect are substantially larger.[46] This is consistent with sponsors designing their plan to merely comply with regulation and minimize liability risk. Current ERISA regulation indeed prescribes that sponsors are not liable for funds' market performance to the extent that their plan includes high quality options compared to the alternative available in the market. Because fees are known whereas performance is uncertain, it is not surprising that sponsors care more about fees rather than gross performance as they cannot be held accountable for the latter.

Plan investors seem to value funds' affiliation less than their sponsors. For the latter, funds' affiliation is a crucial driver of plan inclusion decisions whereas for plan investors the importance of funds' affiliation is more modest although not irrelevant. This is consistent with agency frictions mostly biting at the plan design stage where sponsors

[45] Another way to see this is thinking about G_p as a quadratic interaction of fixed effects where each level of classification is a fixed effect. I provide an illustrative example in Appendix B.

[46] Nevertheless, investors, like sponsors, still care much more about fees relative to returns. Although the coefficient magnitudes are similar, in Table (3) fees are measured in percentage points whereas returns are measured in decimal form.

Plan investors preference parameters					
	OLS			IV	ME
Expense ratio ($\tilde{\gamma}$)	-0.004	-0.011	-0.012	-0.044	-0.098
	(0.001)	(0.001)	(0.001)	(0.004)	
Affiliated ($\tilde{\beta}_1$)	0.001	0.004	0.017	0.015	0.034
	(0.000)	(0.000)	(0.001)	(0.001)	
Gross returns ($\tilde{\beta}_2$)	0.002	0.012	0.013	0.020	0.043
	(0.001)	(0.001)	(0.001)	(0.001)	
Fraction inactive (δ)	0.267	0.290	0.407	0.371	-
	(0.006)	(0.006)	(0.008)	(0.009)	
Median fee elasticity	-0.225	-0.645	-0.688	-2.517	-
Median fee elasticity (active)	-0.306	-0.909	-1.162	-4.002	-
Fstat	-	-	-	82.55	-
R2	0.24	0.25	0.42	0.82	-
Fund brand FE	N	Y	Y	Y	-
Employer FE	N	N	Y	Y	-

Table 3: Estimates of plan investors preferences. All specifications include year, category and passive fixed effects. Expense ratios are in percentage points (pp.). R2 for IV columns is first stage. ME are the (median) marginal effects for portfolio allocations in pp. for a basis point increase in expenses or a pp. increase gross returns.

tend to favor affiliated funds when constructing their retirement plan (Pool, Sialm and Stefanescu (2016)). Agency frictions could potentially spillover to the investment stage if, for instance, recordkeepers also offered advising services to plan investors and were to push investors towards affiliated funds. Although I am not aware of any empirical evidence on mis-advising for the particular context I am considering, a theoretical literature in financial economics contemplates this possibility (Inderst and Ottaviani (2012a), Inderst and Ottaviani (2012b)).

The coefficient on fees increases moving from the first to the third column of the OLS estimates. In particular, it more than doubles when I control for fund brand fixed effects. This is consistent with fund brands potentially capturing a good amount of unobserved heterogeneity in investors' preferences for particular fund brands thereby making funds of the same investment provider (and within the same category) perceived as more substitutable to each other. The same happens to the coefficient on funds' gross returns which is not surprising as we expect investors to care about returns net of fees. Including plan (or equivalently sponsor) fixed effects does not affect too much fees and returns coefficients but triples the coefficient on funds' affiliation suggesting that the extent with which agency frictions matter might depend on sponsor-level unobservables such as sponsors bargaining power in negotiations with the recordkeeper (Bhattacharya and Illannes (2022)). Lastly, after controlling for sponsor fixed effects, the model estimates that nearly two out of five investors do not make an active investment decision.

I report the IV estimates in the fourth column of Table 3. All coefficients, except for the one on funds' fees, are largely unchanged. Conversely, the coefficient on funds' expenses is almost four times larger. The discrepancy between OLS and IV estimates is common in contexts where prices and quantities are determined simultaneously in equilibrium. OLS estimates often imply inelastic demand curves because the observed variation in quantity and prices is also due to shifts in demand. However, after instrumenting for prices, the resulting estimates recover demand curves that are much more elastic. This

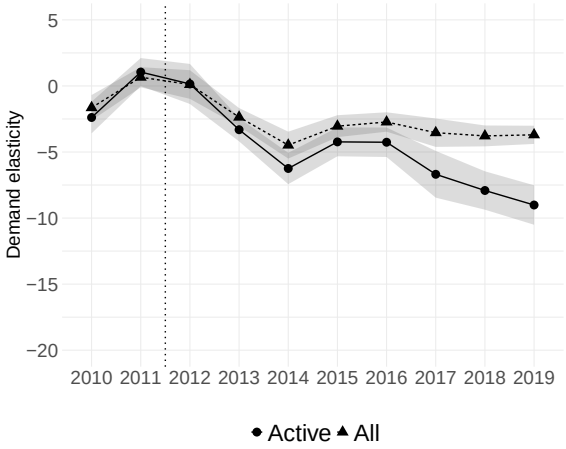

Figure 5: Cross-sectional estimates of plan investors median portfolio elasticity to funds' fees. Dahshed vertical line corresponds to the DOL fee disclosure reform.

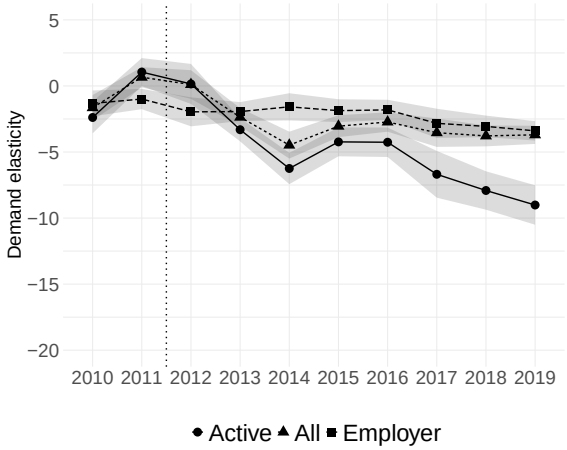

Figure 6: Cross-sectional estimates of plan investors median portfolio elasticity to funds' fees. Dark blue squared-dashed line is the median sponsor elasticity.

pattern emerges clearly when looking at the elasticity to fees of investors' portfolio allocations, which I report in the second part of Table 3. Plan investors' portfolio allocations are inelastic under OLS but become elastic after instrumenting funds' fees with funds' turnover ratios. The median elasticity is around 2.5, almost 50% larger than sponsors' elasticity, indicating that sponsors may not be internalizing investors' preferences when constructing their plan menu. The misalignment in elasticities increases even more if we consider investors who actively form their retirement portfolio. The median fee elasticity for active investors is around 4, over twice larger than sponsors' elasticity to fees. The estimates remain largely unchanged if I use Hausman instruments instead of funds' residual turnover ratios (Table A10).

To get a better sense of the magnitudes the last column of Table 3 reports the marginal change in investors' portfolio allocations implied by a unit change of the corresponding characteristic. The change in allocations is measured in percentage points per basis point increase in expenses and per percentage point increase in gross returns. A 10 basis points increase in fees reduces the corresponding portfolio allocation by nearly 1%. This magnitude is not small if one considers that the average retirement portfolio allocation is of about 3%. Conversely, a 1 percentage point increase in past gross returns increases the corresponding portfolio allocation by a modest 0.04%. Because I am absorbing category fixed effects, the latter should be interpreted as the effect of a 1% increase in funds' performance, measured relative to the corresponding investment category, on their plan-level portfolio allocation. The modest magnitude suggests that plan investors do not chase performance as much as documented for the whole mutual fund industry (Chevalier and Ellison (1997)).

The model estimates provided in Table 3 pool together all cross-sections from 2010 to 2019, although the identifying variation remains cross-sectional because I always include year fixed effects. That being said, nothing prevents me from estimating investors' preferences separately for each observed cross-section of plan menus and assess how such estimates have changed over time.

Figure 5 plots the estimated median fee elasticity for each cross-section from 2010 to 2019. Two broad patterns emerge. First, investors seem to have become more sensitive to

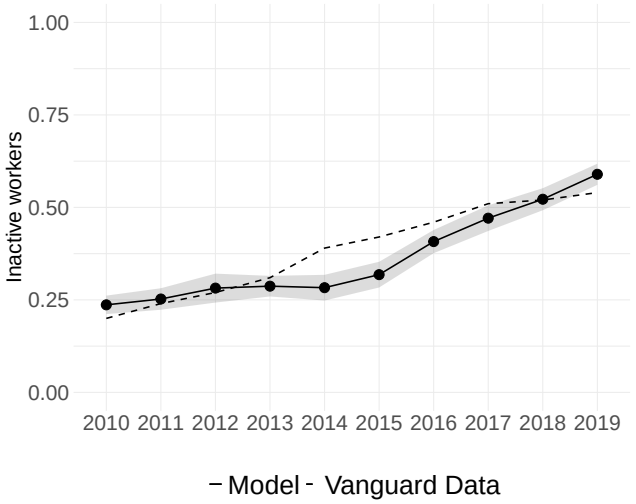

Figure 7: Cross-sectional estimates of the fraction of inactive investors (black). Share of plan investors in Vanguard plans holding a single TDF.

fees over time, with a sharp drop in the estimated elasticity from 2011 to 2013. Second, investors have become more inactive, with the estimated elasticity for active investors diverging from the elasticity of all investors. The first pattern could be a consequence of the regulatory push that required plan sponsors and investment providers to disclose investment fees to plan investors. Specifically, starting from the year 2012 the Department of Labor (DOL) required plan sponsors to disclose information about funds' expenses and performance directly to plan investors and recent empirical evidence suggests that investors have become significantly more attentive to fees as a consequence of that (Kronlund, Pool, Sialm and Stefanescu (2021)). Sponsors' elasticity has also been decreasing over time from around -1.3 in 2010 to about -3.4 in 2019 (Figure 6). Except for the years before the DOL reform, sponsors tend to be less elastic to fees than investors, especially if compared to active investors.

The second pattern is likely a symptom of the growth in the demand and supply of TDFs following the 2006 Pension Protection Act which identified TDFs as one of the qualified default investment alternatives for retirement plans. Since then, TDFs have become a constant component of retirement plan menus with more than 80% of sponsors offering at least one TDF in their plan as of 2019 (Figure A12). At the same time, plan investors have been increasing their TDFs holdings, with the average portfolio share of TDFs across plans growing three-folds from approximately 10% in 2010 to more than 30% as of 2019 (Figure A13).[47] My model attributes this increase in TDFs' portfolio share to the presence of more inactive investors. Indeed, the estimated fraction of inactive investors increases from roughly 25% in 2010 to 60% in 2019 (Figure 7), matching closely the share of investors holding a single TDF as reported by Vanguard (2022). This large increase in the share of investors that do not actively form their portfolio explains why the estimated fee elasticity for active investors diverges from the fee elasticity of all investors and why the latter moves closer to sponsors' estimated elasticity (Figure 6).

[47]Interestingly, this increase in TDFs' market shares has not been accompanied by a reduction in TDFs fees. Relatively to other type of funds TDFs have experienced a much lower decline in fees (Figure A14 and A15)

7 Price-cost Margins and Fee Decomposition

In this section I combine the estimates of sponsors' and investors' preferences together with the Nash-Bertrand first order conditions derived in (19) to recover funds' price-cost margins. After that, I exploit the characterization of equilibrium fees derived in equation (20) to decompose the observed variation in fees into the monopolist margin, the Hotelling markdown and the plan inclusion markdown.

7.1 Recovering funds' price-cost margins

To begin with, I rewrite funds' profit maximization problem making explicit its dependence on the various dimensions of variation I have in the data. I index time periods (i.e., years) by t and recordkeepers by r. Next, I denote by R_{jt} the set of recordkeepers that include fund j in their network of funds and by P_{rt} the set of retirement plans administered by recordkeeper r. I assume that funds set fees simultaneously in each period before sponsors form their retirement menu but knowing which recordkeeper networks they belong to. I rewrite problem (14) as follows:

$$\max_{(f_{jt})_t} \quad \sum_t \sum_{r \in R_{jt}} P_{rt} \cdot (f_{jt} - c_{jt}) \cdot \int \phi^r_{jpt}(\boldsymbol{f}; \boldsymbol{\theta}_p) s^r_{jpt}(\boldsymbol{f}; \boldsymbol{\eta}_p) A_p dF(A_p, \boldsymbol{\theta}_p, \boldsymbol{\eta}_p) \tag{26}$$

where I made explicit the dependence of the inclusion probabilities and portfolio shares on the identity of the plan recordkeeper. This dependence is a consequence of the fact that different recordkeepers have different networks of funds.

The first order condition associated with problem (26) imply the following price-cost margins for fund j:

$$f_{jt} - c_{jt} = -\frac{\sum_{r \in R_{jt}} P_{rt} \cdot \int \phi^r_{jpt} s^r_{jpt} A_p dF_p}{\left(\sum_{r \in R_{jt}} P_{rt} \cdot \int \frac{\partial \phi^r_{jpt}}{\partial f_{jt}} s^r_{jpt} A_p dF_p\right) - \left(\sum_{r \in R_{jt}} P_{rt} \cdot \int \phi^r_{jpt}(1-\delta_p)\gamma_p^{-1}(1-\bar{\kappa}^{r,p,t}_{jj}) A_p dF_p\right)} \tag{27}$$

where the two addends in the denominator represent the revenue loss from the marginal sponsor and the revenue loss from the marginal investor respectively.

By plugging in (27) the estimated distributions of sponsors' and investors' preference parameters one obtains the margins charged by funds in an interior Nash-Bertrand equilibrium. I report the estimated margins and marginal costs in Table 4. The first set of columns presents the estimates for the full sample of funds, whereas the remaining sets of columns focus on active funds, passive funds and TDFs respectively.[48] Starting from the sample of all funds the estimates suggest that the median fund charges a margin of about 14 basis points and a median markup around 20%.

Perhaps not surprisingly things change when looking at passive funds. The median passive fund has a marginal cost more than three times lower than the median fund among all funds and charges a margin that is around 6 basis points. Interestingly, although the absolute margin for the median passive fund is twice smaller than the margin for the median fund among all funds, in relative terms, it charges a markup of about 30%. This is suggestive of the fact that, although passive funds are typically perceived as more

[48] I provide more details on the derivation in Appendix B

	All Funds			Active Funds			Passive Funds			Target Date Funds		
	Fee	MC	PCM	Fee	MC	PCM	Fee	MC	PCM	Fee	MC	PCM
p25	43	27	8	63	45	9	9	2	3	13	2	8
p50	74	58	14	85	70	15	22	13	6	32	19	14
p75	101	87	19	109	95	21	45	36	10	52	39	15
Mean	73	59	14	85	70	15	29	21	8	37	26	11

Table 4: Price cost margins and marginal costs implied by the Nash-Bertrand first order conditions. Magnitudes are in basis points.

homogeneous products, they still enjoy substantial market power and do not pass all their cost efficiency down to investors (Hortaçsu and Syverson (2004)). Compared to all funds taken together, the distribution of costs for passive funds is more skewed, with the average fund bearing a marginal cost of about 21 basis points. However, the absolute margin charged by the average fund is similar to the median suggesting that the skewness is mostly driven by the cost structure.

The estimated margins and costs for Target Date Funds are reported in the last set of columns of Table 4. Starting from the fees, we can see that TDFs tend to be more expensive than passive funds but cheaper than all funds taken together, with the median and average TDFs charging an expense ratio of about 32 and 37 basis points respectively. On the other hand, the estimated cost structure for TDFs is not too dissimilar to the one for passive funds, suggesting that TDFs, although being cost-efficient investment vehicles, charge higher margins to investors. This is particularly evident for the most efficient TDFs. In fact, the margins for the TDFs in the 50th and 25th percentiles are of 14 and 8 basis points respectively, with a median markup of about 39%.

TDFs' pricing power comes from two sources. First, most TDFs are funds affiliated with the plan recordkeeper. Because sponsors value funds' affiliation, inclusion probabilities will be elastic to TDFs' fees. Second, TDFs are the default option in the vast majority of plans, allowing them to capture assets from inactive investors who do not respond to fees. Recent empirical evidence suggests that TDFs charge excessive fees because they are structured as funds of funds and, as such, their expenses reflect multiple layers of fees. Moreover, the vast majority of a TDF's holdings are funds that belong to same fund family as the TDF itself and, some TDFs tend not to include the cheapest share classes of such funds (Brown and Davies (2021), Sandhya (2011)).[49]

7.2 Decomposition of equilibrium fees

In this section I decompose the observed fees exploiting the decomposition derived in equation (20). Specifically, I use the estimated preference parameters and the estimated marginal costs to decompose the observed fees into (1) monopolistic fee, (2) Hotelling markdown and (3) plan inclusion markdown. I perform this decomposition for each year of the sample from 2010 to 2019 and obtain the following decomposition for each fund j

[49] Agency frictions not only impact TDFs fee setting behavior but also their risk-taking incentives as recently documented by Balduzzi and Reuter (2018).

in each year t

$$f_{jt} = \frac{\tilde{\mu}_{jt} + c_{jt}}{2} - h_{jt} - \bar{\iota}_{jt} \tag{28}$$

where $\bar{\iota}_{jt}$ is the jth component of the vector

$$\left(I - \frac{diag(\tilde{\mathcal{K}}_t)}{2} - \frac{\tilde{\mathcal{K}}_t}{2}\right)^{-1} \frac{\boldsymbol{\iota_t}}{2}. \tag{29}$$

To get a sense of the magnitudes, Table A11 shows the average of each of the three components across all funds and cross-sections. The first column reports the average observed fee which is about 66 basis points.[50] The last three columns instead show the averages for each of the three components respectively. A monopolist, on average would charge a fee of about 120 basis points but, because of competition, it needs to give up about 44% of such fee. On average, the two competitive markdowns reduce the monopolist fee by nearly 54 basis points. The contribution of each of those is similar, suggesting that competition for entering investors' choice sets and competition in terms of product characteristics are equally important. On average, the Hotelling and inclusion markdowns each erode more than 20% of the fee a monopolist would be able to charge to its consumers.

Figure A16 repeats the same exercise for each cross-section from 2010 to 2019. The black solid line represents the average fee and shows its declining trend from around 80 basis points in 2010 to nearly 50 basis points in 2019. The decomposition sheds light on the sources of this decline. The blue bars suggest that the decline in fees is not a consequence of changes in investors willingness to pay, as captured by $\tilde{\mu}$, nor of changes in technological primitives as captured by funds' marginal costs c. The monopolist fee $(\tilde{\mu}+c)/2$ indeed has been roughly stable over time fluctuating between 100 and 130 basis points. On the other hand, the two markdowns seem to be the driver of such declining trend in fees. In absolute terms, they went from reducing the monopolist fee by about 29 basis points in 2010 to nearly 78 basis points in 2019. In relative terms, they accounted for a 27% reduction of the monopolist fee in 2010 which has more than doubled over time, accounting for a 59% reduction in 2019. Overall, these patterns are consistent with both sponsors and investors becoming more sensitive to funds' expenses.

8 Counterfactuals

In this section I evaluate the effects of three policy counterfactuals regulating the design of retirement plans. First, I consider the elimination of agency frictions whereby sponsors favor funds affiliated with their plan recordkeeper. Second, I consider the effects of a policy that mandates the inclusion of low-cost options such as low-cost S&P 500 index funds trackers or low-cost TDFs. Lastly, I consider a policy that caps funds' expenses.

For all counterfactuals, I quantify how the policy in question impacts plan investors welfare and plan expenses relative to the status quo. The latter corresponds to the welfare

[50] This fee is slightly lower than the overall average fee because, to reduce the computational burden, I performed the decomposition only including the 200 largest funds in each recordkeeper network of funds. On average the 200 largest funds accounted for more than 80% of the AUM managed by the recordkeeper.

and expenses computed for the observed plan menus, plan expenses and portfolio allocations. The main takeaway is that mandating the inclusion of low-cost default options and imposing expense ratio caps are the most effective policies. Assuming that sponsors do not value funds' affiliation does not improve investors' outcomes because nothing prevents them to include expensive options that are not affiliated. Similarly, mandating the inclusion of low-cost index funds has only a modest effect on welfare and expenses. The reason is that sponsors will still be including expensive funds and investors will still be investing in those either because they are inactive or for diversification purposes. Mandating the inclusion of low-cost TDFs improves outcomes because inactive investors benefit from having access to cheaper default options.

To measure investors surplus I rely on the quadratic specification of investors' portfolio problem defined in (8). At the optimal portfolio allocation, the surplus for active investor i in plan p can be written as,[51]

$$IS_i \equiv \frac{1}{2}\sum_{j=1}^{J_p} a_{ji}(\boldsymbol{f};\boldsymbol{\eta}_p)(\mu_{jp} - f_j)$$

where $\mu_{jp} = \boldsymbol{w}'_{jp}\boldsymbol{\beta} + \xi_{jp}$. Investors' surplus is the sum of the areas below the demand curves of each asset. Because preferences are quadratic and, in turn, the demand for each asset is linear, this area corresponds to a rectangular triangle with height given by $(\mu_{jp} - f_j)$ and base given by given by a_j. Integrating over all active investors we obtain the average surplus for a plan p active investor:

$$IS_p^{\text{active}} = \frac{1}{2}\sum_{j=1}^{J_p} s_{jp}^{\text{active}}(\boldsymbol{f};\boldsymbol{\eta}_p)(\mu_{jp} - f_j)$$

where

$$s_{jp}^{\text{active}} \equiv \sum_{i \in I_{p,\text{active}}} \frac{A}{(1-\delta)A_p} a_{ji}.$$

To obtain a complete welfare measure I need to incorporate the surplus of inactive investors. Because, in the model I do not specify any preference for these investors, I define their surplus as

$$IS_p^{\text{inactive}} = \frac{1}{2}(\mu_d - f_d)$$

where d is plan p's default option. The surplus for a plan p investor is given by

$$IS_p = \delta \cdot IS_p^{\text{inactive}} + (1-\delta) \cdot IS_p^{\text{active}}$$

and, the overall surplus is then

$$\int IS_p(\boldsymbol{f};\boldsymbol{\eta}_p) dF(\boldsymbol{\eta}_p). \tag{30}$$

[51] See Appendix B for a derivation.

Each counterfactual amounts to (1) imposing the policy, (2) solve for funds' counterfactual equilibrium fees, (3) simulate sponsors' plan menus under the counterfactual policy, (4) compute investors' counterfactual portfolios and compute their surplus from equation (30) . Because investors have preferences over their portfolio allocations, their surplus is measured in units of (subjective) excess returns net of fees (e.g., $\mu_{jp} - f_j$). For active investors, I recover μ_{jp} from their estimated preference parameters by computing $\mu_{jp} = \boldsymbol{w}'_{jp}\hat{\boldsymbol{\beta}} + \hat{\xi}_{jp}$, where $\hat{\xi}_{jp}$ are obtained from the residuals of the linear regression in (11) multiplied by the estimated $\tilde{\gamma}$. For inactive investors, I use the average annual (excess) return to measure μ_d.

8.1 Eliminating preference for affiliated funds

The first counterfactual I consider restricts sponsors' preferences by forcing them not to value funds' affiliation when constructing their plan menu. Under this restriction, holding every other characteristic constant, an affiliated fund and a non-affiliated fund will have the same likelihood of being included in a given plan. A practical way to implement such policy would be issuing penalties for sponsors that are found to be favoring affiliated funds even though they exhibit worse performance than otherwise similar alternatives (Pool, Sialm and Stefanescu (2016)).

The second row of Table A12 presents the results from this counterfactual exercise and shows that such policy would be ineffective, leaving investor surplus and plan expenses almost unchanged. In fact, plan expenses decrease by 1 basis point and investor surplus decreases by 3 basis points.

This policy is ineffective because removing sponsors' preference for affiliated funds does not prevent them from including expensive funds that are not affiliated with their recordkeeper. For this reason, counterfactual plan expenses remain unchanged as well as sponsors' sensitivity to funds' fees.

Investors' surplus is also almost unaffected. The slight decrease in their surplus could be either because investors have a small preference for affiliated funds or because affiliated funds are not systematically worse than other alternatives. For example, in many cases TDFs are affiliated with the plan recordkeeper and, we know that their expenses are typically below average (Table 4) and are particularly valued by inactive investors.

8.2 Mandating the inclusion of low-cost options

The second set of counterfactuals studies the effect of policies mandating the inclusion of low-cost investment options. I consider both the inclusion of low-cost index funds that track the S&P 500 and the inclusion of low-cost TDFs. Both policies improve investors' outcome and lead to a reduction of the average plan expenses relative to the status quo.

Mandating the inclusion of at least one low-cost index fund increases investors' surplus by 2% and decreases the average plan expense by 10% relative to the status quo. In magnitudes, the surplus for an investor with a $35,000 account balance increases by about $20 per year (Figure 8). To add perspective, in the last column of Table A12 I consider the dollar savings over 40 years for an household receiving an annual income of $70,000 and contributing 10% to its 401(k) every year.[52] Such household would save approximately $12,000 in fees after the implementation of this policy.

[52]For such computation I assume an annual return of 6%.

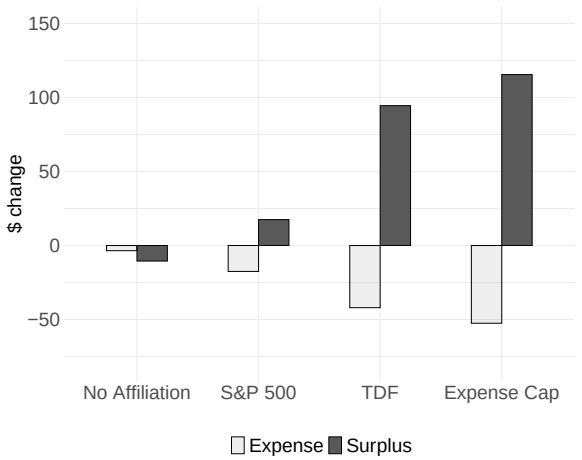

Figure 8: Dollar change in surplus and average plan expense relative to the status quo for an investor holding a retirement account of $35,000 under different plan design policies. Magnitudes are in dollars-per-year.

While one might have expected significant impact from such a policy, its actual effects are relatively modest. This is noteworthy because I selected low-cost funds renowned for being both affordable. Each comes with an expense ratio well under 10 basis points, complemented by a 5-star Morningstar rating. My selection includes the Vanguard 500 Index fund (VFIAX) and ETF (VOO), Fidelity 500 Index Fund (FXAIX), Schwab S&P 500 Index Fund (SWPPX), Blackrock iShare Core S&P 500 ETF (IVV), and SPDR S&P 500 ETF (SPY).

Despite this, there are a couple of key reasons why this policy has had only a modest effect on investors' welfare. Firstly, although investors value lower fees, they also want to diversify across all assets available. As such, they will substitute toward the low-cost index only up to the point that does not hurt their diversification needs. This in practice requires maintaining part of the holdings into more expensive funds. Secondly, a significant segment of investors remain inactive in their investment approach. Instead of actively selecting funds, they default their contributions to the available TDF. The addition of a low-cost index fund does not alter the behavior of this group. Furthermore, it's worth noting that nearly half of sponsors already incorporate these type of funds in their existing offerings, implying that the potential welfare gains only come from the half the sponsors not offering those type of funds as part of their plan menu.

Investors' outcome improves if, instead of mandating low-cost index funds, the policy mandates the inclusion of low-cost TDF. In this case, investor surplus increases by 11% and the average plan expense decreases by 23%. In magnitudes, the surplus for an investor with a $35,000 account balance increases by nearly $100 per year, an increase five times larger than the one obtained by mandating a low-cost S&P 500 tracker. Similarly, an household with a $70,000 income who contributes 10% to its 401(k) account would be saving $28,832 in fees over a 40 years period, an amount twice larger than the savings under a policy mandating the inclusion of low-cost S&P 500 trackers.

Why is it more effective to mandate the inclusion of low-cost TDFs than simply focusing on low-cost index funds? The primary reason lies in the benefit distribution: low-cost TDFs serve both active and, especially, inactive investors because they are used as qualified default options. In contrast, mandating low-cost index funds predominantly benefits active investors, leaving inactive ones with no benefit in terms of reduced fees.

Additionally, while TDFs often carry higher fees compared to index funds, the most affordable TDFs have expense ratios closely aligned with the ones charged by low-cost index funds. Take, for instance, the Fidelity Freedom Index series and the Vanguard Target Retirement series — both offer TDFs with expense ratios under 10 basis points.

8.3 Capping funds' expenses

The last counterfactual I consider studies the effect of a 50 basis point expense ratio cap. Under this policy sponsors are allowed to include in their menu only funds with an expense ratio below 50 basis points. As a consequence, all funds whose marginal cost is higher than 50 basis point will exit the market. The latter are in most cases active funds, as more than 3/4 of passive funds and TDFs have an expense ratio below 50 basis points (Table 4).

This policy increases investor surplus by 14% and decreases the average plan expenses by 30%, corresponding to an increase in surplus of about 33 basis points and a decrease in plan expenses of about 15 basis points. A plan investor with a balance account of $35,000 enjoys an increase in surplus of about $120 dollars per-year whereas an investor contributing 10% of its $70,000 income is expected to save about $36,000 in fees over 40 years.

Perhaps not surprisingly, this policy is the most effective among the ones I considered thus far. On the one hand, it benefits inactive investors by eliminating the right-tail of expensive TDFs. On the other, it benefits active investors by ensuring that all investment options available are not excessively expensive.

Before concluding, one remark is in order. My analysis so far abstracted away from extensive margin considerations and implicitly assumed that sponsors would be always willing to provide a retirement plan to their employees. In practice, plan provision could be affected by these type of policies. For example, under a 50 basis point expense ratio cap, it is likely that recordkeepers would lose revenues from revenue-sharing fees unless sponsors themselves compensate such loss by increasing their direct payments to their recordkeepers (Bhattacharya and Illannes (2022)). Some sponsors may be unwilling or might not have the resources to bear such costs and, consequently, might decide not to offer a retirement plan to their workers in the first place.

A plausible solution to minimize the extensive margin repercussions of these type of policies would be implementing such policies while at the same time subsidizing plan sponsors to incentivize plan provision. In practice, these type of subsidies have been already introduced in the 2019 SECURE Act to push small business to offer a retirement plan to their employees.

9 Conclusions

This paper proposes an equilibrium model of retirement plan menu choice, portfolio choice and fee competition between investment providers to uncover the factors contributing to the design high-cost employer-sponsored retirement plans and quantify the welfare effects of policies regulating plan design.

The model features a two-layer demand system where, in the first layer, sponsors choose their retirement plan and, in the second layer, plan investors form their retirement portfolio from the options available in their menu. On the supply side, investment funds

compete by setting fees simultaneously while accounting for the two layers of demand. Funds compete for being included in a plan menu and for the plan assets.

I estimate the model using comprehensive data on retirement plan menus. Model estimates suggest that plan sponsors are less responsive to funds' fees than plan investors and value other fund characteristics, such as funds' affiliation with the plan recordkeeper. I use the estimated demand parameters to recover funds' price-cost margins and marginal costs from the implied Nash equilibrium conditions. Funds enjoy significant pricing power. This is particularly evident for Target-Date-Funds (TDFs), who, although almost as cost-efficient as index funds, charge double the margins, with an implied median markup of about 34%.

In the last part of the paper, I consider four policy counterfactuals that regulate the design of plan menus and quantify their effect on plan investors' welfare. The first counterfactual shuts down sponsors' preferences for funds' affiliation. The second set of counterfactuals considers mandating the inclusion of low-cost options. The last counterfactual instead imposes a 50 basis point cap on funds' expense ratios. Among those, requiring the inclusion of low-cost TDFs and capping expense ratios are the most effective in improving investors' outcomes. Specifically, mandating the inclusion of low-cost TDFs increases investors' surplus by 11%, whereas a 50 basis points expense ratio cap leads to a 14% increase. Both policies also significantly reduce average plan expenses by 23% and 30%, respectively.

An important caveat of my analysis is that it abstracts from extensive margin considerations. In particular, it assumes that these types of regulations do not affect sponsors' incentives to offer a retirement plan in the first place. In practice, imposing expense ratio caps might reduce plan provision (Bhattacharya and Illannes (2022)). A practical solution would be pairing these policies with plan provision subsidies. Quantifying the optimal subsidy scheme and the effect of this combination of policies on plan investors' welfare is an important direction for future research.

References

Agarwal, S., Grigsby, J., Hortaçsu, A., Matvos, G., Seru, A. and Yao, V. (2021), 'Searching for approval', working paper.

Allen, J., Clark, R. and Houde, J.-F. (2014), 'The effect of mergers in search markets: Evidence from the canadian mortgage industry', *American Economic Review* 104(10), 3365–3396.

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.

Baker, M., Egan, M. and Sarkar, S. (2022), 'How do investors value esg?', NBER working paper.

Balduzzi, P. and Reuter, J. (2018), 'Heterogeneity in target date funds: Strategic risk-taking or risk matching?', *The Review of Financial Studies* 32(1), 300–337.

Ballester, C., Calvó-Armenagol, A. and Zenou, Y. (2006), 'Who's who in networks. wanted: The key player', *Econometrica* 74(5), 1403–1417.

Barberis, N. and Shleifer, A. (2003), 'Style investing', *Journal of Financial Economics* 68(2), 161–199.

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.

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

Benetton, M., Buchak, G. and Robles-Garcia, C. (2022), 'Wide or narrow? competition and scope in financial intermediation', working paper.

Berk, J. and Green, R. (2004), 'Mutual fund flows and performance in rational markets', *Journal of Political Economy* 112(6), 1269–1295.

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

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

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.

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.

Bonacich, P. (1987), 'Power and centrality: A family of measures', *American Journal of Sociology* 92(5), 1170–1182.

Brancaccio, G., Li, D. and Schüroff, N. (2020), 'Learning by trading: The case of the us market for municipal bonds', working paper.

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

Brown, Z., Egan, M., Jeon, J., Jin, C. and Wu, A. (2023), 'Why do index funds have market power? quantifying frictions in the index fund market', working paper.

Buchack, G., Matvos, G., Piskorski, T. and Seru, A. (2022), 'Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy', NBER working paper.

Buchak, G., Matvos, G., Piskorski, T. and Seru, A. (2018), 'Fintech, regulatory arbitrage, and the rise of shadow banks', *Journal of Financial Economics* 130(3), 453–483.

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.

Chevalier, J. and Ellison, G. (1997), 'Risk taking by mutual funds as a response to incentives', *Journal of Political Economy* 105(6), 1167–1200.

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

Choukhmane, T. (2021), 'Default options and retirement saving dynamics', working paper.

Cuesta, I. and Sepúlveda, A. (2020), 'Price regulation in credit markets: A trade-off between consumer protection and credit access', working paper.

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

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

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

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.

Fama, E. and French, K. (1992), 'The cross-section of expected stock returns', *The Journal of Finance* 47(2), 427–465.

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.

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.

Goeree, M. (2008), 'Limited information and advertising in the u.s. personal computer industry', *Econometrica* 16, 1017–1074.

Grice, R. (2023), 'Chain-linked markets', working paper.

Grice, R. and Gueciloueur, A. (2023), 'Mutual fund market structure and company fee competition: Theory and evidence', SSRN working paper.

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.

Hausman, J. (1996), 'Valuation of new goods under perfect and imperfect competition.', *The Economics of New Goods* pp. 207–248.

Hortaçsu, A. and McAdams, D. (2010), 'Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market', *Journal of Political Economy* 118(5), 833–865.

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

Hotelling, H. (1929), 'Stability in competition', *The Economic Journal* 39(153), 41–57.

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

Inderst, R. and Ottaviani, M. (2012a), 'Competition through commissions and kickbacks', *American Economic Review* 102(2), 780–809.

Inderst, R. and Ottaviani, M. (2012b), 'How (not) to pay for advice?', *Journal of Financial Economics* 105, 393–411.

Jackson, M. (2008), *Social and Economic Networks*, Princeton University Press.

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

Kacperczyk, Sialm, C. and Zheng, L. (2005), 'On the industry concentration of actively managed equity mutual funds', *The Journal of Finance* 60(4), 1983–2011.

Kastl, J. (2011), 'Discrete bids and empirical inference in divisible good auctions', *The Review of Economic Studies* 78(3), 974–1014.

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

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

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

Loseto, M. (2023), 'Network games of imperfect competition: An empirical framework', SSRN working paper.

Ma, L., Tang, Y. and Gomez, J. (2019), 'Portfolio manager compensation in the u.s. mutual fund industry', *Journal of Finance* 74(2), 587–638.

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.

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

Massa, M. (2003), 'How do family strategies affect fund performance? when performance-maximization is not the only game in town', *Journal of Financial Economics* 67, 249–304.

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

Nelson, S. (2020), 'Private information and price regulation in the us credit card market', working paper.

Pástor, P., Stambaugh, R. and Taylor, L. (2017), 'Do funds make more when they trade more?', *Journal of Finance* 72(4), 1483–1528.

Pástor, P., Stambaugh, R. and Taylor, L. (2020), 'Fund tradeoffs', *Journal of Financial Economics* 138, 614–634.

Pellegrino, B. (2023), 'Product differentiation and oligopoly: A network approach', SSRN working paper.

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. and Richardson, D. (2022), 'New evidence on the demand for advice within retirement plans', working paper.

Richert, E. (2021), 'Quantity commitment in multiunit auctions: Evidence from credit event auctions', working paper.

Robles-Garcia, C. (2021), 'Competition and incentives in mortgage markets: The role of brokers', working paper.

Roussanov, N., Ruan, H. and Wei, Y. (2021), 'Marketing mutual funds', *Review of Financial Studies* 34, 3045–3094.

Sandhya, V. (2011), 'Agency problems in target date funds.', working paper.

Stigler, G. (1961), 'The economics of information', *Journal of Political Economy* 69(3), 213–225.

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.

Ushchev, P. and Zenou, Y. (2018), 'Price competition in product variety networks', *Games and Economic Behavior* 110(2), 226–247.

Vanguard (2022), 'How america saves.'.

Wermers, R. (2000), 'Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses', *The Journal of Finance* 55(4), 1655–1695.

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

A Additional Figures and Tables

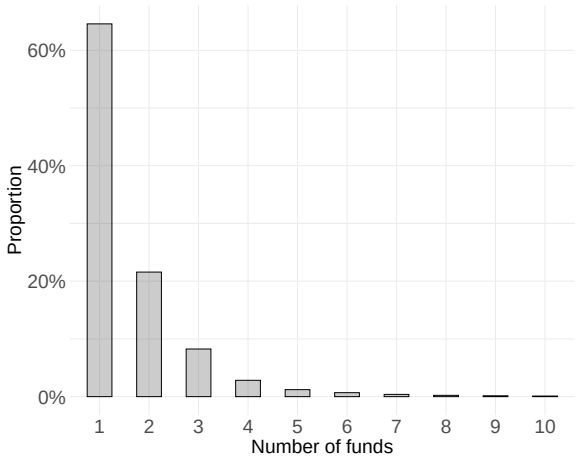

Figure A1: Distribution of number of options offered within investment category.

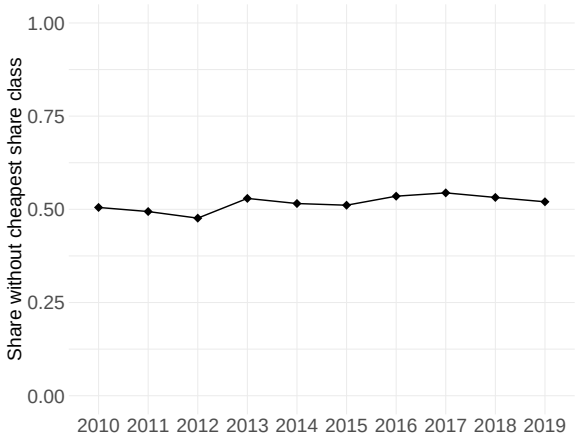

Figure A2: Within in fund × year share of employers who meet minimum investment required for cheapest share class but offer a more expensive one. The black line is the average share of employers without cheapest share class.

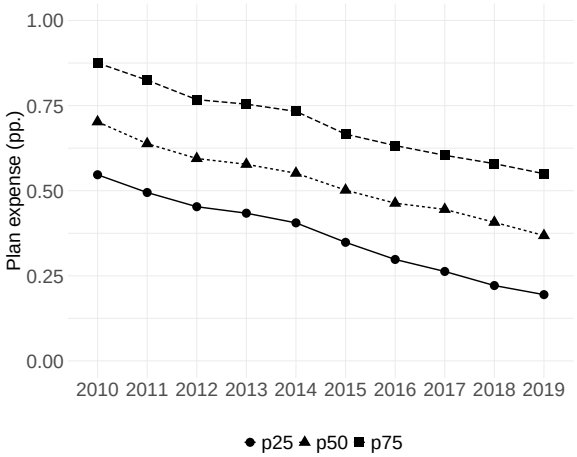

Figure A3: Distribution of average asset-weighted plan expense over time. Expenses are measured in percentage points

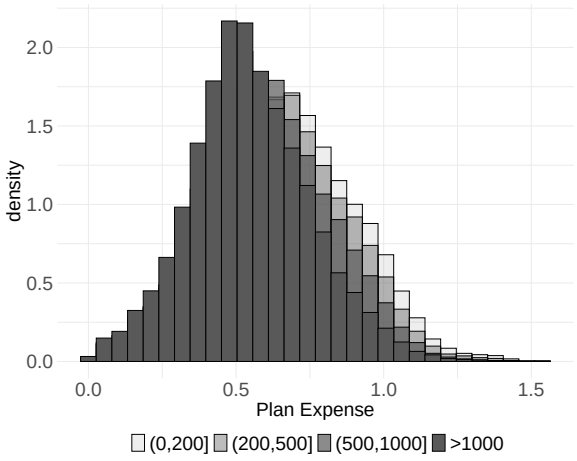

Figure A4: Distribution of plan expenses by plan size groups. Plan size is measured in number of participants.

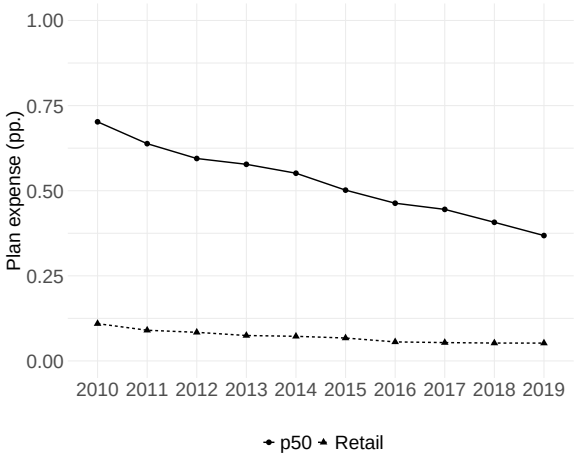

Figure A5: Median asset-weighted plan expense (dot-solid). Average expense ratio for a portfolio of Vanguard retail index funds (triangle-dashed). Expenses are measured in percentage points.

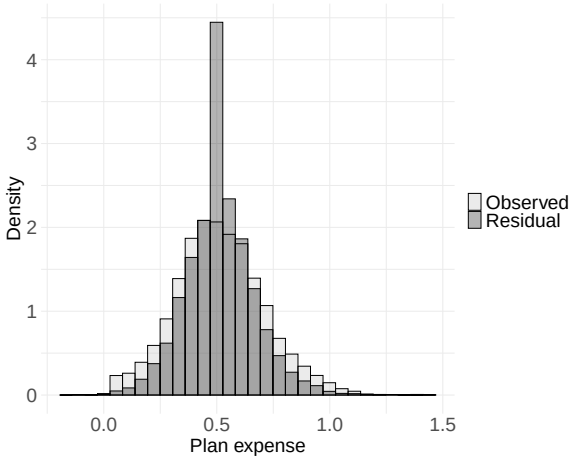

Figure A6: Within recordkeeper × six-digit NAICS dispersion in expenses. Other controls include plan assets, number of participants, and number of options.

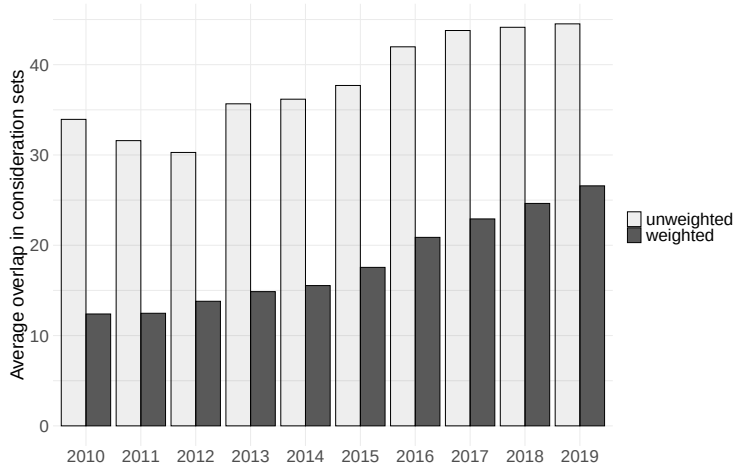

Figure A7: Average overlap in recordkeepers' network of funds. A fund belongs to a recordkeeper's network if it is offered in a plan managed by that same recordkeeper. The red bars represent the average fraction of funds that belong to the network of any two of the 10 largest recordkeepers. The turquoise bars represent the asset-weighted overlap.

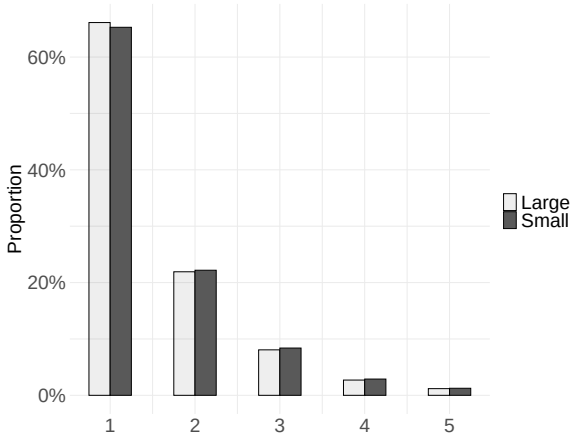

Figure A8: Distribution of number of options within category by plan size.

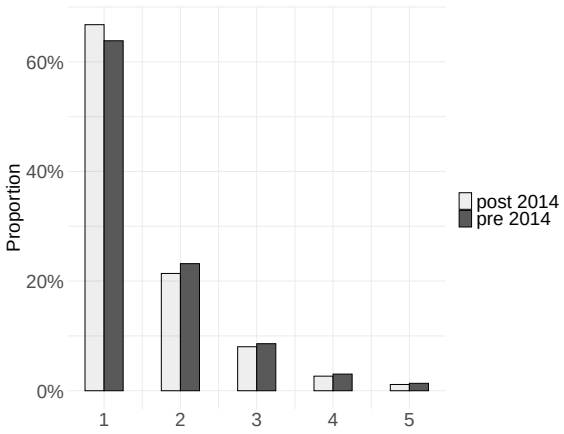

Figure A9: Distribution of number of options within category pre and post 2014.

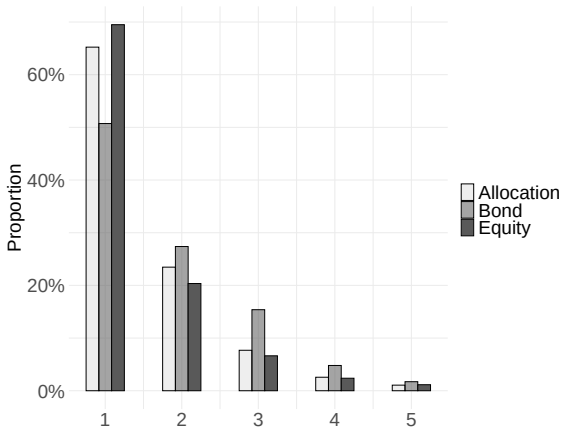

Figure A10: Distribution of number of options within category by asset class.

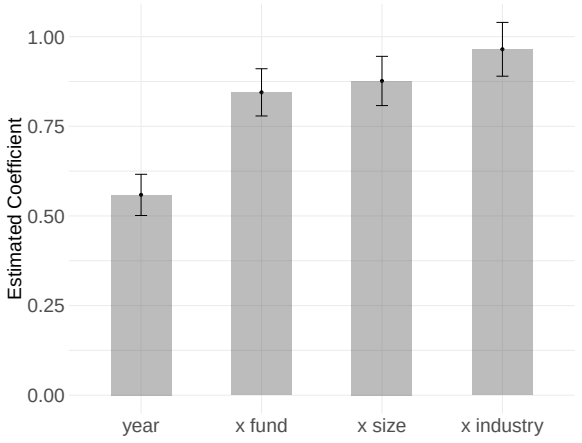

Figure A11: Coefficient from regressing log(inclusion probability) on affiliation dummy. Inclusion probability is the share of 401(k) plans offering a given fund. Inclusion probabilities are computed at the (year × size group × industry × recordkeeper) level for each fund. Size groups are based on the number of plan participants. Industry is the 2-digit NAICS.

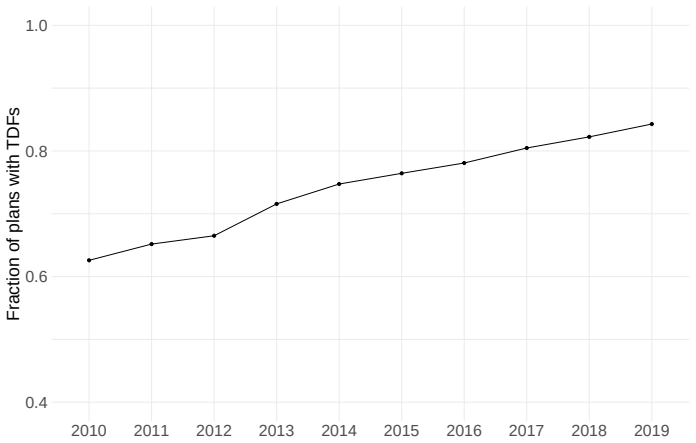

Figure A12: Share of retirement plans that offer at least one Target-Date-Fund (TDF).

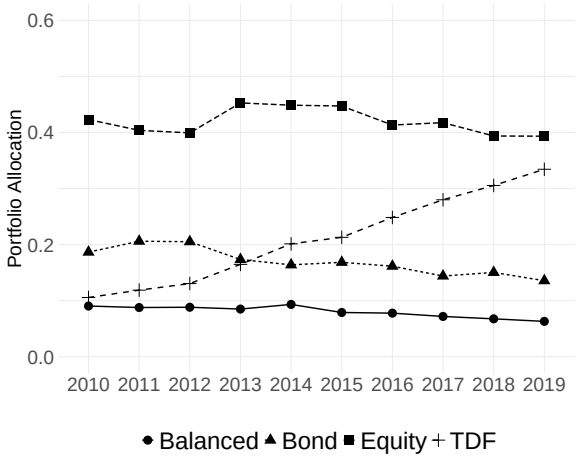

Figure A13: Average portfolio share across plan menus by asset class. Equity includes both US and International Equity funds. Balanced includes aggressive, moderate and conservative allocation funds that are not Target Date Funds (TDFs).

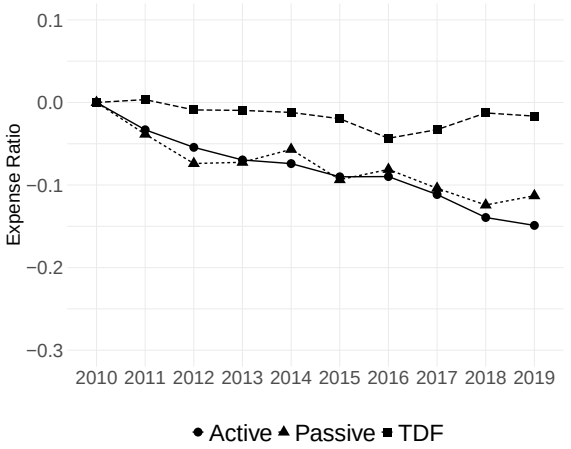

Figure A14: Secular decline in fees by fund type. The sample includes only funds that were available since 2010. The series for each type of fund has been shifted by the average fee as of 2010.

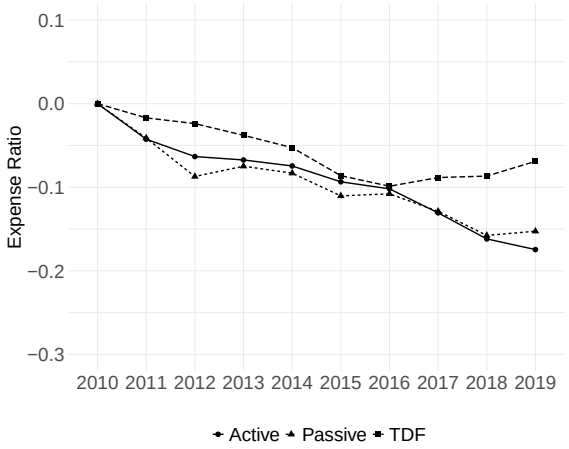

Figure A15: Secular decline in fees by fund type. The sample also includes funds introduced after 2010. The series for each type of fund has been shifted by the average fee as of 2010.

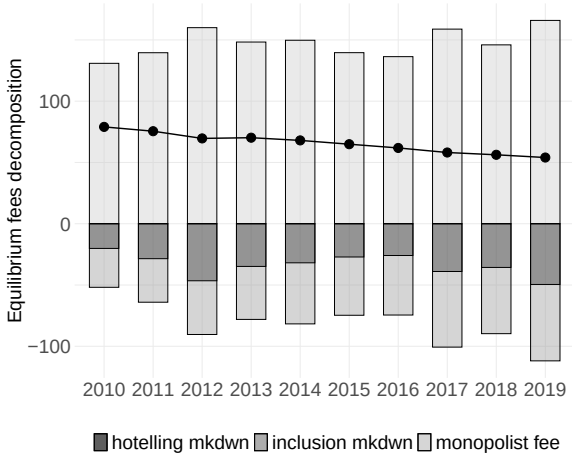

Figure A16: Cross-sectional decomposition of the average expense ratio into monopolist fee, hotelling markdown and plan inclusion markdown as defined in equation (20). Magnitudes are in basis points.

	Plan performance	Plan performance	Plan performance	Plan performance
Plan expenses	-0.76	-0.38	-0.44	-0.06
	(0.01)	(0.01)	(0.01)	(0.01)
Weighted	Y	Y	N	N
Year FE	N	Y	N	Y
R2	0.02	0.19	0.01	0.10

Table A1: Fund performance is the difference between fund return and the average category return. Plan performance is the average performance (possibly asset weighted) of all funds in the plan. Returns are gross of fees. Returns and fees are in percentage points.

	Exp. Ratio	Exp. Ratio	Exp. Ratio
log(# opt. in cat)	-0.047	-0.077	-0.051
	(0.001)	(0.001)	(0.002)
R2	0.07	0.27	0.62
Year FE	Y	Y	Y
Sponsor FE	N	Y	Y
Fund brand FE	N	N	Y

Table A2: Dependent variable is funds' expense ratio. Independent variable is the number of funds within an investment category. Expense ratios are in percentage points.

Year	Audited Plan Count	Potential Plan Count	Plan Coverage (%)	Audited Assets (bln)	Potential Assets (bln)	Asset Coverage (%)
2010	54626	66708	81.9	2,940	3,064	96.0
2011	52504	67394	77.9	3,023	3,134	96.4
2012	42372	67085	63.2	3,300	3,519	93.8
2013	42741	67514	64.8	4,086	4,158	98.3
2014	39360	68142	57.8	4,105	4,439	92.5
2015	40632	69038	58.9	4,225	4,437	95.2
2016	63464	70577	89.9	4,558	4,799	95.0
2017	68600	71855	95.5	5,504	5,553	99.1
2018	69424	73273	94.7	5,243	5,333	98.3
2019	71341	75057	95.0	6,289	6,372	98.7

Table A3: BrightScope Beacon data coverage. Assets are in billions.

Variable	N	Mean	SD	p5	p25	p50	p75	p95
Total Assets (mln.)	5615	190.736	1062.082	0.034	0.77	6.746	46.997	683.457
Portfolio share (avg.)	5615	2.896	3.821	0.119	0.887	1.795	3.279	10.039
Portfolio share (sd.)	5009	3.016	3.201	0.215	1.151	2.089	3.559	10.284
Fund-Plan turnover	5615	48.265	19.1	20.073	34.777	46.875	61.334	82.344
N. of share classes	5615	2.644	2.027	1	1	2	4	7

Table A4: Fund level summary statistics for the years 2010 to 2019. Each variable is first averaged (or summed in the case of 'total assets') within fund-year across plans, then within plan across years and tabulated across funds. The variable 'N' is the number of funds, excluding cash accounts and company stocks. Portfolio share (sd.) is the within fund-year standard deviation of the fund portfolio share across plans, which is then averaged within fund across years.

Employers preference parameters				
	Homogeneous preferences		Heterogeneous q	
	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)
Expense Ratio (bp.)	-0.021	-0.008	-0.023	-0.009
	(0.002)	-	(0.002)	-
Affiliated (dummy)	0.823	0.328	0.852	0.330
	(0.044)	-	(0.044)	-
Target (dummy)	-0.414	-0.165	-0.463	-0.179
	(0.084)	-	(0.084)	-
Target × Affiliated	0.242	0.097	0.403	0.156
	(0.104)	-	(0.101)	-
Gross returns (pp.)	0.004	0.002	0.004	0.002
	(0.001)	-	(0.001)	-
Median fee elasticity	-1.83		-2.03	
q (Calibrated)	0.70		0.66	
GMM objective (df)	4.57 (2)		4.33 (2)	

Table A5: Two-step GMM estimates of plan sponsor preferences. Robust standard errors are reported in parentheses. Year, category, passive and fund brand fixed effects are included. For the marginal effects, inclusion probabilities are in percentage points. For the heterogeneous q specification, q varies at the year-recordkeeper-category level.

Employers preference parameters								
	$(0, 200]$		$(200, 500]$		$(500, 1000]$		> 1000	
	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)
Expense Ratio (bp.)	-0.019	-0.011	-0.017	-0.011	-0.024	-0.016	-0.026	-0.018
	(0.004)	-	(0.003)	-	(0.005)	-	(0.005)	-
Affiliated (dummy)	1.173	0.678	0.903	0.549	0.914	0.615	0.732	0.497
	(0.074)	-	(0.055)	-	(0.075)	-	(0.075)	-
Target (dummy)	-0.228	-0.132	-0.249	-0.151	-0.340	-0.229	-0.451	-0.307
	(0.133)	-	(0.105)	-	(0.160)	-	(0.176)	-
Target × Affiliated	-0.096	-0.055	0.147	0.090	-0.111	-0.075	-0.217	-0.148
	(0.142)	-	(0.122)	-	(0.161)	-	(0.170)	-
Gross returns (pp.)	0.002	0.001	0.004	0.003	0.005	0.003	0.003	0.002
	(0.002)	-	(0.002)	-	(0.002)	-	(0.003)	-
Median fee elasticity	-1.65		-1.50		-1.98		-2.08	
q (Calibrated)	0.70		0.70		0.70		0.70	
GMM objective (df)	3.91 (2)		2.91 (2)		1.56 (2)		2.76 (2)	

Table A6: Two-step GMM estimates of plan sponsor preferences for plans with number of participants below the median (small) and above the median (large). Robust standard errors are reported in parentheses. Year, category, passive and fund brand fixed effects are included. For the marginal effects, inclusion probabilities are in percentage points.

	Employers preference parameters			
	Before 2014		After 2014	
	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)
Expense Ratio (bp.)	-0.014	-0.007	-0.029	-0.010
	(0.003)	-	(0.004)	-
Affiliated (dummy)	0.788	0.383	0.875	0.287
	(0.064)	-	(0.062)	-
Target (dummy)	-0.399	-0.194	-0.477	-0.157
	(0.126)	-	(0.117)	-
Target × Affiliated	0.381	0.185	0.158	0.052
	(0.156)	-	(0.142)	-
Gross returns (pp.)	0.003	0.001	0.004	0.001
	(0.002)	-	(0.002)	-
Median fee elasticity	-1.37		-2.40	
q (Calibrated)	0.70		0.70	
GMM objective (df)	4.88 (2)		6.38 (2)	

Table A7: Two-step GMM estimates of plan sponsor preferences for the pre 2014 and post 2014 subsamples. Robust standard errors are reported in parentheses. Year, category, passive and fund brand fixed effects are included. For the marginal effects, inclusion probabilities are in percentage points.

	Sponsors preference parameters			
	No inertia		Yes inertia	
	Mean	Marg. Effect (pp.)	Mean	Marg. Effect (pp.)
Expense Ratio (bp.)	-0.023	-0.009	-0.015	-0.006
	(0.003)	-	(0.003)	-
Affiliated (dummy)	0.828	0.338	0.510	0.208
	(0.045)	-	(0.045)	-
Target (dummy)	-0.453	-0.185	-0.097	-0.040
	(0.090)	-	(0.088)	-
Target × Affiliated	0.236	0.096	0.265	0.108
	(0.101)	-	(0.106)	-
Gross returns (pp.)	0.003	0.001	0.004	0.002
	(0.002)	-	(0.002)	-
log(Lag inclusion prob.)	-	-	0.259	0.106
	-	-	(0.003)	-
Median fee elasticity	-1.97		-1.30	
q (Calibrated)	0.70		0.70	
GMM objective (df)	6.68 (2)		3.41 (2)	

Table A8: Two-step GMM estimates of plan sponsor preferences accounting for inertia in menu choices. Robust standard errors are reported in parentheses. Year, category, passive and fund brand fixed effects are included. For the marginal effects, inclusion probabilities are in percentage points. Sample is restricted to plans observed for at least two consecutive years.

Plan investors preference parameters						
	OLS			IV-Turnover	IV-Hausmann	ME
Expense ratio ($\tilde{\gamma}$)	-0.004	-0.011	-0.012	-0.044	-0.040	-0.098
	(0.001)	(0.001)	(0.001)	(0.004)	(0.003)	
Affiliated ($\tilde{\beta}_1$)	0.001	0.004	0.017	0.015	0.016	0.034
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	
Gross returns ($\tilde{\beta}_2$)	0.002	0.012	0.013	0.020	0.020	0.043
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	
Fraction inactive (δ)	0.267	0.290	0.407	0.371	0.375	-
	(0.006)	(0.006)	(0.008)	(0.009)	(0.008)	
				1st Stage	1st stage	
Turnover ratio	-	-	-	0.028	-	-
				(0.000)		
Hausman IV	-	-	-	-	-0.208	-
					(0.002)	
Median fee elasticity	-0.225	-0.645	-0.688	-2.517	-2.263	-
Median fee elasticity (active investors)	-0.306	-0.909	-1.162	-4.002	-3.621	-
Fstat	-	-	-	82.55	86.69	-
R2	0.24	0.25	0.42	0.82	0.82	-
Fund brand FE	N	Y	Y	Y	Y	-
Employer FE	N	N	Y	Y	Y	-

Table A10: Estimates of plan investors preferences. All specifications include year, category and passive fixed effects. Expense ratios are in percentage points (pp.). R2 for IV columns is first stage. ME are the (median) marginal effects for portfolio allocations in pp. for a basis point increase in expenses or a pp. increase gross returns. Turnover and Hausman IV are standardized.

	Projected R2	R2
alpha	0.006	0.051
beta	0.039	0.083
category	0.143	0.182
alpha | category	0.000	0.182
beta | category	0.002	0.184
category | fund provider	0.117	0.203

Table A9: Dependent variable is plan-level portfolio allocations. All specification include plan $\times$ year fixed effects. Beta are 3 Fama-French plus Momentum and 3 bond factors.

Equilibrium decomposition of observed fees			
Fee	Monopolist fee	Hotelling markdown	Plan inclusion markdown
65.75	119.44	25.15	28.54

Table A11: Decomposition of fees following equation (20). All magnitudes are in basis points. The figures shown are averages across time and funds.

	Investor Surplus (bp.)	Average plan expense (bp.)	Fee savings/year ($)	Fee savings/40 years ($)
Status Quo	238	51	-	-
No Affiliation Preference	235	50	-	-
Low-cost Index Fund	243	46	18	11,897
Low-cost TDF	265	39	42	28,832
Expense cap (50 bp.)	271	36	53	36,192

Table A12: Investor surplus and average plan expense under different counterfactual policies. Magnitudes are in basis points. Savings are relative to the status quo. Fee savings per-year assumes a retirement account balance of $35,000. Fee savings over 40 years assumes an annual income of $70,000, contribution rate of 10% and an annual return of 6%. Expense ratio cap is at 60 basis points.

B Model derivations

In this Appendix I provide more formal derivations of the results introduced in the main text.

Derivation of ranking probabilities. Consider the simple example in the main text where we have four options $\{j, k, l, m\}$ and we want to compute the probability that option j is ranked 2nd in terms of sponsors' indirect utilities. 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:

$$\begin{aligned}\phi_j^2 = &\Pr\{u_k > u_j > u_l > u_m\} + \Pr\{u_k > u_j > u_m > u_l\} \\ &+ \Pr\{u_l > u_j > u_k > u_m\} + \Pr\{u_l > u_j > u_m > u_k\} \\ &+ \Pr\{u_m > u_j > u_k > u_l\} + \Pr\{u_m > u_j > u_l > u_k\}. \end{aligned} \tag{31}$$

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

$$\Pr\{u_k > u_j > u_l > u_m\} = \frac{\exp(V_k)}{\sum_{s\in\{j,k,l,m\}}\exp(V_s)} \cdot \frac{\exp(V_j)}{\sum_{s'\in\{j,l,m\}}\exp(V_{s'})} \cdot \frac{\exp(V_l)}{\sum_{s''\in\{l,m\}}\exp(V_{s''})} \tag{32}$$

Expression (32) can be derived analytically by integrating over the T1EV extreme value shocks and is known as ranked-ordered-logit (ROL),[53] which can be interpreted as a sequential multinomial logit decision problem.

Expression (32) applies analogously to all six terms in (31) 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 (31) and note that from (32) we can factor out the first two factors of each addend and that the sum of the last factors equals. 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 (31), the three term expression in the main text obtains.

Derivation of unconditional plan inclusion probability. Letting λ_{gp} the probability that p chooses to offer category g, $q(1-q)^{n-1}$ the probability that p includes n funds from category g and $\phi_j^{1:n}$ the probability that j is chosen conditional on n options and g being offered, the unconditional probability that j ends up being in p's retirement plan

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

can be rearranged as follows

$$\begin{aligned}
\phi_{jp} &= \lambda_{gp} \cdot \left(\sum_{n=1}^{\infty} q(1-q)^{n-1} \phi_{jp}^{1:n} \right) \\
&= \lambda_{gp} \cdot \left(\sum_{n=1}^{\infty} q(1-q)^{n-1} \sum_{z=1}^{n} \phi_{jp}^{z} \right) \\
&= \lambda_{gp} \cdot \sum_{z=1}^{\infty} \sum_{n=z}^{\infty} q(1-q)^{n-1} \phi_{jp}^{z} \\
&= \lambda_{gp} \cdot \sum_{z=1}^{\infty} \phi_{jp}^{z} (1-q)^{z-1} \sum_{n=z}^{\infty} q(1-q)^{n-z} \\
&= \lambda_{gp} \cdot \sum_{z=1}^{\infty} \phi_{jp}^{z} (1-q)^{z-1}
\end{aligned}$$

where the latter coincides with expression (6) provided in the main text.

Heterogeneous individual investors. Let us consider the case in which individual investors have heterogeneous preferences. Formally, let A_i, δ_i, $\boldsymbol{\beta}_i$ and γ_i be individual specific. Moreover, assume that there is a subset $D \subset J_p$ of default funds and denote by d_i investor i's default fund. Under these assumptions i's demand system is given by

$$\boldsymbol{a}_i(\boldsymbol{f}) = \begin{cases} \boldsymbol{e}_{d_i} & \text{if } i \text{ defaults} \\ \frac{1}{\gamma_i}(I+G)^{-1}(\boldsymbol{\mu}_i - \boldsymbol{f}) & \text{o.w} \end{cases} \tag{33}$$

where to save on notation I defined $\boldsymbol{\mu}_i \equiv W\boldsymbol{\beta}_i + \boldsymbol{\xi}_i$. To obtain the aggregate demand system let us define the following weighted averages

$$\begin{aligned}
\boldsymbol{\mu}_p &\equiv \sum_{i \in I_p} \frac{(1-\delta_i)\gamma_i^{-1}}{\sum_{i \in I_p}(1-\delta_i)\gamma_i^{-1}} \boldsymbol{\mu}_i \\
\gamma_p &\equiv \left(\sum_{i \in I_p} \frac{A_i}{A_p} \frac{1-\delta_i}{\gamma_i} \right)^{-1} \\
\delta_{dp} &\equiv \left(\frac{\sum_{i \in I_{dp}} A_i}{A_p} \right) \sum_{i \in I_{dp}} \frac{A_i}{\sum_{i \in I_{dp}} A_i} \delta_i
\end{aligned}$$

where I_p is the set of plan p investors and I_{dp} is the set of investors that have fund d as default option. Then taking the horizontal sum of the $\boldsymbol{a}_i$ it is easy to check that the plan level demand system is given by

$$\boldsymbol{s}_p(\boldsymbol{f}; \boldsymbol{\eta}_p) = \sum_{d \in D} \delta_{dp} \mathbf{e}_d + \frac{1}{\gamma_p}(I + G_p)^{-1}(\boldsymbol{\mu}_p - \boldsymbol{f}).$$

Overall, the estimated parameters from the aggregate demand system $\boldsymbol{\eta}_p = (\delta_{dp}, \gamma_p, \boldsymbol{\mu}_p)$ are weighted averages of the underlying individual investors parameters.

Derivation of aggregate demand system in (12). To show that equations (10) and

(12) are equivalent it is enough to note that

$$(I - \tilde{G}(I + \tilde{G}'\tilde{G})^{-1}\tilde{G}')(I + G) = \tag{34}$$
$$= (I - \tilde{G}(I + \tilde{G}'\tilde{G})^{-1}\tilde{G}')(I + \tilde{G}\tilde{G}') = \tag{35}$$
$$= I - \tilde{G}(I + \tilde{G}'\tilde{G})^{-1}\tilde{G}' + \tilde{G}\tilde{G}' - \tilde{G}(I + \tilde{G}'\tilde{G})^{-1}\tilde{G}'\tilde{G}\tilde{G}' \tag{36}$$
$$= I + \tilde{G}(I - (I + \tilde{G}'\tilde{G})^{-1} - (I + \tilde{G}'\tilde{G})^{-1}\tilde{G}'\tilde{G})\tilde{G}' \tag{37}$$
$$= I + \tilde{G}(I + \tilde{G}'\tilde{G})^{-1}(I + \tilde{G}'\tilde{G} - I - \tilde{G}'\tilde{G})\tilde{G}' = I \tag{38}$$

Next, I show that $\kappa_{jj} \in (0, 1)$ and $\kappa_{jl} \in (-1, 1)$ for all j and l. To see this, consider note that by construction $\mathcal{K}_x$ and $I - \mathcal{K}_x$ are positive definite matrices. The let e_j be the jth unit vector and note the definition of positive definite matrix implies that

$$\kappa_{jj} = e_j'\mathcal{K}_x e_j > 0 \quad \text{and} \quad 1 - \kappa_{jj} = e_j'(I - \mathcal{K}_x)e_j > 0. \tag{39}$$

Next, take any (j, l) pair with $j \neq l$ and note that

$$\kappa_{jj} + \kappa_{ll} - 2\kappa_{jl} = (e_j - e_l)'\mathcal{K}_x(e_j - e_l) > 0 \tag{40}$$

where the first equality exploits the fact that $\mathcal{K}_x$ is symmetric. From (40) and the fact that $\kappa_{jj} < 1$ for all j, we can conclude that $\kappa_{jl} < 1$. To show that $\kappa_{jl} > -1$ it is enough to repeat the previous argument using $\kappa_{jj} + \kappa_{ll} + 2\kappa_{jl}$.

Derivation of expected κ_{jl} under biased beliefs about q. Suppose that fund j believes that sponsors will include at most one fund per investment category. This means that funds evaluate inclusion probabilities assuming that $q = 1$ and will assign positive probabilities only to menus S that include at most one fund per category.

Denoting by G_S the set of categories included in menu S, by j_g a generic option from category g and by g_j the investment category fund j belongs to, the probability that such menu $S \in \mathcal{S}_{jp}$ is chosen by sponsor p can be factored us

$$\phi_p(S) = \phi_{jp} \prod_{g \in G_S/g_j} \phi_{j_g p} \prod_{j_{g'p}} (1 - \lambda_{j_{g'}p}) \tag{41}$$

where the factorization is a consequence of the fact that inclusion decision are made independently across investment categories.

Next, consider fund j and fund l and note that

$$\bar{\kappa}_{jl} = \sum_{S \in \mathcal{S}_{jp}} \frac{\phi_p(S)}{\phi_{jp}} \kappa_{jl}^S \tag{42}$$
$$= \sum_{S \in \mathcal{S}_{jp}} \left(\prod_{g \in G_S/g_j} \phi_{j_g p} \prod_{j_{g'p}} (1 - \lambda_{j_{g'}p}) \right) \kappa_{jl}^S \tag{43}$$

which does not depend on f_j because inclusion probabilities of competitors funds belonging to different categories do not depend on f_j.

Derivation of equilibrium fees decomposition. Recall the system of Bertrand FOCs'

derived in the main text

$$\bar{\delta}\boldsymbol{e}_d + (I - \tilde{\mathcal{K}})(\tilde{\boldsymbol{\mu}} - \boldsymbol{f}) - \boldsymbol{\iota} - (I - diag(\tilde{\mathcal{K}}))(\boldsymbol{f} - \boldsymbol{c}) = 0 \tag{44}$$

which can be rearranged as

$$2\left(I - \frac{diag(\tilde{\mathcal{K}})}{2} - \frac{\tilde{\mathcal{K}}}{2}\right)\boldsymbol{f} = (I - \tilde{\mathcal{K}})\tilde{\boldsymbol{\mu}} + (I - diag(\tilde{\mathcal{K}}))\boldsymbol{c} - \boldsymbol{\iota}. \tag{45}$$

Next define

$$\tilde{\boldsymbol{\iota}} \equiv \left(I - \frac{diag(\tilde{\mathcal{K}})}{2} - \frac{\tilde{\mathcal{K}}}{2}\right)^{-1}\frac{\boldsymbol{\iota}}{2} \tag{46}$$

and rewrite the system of FOCs as

$$\begin{aligned}
\boldsymbol{f} &= \frac{1}{2}\left(I - \frac{diag(\tilde{\mathcal{K}})}{2} - \frac{\tilde{\mathcal{K}}}{2}\right)^{-1}\left[(I - \tilde{\mathcal{K}})\tilde{\boldsymbol{\mu}} + (I - diag(\tilde{\mathcal{K}}))\boldsymbol{c}\right] - \tilde{\boldsymbol{\iota}} \\
&= \boldsymbol{c} + \frac{1}{2}\left(I - \frac{diag(\tilde{\mathcal{K}})}{2} - \frac{\tilde{\mathcal{K}}}{2}\right)^{-1}(I - \tilde{\mathcal{K}})(\tilde{\boldsymbol{\mu}} - \boldsymbol{c}) - \tilde{\boldsymbol{\iota}} \\
&= \boldsymbol{c} + \frac{1}{2}\left(I - diag(\tilde{\mathcal{K}}) - \frac{\tilde{\mathcal{K}} - diag(\tilde{\mathcal{K}})}{2}\right)^{-1}(I - \tilde{\mathcal{K}})(\tilde{\boldsymbol{\mu}} - \boldsymbol{c}) - \tilde{\boldsymbol{\iota}} \\
&= \frac{\tilde{\boldsymbol{\mu}} + \boldsymbol{c}}{2} - \left(I - diag(\tilde{\mathcal{K}}) - \frac{\tilde{\mathcal{K}} - diag(\tilde{\mathcal{K}})}{2}\right)^{-1}\frac{\tilde{\mathcal{K}} - diag(\tilde{\mathcal{K}})}{2}\frac{\tilde{\boldsymbol{\mu}} - \boldsymbol{c}}{2} - \tilde{\boldsymbol{\iota}} \\
&= \frac{\tilde{\boldsymbol{\mu}} + \boldsymbol{c}}{2} - (I - diag(\tilde{\mathcal{K}}))^{-1/2}\left(I - \frac{G(\tilde{\mathcal{K}})}{2}\right)^{-1}\frac{G(\tilde{\mathcal{K}})}{2}(I - diag(\tilde{\mathcal{K}}))^{1/2}\frac{\tilde{\boldsymbol{\mu}} - \boldsymbol{c}}{2} - \tilde{\boldsymbol{\iota}}
\end{aligned}$$

where

$$G(\tilde{\mathcal{K}}) \equiv (I - diag(\tilde{\mathcal{K}}))^{-1/2}(\tilde{\mathcal{K}} - diag(\tilde{\mathcal{K}}))(I - diag(\tilde{\mathcal{K}}))^{-1/2}. \tag{47}$$

Expression (21) in the main text obtains by setting $diag(\tilde{\mathcal{K}}) = k_0 I$. For the case in which there is a default fund the same steps apply after redefining $\tilde{\boldsymbol{\mu}} = \tilde{\boldsymbol{\mu}} + \bar{\delta}(I - \tilde{\mathcal{K}})^{-1}\boldsymbol{e}_d$.

Estimation algorithm. Before starting the estimation, I draw a vector $\boldsymbol{\nu}_s$ of random taste parameters for $s = 1, ..., S$ simulated sponsors from a normal $N(\mathbf{0}, I)$ and store it. Then the algorithm proceeds as follows.

Step 0: Guess Γ_θ

Step 1: For a given guess of the vector of the mean utility mean utility $\bar{\boldsymbol{v}}^{(k)}$. Compute the following variables for each fund j, market t

1.1 for each simulated sponsor s calculate the following objects

- the probability that j's category g is included by sponsor s

$$\lambda_{gst}(\Gamma_\theta, \bar{\boldsymbol{v}}_t^{(k)}) = \frac{\sum_{l \in g}\exp(\bar{v}_{lt}^{(k)} + \boldsymbol{w}_{lt}'\Gamma_\theta\boldsymbol{\nu}_s)}{1 + \sum_{l \in g}\exp(\bar{v}_{lt}^{(k)} + \boldsymbol{w}_{lt}'\Gamma_\theta\boldsymbol{\nu}_s)}$$

- for $n = \{1, 2, 3\}$ calculate the ranking probabilities

$$\phi^n_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) = \sum_{(j_1,\dots,j_{n-1})\in g/\{j\}} \prod_{n'=1}^{n-1} \frac{\exp(\bar{v}^{(k)}_{j_{n'}t} + \boldsymbol{w}'_{j_{n'}t}\Gamma_\theta \boldsymbol{\nu}_s)}{\sum_{n''=n'}^{N_{gp}} \exp(\bar{v}^{(k)}_{j_{n''}t} + \boldsymbol{w}'_{j_{n''}t}\Gamma_\theta \boldsymbol{\nu}_s)} \cdot \frac{\exp(\bar{v}^{(k)}_{jt} + \boldsymbol{w}'_{jt}\Gamma_\theta \boldsymbol{\nu}_s)}{\sum_{n'''=n}^{N_{gp}} \exp(\bar{v}^{(k)}_{j_{n'''}t} + \boldsymbol{w}'_{j_{n'''}t}\Gamma_\theta \boldsymbol{\nu}_s)}$$

- compute the probability that sponsor s includes fund j

$$\phi^{1:3}_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) = \phi^1_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) + (1-q_t)\phi^2_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) + (1-q_t)^2\phi^3_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t)$$

where q_t is calibrated to match the empirical distribution of the number of options within investment category in market t. The typical values for q_t are of 0.7 or higher so that $(1-q_t)^{n-1}$ decays quite fast in n. In simulations, I find that stopping at $n = 3$ works well in recovering the true parameters. Moreover the computational burden of calculating ϕ^n_{jst} for $n \geq 4$ is non negligible.

1.2 approximate the RHS of (22) with

$$\phi^S_{jt}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) = \frac{1}{S}\sum_{s=1}^{S} \lambda_{gst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t) \cdot \phi^{1:3}_{jst}(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t)$$

Step 2: For each t update the mean utility vector by computing $\bar{\boldsymbol{v}}^{(k+1)}_t$ as

$$\bar{\boldsymbol{v}}^{(k+1)}_t = \bar{\boldsymbol{v}}^{(k)}_t + \log(\hat{\boldsymbol{\phi}}_t(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t)) - \log(\boldsymbol{\phi}^S_t(\Gamma_\theta, \bar{\boldsymbol{v}}^{(k)}_t)) \tag{48}$$

Step 3: For each t Repeat Step 1 and Step 2 until

$$||\bar{\boldsymbol{v}}^{(k+1)}_t - \bar{\boldsymbol{v}}^{(k)}_t||_\infty < \epsilon \tag{49}$$

for some tolerance level ϵ.

Step 4: Recover sponsors' preference parameters $\boldsymbol{\mu}_\theta$ that enter sponsors' utility linearly

$$\boldsymbol{\mu}_\theta(\Gamma_\theta) = (W'Z(Z'Z)^{-1}Z'W)^{-1}W'Z(Z'Z)^{-1}Z'\boldsymbol{v}(\Gamma_\theta) \tag{50}$$

where W is a matrix of funds' characteristics with number of rows equal to the total number of observations $\bar{N} = \sum_t N_t$ and number of columns equal to the number of characteristics and Z is including both excluded and included instruments.

Step 5: Recover demand residuals

$$\boldsymbol{\zeta}(\Gamma_\theta) = \boldsymbol{v}(\Gamma_\theta) - W\boldsymbol{\mu}_\theta \tag{51}$$

and compute the GMM norm

$$\boldsymbol{\zeta}(\Gamma_\theta)'Z\Omega(\Gamma_\theta)Z'\boldsymbol{\zeta}(\Gamma_\theta) \tag{52}$$

where $\Omega(\Gamma_\theta) = (Z'Z)^{-1}$ in the first GMM estimation step and then is updated to $\Omega(\Gamma_\theta) = Z'diag(\boldsymbol{\zeta}(\Gamma_\theta)^2)Z = \sum_{j,t}\zeta^2_{jt}(\Gamma_\theta)\boldsymbol{Z}_{jt}\boldsymbol{Z}'_{jt}$.

Interior equilibrium existence and uniqueness. Consider first the case in which funds know with certainty which plan menu will include them. Formally, this means that $\phi_{jp} = 1$ when p includes fund j and zero otherwise which further implies that the $\boldsymbol{\iota}$ in equation (21) equals zero. Given this, the system of funds' best replies becomes linear in $\boldsymbol{f}$:

$$(I - \tilde{\mathcal{K}})\boldsymbol{f} + (I - diag(\tilde{\mathcal{K}}))\boldsymbol{f} = (I - \tilde{\mathcal{K}})\tilde{\boldsymbol{\mu}} + (I - diag(\tilde{\mathcal{K}}))\boldsymbol{c} \tag{53}$$

because inclusion probabilities do not depend on fees anymore and with that also $\mathcal{K}$, $\tilde{\boldsymbol{\mu}}$ do not depend on fees. A well-defined solution is then guaranteed to exist as long as $(I - \tilde{\mathcal{K}})$ is invertible. Moreover, linearity would also imply that such solution is unique.

What we do not know is whether this solution is such that equilibrium fees are non-negative. In what follows I show that the following dominance-diagonal condition

$$(1 - \tilde{\kappa}_{jj})(\tilde{\mu}_j - c_j) > \sum_{k \neq j} |\tilde{\kappa}_{jk}|(\tilde{\mu}_k - c_k) \quad \text{all } j \tag{54}$$

implies that the system of best replies is a self-map over the interior of the set $\times_{j\in\{1,\dots,J\}}[c_j, \tilde{\mu}_j]$ which ensures that equilibrium fees are positive and above marginal costs. I start by defining the following linear operator $T : \mathbb{R}^J \to \mathbb{R}^J$ whose j-th component is fund j's best reply

$$T_j(\boldsymbol{f}) \equiv \frac{1}{2}\left[\tilde{\mu}_j - \sum_{k\neq j} \frac{\tilde{\kappa}_{jk}}{1 - \tilde{\kappa}_{jj}}\mu_k + c_j + \sum_{k\neq j} \frac{\tilde{\kappa}_{jk}}{1 - \tilde{\kappa}_{jj}} f_k\right]. \tag{55}$$

Assumption (54) implies that T is a self-map in the interior of $\times_{j\in\{1,\dots,J\}}[c_j, \tilde{\mu}_j]$. To see this take any $\boldsymbol{f} \in \times_{j\in\{1,\dots,J\}}[c_j, \tilde{\mu}_j]$ and note that

$$c_j < T_j(\boldsymbol{f}) < \tilde{\mu}_j \tag{56}$$

$$\Leftrightarrow \left|\sum_{k\neq j} \frac{\tilde{\kappa}_{jk}}{1 - \tilde{\kappa}_{jj}} \frac{\tilde{\mu}_k - c_k}{\tilde{\mu}_j - c_j} \frac{\tilde{\mu}_k - f_k}{\tilde{\mu}_k - c_k}\right| < 1 \tag{57}$$

$$\Leftrightarrow \sum_{k\neq j} \frac{|\kappa_{jk}|}{1 - \tilde{\kappa}_{jj}} \frac{\tilde{\mu}_k - c_k}{\tilde{\mu}_j - c_j} \frac{\tilde{\mu}_k - f_k}{\tilde{\mu}_k - c_k} < 1 \tag{58}$$

is always satisfied when (54) holds and $\boldsymbol{f} \in \times_{j\in\{1,\dots,J\}}[c_j, \tilde{\mu}_j]$. Because T maps a closed and bounded set into itself, Bower fixed point theorem implies that there exists an $\boldsymbol{f}^* \in \times_{j\in\{1,\dots,J_p\}}[c_j, \mu_j]$ such that $T(\boldsymbol{f}^*) = \boldsymbol{f}^*$. Moreover, from (58) we know that such fixed point is interior and unique, because T is linear.

We can also show that in this interior equilibrium each fund manages a positive amount of asset. To see this, consider fund j's first order condition evaluated at the optimum

$$s_j(\boldsymbol{f}^*) - (1 - \tilde{\kappa}_{jj})(f_j^* - c_j) = 0 \tag{59}$$

which implies that

$$s_j(\boldsymbol{f}^*) = (1 - \tilde{\kappa}_{jj})(f_j^* - c_j) > 0 \tag{60}$$

where the latter inequality holds because we just proved that $f_j^* > c_j$ for all j and

$(1-\tilde{\kappa}_{jj})>0$ follows from the fact that $I-\mathcal{K}$ is positive definite and that $\tilde{\kappa}_{jj}$ is just an average of the κ_{jj} across plans:

$$\tilde{\kappa}_{jj}=\bar{\phi}_j^{-1}\int \mathbf{1}\{j\in S_p\}(1-\delta_p)\gamma_p^{-1}A_p\kappa_{jj}^{S_p}dF_p$$

with

$$\bar{\phi}_j=\int \mathbf{1}\{j\in S_p\}(1-\delta_p)\gamma_p^{-1}A_p dF_p. \tag{61}$$

Lastly I show that the equilibrium is stable. To see this define the following variable

$$\hat{f}_j\equiv\frac{f_j-c_j}{\tilde{\mu}_j-c_j} \tag{62}$$

and note that fund j best reply can be rewritten as

$$\hat{f}_j=\frac{1}{2}\left[1-\sum_{k\neq j}\frac{\tilde{\kappa}_{jk}}{1-\tilde{\kappa}_{jj}}\frac{\tilde{\mu}_k-c_k}{\tilde{\mu}_j-c_j}\left(1-\frac{f_k-c_k}{\tilde{\mu}_k-c_k}\right)\right] \tag{63}$$

Defining the linear mapping on the above RHS as $\hat{T}:R^J\to R^J$, it can be shown that this mapping is a self-map into $[0,1]^J$ under assumption (54). Additionally, this mapping is a contraction in the L_∞ norm.

$$\begin{aligned}
||\hat{T}(\hat{\boldsymbol{f}}_1)-\hat{T}(\hat{\boldsymbol{f}}_0)||_\infty &= \max_j|\hat{T}_j(\hat{\boldsymbol{f}}_1)-\hat{T}_j(\hat{\boldsymbol{f}}_0)| && (64)\\
&= \max_j\left|\sum_{k\neq j}\frac{\kappa_{jk}}{1-\kappa_{jj}}\frac{\mu_k-c_k}{\tilde{\mu}_j-c_j}(\hat{f}_{1k}-\hat{f}_{0k})\right| && (65)\\
&\leq \max_j\sum_{k\neq j}\frac{|\kappa_{jk}|}{1-\kappa_{jj}}\frac{\mu_k-c_k}{\mu_j-c_j}|\hat{f}_{1k}-\hat{f}_{0k}| && (66)\\
&< \max_k|f_{1k}-f_{0k}| && (67)
\end{aligned}$$

which ensures that the unique equilibrium is also stable.

Next, I consider the case in which sponsor preferences are homogeneous $\theta_p\equiv\theta$ and funds do not know with certainty whether or not they will be included in sponsors' retirement menus (e.g., $\phi_j\in(0,1)$). For simplicity, I also assume that there is only one recordkeeper or equivalently that all recordkeerpers have the same network of funds. Under this assumptions, I will show that a Nash-Bertrand equilibrium exists when funds believe that sponsors include at most one fund per category and the previous dominance diagonal condition holds.

To start with, note that fund j problem in any given period simplifies to

$$\max_{f_j}\quad P\cdot(f_j-c_j)\cdot\phi_j(\boldsymbol{f}_j;\boldsymbol{\theta})\int s_{jp}(\boldsymbol{f};\boldsymbol{\eta}_p)A_p dF(\boldsymbol{\eta}_p,A_p) \tag{68}$$

where fund j expected portfolio share is given by

$$\begin{aligned} s_{jp}(\boldsymbol{f};\eta_p) &= \sum_{S\in\mathcal{S}_j} \tilde{\gamma}_p \frac{\phi(S;\boldsymbol{\theta})}{\phi_j(\boldsymbol{f};\boldsymbol{\theta})} \left[(1-\kappa^S_{jj})(\mu_{jp}-f_j) - \sum_{l\neq j, l\in S} \kappa^S_{jl}(\mu_{lp}-f_l)\right] \\ &= \tilde{\gamma}_p(1-\bar{\kappa}_{jj})(\mu_{jp}-f_j) - \tilde{\gamma}_p \sum_{l\neq j}\sum_{S\in\mathcal{S}_j} \kappa^S_{jl}\mathbf{1}\{l\in S\}\frac{\phi(S;\boldsymbol{\theta})}{\phi_j(\boldsymbol{\theta})}(\mu_{lp}-f_l) \\ &= \tilde{\gamma}_p(1-\bar{\kappa}_{jj})(\mu_{jp}-f_j) - \tilde{\gamma}_p \sum_{l\neq j} \phi_l \mathbb{E}[\bar{\kappa}^S_{jl}|j,l\in S](\mu_{lp}-f_l) \\ &= \tilde{\gamma}[I-\bar{\mathcal{K}}]'_j(\boldsymbol{\mu}_p-\boldsymbol{f}) \end{aligned}$$

with

$$\bar{\kappa}_{jl} \equiv \begin{cases} \mathbb{E}[\bar{\kappa}^S_{jj}|j\in S] & \text{if } j=l \\ \phi_l(\boldsymbol{f};\boldsymbol{\theta})\mathbb{E}[\bar{\kappa}^S_{jl}|j,l\in S] & \text{if } j\neq l. \end{cases}$$

Overall, fund j pricing problem can be written more compactly as

$$\max_{f_j} \quad P\cdot(f_j-c_j)\cdot\phi_j(\boldsymbol{f}_j;\boldsymbol{\theta})\cdot[I-\bar{\mathcal{K}}]'_j(\boldsymbol{\mu}-\boldsymbol{f})\cdot\tilde{\gamma}\bar{A} \tag{69}$$

where for simplicity I assumed that investors preferences are homogeneous too.[54]

In what follows, I first prove a lemma that provides conditions ensuring that funds' problem is concave and then show that funds' objective satisfies such conditions under the previous assumptions. Equilibrium existence will then follow directly from Kakutani fixed point theorem.

Lemma 1 *Let $\pi(f)$ be a continuous and strictly concave function $\pi''(f)<0$ that is uniquely maximized at f^* and is such that $\pi(f^*)>0$. Let $\phi(f)$ a continuous and decreasing function such that*

$$\phi(f)\in(0,1) \tag{70}$$
$$\phi'(f) = -\phi(f)(1-\phi(f)) < 0 \tag{71}$$

then the function $\phi(f)\pi(f)$ is concave on a compact set $[\underline{f},\overline{f}]$ with unique interior maximizer $f^{**}\leq f^*$.

Proof: Because π is continuous and $\pi(f^*)>0$, we can construct $[\underline{f},\overline{f}]$ such that $\pi(f)>0$ for all $f\in[\underline{f},\overline{f}]$ and $f^*\in[\underline{f},\overline{f}]$.

Next suppose there exists a $f^{**}\in(\underline{f},\overline{f})$ such that

$$\phi'(f^{**})\pi(f^{**}) + \phi(f^{**})\pi'(f^{**}) = 0 \tag{72}$$

which can be rearranged as

$$-\frac{\phi'(f^{**})}{\phi(f^{**})} = \frac{\pi'(f^{**})}{\pi(f^{**})} \tag{73}$$

[54] This assumption is irrelevant for the existence result.

Taking the second order condition

$$\phi''(f^{**})\pi(f^{**}) + 2\phi'(f^{*})\pi'(f^{*}) + \phi(f^{**})\pi''(f^{**}) < 0 \tag{74}$$

The last term of the above is negative by assumption. The sum of the first two terms is also negative:

$$\phi''(f^{**})\pi(f^{**}) + 2\phi'(f^{*})\pi'(f^{*}) < 0 \tag{75}$$

$$\Leftrightarrow \frac{\phi''}{\phi'} + 2\frac{\pi'}{\pi} > 0 \tag{76}$$

$$\Leftrightarrow \frac{\phi''}{\phi'} - 2\frac{\phi'}{\phi} > 0 \tag{77}$$

$$\Leftrightarrow -(1-2\phi) + 2(1-\phi) = 1 > 0 \tag{78}$$

where the latter equivalence uses the fact that $\phi'' = -\phi'(1-2\phi)$. This shows that if an interior f^{**} that satisfies the necessary FOC exists then it is always a maximum which implies that $\phi(f)\pi(f)$ is concave on $[\underline{f}, \overline{f}]$.

Lastly, we can show that such p^{**} indeed exists and is interior. To see this evaluate the FOC at $\underline{f}$ and note that under the above assumptions the following

$$\phi'(\underline{f})\pi(\underline{f}) + \phi(\underline{f})\pi'(\underline{f}) > 0 \tag{79}$$

holds by choosing $\underline{f}$ sufficiently low (for instance choosing $\underline{f}$ = marginal cost such that $\pi(\underline{f}) = 0$) and by noting that $\pi'(\underline{f}) > 0$ because $\underline{f} < f^{*}$. Then evaluate the FOC at $\overline{f}$ and note that

$$\phi'(\overline{f})\pi(\overline{f}) + \phi(\overline{f})\pi'(\overline{f}) < 0 \tag{80}$$

which implies, by continuity that there exists an interior f^{**} that satisfies the FOC. Moreover, it must be the case that $f^{**} < f^{*}$.

Overall, conditions all conditions of the lemma hold and we can be assured that fund j objective in (69) is concave in f_j which ensures that funds' best replies $f_j(f_{-j})$ are proper functions. Thus, all conditions of Kakutani fixed point theorem are satisfied and a Nash equilbrium exists (see MWG chp. 8 Proposition 8.D.3).

Funds' maximization problem in (69) satisfies the conditions in the previous lemma after defining $\pi(f_j) \equiv (f_j - c_j) \cdot [I - \bar{\mathcal{K}}]'_j(\boldsymbol{\mu} - \boldsymbol{f})$. First note that, because funds' believe that sponsors include at most one fund per category we have that

$$\phi_j(f_j) = \lambda_g \phi_j^1 = \frac{\exp(V_j(\boldsymbol{\theta}))}{1 + \sum_l \exp(V_l(\boldsymbol{\theta}))}$$

which is decreasing in f_j and such that

$$\phi'_j = -\theta_f \phi_j (1 - \phi_j) < 0.$$

Next, note that because funds' believe that $q = 1$, the matrix $\mathcal{K}$ does not depend on f_j as I showed in previous derivations. This in turn implies that $\pi(f_j)$ is quadratic in f_j

and thus concave f_j if and only if

$$(1 - \bar{\kappa}_{jj}) > 0.$$

The latter holds because $\bar{\kappa}_{jj} = \mathbb{E}[\kappa^S_{jj}|j, l \in S]$ and $\kappa^S_{jj} \in (0, 1)$ for any S as I showed in previous derivations. Because π is globally concave in f_j it admits a unique maximum, f^*_j. Moreover we have $f^*_j > c_j$ whenever the previous dominance diagonal holds:

$$(1 - \bar{\kappa}_j)(\mu_j - c_j) > \sum_{l \neq j} |\bar{\kappa}_{jl}|(\mu_l - c_l) \tag{81}$$

and $f_l \in [c_l, \mu_l]$ for all l. To see this note that the above condition implies that

$$\begin{aligned}(1 - \bar{\kappa}_j)(\mu_j - c_j) &> \sum_{l \neq j} |\bar{\kappa}_{jl}|(\mu_l - c_l) \\ &\geq \sum_{l \neq j} |\bar{\kappa}_{jl}|(\mu_l - f_l) \\ &> \sum_{l \neq j} \bar{\kappa}_{jl}(\mu_l - f_l)\end{aligned}$$

where the last two inequalities holds whenever $f_l \in [c_l, \mu_l]$. But note that the previous inequality corresponds to fund j FOC (when maximizing π) evaluated at $f_j = c_j$

$$\begin{aligned}\frac{\partial \pi(f_j)}{\partial f_j} &= (1 - \bar{\kappa}_{jj})(\mu_j - f_j) - \sum_{l \neq j} \bar{\kappa}_{jl}(\mu_l - f_l) - (1 - \bar{\kappa}_{jj})(f_j - c_j)\Big|_{f_j = c_j} \\ &= (1 - \bar{\kappa}_j)(\mu_j - c_j) - \sum_{l \neq j} \bar{\kappa}_{jl}(\mu_l - f_l) > 0.\end{aligned}$$

Then it must also be the case that $\pi(f^*_j) > 0$, if not setting $f_j = c_j$ would lead to higher π which would be a contradiction.

Microfoundation of the distribution of the number of options. In what follows I offer a simple microfoundation for the distribution of the number of options sponsors include in any given category building on the Stigler (1961) simultaneous search model.

Sponsor p first commits to include n investment options in investment category g and then conditional on n, selects the n options providing her with the n highest utilities. I assume that sponsors choose n before observing their random utility shocks ε_{jp} but knowing the mean utility of each option $V_j(\boldsymbol{\theta}_p)$. From this perspective the utility sponsor p derives from option j is distributed as

$$u_{jp} \sim T1EV(V_j). \tag{82}$$

I assume that sponsors incur a cost $c(n)$ for including n options which is increasing, convex in n and is such that $c(1) < \mathbb{E}[u_{j_1 p}]$. The benefit from choosing n options is given by

$$\sum_{s=1}^{n} \mathbb{E}[u_{j_s p}] \tag{83}$$

where j_s is the option that provides the sth highest utility. Overall, sponsor p problem

becomes

$$\max_{n\geq 1} \sum_{s=1}^{n} \mathbb{E}[u_{j_s p}] - c(n). \tag{84}$$

The are two main differences between this model and the Stigler (1961) model. First, in this case the agent commits to consume n options whereas in Stigler (1961) the agent commits to sample n options and among those to consume the one with the highest utility. Second, in this model after committing to n, sponsor p observes the realized utility of all options available but has chosen to only consume the n highest whereas in the Stigler (1961) model the agent observes the utility realizations of the searched options only, and among those selects the highest.

The solution to the above problem is given by the n^* such that the marginal benefit from chosing to include $n^* + 1$ options is lower than the change in the cost

$$\mathbb{E}[u_{j_{n^*+1}}] \leq c(n^* + 1) - c(n^*). \tag{85}$$

Heterogeneity in the cost of adding options c_p or in the benefits u_{jp} across sponsors would produce a different n_p^* for each sponsors. In the data such distribution of n^* corresponds to the one plotted in Figure A1 suggesting that most sponsors do not include more than one option and that the likelihood decreases geometrically in the number of options.

In estimation I do not attempt to estimate the distribution of costs c_p that matches the observed distribution in Figure A1 because it would complicate substantially the estimation of sponsor preferences. First, it would require estimating such distribution at each iteration of the estimation algorithm because the optimal number of options n^* itself depends on sponsors preference parameters. Second, it would require finding a solution to problem (84) which is non-convex without further restrictions.

To keep estimation tractable I instead model sponsors' choice of n as a random draw from the empirical distribution which I parametrize as geometric with parameter q. In estimation I allow for such distribution to be heterogeneous at the recordkeeper-year-category level. In general, the distribution of the number of options included within each category looks similar to the one in Figure A1 for many cuts of the data I have considered.

Derivation of investor surplus. Consider active investor i in plan p. Investor i has preferences over its retirement portfolio allocation $\boldsymbol{a}_i$ given by

$$u_i(\boldsymbol{a}_i) = \boldsymbol{a}_i'(\boldsymbol{\mu} - \boldsymbol{f}) - \frac{\gamma}{2}\boldsymbol{a}_i' V \boldsymbol{a}_i$$

where

$$V \equiv I + X_{(2)}X_{(2)}'.$$

Investor i's demand implied by the previous problem is

$$\boldsymbol{a}_i(\boldsymbol{f}) = \frac{1}{\gamma} V^{-1}(\boldsymbol{\mu} - \boldsymbol{f}).$$

Next, combine the expression for $\boldsymbol{a}_i$ with the expression for $u_i(\boldsymbol{a}_i)$ as follows

$$\begin{aligned} u_i(\boldsymbol{a}_i) &= \boldsymbol{a}_i'(\boldsymbol{\mu} - \boldsymbol{f}) - \frac{\gamma}{2}\boldsymbol{a}_i' V \boldsymbol{a}_i \\ &= \boldsymbol{a}_i(\boldsymbol{f})'(\boldsymbol{\mu} - \boldsymbol{f}) - \frac{1}{2}\boldsymbol{a}_i(\boldsymbol{f})'(\boldsymbol{\mu} - \boldsymbol{f}) \\ &= \frac{1}{2}\boldsymbol{a}_i(\boldsymbol{f})'(\boldsymbol{\mu} - \boldsymbol{f}), \end{aligned}$$

which is the measure of investors' surplus provided in the main text.

Example of category-based correlation structure. Consider a plan menu with 8 assets classified in the following four investment categories 'Equity-Growth', 'Equity-Value', 'Bond-Government' and 'Bond-Corporate'. Also assume that there are two assets for each of the four categories.

The vector of characteristic for a given asset j, $\tilde{\boldsymbol{g}}_j$, has six elements corresponding to the 1st level characteristics (Equity, Bond) and 2nd level characteristics (Equity-Growth, Equity-Value, Bond-Gov, Bond-Corp). If asset j is an Equity-Value fund its vector of characteristics is given by:

$$\tilde{\boldsymbol{g}}_j = (1, 0, 0, 1, 0, 0)'. \tag{86}$$

With the assumption that there are two assets within each category the 8×8 outer-product matrix G_p is given by:

$$G_p = \tilde{G}_p \tilde{G}_p' = \begin{bmatrix} 2 & 2 & 1 & 1 & 0 & 0 & 0 & 0 \\ 2 & 2 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 2 & 2 & 0 & 0 & 0 & 0 \\ 1 & 1 & 2 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 2 & 1 & 1 \\ 0 & 0 & 0 & 0 & 2 & 2 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 2 & 2 \\ 0 & 0 & 0 & 0 & 1 & 1 & 2 & 2 \end{bmatrix}$$

Funds' investment category classifications capture cross-substitution patterns between assets. In this example investors treat Equity and Bond assets as independent and within Equity, Growth and Value assets as less substitutable than the two Growth or the two Value assets.

C Turnover ratios & identification

To consistently estimate sponsors' and investors' preferences I instrument funds' fees with funds' turnover ratios which capture trading costs that are typically pass on to investors through fees. In this appendix, I first provide more details on how a fund's turnover is computed and, after that, I discuss a possible threat to identification. Lastly, to motivate the relevance of the instrument, I provide an illustrative example of how funds' turnover might affect funds' fees.

C.1 Funds' turnover ratios.

I obtain data on funds' turnover from CRSP, which reports it at fiscal year frequency. Turnover for fund j in year t is defined as

$$\text{turnover}_{jt} = \frac{\min(\text{buys}_{jt}, \text{sells}_{jt})}{\text{Average TNA}_{jt}}$$

where the numerator is the smaller of the funds' total purchases and sales over fiscal year t and the denominator is the average total net asset value (TNA) in year t. This measure is the one that the SEC requires funds' to report each year.

A key advantage of this measure is that by taking the minimum between sales and purchases it excludes turnover arising from trading activities triggered by persistent inflows or outflows. For example, if a fund experiences substantial inflows most of the trading activity will be implemented to buy more securities or increase current portfolio positions. In this case though, because of the min() in the numerator, the turnover reported will capture the sales and not the purchases. Similarly, if a fund experiences substantial outflows, the turnover measure will likely pick up the fund's purchases. Overall, because flows are notoriously persistent, this measure of turnover is largely immune to flows and instead will capture discretionary trading decision from funds' managers (Pástor, Stambaugh and Taylor (2017)).

From an identification perspective, this property of the turnover measure is particularly appealing because it makes it mechanically less dependent of persistent demand shocks. Nonetheless, the measure is not completely independent of demand shocks that generate non-persistent inflows or outflow and, as such, may be a non valid instrument if it correlates with some driver of funds' flows.

C.2 Identification: turnover vs. performance

Funds' performance is a well-known driver of funds' flows (Chevalier and Ellison (1997)) and recent research suggests that active funds achieve better performance when they trade more and have higher turnover ratio (Pástor, Stambaugh and Taylor (2017)).[55] If investors' chase performance and turnover determines or is correlated with performance then the exclusion restriction I employ to identify sponsors' preferences would be violated.

To reduce the concern about this potential identification threat I check whether my instrument i.e., the residual turnover after absorbing funds' brand, year, category and

[55]The turnover-performance relationship is non-significant for passive funds. Also, for active funds, the relationship is stronger over time rather than across funds.

	turnover ratio	expense ratio	alpha	MS adj. return	BS adj. return
turnover ratio	1.00	0.16	-0.00	0.00	-0.02
expense ratio	0.16	1.00	0.01	-0.01	0.01
alpha	-0.00	0.01	1.00	0.49	0.54
MS adj. return	0.00	-0.01	0.49	1.00	0.69
BS adj. return	-0.02	0.01	0.54	0.69	1.00

Table C1: Correlation table of instrument (turnover ratio) with fund performance measures. Turnover ratio and expense ratios are residuals after absorbing funds' brand, year and category fixed effects. Performance measures are yearly-demeaned.

passive fixed effects, shows significant correlation with some measures of investment performance and find that this is not the case.

Table C1 presents a correlation table between the residualized turnover (i.e., the instrument), the residualized expense ratio (i.e., the endogenous variable) and three measures of investment performance; (i) a fund's (gross of fees) alpha from a three Fama-French factor regression plus Momentum, (ii) a BrightScope-category-adjusted measure of performance computed as the difference between a fund's gross return and the return of the corresponding BrightScope category and (iii) a Morningstar-category-adjusted measure of performance computed as the difference between a fund's gross return and the return of the corresponding Morningstar-category-adjusted.

Turnover and expense ratio exhibit a positive correlation confirming that the instrument has a strong first stage. Conversely, turnover does not seem to be correlated with any of the performance measure considered, alleviating the concern of a possible violation of the exclusion restriction.

To check the statistical significance of these correlations I present a series of binscatter plots where I regress the instrument on the above mentioned measures of performance and on the endogenous variable. Figure C1 shows the strong and significant correlation between turnover and expense ratio. The other set of figures (C2, C3, C4) instead show that there is no significant correlation between the instrument and investment performance. Lastly, in Figures (C5, C6, C7) I check whether turnover correlates with performance in the following year and still find no evidence of a significant correlation, again alleviating the concern of a possible violation of the exclusion restriction.

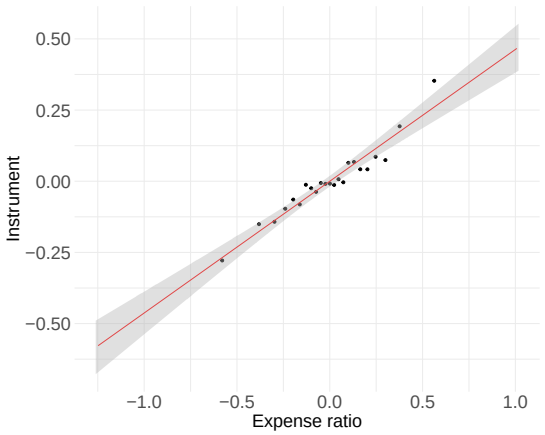

Figure C1: Expense ratio is residual expense ratio after absorbing funds' brand, year and category fixed effects. Instrument is residual turnover ratio.

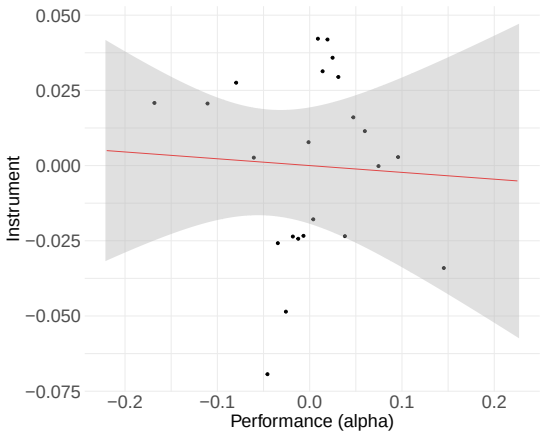

Figure C2: Performance is yearly-demeaned alpha from 3 FF factors plus Momentum. Instrument is residual turnover.

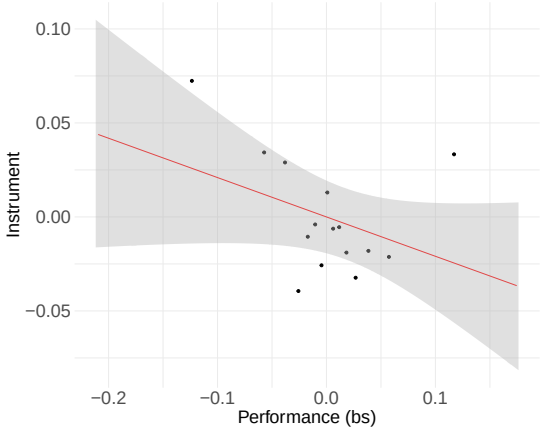

Figure C3: Performance is yearly-demeaned (gross) return relative to BrightScope category return. Instrument is residual turnover.

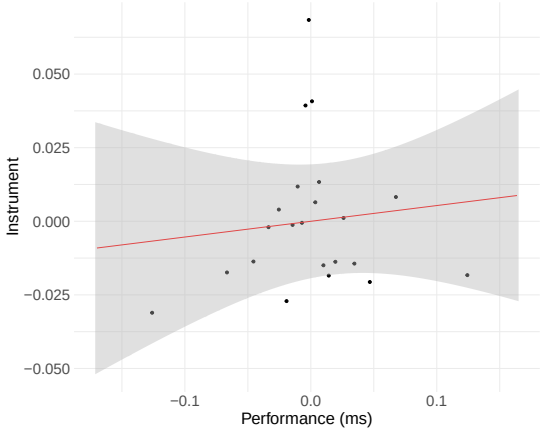

Figure C4: Performance is yearly-demeaned (gross) return relative to Morning Star category return. Instrument is residual turnover.

C.3 How does turnover affect fees?

Figure (C1) provides evidence of a strong relationship between turnover (i.e., the instrument) and fees (i.e., the endogenous variable). In this section, I show how a simple model in which funds' maximize their fee revenue requires fund managers to pass trading costs onto investors via higher fees.

Following Pástor, Stambaugh and Taylor (2020), consider fund j who optimally chooses its fees f_j to maximize its dollar profit:

$$\max_{f_j}(f_j - c_j(t_j))A_j(f_j, t_j)$$

where A_j is the dollar AUM of the fund, c_j is the marginal cost of operating the fund and t_j is the fund turnover. In a world in which investors chase performance it is reasonable to assume that a fund's AUM depend on the level of fees and turnover. For example,

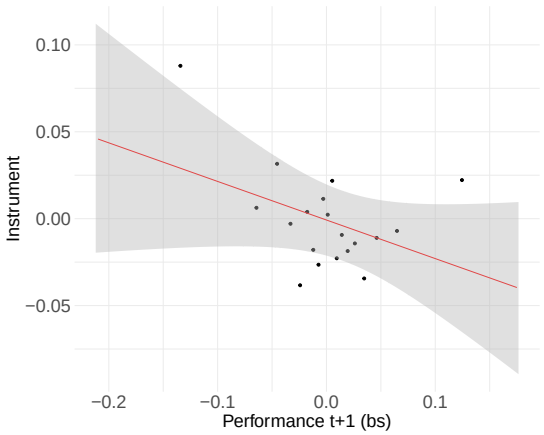

Figure C5: Performance is next period yearly-demeaned (gross) return relative to BrightScope category return. Instrument is residual turnover.

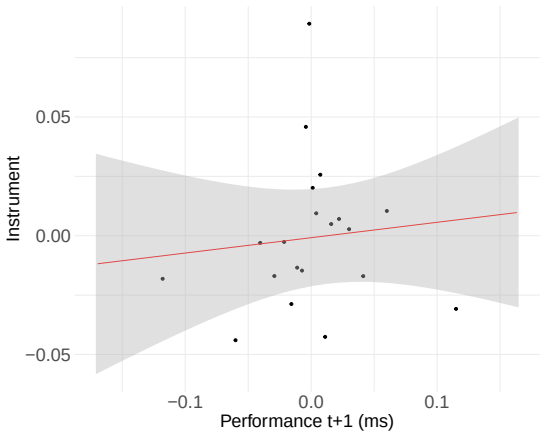

Figure C6: Performance is next period yearly-demeaned (gross) return relative to Morning Star category return. Instrument is residual turnover.

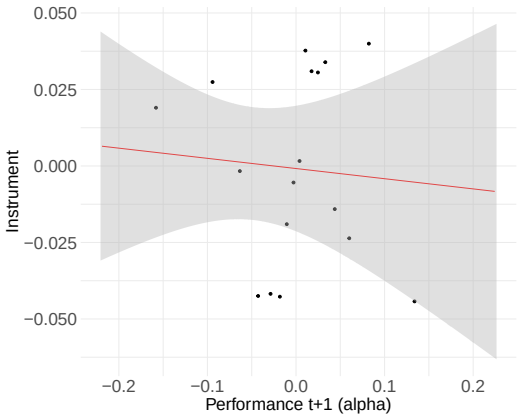

Figure C7: Performance is next period yearly-demeaned alpha from 3 FF factors plus Momentum. Instrument is residual turnover.

when investors are rational and supply capital perfectly elastically a la Berk and Green (2004), the AUM of fund j must be such that investors expect zero returns net of fees and trading costs:

$$\mu_j(t_j) - f_j - q_j(A_j, t_j) = 0 \tag{87}$$

where μ_j is fund j expected return gross of fees and q_j is a per-dollar trading cost. The return generated by the fund may depend on its turnover if, for instance, higher skilled managers trade more (Pástor, Stambaugh and Taylor (2017)). A fund trading cost might depend on its size (Berk and Green (2004)) and on its portfolio turnover (Pástor, Stambaugh and Taylor (2020)). Overall, investors' behaviour (i.e., equation (87)) will determine fund j's demand $A_j(f_j, t_j)$ as a function of fees and turnover.

Funds' turnover ratio can also affect its operating costs. For example, funds that trade more might incur in higher operating costs (e.g., higher more research analysts) and may need to implement more rewarding compensation structures for their portfolio managers which force the fund to charge higher advisory fees to its investors (Ma, Tang and Gomez (2019)).

The optimal fee charged by fund j will then depend on a fund turnover or expected turnover t_j

$$f_j = c_j(t_j) + \frac{A_j(f_j, t_j)}{\partial A_j(f_j, t_j)/\partial f_j}.$$

Overall, funds' optimizing behaviour explains why we observe a positive relationship between trading costs and fees (Pástor, Stambaugh and Taylor (2020)). Funds' managers pass trading costs to investors to maximize dollar profits which provides a rationale for using turnover as a cost-shifter instrument for fees.

D Nesting investors & sponsors preferences

In Section 4 I assumed that sponsors have their own preference parameters over funds' attributes. When making their plan inclusion decisions sponsors evaluate the suitability of each fund based of their preference parameters which I denoted by $\boldsymbol{\theta}_p$. In this appendix I illustrate how sponsors' estimated parameters can be interpreted as a weighted average of what I refer to as sponsors' 'true' preference parameters, denoted by $\boldsymbol{\theta}_p^s$ and investors preference parameters denoted by $\boldsymbol{\theta}_p^i$.

To this end, suppose sponsor p random utility from including fund j is given by

$$u_{jp} = \chi V_{jp}^{\text{sponsor}}(\boldsymbol{\theta}_p^s) + (1-\chi)V_{jp}^{\text{investor}}(\boldsymbol{\theta}_p^i) + \zeta_j + \varepsilon_{jp}$$

where the weight $(1-\chi) \in (0,1)$ captures how much sponsors account for investors preferences when making their plan inclusion decisions.

As before, let us assume that sponsors mean utility $V_{jp}^{\text{sponsor}}(\boldsymbol{\theta}_p^s)$ is linear in funds' characteristics, formally $V_{jp}^{\text{sponsor}} = \boldsymbol{w}_{jp}'\boldsymbol{\theta}_p^s$. The parameter vector $\boldsymbol{\theta}_p^s$ captures how much a non-altruistic sponsor (e.g., $\chi = 1$) would value funds' attributes. For example, a non-benevolent sponsor might prefer to include expensive funds (i.e., its 'true' fee parameter is positive $\theta_{pf}^s > 0$) to maximize its recordkeeper revenues from indirect compensations (Bhattacharya and Illannes (2022)). If sponsors internalize their investors utility then sponsors' parameters estimated starting from the preference specification in (1) will also reflect part of investors preferences. To see this suppose that investors mean utility $V_{jp}^{\text{investor}}(\boldsymbol{\theta}_p^i)$ is also linear in funds' characteristics i.e., $V_{jp}^{\text{investor}}(\boldsymbol{\theta}_p^i) = \boldsymbol{w}_{jp}'\boldsymbol{\theta}_p^i$. Then, it is straightforward to see that the parameter $\boldsymbol{\theta}_p$ we estimated in the main text is a weighted average of sponsors' and investors' preferences

$$u_{jp} = \boldsymbol{w}_{jp}' \underbrace{(\chi\boldsymbol{\theta}_p^s + (1-\chi)\boldsymbol{\theta}_p^i)}_{\equiv\boldsymbol{\theta}_p} + \zeta_j + \varepsilon_{jp}. \tag{88}$$

At this point there are two challenges. First, investors' indirect utility is more complex than the simple linear specification I used just above. Second, it is unclear how we can identify and estimate χ. In what follows, I explain why I abstract from the first challenge and, under the assumption that investors utility is linear, I will illustrate what restrictions are needed to identify χ and I will provide an estimate of it.

Assuming that investors indirect utility is linear in funds' characteristics is not consistent with the quadratic type of preferences I specified when defining investors' portfolio problem. Under those type of preferences investors indirect utility depends on the characteristics of all the funds included in the plan menu because investors optimally diversify across the funds available. In other words, investors indirect utility depends on the menu sponsors choose on their behalf. For this reason, nesting investors' utility into sponsors random utility is not straightforward because it would require sponsors taking expectation over all possible menus that could be chosen. Thus, I will take a different approach and specify the utility investors derive from having fund j in their menu as a linear function of j characteristics. This would be consistent with a demand model, like Berry (1994), in which plan investors make a discrete choice among the options available in their plan, which under the appropriate coefficient restrictions, would be observationally equivalent to the asset demand derived from mean-variance preferences (Koijen and Yogo (2019)).

With the assumption that V_{jp}^{investor} is linear in funds characteristics we can identify

	Preference parameters	
	Active Investors	All Investors
χ	0.749	0.749
Expense ratio (f, bp.)		
θ_{pf}	-0.021	-0.021
θ^i_{pf}	-0.039	-0.025
θ^s_{pf}	-0.019	-0.019
Affiliation (a)		
θ_{pa}	0.823	0.823
θ^i_{pa}	0.276	0.174
θ^s_{pa}	1.040	1.040
Return (r, pp.)		
θ_{pr}	0.004	0.004
θ^i_{pr}	0.025	0.016
θ^s_{pr}	0	0

Table D1: Estimated investors and sponsors parameters assuming investors make their portfolio choice according to a discrete choice demand with linear random utility.

and estimate χ as follows. First, we need to obtain an estimate of investors' preferences $\hat{\boldsymbol{\theta}}^i_p$ and an estimate of $\hat{\boldsymbol{\theta}}_p$. Second we need to assume that at least one fund characteristic enters investors' preferences but is excluded from sponsors' preferences. For example, if investors care about past returns but sponsors do not (i.e., $\theta^s_{pr} = 0$) then χ will be determined by

$$\chi = 1 - \frac{\hat{\theta}_{pr}}{\hat{\theta}^i_{pr}}.$$

Table D1 presents the estimates of sponsors and investors preferences assuming that active investors solve a discrete choice portfolio problem and that past returns gross of fees enter investors preferences but not sponsor preferences. Under this assumption investors' preference parameters $\boldsymbol{\theta}^i$ can be estimated from the following linear regression:

$$\log(s^{\text{active}}_{jpt}) = \boldsymbol{w}'_{jpt}\boldsymbol{\theta}^{i,\text{active}} + \psi_{pt} + \xi_{jpt} \tag{89}$$

where s^{active}_{jpt} is the active investors portfolio share of fund j in plan p and year t, ξ_{jpt} is an unobserved demand shock possibly correlated with fees and ψ_{pt} are plan-by-year fixed effects which are needed to absorb the inclusive value component of investors discrete choice problem.[56] The data does not distinguish between portfolio shares of inactive v. active investors. To recover the latter I use the estimated fraction of inactive investors

[56] Plan-by-year fixed effect absorb the following term $\ln(1 + \sum_{j \in S_{pt}} \exp(V^{\text{investors}}_{jpt}))$ which is plan-year specific. Logit demand implies is given by $s_{jpt} = \frac{\exp(V^{\text{investors}}_{jpt})}{1+\sum_{j' \in S_{pt}} \exp(V^{\text{investors}}_{j'pt})}$.

and obtain the portfolio share of the active as

$$s_{jp}^{\text{active}} = s_{jp}\mathbf{1}\{j \neq d\} + \frac{s_{jd} - \delta_d}{1 - \delta_d}\mathbf{1}\{j = d\}.$$

To account for the endogeneity of fees I use funds' turnover ratio as instrument. The parameters $\boldsymbol{\theta}_p$ are estimated from the menu choice problem I developed in Section 6. Specifically, I use the estimates reported in the left column of Table (A5). Assuming inactive investor preference parameters is equal to 0 I can recover the preference weighting as

$$\chi = 1 - \frac{\hat{\theta}_{pr}}{\hat{\theta}^i_{pr}(1 - \delta)} \tag{90}$$

The estimates suggest that preference misalignment are important in this market. Sponsors weight their own preferences roughly three times more (0.75/0.25) than their investors preferences. Looking at the estimated coefficients, it appears that the largest misalignment is in the preference for fund affiliation. Furthermore, consistent with previous results, sponsors tend to tolerate higher fee more than their investors and particularly so if compared to active investors. Active investors marginal dis-utility from higher fees is twice larger than the one of their sponsors.

E 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.

F Differentiated Bertrand Network Game

In this Appendix I describe a simple differentiated Bertrand Network game and show how firms Nash equilibrium prices and margins relate to firms' network centrality. The discussion is based on Loseto (2023). In what follows, I frame everything in terms of products or firms, instead of calling them investment funds. I will assume that consumers have quadratic preferences with a taste for variety which makes the consumers' problem mathematically equivalent to an investor mean-variance portfolio problem. The same type of preferences are considered in Pellegrino (2023) who instead instead studies Cournot competition.

Consider a market with $j \in \{1, ..., J\}$ products available. Each product j is characterized by a set of K attributes whose values are collected in the K dimensional real-valued vector $x_j = (x_{jk})_{k=1}^K$ where x_{jk} is measured in units of quantity consumed. Characteristic x_{jk} tells you how much of attribute k you would get if you consume one unit of product j.

I assume there is a representative consumer who takes product prices $p = (p_j)_{j=1}^J$ as given, and chooses how much to consume of each product available. I denote by $q = (q_j)_{j=1}^J$ their consumption vector and define their preference as

$$u(q, X) = q_0 + q'\mu - \frac{\gamma}{2} q' \left(I + XX'\right) q \tag{91}$$

where q_0 is an outside good, X the $J \times K$ matrix of products attributes, μ is a J-vector parameters determining the marginal utility that comes from the linear term in (91). Finally, γ captures consumer's taste for variety.

The representative consumer takes prices p as given and maximizes (91) subject to

$$q_0 + q'p \leq y \tag{92}$$

where y income. After substituting for the budget constraint in (92), the demand system is given by

$$q(p) = \frac{1}{\gamma} \left(I + XX'\right)^{-1} (\mu - p) \tag{93}$$

$$= \frac{1}{\gamma} \left(I - \Theta\right)^{-1} (\mu - p) \tag{94}$$

which is always well defined because $(I + XX')$ is positive definite and therefore non-singular and $\Theta \equiv X(I + X'X)^{-1}X'$

Next consider assume that J single-product firms producing the J products with constant marginal costs. Firm j takes the vector of competitor prices p_{-j} as given and solves

$$\max_{p_j} \quad (p_j - c_j) q_j(p_j, p_{-j}) \tag{95}$$

$$\text{s.t.} \quad q_j(p_j, p_{-j}) = a_j - \frac{1}{\gamma}(1 - \theta_{jj}) p_j + \frac{1}{\gamma} \sum_{l \neq j} \theta_{jl} p_l \tag{96}$$

which is equivalent to

$$\max_{p_j} \left(a_j + \frac{c_j}{\gamma}(1-\theta_{jj}) \right) p_j - \frac{1}{\gamma}(1-\theta_{jj})\, p_j^2 + \frac{1}{\gamma} \sum_{l \neq j} \theta_{jl} p_j p_l. \tag{97}$$

The payoff function in equation (97) is analogous to the linear-quadratic utility functions considered in Ballester, Calvó-Armenagol and Zenou (2006) and, as such, defines a linear-quadratic network game in which each product is a node and the $J \times J$ matrix

$$A(\Theta) \equiv \Theta - diag(\Theta) \tag{98}$$

is the weighted and undirected adjacency matrix of the network.

Network game interpretation. The adjacency matrix defined in (98) shows that network connections and products' substitution patterns are isomorphic to each other. We know that an off-diagonal element θ_{jl} of the matrix Θ captures the degree of substitution between product j and product l because it is defined as the (j,l) element of the demand jacobian. From equation (98), we can interpret θ_{jl} as a network link between product j and product l and therefore, we can think of the product differentiation space as being a network whose nodes are the products and whose links tell us how close, or equivalently how substitutable, are any two products.

Framing the product differentiation space as a network enables us to learn how product differentiation affects equilibrium outcomes by studying the topological properties of the competitive network. Loseto (2023) shows that equilibrium Bertrand price-cost margins depend negatively on a product's Bonacich centrality, which, following Jackson (2008), is defined as

Definition 1 *Let (A, J) be a network with J nodes and adjacency matrix A. The J-vector of (weighted) Bonacich centralities $\mathbf{b}(A, \delta, u)$ is given by*

$$\mathbf{b}(A, \delta, u) \equiv (I - \delta A)^{-1} \delta A u = \sum_{k=1}^{\infty} \delta^k A^k u, \tag{99}$$

where $\delta > 0$ is a scalar and $u > 0$ is J-vector.

The j-th element of $\mathbf{b}(A, \delta, u)$ summarizes how central node j is in the network. This measure of centrality is widely used in social networks because it captures a node's importance in terms of how close/connected this node is to others and how close/connected the nodes it is connected to. According to the definition of Bonacich centrality, a node's importance is a weighted sum of the walks that emanate from it. Moreover, if $\delta \in (0,1)$, walks of shorter length are weighted more.[57]

In an interior Betrand-Nash equilibrium, firms' equilibrium price-cost margins can be decomposed as

$$p^* - c = \underbrace{\frac{\mu - c}{2}}_{\text{monopolist margin}} \underbrace{- \mathbf{b}\left(A(\Theta), \frac{1}{2(1-\theta)}, \frac{\mu - c}{2} \right)}_{\text{Bonacich centrality}}, \tag{100}$$

[57] This interpretation is motivated by the fact that when A is a binary $\{0,1\}$ it k-th power A^k counts how many walks of length k are between any two nodes.

or equivalently, firms equilibrium fees can be decomposed as

$$p^* = \underbrace{\frac{\mu + c}{2}}_{\text{monopolist fee}} \underbrace{- \mathbf{b}\left(A(\Theta), \frac{1}{2(1-\theta)}, \frac{\mu - c}{2}\right)}_{\text{Bonacich centrality}}, \tag{101}$$

The key insight is that the more central a product is in the competitive network, the lower its equilibrium price-cost margins. What does this mean in practice? From Definition 1, we can see that the higher any of the entries of the j-th row of A, the more central node j is. In this setting, product j is more central the higher its substitutability with any other product (i.e., the higher the elements $(\theta_{jl})_{l \neq j}$ of the j-th row of Θ). Overall, the expression for the Bertrand price-cost margins in (101) tells us two things. First, a less central or, equivalently, more differentiated product will be able to charge higher markups. Second, a product's Bonacich centrality is a sufficient statistic to measure how product differentiation allows firms to price above marginal costs.

Next, I perform a simple simulation exercise to summarize and visualize how the Bertrand Network model works. Table F1 describes the parameters used in the simulation. There is a single market with $J = 30$ products and $K = 7$ characteristics whose values are drawn from a uniform distribution in between $[0, 1]$. The demand intercept μ is the same across all products and set to 0.15 whereas marginal costs are heterogeneous across products and drawn from a $[.01, .03]$ uniform distribution.

Parameter	Value
J	30
K	7
μ	0.15
c	U[0.01, 0.03]
x_{jk}	U[0, 1]

Table F1: Parameters for simulation of Bertrand network game

Given this parameters, Figure F1 plots the underlying Bertrand network. Each product is a node and the edges capture the degree of substitution between any two products/nodes; the longer the edge the less substitute are the two products. The location of dots and edges is exogenous and entirely determined by the realization of the draws of product characteristics. Conversely, the size of the dots is endogenous and it is proportional to the equilibrium price-cost margins. The plot shows that nodes that are more peripheral tend to have larger dot sizes whereas dots that are more central are smaller. The intuition for this result is the following: peripheral products are more unique or equivalently less central and, per equation (101) will charge higher margins in equilibrium. On the other hand, more central nodes face more intense competition and must lower their margins.

Figure F2 instead visualizes the previous decomposition and plots the equilibrium price-cost margins on the y-axis against the Bonacich product centrality on the x-axis. It should be clear by now why the relationship is decreasing; higher centrality implies lower equilibrium markups. The noise around the downward sloping relationship is due to the fact that marginal costs are heterogeneous. By increasing the variance of the distribution

of costs, Figure F2 would start looking noisier and the resulting relationship between centrality and margins might not look as clear. This highlights how empirically it is important to control for the unobserved costs in order to recover the downward sloping relationship. The same would be true if we were to introduce heterogeneity in the demand intercept μ.

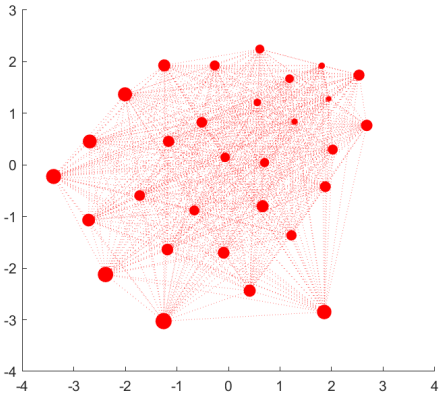

Figure F1: Simulated Network. Location is exogenous. Node size is proportional to markups.

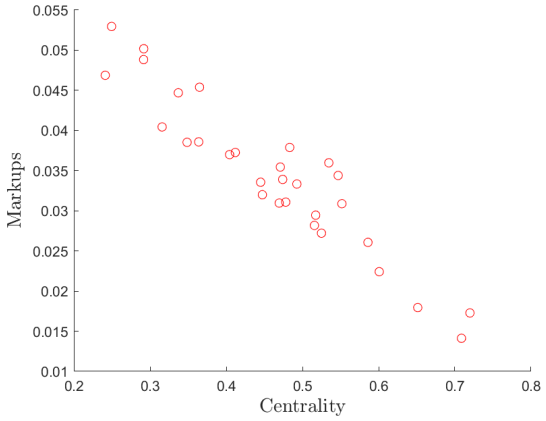

Figure F2: Simulated Network. Price-cost margins (y-axis) against Bonacich centrality (x-axis).