Bundling and Exclusion in the Defined Contribution Plan Market

Lin Zhao *

Duke University, Fuqua School of Business

November 21, 2025

Click here for the latest version

Abstract

Employers work with recordkeepers to provide defined contribution plans for their employees. Yet access to high-quality plans remains unequal, particularly for employees at small firms. I show that this disparity is driven by a supply-side friction: recordkeepers exercise their market power by strategically excluding small employers through minimum revenue thresholds. Because most major recordkeepers are vertically integrated and bundle recordkeeping with their own investment products, exclusion enables them to raise recordkeeping fees and induce employers towards plans with more affiliated and expensive investment options. I build and estimate a structural model that quantifies the extent and consequences of exclusions. Counterfactual analyses show that policies designed to relax the exclusion, such as pooled employer plans, can expand access to high-quality plans and improve participant welfare, although these gains can be mitigated by recordkeepers' strategic adjustment of exclusions in equilibrium.

Keywords: Bundling, Exclusion, Defined Contribution Plan, Nash Bargaining

*I thank my committee members Alon Brav, Simon Gervais, John Graham, Arthur Taburet, and S. "Vish" Viswanathan for their guidance on this project. I am also grateful to Manuel Adelino, John Barry, Xinyue Bei, Anna Cieślak, Nuno Clara, Allan Collard-Wexler, Paymon Khorrami, Steven Medema, Adriano Rampini, João Thereze, Melanie Wallskog, Daniel Yi Xu, and participants in seminars at the Fuqua School of Business and the Department of Economics at Duke University for their comments and suggestions. The results are subject to change. All errors are my own.

1 Introduction

Defined contribution (DC) plans have become the dominant channel of retirement saving in the United States, holding over $12 trillion in assets and covering more than 120 million participants. The design and pricing of these plans directly influence the accumulation of retirement wealth for a large share of the labor force. Yet access to high-quality, low-cost plans remains uneven across employers: employees at small firms systematically face higher administrative and investment fees than those at large firms.

I show that these disparities arise from a supply-side friction rooted in the behavior of recordkeepers, the service providers that administer DC plans and design their investment menus. I argue that major recordkeepers impose minimum revenue thresholds that effectively exclude small plan sponsors (typically employers) from their platforms.[1] Because most large recordkeepers are vertically integrated with asset-management affiliates and are able to bundle their own investment products, exclusion enables them to charge higher recordkeeping fees and steer plans toward more expensive affiliated investment products. I build and estimate a structural model of the DC plan market to quantify the extent and welfare consequences of these exclusions and to evaluate how market structure changes on both sides shape plan access, pricing, and participant welfare.

The model captures the strategic interaction between plan sponsors and recordkeepers. On the demand side, each employer chooses among recordkeepers and plan menus, taking into account investment-menu characteristics and recordkeeping fees. On the supply side, recordkeepers decide which sponsors to serve and set exclusion thresholds that determine the minimum revenue required from a client. Once a sponsor–recordkeeper match occurs, the two parties bargain over transfers, reflecting the outside options available to them. Recordkeepers' vertical integration with asset-management affiliates links these bargaining outcomes to the profitability of their investment products, providing an additional incentive to exclude smaller, less profitable clients.

At the core of the analysis is the idea that recordkeepers exercise market power through strategic exclusion. In the Nash bargaining framework, the recordkeeper commits the exclusion threshold as its disagreement payoff. Importantly, the same exclusion threshold applies to bilateral Nash bargaining in every potential match. Exclusion therefore plays a role analogous to price in standard models of price competition: a higher threshold increases revenue from each client but reduces the set of sponsors served. Each recordkeeper optimally chooses its exclusion threshold to maximize total profit given the choices of its competitors, yielding a Bertrand-Nash equilibrium in exclusion and pricing behavior.

[1] I use sponsor to refer to the entity that legally establishes and maintains the DC plan. This is most often the employer, but not always: in some cases the sponsor can be a union, a professional or trade association, or a pooled or multiple-employer arrangement.

Exclusion has asymmetric effects across employers of different sizes. For large employers, the per-participant exclusion is close to zero, so exclusion is largely irrelevant for their recordkeeper and menu choices. Small employers, by contrast, face binding exclusion constraints and are often forced to rely on low-exclusion but lower-quality providers. Employers in the middle of the size distribution are disciplined by exclusion to select arrangements with higher recordkeeping fees and menus containing a greater share of affiliated and higher-fee investment options.

The optimal exclusion is also affected by factors beyond the trade-off between revenue and market share. On the one hand, recordkeepers face limited operational capacity: expanding their client base requires additional investment in technology infrastructure and labor. These capacity constraints create an opportunity cost of serving more plans, pushing recordkeepers to set higher exclusion thresholds. On the other hand, a larger footprint in the DC plan market can generate synergies with a recordkeeper's other lines of business. For example, insurance companies and payroll service providers can cross-sell products to the same employers, even when these sales are not formally bundled into the DC plan contract and therefore fall outside the bargaining process. Such synergies make it valuable to serve more plans, which tends to lower the optimal exclusion threshold. Because both forces can operate simultaneously and are tied to the recordkeeper's effective service capacity, I refer to their combined effect as the recordkeeper's shadow value of capacity and estimate it in the model.

This supply-side perspective complements a growing literature that studies demand-side frictions in the defined contribution plan market. Prior research emphasizes that employers and employees often fail to select optimal plans because of misaligned incentives, limited attention, or inertia. Bhattacharya and Illanes (2022) and Loseto (2023) model the principal–agent problem between employers and plan participants, showing that sponsors may tolerate high-fee investment options when participants bear most of the costs. Yang (2023) highlights switching costs that discourage employers from renegotiating or changing recordkeepers even when better alternatives exist. In these studies, the central friction is on the employer/participant side, and recordkeepers appear mainly as service providers that supply plans at given terms.

My framework departs from prior demand-side models by explicitly modeling the profit-maximizing behavior of the recordkeepers. This structural perspective yields three insights extending the demand-side framework. First, it explains the observed matching pattern between sponsors and recordkeepers, that large firms pair with high-quality, high-exclusion providers while small firms rely on low-exclusion alternatives. Second, it rationalizes the empirical finding that, for a given sponsor, average plan fees decline as plan size expands, a prediction unique to the supply-side mechanism of exclusion. Third, it delivers a different prediction for pooled employer plans (PEPs): demand-side mod-

els suggest such pooling would have a limited effect because the friction roots in plan sponsors' preferences, whereas in my analysis, the friction comes from the supply-side exclusions, thus PEPs can meaningfully expand access and improve welfare. Together, these distinctions highlight how incorporating strategic supply-side behavior complements and extends the insights from existing demand-side studies.

The model estimates yield several main findings about the structure of the DC plan market. Recordkeepers impose sizable and heterogeneous exclusion thresholds. Fidelity, the largest recordkeeper, has an estimated exclusion of $51,800, given its median client size of 345 participants, this translates to about $150 per participant, nearly matching its median annual revenue per participant. Vanguard's exclusion of $29,700 corresponds to roughly $6 million in plan assets if revenue derives solely from management fees, consistent with anecdotal evidence that Vanguard historically required several million dollars in plan balances to initiate service. In contrast, ADP shows no evidence of exclusion, consistent with its focus on small plans and integration with payroll services.

Second, the estimated exclusions are strongly increasing in each recordkeeper's mean utility, indicating that high-quality providers set higher thresholds and attract larger employers. This pattern reflects the sorting mechanism central to the model: recordkeepers offering superior services use high exclusions to screen out smaller, less profitable sponsors, while low-exclusion providers compete for small plans. Together, these results provide direct empirical evidence that strategic exclusion drives the observed link between employer size and plan quality in the data.

Third, differences in recordkeepers' shadow value of capacity play an important role in determining exclusion levels. Most top recordkeepers are constrained by the opportunity cost of expanding coverage, with estimated shadow costs ranging from $6 to $103 per participant, about one-quarter of average per-participant profits. By contrast, ADP exhibits a positive shadow value of scale, consistent with complementarities between its recordkeeping and payroll businesses. That is, by expanding coverage it increases its overall profitability. This result is consistent with ADP's business model: as a payroll service provider, it benefits from bundling DC plan administration with payroll processing and other employer services. These complementarities make serving additional small plans profitable, explaining why I do not find evidence of a positive exclusion level for ADP.

Fourth, exclusion has first-order welfare consequences. Removing thresholds in the model substantially expands access to high-quality recordkeepers and improves plan quality. The overall provision rate of DC plans rises from 97 percent to nearly 100 percent, and the share of employers choosing Fidelity increases from 15.3 percent to 20.7 percent. For small plans with 100 participants, the expected management fee declines by 8 basis points and

the recordkeeping fee by about $30 per participant, generating meaningful welfare gains for employees at small firms.

Fifth, removing exclusions also has large effects on recordkeepers' profitability. When all exclusion thresholds are eliminated, total profits decline sharply across all major recordkeepers, with larger initial exclusions leading to larger losses. For example, Fidelity's total profit falls by about 30 percent and its per-participant profit by 48 percent. Overall, the results show that exclusion raises recordkeepers' profitability primarily by restricting access for small plans, and that removing these thresholds redistributes surplus from recordkeepers to employers and participants.

Having established the magnitude and determinants of exclusion, I next use the model to examine how policy and market structure shape equilibrium outcomes through a series of counterfactual analyses. First, I evaluate the potential effects of pooled employer plans (PEPs), a demand-side policy introduced under the 2019 SECURE Act, assuming recordkeepers' exclusion behavior remains unchanged. Pooling small employers into larger plans increases the provision rate from 80 to 90 percent and raises per-participant welfare by approximately $200, as pooled plans gain access to higher-quality recordkeepers and lower-fee investment menus.

I then incorporate recordkeepers' strategic response to pooling. When I consider the possibility that all plans with fewer than 300 participants in the same ZIP code merge into a single pooled plan, I find that the resulting increase in large plans induces recordkeepers to raise their exclusion thresholds. This response dampens the benefits of PEPs: while participating firms enjoy lower costs, firms that remain outside pooled arrangements face higher equilibrium exclusions and welfare losses. The overall welfare effect depends on firm-size distribution and PEP adoption rates.

Next, I consider a scenario in which Fidelity introduces its own pooled plan, which is a lower-quality plan with no exclusion requirement, alongside its standard plan. Employers can choose between the two. Fidelity's standard plan loses some market share, but its aggregate share increases as it attracts smaller employers that previously relied on low-quality recordkeepers. Aggregate welfare rises because more participants gain access to a plan administered by a high-quality provider.

Finally, I examine a supply-side counterfactual reflecting the industry's ongoing consolidation. When I assume that all non-top-ten recordkeepers merge into four large firms that also set exclusion thresholds strategically, the new entities gain market power and raise their exclusions, making them less attractive to small employers. In response, the existing top-ten recordkeepers lower their thresholds to capture these displaced clients. In equilibrium, small employers face slightly higher costs but gain access to higher-quality providers, while overall market concentration increases.

Together, the findings show that supply-side behavior plays a central role in shaping the quality and cost of retirement plans. Recordkeepers' strategic use of exclusion not only limits access for small employers but also amplifies the effects of vertical integration on fees and investment choices. Recognizing these supply-side frictions is essential for designing policies that expand retirement plan coverage and improve participant outcomes. Interventions such as pooled employer plans can mitigate disparities, but their effectiveness depends on how recordkeepers adjust in equilibrium. More broadly, the paper highlights how market power and product bundling in financial intermediation can create persistent inequality in access to high-quality financial services. The framework developed here can inform the evaluation of future policies, both in retirement markets and in other settings where intermediaries balance profitability against access and competition.

The remainder of the paper proceeds as follows. Section 2 describes the DC plan market and data. Section 3 presents descriptive evidence on bundling and exclusion. Section 4 introduces a simplified model to build intuition, and Section 5 develops the full structural model. Section 6 reports estimation results, Section 7 conducts counterfactual and policy analyses, and Section 8 concludes.

Related literature

This paper is closely related to recent structural work in the defined contribution plan market to address the issue of low-quality plans, including Bhattacharya and Illanes (2022), Loseto (2023), and Yang (2023). These studies primarily focus on the demand-side frictions, including the principal-agent problem between employer and employees (Bhattacharya and Illanes (2022) and Loseto (2023)), and switching costs in Yang (2023).

Under ERISA, it is the plan sponsor's duty to act solely in the interest of the plan participants. The fact that plan sponsors breached their duties to participants would suggest the existence of the principal-agent problem. The agency problem can occur intentionally or unintentionally. Bhattacharya and Illanes (2022) assume that the plan sponsor pays the recordkeeping fees and plan participants pay the management fees, so the sponsor has the incentive to shift the costs to participants by having them pay higher management fees. Loseto (2023) finds that plan sponsors on average are only half as sensitive to management fees as plan participants, so they design plans with higher fees than what would be under the participants' preference.

The principal-agent problem definitely matters in many cases while it might not exist in other cases. First, many DC plans have a plan committee including representatives of the participants.[2] The agency problem should be less of a concern in those cases. Second, it is

[2] For example, Purdue University's Defined Contribution Retirement Plan Committee contains the Faculty Senate representative, the Clerical and Service Staff Advisory Committee (CSSAC) representative, the Management and Professional Staff Advisory Committee (MaPSAC) representative, and the Purdue University Retirees Association (PURA) representative. See

not clear which party pays the recordkeeping fees. A report[3] by DOL suggests that in the case of a 401(k) plan or other types of DC plans with individual accounts, administrative fees are allocated to each individual account either in proportion to account balance or passed as a flat fee to each account. If the administrative fees are eventually paid by the plan participants as well, the sponsor should not have a strong incentive to choose a plan with expensive investment options.

This paper also contributes to the large empirical literature on the DC plan market by addressing the pricing behavior of recordkeepers on the supply side. Reuter (2024) provides a comprehensive literature review of this literature, which has largely emphasized the demand side of retirement saving behavior, treating the supply of plan services as passive. The literature has addressed how default options (Madrian and Shea (2001)), participant inertia (Choi, Laibson, Madrian, and Metrick (2004), Choi, Laibson, and Madrian (2011)), and limited attention (Beshears, Choi, Laibson, Madrian, and Milkman (2015)) shape participation and investment choices. Other work explores how plan design and menu complexity influence participants' asset allocations (Iyengar, Jiang, and Huberman (2003); Sialm, Starks, and Zhang (2015); Reuter and Richardson (2022)). Collectively, this literature provides a rich view of how plan participants respond to the options presented to them, while neglecting the strategic behavior of recordkeepers. The paper helps to explain plan choice, fee dispersion, and bundling in an equilibrium setting that addresses the behavior of both sides of the market.

This paper complements the literature examining how regulatory interventions influence the DC plan market. For example, Badoer, Costello, and James (2020) show that fee-disclosure reforms reduced indirect compensation, and Kronlund, Pool, Sialm, and Stefanescu (2021) find that the 2012 Department of Labor disclosure rule increased participants' sensitivity to fees. My work complement this literature by studying the impact of the pooled employer plan that was introduced in 2021, taking into consideration the equilibrium effect on both sides of the market.

The paper also contributes to the bundling literature. Bundling has been extensively studied as a mechanism for price discrimination, market segmentation, and foreclosure. Foundational theoretical work, beginning with Adams and Yellen (1976) and extended by McAfee, McMillan, and Whinston (1989) shows that bundling allows firms to segment heterogeneous consumers and thereby extract more surplus from them. Subsequent studies such as Bernheim and Whinston (1990), Matutes and Regibeau (1992), and Nalebuff (2004) highlight its potential strategic role in deterring entry and raising rivals' costs, while Armstrong and Vickers (2010), Zhou (2017), and Armstrong and Vickers (2018),

https://www.purdue.edu/hr//Benefits//retirees/Retirment%20Resources/retirementplanCommittee.php

[3] https://www.dol.gov/sites/dolgov/files/ebsa/about-ebsa/our-activities/resource-center/publications/retirement-plan-fees-expenses.pdf

develop models of competitive and nonlinear pricing that formalize how bundling alters equilibrium conduct. Song, Nicholson, and Lucarelli (2017) further shows how mergers and inter-firm bundling reshape product competition and welfare. Building on this theoretical foundation, recent structural analyses, such as Crawford and Yurukoglu (2012) and Crawford, Lee, Whinston, and Yurukoglu (2018) in multichannel television, and Cao and Chatterjee (2022) in pharmaceuticals—quantify how vertical integration and bundling affect prices, access, and welfare in concentrated industries. My paper contributes to this structural literature by applying these insights to the defined contribution (DC) plan market, where recordkeepers vertically integrate recordkeeping and asset management services. By modeling bundling and exclusion thresholds in a Nash-in-Nash framework, I show how bundling in financial intermediation functions both as an efficiency mechanism—internalizing complementarities—and as a strategic tool that limits access and competition.

This paper is also related to the literature on specialization. Despite some unique features related to the nature of the business, the DC plan market exhibits a market structure similar to industries such as retail and banking, with a small number of dominant firms alongside many niche providers. Ellickson (2006) shows that in retail, escalating fixed investments in quality and distribution create barriers to entry, leading large firms to dominate while smaller ones specialize in narrow segments. In banking, Balyuk, Prabhala, and Puri (2021) and Paravisini, Rappoport, and Schnabl (2023) show that specialization arises from expertise and information advantages that segment credit markets and shape competition. My paper extends these insights to the defined contribution plan market, where recordkeepers specialize by plan size and service scope, and relates the specialization to vertical integration and bundling.

Finally, this paper fits into the structural finance literature, including Gavazza (2011) in the mutual fund industry, Taburet, Polo, and Vo (2024) in the lending market, Egan, MacKay, and Yang (2022) in the index funds industry, Egan, MacKay, and Yang (2021) in the 401(k) plan market, and Bhattacharya, Illanes, and Padi (2019) in the financial advice market.

2 Institutional Background and Data

2.1 Institutional Background

The defined contribution (DC) plan is one of the most common types of employee benefit plans. There are different types of defined contribution plans, including 401(k) plans, 403(b) plans, 457(b) plans, employee stock ownership plans, and profit-sharing plans.

The most popular type is the 401(k) plan. As of 2022, out of the $8 trillion assets in DC plans, $6.8 trillion are in 401(k) plans.[4]

A DC plan consists of a pre-specified investment menu with different types of investment options, including mutual funds, stocks, collective investment trusts, and separate accounts, among others. Both the employer and the employees can contribute to the plan and enjoy the tax advantages of their contributions. On the employee side, the contributions are deducted from the pre-tax income, and the taxes on contributions can be deferred until withdrawals are made. On the employer side, employer contributions are tax-deductible.

The key feature of the DC plan is that the contributions are defined, but the eventual benefits or retirement income are not guaranteed. This differentiates it from the defined benefit (DB) plan that guarantees a predetermined payout after retirement. In a DC plan, employees choose investment options from the menu, whereas in a DB plan, employers are responsible for managing or outsourcing the investments.

A DC plan will have a main service provider, the recordkeeper, who provides services including administering the plan, managing investments, and educating plan participants. To find a recordkeeper, the employer will go through a request-for-proposal (RFP) process, in which the employer will contact a few potential recordkeepers and compare offers from each recordkeeper.

Figure 1 shows the market shares of the top-10 recordkeepers[5] in 2019 among plans with more than 100 participants.[6] The market is not very concentrated, with an HHI of 0.05. The largest recordkeeper Fidelity serves 10,017 plans out of a total of 61,630 plans.

[Insert Figure 1 here]

There are different types of recordkeepers, including asset management companies, insurance companies, payroll service companies, banks, and non-integrated recordkeepers. For the purpose of this paper, I categorize them into two main categories depending on whether they have their own investment products: integrated recordkeepers and non-integrated recordkeepers. Among the top 10 recordkeepers, only ADP does not have its own investment products. When a plan has an integrated recordkeeper, the recordkeeper

[4] https://www.ici.org/system/files/2025-03/25-rpt-dcplan-profile22-401k.pdf

[5] Precisely speaking, the top-10 recordkeepers are the top 10 recordkeepers of my focus in this paper. They are ranked by their market shares except for Charles Schwab, which is the 11th largest recordkeeper in terms of market share and the 9th largest recordkeeper in terms of total number of participants served. In Appendix D, I show the ranks by the number of plans served, the number of participants served, and AUM.

[6] As noted in the next subsection, only plans with more than 100 participants are required to provide information on their recordkeepers. PLANSPONSOR magazine named Paychex as the largest recordkeeper in 2023 in terms of total number of plans, serving more than 110,000 plans.

can sell a bundle of two products at the same time: the recordkeeping service and the affiliated investment options.

A DC plan can generate two main types of fees: recordkeeping fees for administrative services and management fees for asset management services. Recordkeeping fees are typically on a per-participant basis, while asset management fees of each option equal the expense ratio multiplied by the asset under management. Figure 2 shows the structure of revenues in a DC plan. Including the two types of fees, a recordkeeper could receive three sources of revenue: recordkeeping fees, management fees, and revenue sharing from other investment option providers for including their options on the menu. The first and the third are the revenues from the recordkeeping business, and the second is the revenue from the asset management business.

[Insert Figure 2 here]

Table 1 gives the summary statistics of the three sources of revenue. The median plan has around 300 participants. The median per-participant recordkeeping fee is $38, while the medians of per-participant revenue sharing and per-participant management fees from affiliated options are $15 and $49. Among all plans with positive management fees, the ratio of the management fees to the total revenue has a median of 0.5, with the 90th percentile equal to 0.98. This shows that management fees from affiliated investment options constitute an important portion of the recordkeeper's total revenue.

[Insert Table 1 here]

2.2 Data

The primary data source is Form 5500, publicly available from the Department of Labor (DOL). The Employee Retirement Income Security Act of 1974 (ERISA) requires administrators of DC plans to file Form 5500 annually. Small plans with fewer than 100 participants at the beginning of the plan year are eligible to file Form 5500-SF instead of Form 5500. Form 5500-SF does not contain plan details, including service providers and the investment menu. Therefore, in this paper, the analysis is restricted to plans with more than 100 participants in 2019.

Most of the plan information can be directly obtained from the digitized version of Form 5500, except for the plan menu. Due to the complicated and non-standardized format, DOL did not digitize the plan menu. I accessed the digitized menu data from BrightScope.

3 Descriptive Analysis

This section documents the key empirical facts about bundling and exclusion in the DC plan market.

Fact 1: Bundling. Integrated recordkeepers bundle their own investment products in the menu.

[Insert Figure 3 here]

Integrated recordkeepers can include their investment products in the menu. The practice of bundling is very common in the DC plan market. Figure 3 plots the percentages of DC plans with their recordkeeper's own investment products for the top 10 recordkeepers, as well as the rest as an aggregate. The dashed line is the average among all plans. Among 47,656 plans with an integrated recordkeeper, 81% of them have at least one investment option for the recordkeeper. Fidelity, the largest recordkeeper, and Vanguard, the largest asset management company, have the highest bundle sale percentages, which are 96% and 99.7%. The insurance company recordkeepers, Empower, John Hancock, and MassMutual, have slightly lower proportions. ADP is missing because it does not have its own investment products. While bundling is common, Figure 3 also shows that it is not compulsory.

Bundling practices vary across plan sizes, both at the intensive margin, i.e., the proportion of affiliated investment products in the menu, and the extensive margin, i.e., the average expense ratio of affiliated investment products. Large plans tend to have few affiliated investment options while those affiliated investment options tend to have lower expense ratios. In Appendix D, I describe these patterns in the data.

Fact 2: Specialization. Different recordkeepers specialize in serving plans of different sizes.

[Insert Figure 4 here]

Plans of different sizes choose different recordkeepers. Figure 4 plots each of the top 10 recordkeepers' market share against their average client size. It is clear that among recordkeepers with the largest market shares, there is substantial variation in the distribution of their client sizes. In Appendix D, I plot the client size distribution of each of the top-10 recordkeepers. Recordkeepers such as Fidelity, Vanguard, AllSpring, and Charles Schwab, primarily serve large plans, whereas John Hancock and ADP specialize in small plans. Importantly, a large market share does not necessarily correspond to a large average client size. If, following the standard argument from the industrial

organization literature, market share reflects the mean utility of each recordkeeper, then large plans should also favor ADP, the fifth largest recordkeeper. As a result, the client size distributions across recordkeepers would be expected to look similar. The observed differences in these client size distributions indicate that factors other than mean utility affect sponsors' choice of recordkeeper, and these factors are related to plan size.

Fact 3: Plan size matters. Large plans tend to have lower fees.

[Insert Figure 5 here]

Figure 5 illustrates the well-documented pattern in the DC plan market: large plans tend to have lower recordkeeping fees and choose menus with low management fees. This pattern could have a demand-side explanation, whereby sponsors of large plans have different preferences over fees compared to sponsors of small plans (Bhattacharya and Illanes (2022)).

[Insert Figure 6 here]

To isolate the effect of plan sponsors' preferences, I examine how plan size influences recordkeeper and menu choices by focusing on cases in which the same plan experiences a substantial increase in size. Specifically, I identify all events where a plan's total number of participants rises by more than 50% within a year. I then implement an event-study design exploiting these sudden expansions in plan size using the following regression specification:

$$y_{it} = \sum_{t \neq -1} \theta_s 1(s = t) + \tau_t + \lambda_i + \epsilon_{it}. \tag{1}$$

In the first event study, the dependent variable y_{it} is the plan's average fee. In Figure 6, Panel (a) shows that the average fee declines by about 1.5 basis points in the year following the sudden expansion of the plan. The reduction in fees could occur because the sponsor switches to a different recordkeeper following the expansion. Yang (2023) finds that management fees significantly drop following a change of recordkeeper. Therefore, in the second event study, the dependent variable is the indicator for whether the sponsor switches the recordkeeper in the given year. Panel (b) supports the hypothesis and shows that the probability of switching a recordkeeper significantly increases by 17.5% following a sudden plan size expansion. The switching probability remains significantly higher in the first three years after the expansion. These results highlight the importance of plan size in recordkeeper choice.

Furthermore, to isolate the effect of recordkeeper switching, in the third event study, I repeat the first event study but exclude all events where the sudden expansion coincides

with a recordkeeper change. Panel (c) shows that fees still decline following the expansion even when the sponsor retains the same recordkeeper, suggesting that larger plans obtain lower fees even absent switching. Putting together, these event studies show that plan size matters in a sponsor's recordkeeper and plan choice even in the absence of idiosyncratic preferences.

4 Illustrative Model

In this section, I present the setting of my structural model. To illustrate the economics, I consider a simplified version of the model with only one monopolistic recordkeeper. The next section adds additional components to the model and takes it to the data.

4.1 Baseline Setting

Market. There is a market for recordkeeping services. The recordkeeper provides the recordkeeping service to the sponsor. They bargain over the price of the service, i.e., the recordkeeping fee, denoted by T, *à la* Nash bargaining.

The recordkeeper can bundle other businesses while providing the recordkeeping service to the sponsor. For example, the recordkeeper can include its own investment options on the investment menu and receive management fees.

Demand. There is a continuum of sponsors differing in their size N. A sponsor of size N is an employer with N participants. Let δ be the per-participant utility of the recordkeeping service. The utility depends on the quality of the recordkeeping service as well as other businesses bundled together. The per-participant net utility is therefore

$$u = \delta - T. \tag{2}$$

The sponsor of size N gets a total utility of $U = N \cdot u$.

Supply. A monopolist recordkeeper provides the recordkeeping service with the per-participant marginal cost mc. The recordkeeper can bundle other businesses at the same time and receive an additional revenue M from the bundled business. Assume $M = 0$ if there is no other business bundled together. The recordkeeper's per-participant profit is

$$\pi = T + M - mc. \tag{3}$$

Pricing. The sponsor bargains with the recordkeeper over the recordkeeping fee T

$$\max_{T \geq 0} (T + M - mc)^{1-\eta} \cdot (\delta - T)^{\eta} \tag{4}$$

where η is the sponsor's bargaining power. The solution is

$$T = \max\{0, (1-\eta)\delta - \eta(M - mc)\} \tag{5}$$

The condition $T \geq 0$ requires that the transfer can only be from the sponsor to the recordkeeper.

4.2 Sponsor's Bundle Choice

Suppose the recordkeeper provides multiple bundles and the sponsor can choose one of the bundles. Each bundle is associated with a different M, and each such bundle M leads to a different per-participant utility of $\delta(M)$. Write δ as $\delta(M)$, and plug in equation (5) into equation (2), the sponsor's utility given M is therefore

$$u(M) = \begin{cases} \delta(M) & \text{if } T = 0 \\ \eta(\delta(M) + M - mc) & \text{if } T > 0. \end{cases} \tag{6}$$

The question of interest is, which bundle would the sponsor choose? I leave the complete discussion of the question to Appendix B, and I will show the intuition with a few numerical examples.

[Insert Figure 7 here]

Let $mc = 0$ and $\eta = 0.5$. Assume that $\delta(M) = 100 - kM$. With this parameterization, the change in the net utility $u(M)$ given a dollar increase in M is

$$\Delta u(M) = \begin{cases} -k & -0 & \text{if } T = 0 \\ \underbrace{\eta(1-k)}_{\text{Change in } \delta} & -\underbrace{(-k(1-\eta) - \eta)}_{\text{Change in } T} & \text{if } T > 0. \end{cases} \tag{7}$$

Example 4.2.1. Suppose $k < 0$. Let $k = -0.1$. Panel (a) of Figure 7 plots the transfer T, the sponsor's net utility u, and the recordkeeper's profit π as functions of M. This is the case where the sponsor likes bundles with high M. Higher M results in higher utility and lower recordkeeping fee; thus, the sponsor will always choose the bundle with the highest M.

Example 4.2.2. Suppose $k = 0$. This is the case when the sponsor's utility of the service δ does not directly relate to the bundled business. Indirectly, since the recordkeeper receives higher revenue, the recordkeeping fee T decreases in M until being bounded below by 0. Therefore, the sponsor's net utility u increases in M and remains flat when $T = 0$. Panel (b) of Figure 7 plots T, u, and π against M in this case.

Example 4.2.3. Suppose $0 < k < 1$. Let $k = 0.5$. In this case, the sponsor dislikes bundles with high M. For example, an investment menu with high-fee investment options could bring high management fees to the recordkeeper but harms the plan participants. The direct effect of an additional dollar of M is reducing the sponsor's utility δ by 0.5 dollars. On the other hand, when $T > 0$, one additional dollar of M reduces recordkeeping fee T by 0.75 dollars. Therefore, the sponsor prefers a bundle with a higher M if $T > 0$. When M is large enough so that T is bounded below by 0, higher M still lowers the sponsor's utility, but the recordkeeping fee cannot be reduced further. In this range, a higher M translates to a lower net utility.

Panel (c) Figure 7 displays the transfer, net utility, and profit in this case. As discussed above, the sponsor's net utility first increases in M before eventually decreasing. This hump-shaped relationship implies that the recordkeeper could be indifferent to two bundles with different M.

Example 4.2.4. Suppose $k = 1$. In this case, in the range where $T > 0$, one additional dollar of M decreases both the utility δ and the recordkeeper fee T by 1 dollar, therefore, the sponsor is indifferent to bundles with different M. While M is large enough so that $T = 0$, a larger M only reduces the utility, so the sponsor prefers a bundle with a lower M. As a result, the sponsor is indifferent to bundles whose M are not so large. Panel (d) of Figure 7 displays the relationship between the net utility u and M in this case.

Example 4.2.5. Suppose $k > 1$. Let $k = 1.5$. In this case, while $T > 0$, one dollar increase in M reduces utility δ by 1.5 dollars while only lowers recordkeeping fee by 1.25 dollars; while $T = 0$, a larger M only reduces the utility. Putting together, the net utility strictly decreases in M, so the sponsor always prefers the bundle with the lowest M. Panel (e) of Figure 7 displays this monotone relationship.

4.3 Recordkeeper's Exclusion

The monopolist recordkeeper can increase its profit by exercising strategic exclusion. In the Nash bargaining, each side can extract a higher proportion of the joint surplus if it can credibly increase the value of its outside option. In this baseline setting, there is only one monopolist recordkeeper; therefore, the sponsor's only outside option is to forgo service entirely, yielding zero net utility. The recordkeeper, by contrast, faces a

continuum of potential sponsors and can credibly commit to serving only a subset of the market, thereby improving its bargaining position and increasing profits.

The recordkeeper can require a minimum amount of profit in each pair of Nash bargaining with sponsors. Denote this minimum profit requirement by E. For a sponsor of size N, the effective per-participant minimum profit requirement is $\frac{E}{N}$. The Nash bargaining becomes

$$\max_{T \geq 0} \left(T + M - mc - \frac{E}{N} \right)^{1-\eta} \cdot (\delta - T)^{\eta} \tag{8}$$

The Nash bargaining is valid if $M - mc - \frac{E}{N} + \delta \geq 0$, therefore, a high minimum profit requirement E excludes sponsors that cannot generate enough profit. For this reason, I call E the recordkeeper's exclusion threshold.

The direct effect of exclusion. With the exclusion, the recordkeeping fee becomes

$$T = \max \left\{ 0, (1-\eta)\delta - \eta \left(M - mc - \frac{E}{N} \right) \right\} \tag{9}$$

The recordkeeping fee T increases in the exclusion E given that the Nash bargaining is valid. Notice that the exclusion E is not a real cost; therefore, the joint surplus of the Nash bargaining is still $M - mc + \delta$, the same as in the baseline model. By exercising exclusion, the recordkeeper raises the negotiated recordkeeping fee and thereby captures a larger share of the total surplus.

The indirect effect of exclusion. Exclusion can also induce the sponsor to make a different bundle choice. The following example demonstrates the intuition.

[Insert Figure 8 here]

Example 4.3.1. Continue with Example 4.2.3. Let $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - 0.5M$. Assume that $E = 5,000$ and the sponsor has 100 participants. Panel (c) of Figure 8 compares the recordkeeping fee T, the sponsor's net utility u, and the recordkeeper's profit π with and without the exclusion.

Suppose the recordkeeper provides two bundles A and B. Bundle A has $M = 50$ and bundle B has $M = 100$. The table below shows the utility δ, recordkeeping fee T, and the net utility u of each bundle. As shown in the table, without the exclusion, the net utility of the bundle A is 62.5 dollars and the net utility of bundle B is 50 dollars; with the exclusion, the net utility of bundle A is 37.5 dollars, and the net utility of bundle B is 50 dollars. Therefore, the sponsor chooses bundle A, the bundle that yields a higher utility δ when there is no exclusion, and chooses bundle B vice versa.

			Without exclusion		With exclusion	
	M	δ	T	u	T	u
A	50	75	12.5	62.5	37.5	37.5
B	100	50	0	50	0	50

The exclusion does not necessarily result in the change of bundle choice, as illustrated in Example 4.3.2.

Example 4.3.2. Continue with Example 4.2.5. Let $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - 1.5M$. Assume that $E = 5,000$ and the sponsor has 100 participants. Panel (e) of Figure 8 compares T, u, and π with and without exclusion in this case. Since the sponsor always prefers the bundle with a lower M, the exclusion cannot shift the sponsor's bundle choice.

4.4 Effect of Exclusion on Sponsors of Different Sizes

Focus on the case where the exclusion is able to affect the sponsor's bundle choice. How would this effect differ by sponsor size? Intuitively, sponsors that are sufficiently large will not be affected by the exclusion, because the per-participant exclusion $\frac{E}{N}$ approaches zero as N goes large enough. Meanwhile, sponsors that are small enough cannot meet the requirement of exclusion, because even if the small sponsor transfers all the surplus to the recordkeeper, the total profit of the recordkeeper is still below the exclusion.

[Insert Figure 9 here]

Example 4.4.1 Continue with the setting in Example 4.3.1. Given the recordkeeper's exclusion $E = 5000$. Figure 9 plots the utility u as a function of M with and without exclusion for sponsors of different sizes.

Comparing Panel (a) and Panel (b) of Figure 9, as the sponsor's size N increases, the curve with exclusion gradually approaches the curve without exclusion, meaning that the exclusion has a smaller impact on the sponsor. Meanwhile, if the sponsor is small enough, Nash bargaining always fails, and so the sponsor will always choose the outside option and receive a zero utility, as shown in Panel (c) of Figure 9. In this sense, the small sponsor is excluded from the recordkeeper's client pool.

Consider again the two bundles where bundle A has $M = 50$ and $M = 100$. Panel (c) of Figure 9 shows that a large sponsor with $N = 100$ still chooses bundle A with exclusion; a mid-sized sponsor with $N = 50$ is induced to choose bundle B, while a small sponsor

with $N = 20$ is excluded from the market. Figure 10 shows how the bundle choices differ by sponsor size given the same exclusion.

[Insert Figure 10 here]

4.5 Recordkeeper's Optimal Exclusion

The recordkeeper can increase its profit by setting an exclusion threshold. What should be the optimal exclusion chosen by a recordkeeper?

[Insert Figure 11 here]

Once a recordkeeper commits to a particular exclusion threshold, the exclusion threshold applies to the entire market. As discussed above, the exclusion has different effects on sponsors of different sizes. In this simple setting, increasing exclusion has two effects: the recordkeeper could extract a higher proportion of surplus from sponsors who still get a positive surplus, while small sponsors will step away. The recordkeeper must balance the two effects. The recordkeeper chooses the exclusion to maximize total profit across the market. Denote by $\pi(E; N, M)$ the recordkeeper's profit from a sponsor of size N when providing the bundle M and setting an exclusion of E, the recordkeeper's total profit is thus

$$\max_E \Pi = \int_N N \cdot \sum_M \pi(E; N, M) dF(N). \tag{10}$$

The optimal exclusion depends on various factors. First, it depends on the size distribution of all sponsors. In Figure 11, I plot the optimal exclusion against the average sponsor size. The sponsor size distribution is modeled as a beta distribution on $[0, 100]$. Figure 11 shows that there are more larger sponsors, the recordkeeper will set a higher exclusion.

Second, the effect of exclusion also depends on how the sponsor's utility responds to changes in M. Exclusion can induce the sponsor to choose a more expensive portfolio only if the sponsor is not overly sensitive to M. In Panel (a) of Figure 11, I plot the optimal exclusions for different values of the sensitivity parameter k. The figure shows that, holding the sponsor size distribution constant, the optimal exclusion decreases as the sponsor's sensitivity to management fees M increases.

Third, it also depends on the set of bundles offered to sponsors. In Panel (b) of Figure 11, I plot the optimal exclusion for different bundles M. When there are more large sponsors, the optimal exclusion is higher for a larger M. In contrast, when most sponsors are small,

the optimal exclusion can be lower with larger M. The intuition is as follows: with a higher M, a larger fraction of small sponsors can afford to participate without paying a direct transfer, making the plan feasible for more of them. Consequently, increasing the exclusion in this case would drive more small sponsors out of the market compared to a setting with a smaller M, where many small sponsors are already excluded due to infeasibility.

5 Structural Model

In this section, I present the structural model for estimation. On top of the illustrative model, in the full structural model I consider also the competition between recordkeepers, sponsors' limited consideration, and recorkdeepers' the benefits and costs from a larger market share.

5.1 Model Setting

Demand. Sponsor s sets a DC plan for N_s participants. The sponsor chooses a recordkeeper $r \in \mathcal{R}_s$ and a plan $p \in \mathcal{P}_{rs}$, where $\mathcal{R}_s$ is the set of potential recordkeepers that sponsor s will consider, and $\mathcal{P}_{rs}$ is the set of plans provided by recordkeeper r for sponsor s to choose from.

If sponsor s chooses recordkeeper r and plan p, given direct transfer of recordkeeping fees T, the per-participant utility is

$$u_{srp}(T) = X_p\beta - T + \delta_r + \epsilon_{srp} \tag{11}$$

which depends on plan characteristics X_p, the direct transfer T, the recordkeeper fixed effect δ_r and the idiosyncratic taste shock ϵ_{srp}. The shock takes a nested logit structure

$$\epsilon_{srp} = \sigma(\nu_{sr} + \rho\omega_{sp}), \tag{12}$$

where $\omega_{sp} \sim EVT1(0,1)$ and $\nu_{sr} + \rho\omega_{sp} \sim EVT1(0,1)$.

Supply. If recordkeeper r provides plan p to sponsor s, the recordkeeper will receive three sources of revenue: a direct transfer T of recordkeeping fees, an indirect transfer R_{rsp} of revenue sharing, and management fees M_{rsp} for all affiliated investment products. Here, T, R, and M are on a per-participant basis. Besides, recordkeeper r has per-participant

marginal cost mc_r. The per-participant profit for providing plan p to sponsor r is thus

$$\pi_{rsp}(T) = T + R_{rsp} + M_{rsp} - mc_r \tag{13}$$

The corresponding total profit is thus

$$\Pi_{rsp}(T) = N_s(T + R_{rsp} + M_{rsp} - mc_r)$$

Bargaining and equilibrium. Given the chosen plan p, the recordkeeper and the sponsor bargain over the direct transfer T:

$$\max_{T \geq 0} \left(\pi_{rsp}(T) - \frac{E_r}{N_s} \right)^{\eta} (u_{srp}(T) - u_s^{sd})^{1-\eta}, \tag{14}$$

Outside options are crucial in Nash bargaining and are determined in the equilibrium. Similar to Bhattacharya and Illanes (2022), I apply the notion of Nash-in-Nash with the threat of replacement as in Ho and Lee (2019) and assume that each sponsor is its own market.

Sponsor's outside option. With this equilibrium notion, the sponsor's outside option is equal to the payoff it can get from the second-best recordkeeper by making the second-best recordkeeper indifferent to providing the service or not, that is,

$$\begin{aligned} u_s^{sd} = & \max_{r' \neq r, p, T \geq 0} u_{sr'p}(T) \\ s.t. \quad & \pi_{r'sp}(T) \geq 0 \end{aligned}$$

In the equilibrium, the transfer that is necessary to make the recordkeeper indifferent is called the reservation transfer. For a given sponsor s, recordkeeper r, and a plan p, the reservation transfer is

$$T_{srp}^{res} = \max \left\{ 0, \frac{E_r}{N_s} + mc_r - R_{rsp} - M_{rsp} \right\} \tag{15}$$

The first condition requires that the transfer should be only from the sponsor to the recordkeeper. The second condition says that the amount of transfer has to at least satisfy the exclusion threshold.

Given the reservation transfer, the maximum extractable utility from recordkeeper r and plan p is

$$u_{srp}^{res} = u_{srp}(T_{srp}^{res}) \tag{16}$$

Recordkeeper's outside option. Each recordkeeper has an exclusion threshold E_r that restricts the total profit Π_{rsp} from a single sponsor to be at least E_r. Importantly, E_r is not a real cost. It is different from the case in Grennan (2013) where the outside option of the supplier is assumed to be its fixed cost. I assume that E_r is a commitment and is optimally set by the recordkeeper to maximize its total expected profit across the whole market. I will discuss the recordkeeper's optimal choice of exclusion later in detail.

Sponsor's recordkeeper and plan choice. In the equilibrium, given a set of potential recordkeepers $\mathcal{R}_s$ and the set of plans provided by each recordkeeper $\{\mathcal{P}_{rs}\}_{r\in\mathcal{R}_s}$, sponsor s chooses the plan p that yields the highest utility under the reservation transfer and the recordkeeper that provides the chosen plan, that is, if $u^{res}_{srp} > u^{res}_{sr'p'}$ for all $(r', p') \neq (r, p)$. Let C_{srp} be the indicator that recordkeeper r and plan p is chosen by sponsor s, with the nested logit shock, the probability that sponsor s chooses recordkeeper r and plan p is

$$Pr(C_{srp}=1|\mathcal{R}_s) = \underbrace{\frac{\left[\sum_{p\in\mathcal{P}_{rs}} \exp(\frac{u^{res}_{srp}}{\rho\sigma})\right]^{\rho}}{1+\sum_{r'\in\mathcal{R}_s}\left(\sum_{p\in\mathcal{P}_{r's}}\left[\exp(\frac{u^{res}_{sr'p}}{\rho\sigma})\right]^{\rho}\right)}}_{\text{Prob. that } r \text{ is chosen}} \cdot \underbrace{\frac{\exp(\frac{u_{srp}}{\rho\sigma})}{\sum_{p\in\mathcal{P}_{rs}}\exp(\frac{u_{srp}}{\rho\sigma})}}_{\substack{\text{Conditional on } r \text{ is chosen,}\\ \text{the prob. that plan } p \text{ is chosen.}}} \tag{17}$$

Consideration set formation. Recordkeepers (and econometricians) do not have precise information about $\mathcal{R}_s$ of each sponsor before the sponsor reaches out to the recordkeeper. Instead, recordkeepers have prior information on the potential set of recordkeepers $\bar{R}_s$ that each sponsor s may reach out to, which I refer to as the prior consideration set, and I assume that these potential sets are common knowledge.

Sponsor s considers a subset $\mathcal{R}_s \subseteq \bar{R}_s$ in the run-for-proposal process. Each recordkeeper $r \in \bar{R}_s$ will be included in $\mathcal{R}_s$ with some probability q_{rs}. I assume that q_{rs} is positively related to the utility, i.e., recordkeepers that potentially could yield higher utility to the sponsor have higher probabilities of being included in the consideration set. I model q_{rs} as follows: first, denote by $E[u^{max}_{sr}]$ the expected maximum utility sponsor s can get from recordkeeper r, that is

$$\begin{aligned} E[u^{max}_{sr}] &= \max_{p\in\mathcal{P}_{rs}} u^{res}_{srp} \\ &= \sigma\left[\rho\log\left(\sum_{p\in\mathcal{P}_{rs}}\left(\frac{u^{res}_{srp}}{\rho\sigma}\right)\right)+\gamma\right] \end{aligned} \tag{18}$$

Then let u^{max}_s be the maximum of $E[u^{max}_{sr}]$ among all recordkeepers in the prior set $\bar{R}_s$,

$$u^{max}_s = \max_{r\in\bar{R}_s} E[u^{max}_{sr}] \tag{19}$$

The probability that recordkeeper r is included in the consideration set $\mathcal{R}_s$ is

$$q_{rs} = Pr\left(\frac{E[u_s^{max}] - u_{sr}^{max}}{u_s^{max}} > \epsilon_c\right) \tag{20}$$

where $\epsilon_c \sim N(\mu_c, \sigma_c^2)$. Therefore, given the prior set $\bar{R}_s$, the probability that $\mathcal{R}_s$ is the actual consideration set is

$$Pr(\mathcal{R}_s) = \prod_{r \in \bar{R}_s} q_{rs}^{1(r \in \mathcal{R}_s)} (1 - q_{rs})^{1-1(r \in \mathcal{R}_s)} \tag{21}$$

Recordkeeper's shadow value of capacity. Besides profit from providing DC plans, I consider also other factors that affect a recordkeeper's total profit. On the one hand, some recordkeepers can provide other services to sponsors. For example, a large proportion of the recordkeepers are insurance companies, and are able to provide different types of insurance products to firms. Meanwhile, recordkeepers such as ADP and Paychex are payroll service companies, they typically bundle their payroll services with the recordkeeping services.

On the other hand, recordkeepers may have additional costs associated with their DC plans services. Expanding the business often requires investment in infrastructure, such as data systems or regional offices to serve clients in different locations. Depending on their business model, there can be an increasing or decreasing return to scale. For instance, they need to open local offices to access clients in the region. Insufficient infrastructure can lead to capacity constraints. When recordkeepers are capacity-constrained, there will be a shadow cost for providing services to the additional sponsor (Wenning (2024)).

Importantly, these additional revenues and costs are not directly tied to the provision of DC plan services and therefore do not enter the bilateral Nash bargaining. Instead, they are related to the overall scale of each recordkeeper's business. Accordingly, I assume that the net profit from these ancillary factors is proportional to the total number of plan participants served by each recordkeeper. Because there can be either additional revenues or costs, I refer to the net profit from the ancillary factors as the shadow value of capacity.

The expected total number of plan participants of recordkeeper r is

$$\begin{aligned} E[N_r(E_r, \mathbf{E}_{-r})] &= \sum_{s \in \mathcal{S}} N_s \cdot E[s_{rs}(E_r, \mathbf{E}_{-r})] \\ &= \sum_{s \in \mathcal{S}} 1(r \in \bar{R}_s) \cdot N_s \cdot \sum_{\mathcal{R}_s \subseteq \bar{R}_s} \sum_{p \in \mathcal{P}_{rs}} Pr(C_{rsp} = 1 | \mathcal{R}_s) Pr(\mathcal{R}_s) \end{aligned} \tag{22}$$

The additional profit is $\alpha_r \cdot E[N_r(E_r, \mathbf{E}_{-r})]$.

Recordkeeper's optimal exclusion. Recordkeepers optimally set their exclusions to maximize their total profit across the market. The expected total profit net of the shadow value of capacity is

$$\begin{aligned} E[\Pi_r(E_r, \mathbf{E}_{-r})] &= \sum_{s\in\mathcal{S}} E[\Pi_{rs}(E_r, \mathbf{E}_{-r})] \\ &= \sum_{s\in\mathcal{S}} 1(r \in \bar{R}_s) \sum_{\mathcal{R}_s \subseteq \bar{R}_s} \sum_{p\in\mathcal{P}_{rs}} E[\Pi_{rsp}|C_{rsp}=1, \mathcal{R}_s] Pr(C_{rsp}=1|\mathcal{R}_s) Pr(\mathcal{R}_s) \end{aligned} \tag{23}$$

The recordkeeper maximizes the net total profit

$$\max_{E_r} E[\Pi_r^*(E_r, \mathbf{E}_{-r})] = E[\Pi_r(E_r, \mathbf{E}_{-r})] - \alpha_r \cdot E[N_r(E_r, \mathbf{E}_{-r})] \tag{24}$$

Notice that each component in the expected total profit depends on both the recordkeeper's exclusion and other recordkeepers' exclusions. Meanwhile, every recordkeeper plays the same game to set its exclusion. So the exclusions of recordkeepers are jointly determined in the equilibrium.

5.2 Analysis of the Model

Discussion of exclusion. I assume that exclusions are commitments made by the recordkeepers. On the one hand, it means that each recordkeeper applies the same threshold to the whole market and is unable to set different exclusion thresholds for different sponsors. This is similar to a two-part tariff, which is a type of second-degree price discrimination.

On the other hand, it is important to note that the exclusion is not a real cost; instead, it can be viewed as an opportunity cost. The DC plan market is a large market, so I assume that there always exists an alternative for sponsors of any size. When the recordkeeper commits to its exclusion threshold, the exclusion threshold is also the minimum profit the recordkeeper could receive from an alternative sponsor, so it can serve as the disagreement payoff of the recordkeeper in the Nash bargaining.

Although in a different setting, treating exclusion as a commitment is similar to the idea in Barry, Carlin, Crane, and Graham (2024), where a firm can commit to a high hurdle rate, thus a high disagreement payoff in the Nash bargaining, in order to gain an advantage over counterparties, at the cost of forgone projects.

Discussion of consideration set formation. In reality, since the RFP process is costly (Yang (2023)), the sponsor will not contact too many recordkeepers. Sponsors do not randomly choose the set of potential recordkeepers. I assume that sponsors contact a

subset of recordkeepers that bring the highest expected utility from a large, uninformative prior set of recordkeepers.

This consideration set formation process has two effects. First, it alters the size distribution of sponsors that include a given recordkeeper in their consideration set. As shown in Section 4, this distribution plays an important role in determining the recordkeeper's optimal exclusion threshold. Intuitively, a recordkeeper is more likely to set a higher exclusion threshold when the sponsors that consider it are, on average, larger. Under an uninformative prior consideration set, the sponsor size distribution would be similar across recordkeepers. However, once sponsors selectively consider only the recordkeepers that provide the highest expected utility, these distributions diverge. For instance, if a recordkeeper sets a high exclusion threshold, its expected utility will be lower for small sponsors, making them less likely to include it in their consideration set. Consequently, the size distribution of sponsors that consider this recordkeeper shifts toward larger sponsors.

Accordingly, exclusion influences recordkeeper selection through two distinct mechanisms. First, it lowers the conditional probability of being chosen given inclusion in the consideration set. Second, it reduces the probability of being included in the consideration set itself. As a result, relative to the setting with an exogenously fixed consideration set, exclusion has a larger overall effect on total profit. An increase in the exclusion threshold allows a recordkeeper to attract larger sponsors on average but simultaneously reduces its market share through the consideration set formation process. The optimal exclusion thus reflects a trade-off between these two opposing forces.

Discussion of the nested logit shock. The sponsor chooses a recordkeeper from a set of recordkeepers and each recordkeeper could provide a set of plan candidates. The choice structure naturally leads to a nested logit model. The nested logit model also helps to alleviate the issue with a logit model that the probability of a particular recordkeeper being chosen increases mechanically with the number of plans it provides. In the nested logit model, if the nested parameter $\rho = 1$, it reduces to the simple logit model; if $\rho = 0$, the utility of each nest (called the inclusive value) is equal to the maximum utility among all options, which is independent of the number of options in the nest. I do not rule out the possibility that some recordkeepers are chosen because they are able to provide a spectrum of plans, and I let the model decide how this matters.

Discussion of the shadow value of capacity. Profits outside the recordkeeping business and investment products play an important role in determining each recordkeeper's optimal exclusion. These revenues and costs are related to the provision of DC plan services but are not directly bundled with them and therefore do not enter the Nash bargaining. In the model, I assume that the shadow value of capacity is linear in the

total number of participants served by each recordkeeper. However, this should not be interpreted as an additional profit realized simultaneously with the DC plan services. When there is a net cost, it may reflect infrastructure investments that are not specific to any particular sponsor. Conversely, ancillary profits may not occur concurrently with DC plan activities. For instance, while ADP often provides recordkeeping jointly with its payroll services, John Hancock, one of the largest insurance providers, typically does not bundle recordkeeping and insurance services at the same time. Hence, these additional profits should be interpreted as average, rather than contemporaneous, returns.

5.3 Identfication and Estimation

5.3.1 Identification

With the nested logit structure, the identification of preference parameters β, recordkeeper fixed effects δ_r, nested logit parameters ρ and σ follow standard arguments in discrete choice models. I normalize the value of the outside option (not to provide a DC plan) to zero, and I normalize the parameter on direct transfer T to 1.

The consideration set parameters μ_c and σ_c are identified via plan choice probabilities. A larger μ_c shifts weights towards consideration sets that contain only the recordkeepers with the highest expected utility. Since the normalized utility is between -1 and 0, I assume that $-1 < \mu_c < 0$. Given μ_c, a smaller σ_c amplifies the effect. If σ_c is close to zero, only recordkeepers with expected utility above μ_c will be incorporated in the consideration set; if σ_c goes to infinity, all recordkeepers will be considered.

The identification of the bargaining parameter η follows the argument in Grennan (2013), that the incremental surplus identifies the bargaining parameter. Larger μ_c shifts more weights to recordkeepers with higher expected utilities.

Finally, the differences between the model-implied size distribution of sponsors that choose a recordkeeper and the actual size distribution help to identify the additional cost parameter α_r. A positive α_r will induce a lower exclusion because the recordkeeper benefits from a large market share, while a negative α_r results in higher exclusion because of the additional cost on top of the DC plan services.

5.3.2 Estimation

The model can be estimated using maximum likelihood. In each iteration step, given the parameters, the algorithm solves for each recordkeeper's optimal exclusion, then calculates the joint likelihood of plan choice probabilities.

In practice, it is computationally costly to solve for optimal exclusions. I apply a three-step estimation procedure. First, notice that given exclusions E, the plan choice probabilities do not depend on η and α. In the first step, given parameters other than η and α's, I implement a fixed-point algorithm to jointly solve for δ's and E's that match both the market shares and the average client size of each recordkeeper. I refer to the exclusions obtained from the fixed-point algorithm as the data-implied exclusions. Solving for the fixed point does not involve solving the optimal exclusions. With the data-implied exclusions, I can calculate the likelihood and solve for parameters other than η and α's by maximum likelihood.

Second, with the estimated parameters and the data-implied exclusions, I can solve for the expected direct transfer for each chosen plan for a given η. I estimate η by matching the observed and expected direct transfers

$$\min_{\eta} \sum_{s} (T_s - E[T_s|(r^*, p^*)])^2 \tag{25}$$

Finally, given the parameters estimated in the first step and the second step, for each set of α's, I can solve for optimal exclusions. I refer to these exclusions as model-implied exclusions. I then choose α's that minimize the differences between data-implied exclusions and model-implied exclusions. In sum, the estimation algorithm matches key data moments inside a maximum likelihood estimation.

6 Estimation Results

6.1 Sample Construction

The main sample uses plans with participants between 100 to 5000 in 2019. I exclude plans if (1) the plan does not report recordkeeper information; (2) BrightScope does not provide the plan's investment menu; (3) the employer has multiple DC plans; (4) the plan sponsor's state or industry information is missing; (5) the plan sponsor is located in the U.S. territories.[7]

I include sponsors providing only DB plans as sponsors who choose outside options as those who choose outside options. These plans do not have information on the recordkeeper and the investment menu; otherwise, I apply the same filters as to the main sample. The final sample contains 51,429 DC plans and 1,885 outside option plans.

To estimate the model, I need information on each sponsor's recordkeeper choice set and

[7] Precisely, I exclude samples in American Samoa (AS), Guam (GU), the Northern Mariana Islands (MP), Puerto Rico (PR), the U.S. Virgin Islands (VI), and Trust Territory of the Pacific Islands (TT)

plan choice set. However, the data only contain the chosen recordkeepers and plans, so it is necessary to construct prior choice sets. Because the sponsor size distribution plays a key role in determining the optimal exclusion, the prior choice sets must be constructed in an uninformative manner such that the size distributions of potential sponsors are approximately the same across recordkeepers.

I construct the prior choice set for each sponsor in the following way: first, similar to Bhattacharya and Illanes (2022), I assume that potential recordkeepers of a sponsor are among those who serve other sponsors in the same state and the same industry (with the same 2-digit NAICS code). Then I rely on geographic information to select the prior set of recordkeepers for each sponsor.[8] First, for each potential recordkeeper, I calculate the distances between the sponsor and other firms served by the recordkeeper in the same state and industry. The smallest distance is defined as the distance between the sponsor and the recordkeeper. Second, I define neighboring plans as those with a distance of less than 100 km and count the number of neighboring plans of each recordkeeper. I sort the recordkeepers by the distance and the number of neighboring plans. I keep the all top 10 recordkeepers, 3 in the non-top-10 integrated recordkeepers, and 3 in the non-top-10 pure recordkeepers, so at maximum, I have 16 recordkeepers in the prior choice set for each sponsor.

I then build the plan choice set for each recordkeeper in the sponsor's choice set. There are two issues with the plan choice set. The first issue is selection bias: I observe only plans that were chosen, while the chosen plans could be biased towards good plans. Among all plans with investment menu information available, 70% have an S&P 500 tracker and 82.4% have target-date funds. It is possible that lower-quality plans were available but not selected, leading to an upward bias in observed plan quality. The second issue is that with the nested logit model, the probability that a recordkeeper is chosen mechanically increases with the total number of plans provided by the recordkeeper. Although the nested-logit model can alleviate the issue, a large variation in the number of plans by each recordkeeper could affect the estimation.

To address these issues, the plan choice set is constructed as follows: given a pair of sponsor and recordkeeper, I group the plans provided by the recordkeeper to sponsors in the same state and industry into quintiles based on management fees; I keep one plan with the closest number of participants in each quintile, making sure that there is enough dispersion in fees. Furthermore, if all selected plans are good plans, meaning that they all have at least one S&P 500 tracker and target-date fund, I add one plan without S&P 500 trackers and one plan without target-date funds from all plans provided by the recordkeeper. Conversely, if all selected plans lack these preferred investment options, I add two plans that contain them. This procedure ensures that each choice set includes

[8] In Appendix D, I plot the geographic distribution of plans served by each recordkeeper.

variation in both management fees and investment menu quality.

The structural model endogenizes the recordkeeping fees and requires information on management fees and revenue sharing that could have been received by the recordkeeper for each plan in the sponsor's choice set. I follow Bhattacharya and Illanes (2022) and predict the management fees and revenue sharing using XGBoost with characteristics of the recordkeeper, the sponsored, and the menu. Appendix A gives details on the sample construction and prediction.

Finally, the plan characteristics in the sponsor's utility function include the number of options in the menu, the number of BrightScope investment categories, whether the plan contains an S&P 500 tracker, whether the plan has target-date funds, the proportion of index funds, and the expected management fees.

6.2 Estimation Details

I pool all plans and estimate the model. The parameters of interest include the preference parameter β, the variance of the nested logit shock σ, the nested parameter ρ, the Nash bargaining parameter η, and the consideration set parameters μ_c and σ_c. Besides, I estimate the recordkeeper-specific marginal costs mc_r for top-10 recordkeepers, and two separate marginal costs for non-top-10 integrated recordkeepers, and non-top-10 pure recordkeepers.

Recordkeepers play a Bertrand game to determine their optimal exclusion thresholds. For computational tractability, only the top 10 recordkeepers are modeled as active players in this game, while the exclusions of non–top-ten and pure recordkeepers are treated as parameters. In the three-step estimation procedure, the parameters δ and α are obtained through the Berry inversion.

To reduce computational burden, I assume that sponsors apply the consideration set formation process only to the top ten recordkeepers. All smaller recordkeepers in the prior set are included with probability one. Moreover, I restrict attention to consideration sets containing between one and five of the top ten recordkeepers, so there are at most $\sum_{k=1}^{5} \binom{10}{k} = 637$ consideration sets.

6.3 Results

[Insert Table 2 here]

Table 2 presents the estimated model parameters. Sponsors prefer plans with diversified investment products, plans with S&P 500 index trackers and target-date funds; on the

other hand, sponsors dislike plans with too many investment options, consistent with the findings in Iyengar et al. (2003) that participants could be overwhelmed with too many investment options. While sponsors dislike plans with high fees, the coefficient on the expected indirect transfer is only -0.14, meaning that one dollar increase in dollar management fees brings about 0.14 dollar loss in utility. This shows that on average plan sponsors are not very sensitive to plan fees. This insensitivity is largely due to the estimation pools small plans and large plans together. From the theory model in Section 4, this relatively low insensitivity is necessary for exclusions to be effective in shifting sponsors' plan choices.

The nested logit model has a standard deviation of 143.4 and a nested parameter of 0.75, indicating that the provider and plan choice involve a lot of randomness. Given that the mean utility of the top 10 recordkeepers ranges from \$200 to \$400, the randomness is not too large. The nested parameter $\rho = 0.5$ indicates that the choice probability does increase in the number of options, but not in an absolute mechanical form. The Nash bargaining power parameter η is around 0.5, meaning that on average, sponsors and recordkeepers roughly have the same bargaining power.

The normal distribution for the consideration set formation has mean $\mu_c = -0.1$ and standard deviation $\sigma_c = 0.30$. Panel (a) of Figure 12 plots how the probability of being included in the consideration changes in the normalized utility. The top recordkeeper has a 64% chance of being selected into the actual consideration set, while the recordkeeper with half of the top provider's normalized utility has only about 10% chances of being included in the consideration set. Panel (b) of Figure 12 plots the empirical probability of being included in the consideration set by the rank of the recordkeepers. Consistent with the theoretical distribution, the top recordkeeper in the data has 64% chance of being considered, while the probability that 5th recordkeeper is considered is 35%, and the 10th recordkeeper has only 1% probability of being considered.

[Insert Figure 12 here]

Regarding the recordkeeper-specific parameters, Figure 13 plots each recordkeeper's exclusion against its mean utility. It's striking that the exclusions are strongly positively correlated with mean utilities. This is consistent with the intuition and the reality that recordkeepers offering high-quality services tend to set high exclusions and attract larger sponsors. In terms of the scale of exclusions, Fidelity, the largest recordkeeper, has an exclusion threshold of \$51,833, while Vanguard, the largest institutional investor, has an exclusion threshold of \$29,734. Given that the average plan expense ratio is roughly 50 basis points, these exclusion thresholds translate into minimum plan balances on the order of several million dollars, consistent with anecdotal evidence of Vanguard's minimum plan size requirement for a single DC plan.

On the other hand, there is no evidence that ADP, the only non-integrated recordkeeper among the top 10, has an exclusion, in line with the fact ADP is a payroll service company that primarily serves small clients. Among the top 10 recordkeepers, Allspring and Charles Schwab are the type of niche recordkeepers. They have high mean utilities and high exclusions, mainly serving large clients and having small market shares.

I also estimate the exclusion thresholds for four types of non-top-10 recordkeepers. In general, both small intergated and small non-integrated recordkeepers have no exclusion, while large integrated recordkeepers have larger exclusions than large non-integrated recordkeepers. This is in accordance with the story that integrated recordkeepers have the incentive to raise theirs exclusion to promote bundle selling of their own investment products.

[Insert Figure 13 here]

The shadow value of capacity α is important in determining the optimal exclusion for the top 10 recordkeepers. These shadow values should be viewed as the joint effect of capacity constraints and related businesses. Among the top 10 recordkeepers, most of them have a negative α, indicating that the shadow costs due to capacity constraints exceed the profits from related businesses. For example, Fidelity has an α of \$103, meaning that the opportunity cost for Fidelity to serve one more participant is 103. If there is no additional profit from additional businesses, given Fidelity's exclusion around \$51,000, it implies that Fidelity's next best alternative client is a plan with about 50 participants. On the other hand, ADP has positive α's, indicating that its profits from payroll services surpass the shadow cost of capacity constraint. This is largely consistent with the fact that ADP is one of the largest payroll service companies that usually bundles the DC plan services with their payroll services.

[Include Table 3 here]

To show the effects of the shadow value of capacity, Table 3 reports the optimal exclusion for each of the top 10 recordkeepers, setting their $\alpha_r = 0$, while holding α for other recordkeepers fixed. Recordkeepers with negative α now set lower exclusions, since they are no longer capacity constrained. The magnitude of the reduction in exclusion is larger for recordkeepers that were more constrained initially. ADP has a slightly higher exclusion as it does not benefit from other businesses.

6.4 Model Fit

I test the model fit by simulating data with the estimated parameters and I compare the key moments of the simulated data and the actual data. Overall, the model matches the

key data moments.

[Include Figure 14 here]

Figure 14 compares market shares, recordkeeping fees, and management fees generated from the simulated data to their counterparts in actual data. Panel (a) of Figure 14 shows that the model almost perfectly matches the market shares of the top 10 recordkeepers and the four types of non-top 10 recordkeepers. This follows the Berry inversion in which the recordkeeper fixed effects are adjusted to match their market shares. Furthermore, Panel (a) of Table 4 presents the average client sizes of the top 10 recordkeepers. The results indicate that the model provides a good fit to the observed distribution of client sizes.

The model also captures the main patterns of direct and indirect transfers. In Panel (b) of Figure 14, I fit the relationship between the average management fees of chosen plans and plan sizes; in Panel (c) of Figure 14, I fit the relationship between the recordkeeping fees and plan sizes in the simulated data. The model generates the pattern that larger sponsors tend to have plans with lower recordkeeping fees and cheaper investment options. Notice that my model does not assume that small and large sponsors differ in their sensitivity to fees; therefore, the pattern is mainly the result of exclusions.

6.5 Decompose the Effects of Exclusion

In the model, exclusion has two effects on recordkeeper choice: it affects both the probability that a recordkeeper is included in a consideration set, and the probability of being chosen conditional on being included in the consideration set. In Table 4, I decompose the two effects.

[Insert Table 4 here]

Panel (a) of Table 4 decomposes the effect of exclusions on determining the average client size of each recordkeeper. The first column displays the average size of all sponsors that include each recordkeeper in their prior consideration set. By construction, these prior means are similar for all recordkeepers, around 570 on average, which confirms that prior sets are uninformative. The second column presents the expected mean of all sponsors that include the recordkeeper in their consideration sets (not the prior set). The distribution shifts towards the actual distribution of client size observed in the data. This is because although recordkeepers with high mean utility before transfer attract all recordkeepers, small sponsors need to pay a high transfer due to exclusion, so their expected

utility after transfer with a large recordkeeper is lower, and hence they put a larger weight on smaller recordkeepers when forming the consideration set. The third column shows the expected average client size of each recordkeeper predicted by the model, which takes into both the effect of consideration sets and the effect of recordkeeper choice given a consideration set. The distribution shifts further towards the actual distribution in the data, which is presented in the last column, where recordkeepers with high mean utility tend to serve large sponsors.

Panel (b) of Table 4 decomposes the effect of exclusions on the market share. The first column reports the percentage of all sponsors that include each recordkeeper in their prior consideration set. These values reflect the sample construction: on average, a larger proportion of sponsors include the top recordkeepers in their prior sets. This pattern likely reflects a brand or visibility effect, as potential sponsors are more familiar with large recordkeepers that already serve a substantial share of the market. However, the differences in these prior inclusion rates are much smaller than the differences in observed market shares, indicating that the model-predicted market share rankings are not mechanically driven by the prior construction.

The second column presents the chance of being included in the consideration set across the market. For example, while 86.5% of sponsors have Fidelity in their prior set, on average, Fidelity shows up in consideration sets with a chance of 57%. From the prior set to consideration, the probability of being considered shrinks by about one-third for Fidelity, while it shrinks about two-thirds for Charles Schwab, highlighting the effect of consideration set formation.

The third column reports the average probability of being chosen conditional on being included in the consideration set. These probabilities do not show a large variation, as sponsors only contact recordkeepers offering relatively high expected utility, which reduces the variance of recordkeeper qualities in the consideration set is smaller compared to the prior set. The last column shows the model-predicted market shares, which equals the product of the probabilities in the first three columns.

7 Counterfactual Analysis

7.1 Effects of Exclusion

How does the exclusion affect the market outcomes? To examine the effects of exclusion, in the first counterfactual analysis, I remove the exclusions of all recordkeepers and re-simulate sponsors' recordkeeper and plan choices.

[Insert Figure 15 here]

Figure 15 compares the market outcomes under no exclusion to those observed in the data. Panel (a) plots the market shares in both cases. Without exclusions, the overall provision of DC plans increases from 96.7% to 97.5%. Recordkeepers that yield high mean-utilities expand in market share. Among the three recordkeepers with the highest exclusion, Fidelity's market share increases from 15.8% to 20%; Allspring's market share increases from 2.0% to 4.0%, and Charles Schwab's market share increases from 1.6% to 3.7%. On the other hand, recordkeepers with low or no exclusions initially experience declines in market share. For example, John Hancock's exclusion drops from 5.3% to 3.3%, while ADP's market share drops from 4.2% to 2.2%. This proves the exclusions are important frictions in this market that prevent sponsors from accessing high-quality recordkeeping services.

Panels (b) and (c) compare the expected management fees and expected recordkeeping fees in the counterfactual analysis against the data. Without exclusions, sponsors on average switch to plans with both lower management fees and lower recordkeeping fees. The effect is especially pronounced for small sponsors: for a sponsor with 100 participants, the management fee drops from 57 basis points to 50 basis points, and the recordkeeping fee drops from about \$100 to less than \$30 per participant. This counterfactual analysis confirms both the direct and indirect effects of exclusion: that exclusion results in both higher recordkeeping fees and induces the small sponsors towards more expensive plans.

Panel (d) further contrasts the expected welfare with and without the exclusions. It's phenomenal that the expected welfare of small sponsors increases far more than large sponsors. For sponsors with more than 1500 participants, there is no significant increase in expected welfare, suggesting that large sponsors are largely unaffected by exclusion thresholds.

Removing exclusions also largely impacts recordkeepers' profits. Panel (e) plots the percentage changes in total profit and per-participant profit if all exclusions are removed. It's striking that total revenues decline for all recordkeepers and the scale is large. For example, if there is no exclusion, Fidelity's total profit will decline by 30% and the per-participant profit will decrease by 48%. In general, the profit drops more for recordkeepers with large exclusions ex ante. The per-participant profit increases for plans with low or no exclusions, for example, John Hancock and Fidelity. This is largely due to the drop in their market shares, since small plans have easier access to recordkeepers with high-quality services.

Overall, removing exclusions increases the provision rate of DC plans, induces sponsors to select plans with lower management and recordkeeping fees, and enhances welfare, particularly for small sponsors. This counterfactual analysis confirms that exclusions

constitute a major friction in the DC plan market.

[Insert Table 5 here]

7.2 Pooled Employer Plans

The ERISA (1974) allows multiple employers to share a single retirement plan, while DOL's advisory opinion[9] restricts the employers involved in a Multi-Employer Plan (MEP) to have some "commonality of interest" or "bona fide group or association". More recently, the SECURE Act of 2019 introduced the Pooled Employer Plan (PEP) that allows unrelated employers to participate in the same plan without having a "common nexus".

According to the demand-side explanation, the friction comes from the agency problem, especially in small plans. Therefore, unless the policy can align the preferences of small sponsors with the preferences of large sponsors, this type of policy will not be effective. In my model, however, the friction is from the supply side. Recordkeepers set a threshold on the total revenue from each plan, therefore, the plan size itself matters in determining recordkeeper and plan choices.

7.2.1 Effect on small plans' choices

In the first counterfactual analysis, I examine how plan size affects sponsors' recordkeeper and plan choices. Specifically, I fix all plan and sponsor characteristics at their observed values and vary the number of participants for every plan in the data to 10, 20, and so forth up to 100, while keeping each recordkeeper's estimated exclusion threshold unchanged. I then simulate recordkeeper and plan choices under each counterfactual plan size to evaluate how market outcomes respond to differences in plan size.

[Insert Figure 16 here]

Figure 16 shows the results of this counterfactual analysis. Panel (a) plots the market shares of the top-10 recordkeepers and the four types of non-top-10 recordkeepers for plans of different sizes. First, the DC plan provision rate increases. The proportion of plans opting for outside options drops from 10% to 5% as the number of participants increases from 10 to 100. Second, as plan size increases, sponsors switch to recordkeepers with high-mean utility and high exclusion. For example, Fidelity's market share among

[9] https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/resource-center/advisory-opinions/2012-04a

sponsors with 10 or 20 participants is almost zero, while about 5% sponsors choose Fidelity when they have 100 participants. On the other hand, small plans mostly choose recordkeepers without exclusion. For plans with 10 participants, 97% of them choose ADP, small intergated, and small non-integrated recordkeepers, because they don't have any exclusion. The combined market share of these recordkeepers steadily decreases as the plan size increases, and drops to 54% among sponsors with 100 participants.

Panels (b) and (c) draw the distributions of the average expense ratio and recordkeeping fees for plans of different sizes. A bit surprisingly, the distributions do not show a significant upward or downward trend as plan size increases. This is driven by the extensive margin, because sponsors switch to recordkeepers with high quality service and high exclusion, so the average expense ratio and the recordkeeping fees do not necessarily decrease. Yet sponsors enjoy have quality of services as they become larger. This is evidenced in Panel (d). I plot the distribution of expected welfare across plans sizes, the per-participant expected utility strictly icnreases in plan size. This proves that pooling small plans into a single large plan can help them get access to high-quality recordkeeping services, and improve the welfare of plan participants.

7.2.2 Equilibrium effects of PEP

The counterfactual analysis above considers only the local effects, that is, when several small plans merge into a large plan, the overall size distribution in the market remains unchanged. This setting can be interpreted as a scenario in which the adoption of pooled employer plans (PEPs) is limited. However, if PEPs become a common practice in the DC plan market, recordkeepers will adjust to the new market structure on the demand side by reoptimizing their exclusion thresholds. In this case, PEPs would generate broader equilibrium effects on market outcomes.

To examine the equilibrium effects of exclusion, I conduct the following counterfactual analysis: I assume that all plans that are in the same zip code area and have fewer than 300 participants merge into a large pooled employer plan. This significantly reduces the total number of sponsors in the market and increases the average size of sponsors.

Facing the change of market structure on the demand side, I consider two cases: in the first case, recordkeepers do not re-optimize their exclusions; in the second case, the recordkeepers re-optimize their exclusions given the new market structure. I compare the markets, expected management fees, expected recordkeeping fees, and the expected welfare of plan participants in each case.

Table 5 column (2) presents the re-optimized exclusions in this counterfactual scenario. All recordkeepers set a higher exclusion because the market houses more large sponsors:

Fidelity's optimal exclusion increases from $51,000 to $93,900, which is an 84% increase. ADP, while having a positive shadow value capacity, also raises its exclusion to $10,710. Other top-10 recordkeepers on average increase their exclusion by more than 80%, highlighting that the equilibrium effect on optimal exclusions cannot be neglected.

[Insert Figure 17 here]

Figure 17 plots the market outcomes under both scenarios and compares them to the baseline case. Panel (a) compares the market shares. Compared to the base case, when small plans in the same zip code area form a single large plan, recordkeepers with high mean utility expand their market shares. For example, Fidelity's market share increases from 15.3% to 18.5%; Vanguard's market share increases from 3.7% to 4.3%. In contrast, recordkeepers that primarily serve small plans experience declines in market share—for example, ADP's share falls from 3.8% to 2.5%.

Panel (b) through Panel (d) of Figure 17 illustrate the changes in management fees, recordkeeping fees, and expected welfare compared to the base case for plans that are enrolled into these local PEPs. By construction, these are plans with fewer than 300 participants. Panel (b) shows that these small sponsors have significantly lower management fees in the pooled employer plan compared to operating as individual plans. If recordkeepers do not re-optimize their exclusions, for plans with 100 participants, enrolling in a PEP can reduce the average management fees by 5 basis points. Meanwhile, small plans benefit more from pooling. However, if recordkeepers re-optimize their exclusions, the reduction in fees will be smaller by about 1 basis point. This is due to the equilibrium effect of the PEP. Panel (c) shows a similar pattern for recordkeeping fees. Compared to the base case, recordkeeping fees decline by about $25 per participant when exclusions are fixed, and by about $22 when exclusions are re-optimized. Panel (d) compares the expected per-participant welfare. For sponsors with 100 participants, the expected welfare improves by $277 under fixed exclusions and by slightly less when exclusions are updated.

Panel (e) to Panel (g) of Figure 17 plot the changes in the market outcomes compared to the base case for plans that are not enrolled into the local PEPs. These are either plans with fewer than 300 participants and no other small plans in the same zip code area, or mid-size to large plans with more than 300 participants. Panel (e) compares the management fees. If recordkeepers do not re-optimize their exclusions, sponsors that do not join the pooled-employer plans are not affected, therefore, there is no change in management fees compared to the base case. However, recordkeepers will raise their exclusions in the equilibrium because of the more concentrated market structure on the demand side. As a result, most plans that are not enrolled in PEPs will face higher exclusions and are induced to choose plans with higher management fees. Tiny plans will

switch recordkeepers with low exclusions and may have lower fees. Panel (f) compares the recordkeeping fees. Again, there is no change in recordkeeping fees for these plans if recordkeepers do not update their exclusions. But sponsors will need to pay higher recordkeeping fees in the new equilibrium. Plans with 100 participants on average will need to pay \$12 per participant. Panel (g) compares the expected welfare. Due to the equilibrium effect, for a plan with 100 participants that does not join a PEP, the per-participant welfare decreases by about 70 dollars.

Putting together, while PEPs can benefit their participants who originally come from small plans, the benefits are mitigated because the recordkeepers will set higher exclusions facing the market structure change on the demand side. At the same time, participants not enrolled in PEPs are adversely affected by the higher exclusions in the new equilibrium. Consequently, the aggregate welfare effect of PEP adoption depends critically on both the prevalence of PEPs and recordkeepers' ability to re-optimizee their exclusion decisions.

7.3 Major Recordkeeper's PEPs

Following the policy, some major recordkeepers have initiated their own pooled employer plans. In practice, a major recordkeeper sets up a pooled employer plan that does not have any exclusion. Sponsors can choose to start an individual plan with the recordkeeper, or join the pooled employer plan.

The pooled employer plan provided by the same recordkeeper could differ from the standard plan in the availability of investment options, fee structure, and quality of recordkeeping services. In this counterfactual analysis, I assume that the top recordkeeper Fidelity starts its own pooled employer plan. The pooled employer plan and the standard DC plan have different quality of services, i.e., they differ in δ. I assume the δ of the standard plan is as estimated, and set the δ of the pooled employer plan to be the estimated δ mius \$100 (from \$248 to \$148). The pooled employer plan does not have an exclusion, and Fidelity chooses the exclusion threshold of the standard plan to maximize the total profit of the pooled employer plan and the standard plans.

Table 5 column (3) reports the exclusions in the new equilibrium when Fidelity introduces its own pooled employer plan. Fidelity sets roughly the same exclusion for its standard plan services, while other recordkeepers reduce their exclusions in response to the competition from Fidelity's pooled employer plan. Panel (a) of Figure 18 compares the market shares. Following the introduction of Fidelity's PEP, the standard Fidelity's market share decreases from 15.3% to 11.9%, while Fidelity PEP captures 10.1% of the market share, so the aggregate market share of Fidelity increases to 25.4%. Fidelity's standard market share declines from 15.3% to 11.9%, while the new Fidelity PEP captures 10.1% of the

market, raising Fidelity's combined market share to 25.4%. The gains for Fidelity's PEP primarily come at the expense of recordkeepers that traditionally serve small plans, as these firms tend to operate without exclusions and thus compete directly with the new PEP.

[Insert Figure 18 here]

Panels (b) and (c) of Figure 18 compare the management and recordkeeping fees following the introduction of Fidelity's pooled employer plan. With Fidelity's pooled employer plan in the market, plans of all sizes pay lower recordkeeper fees and choose plans with low management fees. Panel (d) contrasts the welfare. Both small and large plans experience higher per-participant welfare, driven largely by the decline in recordkeepers' exclusion thresholds in the new equilibrium. At the intensive margin, sponsors benefit from lower direct and indirect transfer payments; at the extensive margin, they gain access to major recordkeepers that would have otherwise excluded them. For example, both Allspring and Charles Schwab expand their market shares as they reduce exclusions in response to the new competitive environment.

7.4 Supply Side Mergers

So far I have focused on the market structure on the demand side. Another salient trend in the DC plan market is the consolidation among recordkeepers. For example, in 2022, Empower struck a deal to acquire Prudential's full-service retirement plan recordkeeping and administration business. In 2024, Vanguard announced a definitive agreement to transfer its Individual 401(k), Multi-SEP, and SIMPLE IRA recordkeeping operations to Ascensus, with the transaction expected to close in the third quarter of 2024.

Using the framework developed in this paper, I am able to examine the effects of these supply-side mergers in the DC plan market. I consider the following counterfactual analysis: each of the four types of the non-top-10 recordkeepers merges into a single large recordkeeper, so there are 14 recordkeepers in the market in total. The four new recordkeepers also decide their optimal exclusions in the game.

I keep the assumption that only the top 10 recordkeepers are involved in the consideration set formation. Meanwhile, I treat the small recordkeepers previously in the sponsor's prior set as different branches of the consolidated recordkeepers, so the consolidated recordkeepers maximize their profits by aggregating profits from small recordkeepers of the same type in order to avoid the choice probability from mechanically decreasing in the number of options.

Since additional profits are important in determining the optimal exclusions, I assign the

counterfactual α's for the four new consolidated recordkeepers in the following way: the large integrated recordkeeper has the α as the average of α's of all top 10 recordkeepers whose average client size is greater than 600; the small integrated recordkeper has the *alpha* of the average of all top 10 recordkeepers whose average client size is smaller or equal to 600; the large non-integrated recordkeeper has the α as the average of all top 10 recordkeepers; the small non-integrated recordkeeper has the α the same as ADP.

Table 5 column (4) reports the optimal exclusions all 14 large recordkeepers in the counterfactual analysis. Relative to the baseline, the top 10 recordkeepers respond to the more concentrated competition by reducing their exclusions. The four consolidated recordkeepers significantly increase their exclusions following the merger. Non-top-10 small-serving recordkeepers on average have no exclusions before the merger, now they have exclusions as high as major recordkeepers. This occurs partly because, in the simulation, the four consolidated recordkeepers are excluded from the consideration set formation process and thus always appear in sponsors' consideration sets, enabling them to profitably set higher exclusions.

[Insert Figure 19 here]

Figure 19 compares the market outcomes upon the consolidation on the supply side. Panel (a) shows that, consistent with the reduction in exclusions among large recordkeepers, the top-10 recordkeepers gain market share, absorbing the outflow of sponsors that previously selected small-serving competitors. Panels (b) and (c) indicate that, post-merger, plans of all sizes pay higher recordkeeping and management fees. As a result, Panel (d) demonstrates a decline in per-participant expected welfare across all plan sizes, although the magnitude is limited. The mild welfare reduction reflects the trade-off between higher fee burdens and a shift toward top-ten recordkeepers that deliver greater mean utilities.

8 Conclusion

In the defined contribution plan market, most major recordkeepers are vertically integrated and are able to bundle recordkeeping services with asset management services. They exercise their market power by setting an exclusion threshold that requires a minimum amount of payment from each client. The exclusion threshold serves to increase the direct payment of recordkeeping fees and induces the sponsors to choose expensive plans, at the cost of excluding small plans from accessing high-quality plans.

In this paper, I develop a structural model where the sponsor chooses a recordkeeper and a plan, and bargains with the chosen recordkeeper over the recordkeeping fees. Recordkeepers optimally set their exclusions to maximize their profit. I show that recordkeepers

that provide high-quality DC plan services tend to have higher exclusions. Therefore, the policy of the pooled employer plan introduced in 2021 can serve as an effective tool, helping small sponsors access high-quality DC plans. However, the change in market structure can induce equilibrium effects, where recordkeepers will re-optimize and set higher exclusions when the average size of sponsors expands. As such, the benefit of the PEP will be mitigated due to the equilibrium effect. Meanwhile, sponsors who do not participate in a PEP will be adversely affected, facing higher exclusions. The aggregate impact on social welfare depends on the prevalence of the PEP and the size distribution of firms. Besides, my model builds a framework to study the effects of market structural changes on both sides.

In this paper, I focus on the quality of DC plans. In reality, employers provide different types of benefits to employees, including retirement savings plans, insurance plans, and most importantly, wages. These different forms of compensation may act as substitutes or complements, but their interactions remain insufficiently understood in the existing literature. Exploring the interdependence between DC plan quality and other components of employee compensation presents an important avenue for future research.

References

Adams, W. J., & Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *The quarterly journal of economics*, *90*(3), 475–498.

Armstrong, M., & Vickers, J. (2010). Competitive non-linear pricing and bundling. *The Review of Economic Studies*, *77*(1), 30–60.

Armstrong, M., & Vickers, J. (2018). Multiproduct pricing made simple. *Journal of Political Economy*, *126*(4), 1444–1471.

Badoer, D. C., Costello, C. P., & James, C. M. (2020). I can see clearly now: The impact of disclosure requirements on 401 (k) fees. *Journal of Financial Economics*, *136*(2), 471–489.

Balyuk, T., Prabhala, N., & Puri, M. (2021). Small bank financing and funding hesitancy in a crisis: Evidence from the paycheck protection program. *Available at SSRN 3717259*.

Barry, J. W., Carlin, B. I., Crane, A. D., & Graham, J. (2024). *Project development with delegated bargaining: The role of elevated hurdle rates* (Tech. Rep.). National Bureau of Economic Research.

Bernheim, B. D., & Whinston, M. D. (1990). Multimarket contact and collusive behavior. *The RAND Journal of Economics*, 1–26.

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., & Milkman, K. L. (2015). The effect of providing peer information on retirement savings decisions. *The Journal of finance*, *70*(3), 1161–1201.

Bhattacharya, V., & Illanes, G. (2022). *The design of defined contribution plans* (Tech. Rep.). National Bureau of Economic Research.

Bhattacharya, V., Illanes, G., & Padi, M. (2019). *Fiduciary duty and the market for financial advice* (Tech. Rep.). National Bureau of Economic Research.

Cao, S., & Chatterjee, C. (2022). *Equilibrium effects of pharmaceutical bundling: Evidence from india*. SIEPR Stanford Institute for Economic Policy Research.

Choi, J. J., Laibson, D., & Madrian, B. C. (2011). $100 bills on the sidewalk: Suboptimal investment in 401 (k) plans. *Review of Economics and statistics*, *93*(3), 748–763.

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2004). For better or for worse: Default effects and 401 (k) savings behavior. In *Perspectives on the economics of aging* (pp. 81–126). University of Chicago Press.

Crawford, G. S., Lee, R. S., Whinston, M. D., & Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*, *86*(3), 891–954.

Crawford, G. S., & Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, *102*(2), 643–685.

Egan, M., MacKay, A., & Yang, H. (2021). *What drives variation in investor portfolios?*

estimating the roles of beliefs and risk preferences (Tech. Rep.). National Bureau of Economic Research.

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

Ellickson, P. B. (2006). Quality competition in retailing: A structural analysis. *International Journal of Industrial Organization*, *24*(3), 521–540.

Gavazza, A. (2011). The role of trading frictions in real asset markets. *American Economic Review*, *101*(4), 1106–1143.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, *103*(1), 145–177.

Ho, K., & Lee, R. S. (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, *109*(2), 473–522.

Iyengar, S. S., Jiang, W., & Huberman, G. (2003). How much choice is too much? contributions to 401 (k) retirement plans.

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

Loseto, M. (2023). Plan menus, retirement portfolios, and investors' welfare. *Retirement Portfolios, and Investors' Welfare (November 1, 2023)*.

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

Matutes, C., & Regibeau, P. (1992). Compatibility and bundling of complementary goods in a duopoly. *The Journal of Industrial Economics*, 37–54.

McAfee, R. P., McMillan, J., & Whinston, M. D. (1989). Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, *104*(2), 371–383.

Nalebuff, B. (2004). Bundling as an entry barrier. *The Quarterly Journal of Economics*, *119*(1), 159–187.

Paravisini, D., Rappoport, V., & Schnabl, P. (2023). Specialization in bank lending: Evidence from exporting firms. *The Journal of Finance*, *78*(4), 2049–2085.

Pool, V. K., Sialm, C., & Stefanescu, I. (2022). *Mutual fund revenue sharing in 401 (k) plans* (Tech. Rep.). National Bureau of Economic Research.

Reuter, J. (2024). Plan design and participant behavior in defined contribution retirement plans: Past, present, and future. In *Pension research council 2024 symposium.*

Reuter, J., & Richardson, D. P. (2022). *New evidence on the demand for advice within retirement plans* (Tech. Rep.). National Bureau of Economic Research.

Sialm, C., Starks, L. T., & Zhang, H. (2015). Defined contribution pension plans: Sticky or discerning money? *The Journal of Finance*, *70*(2), 805–838.

Song, M., Nicholson, S., & Lucarelli, C. (2017). Mergers with interfirm bundling: a case of pharmaceutical cocktails. *The RAND Journal of Economics*, *48*(3), 810–834.

Taburet, A., Polo, A., & Vo, Q.-A. (2024). Screening using a menu of contracts: a structural model of lending markets.

Wenning, D. (2024). National pricing and the geography of us life insurers. *Available at SSRN 4745604*.

Yang, H. (2023). What determines 401 (k) plan fees? a dynamic model of transaction costs and markups. *Unpublished Working Paper*.

Zhou, J. (2017). Competitive bundling. *Econometrica*, *85*(1), 145–172.

Tables and Figures

Table 1: Summary statistics of recordkeeper revenues

This table displays the summary statistics of the recordkeeper's sources of revenue, including recordkeeping fees, revenue sharing, and management fees from affiliated options. Revenues are on a per-participant basis. The statistics in this table are based on Form 5500 data in 2019. The sample uses plans with participants between 100 and 5,000. Plans are excluded if (1) the plan does not report recordkeeper information; (2) BrightScope does not provide the plan's investment menu; (3) the employer has multiple DC plans; (4) the plan sponsor's state or industry information is missing; (5) the plan sponsor is located in the U.S. territories.

	Observations	Average	p10	p25	p50	p75	p90
Plan participants	51,429	1,143.29	134	173	282	611	1697
Recordkeeping fees	48,047	69.17	3.13	10.31	38.44	93.73	170.51
Revenue sharing	13,595	101.60	0.49	3.09	15.45	54.66	126.05
Management fees	31,618	128.91	1.61	10.02	48.82	172.20	359.06
Fees / Total revenue	31,608	0.51	0.03	0.15	0.51	0.88	0.98

Table 2: Estimated parameters

This table presents the estimated parameters. Panel A displays the preference parameters in the sponsor's utility function. Panel B displays the parameters of the nested-logit model, the Nash bargaining model, and the consideration formation. Panel C displays the recordkeeper-specific parameters, including marginal cost mc, recordkeeper fixed effect δ, exclusion threshold E, and recordkeeper shadow value parameter α.

Panel A: Preference parameters β	
Number of options	−1.07
Number of investment categories	1.59
Has S&P 500 tracker	19.45
Number of target-date funds	6.24
Proportion of index funds	31.30
Expected management fees (dollar)	−0.13

Panel B: Model parameters	
Variance of EVT1 shock σ	146.36
Nested parameter ρ	0.58
Bargaining parameter η	0.55
Consideration set distribution mean μ_c	−0.11
Consideration set distribution SD σ_c	0.30

Panel C: Recordkeeper-specific parameters

	mc	δ	E	α
Fidelity	10.85	248.52	51,833	103.76
Principal	12.29	159.01	25,731	64.09
Empower	12.86	161.71	15,312	38.15
John Hancock	11.73	137.12	11,596	6.10
ADP	9.99	95.25	0	-9.16
Vanguard	10.67	165.92	29,734	88.50
MassMutual	8.64	145.98	22,673	61.04
Transamerica	10.85	186.42	27,201	74.77
Allspring	10.67	209.99	46,638	95.50
Charles Schwab	9.39	223.48	50,365	95.02
Large integrated	6.99	135.34	60,114	–
Small integrated	13.30	48.09	0	–
Large non-integrated	8.38	249.56	58,159	–
Small non-integrated	10.36	79.14	0	–

Table 3: Optimal exclusions without shadow value of capacity

This table compares each of the top 10 recordkeepers' optimal exclusion setting the shadow value of capacity to zero. In the counterfactual analysis, for each recordkeeper, I keep the shadow value of capacity of other recordkeepers fixed and set the recordkeeper's own *alpha* to be zero.

	Estimated exclusion	Exclusion with $\alpha = 0$
Fidelity	51,833	48,600
Principal	25,731	17,688
Empower	15,312	11,024
John Hancock	11,596	12,016
ADP	0	2,873
Vanguard	29,734	20,600
MassMutual	22,673	11,240
Transamerica	27,201	15,600
Allspring	46,638	27,600
Charles Schwab	50,365	23,600

Table 4: Effects of consideration set

This table decomposes the effects of the consideration set in the model. Panel (a) reports the average client size for each recordkeeper. Column (1) shows the average client size in the prior consideration set. Column (2) reports the average client size after sponsors form their consideration sets, Column (3) reports the average client size after sponsors make their final choices, and Column (4) reports the average client size in the data. Panel (b) decomposes the choice probabilities. Column (1) shows the probability of being in the prior consideration set. Column (2) reports the probability of being included in the consideration set. Column (3) presents the probability of being chosen conditional on inclusion in the consideration set. Column (4) shows the actual market share, which is the product of the probabilities in the first three columns.

Panel (a): Average client size

	(1) Initial	(2) Before choice	(3) After choice	(4) Data
Fidelity	568	626	688	702
Principal	564	504	471	486
Empower	564	504	477	486
John Hancock	562	454	397	411
ADP	569	428	355	364
Vanguard	573	617	648	656
MassMutual	568	522	497	508
Transamerica	572	613	653	662
Allspring	573	804	900	909
Charles Schwab	576	870	997	1006

Panel (b): Choice probability

	(1) Prior	(2) Consider.	(3) Choice | consider.	(4) Market Share
Fidelity	0.866	0.572	0.312	0.154
Principal	0.806	0.460	0.219	0.081
Empower	0.801	0.466	0.214	0.080
John Hancock	0.747	0.392	0.192	0.056
ADP	0.651	0.341	0.171	0.038
Vanguard	0.645	0.332	0.171	0.037
MassMutual	0.687	0.339	0.167	0.039
Transamerica	0.682	0.330	0.167	0.037
Allspring	0.563	0.227	0.159	0.020
Charles Schwab	0.489	0.212	0.154	0.016

Notes: "Prior" = probability in prior set; "Consider." = probability in consideration set; "Choice | consider." = choice probability conditional on being in consideration set.

Table 5: Counterfactual optimal exclusions

This table compares the estimated exclusion thresholds from the data with several counterfactual scenarios. Column (1) reports the baseline estimated exclusions. Column (2) examines a scenario with local pooled employer plans, where sponsors with fewer than 300 participants in the same ZIP code are pooled into a single plan. Column (3) considers a counterfactual in which Fidelity introduces its own pooled employer plan without any exclusion threshold. Column (4) analyzes a consolidation scenario in which all non–top-10 recordkeepers of each type merge into a single recordkeeper.

	(1) Base	(2) Local PEP	(3) Fidelity's PEP	(4) Supply-side merger
Fidelity	51,833	93,900	51,200	43,000
Principal	25,731	44,700	25,600	24,000
Empower	15,312	31,800	15,600	14,000
John Hancock	11,596	19,800	11,000	10,000
ADP	0	10,710	660	0
Vanguard	29,734	53,100	27,600	24,000
MassMutual	22,673	36,900	18,600	18,000
Transamerica	27,201	49,800	25,200	22,000
Allspring	46,638	70,500	35,800	30,000
Charles Schwab	50,365	75,900	36,200	30,000
Large integrated	60,114	60,114	60,114	70,000
Small integrated	0	0	0	46,000
Large non-integrated	58,159	58,159	58,159	92,000
Small non-integrated	0	0	0	28,000

Figure 1: Market shares of top 10 recordkeepers

This figure shows the market share of the top 10 recordkeepers in 2019 in terms of the total number of plans they serve. All plans in the sample have more than 100 participants.

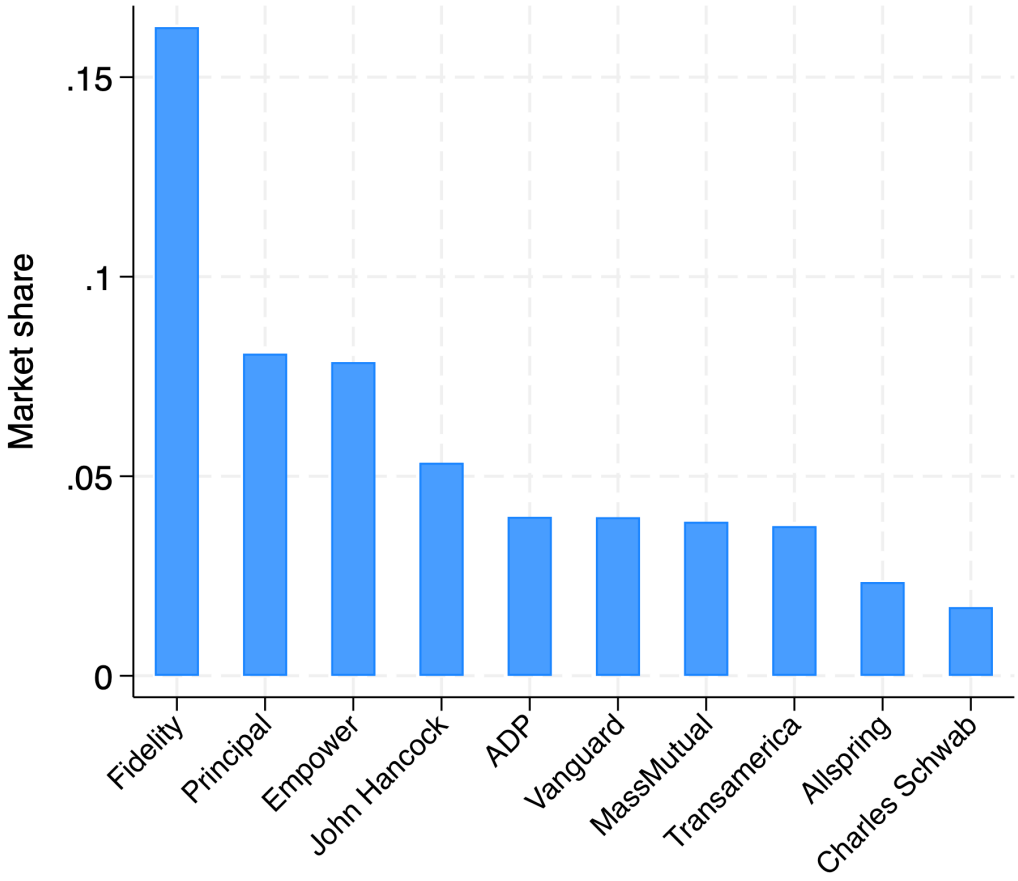

Figure 2: Revenue structure in a DC Plan

The figure shows the relationships among parties involved in a DC plan and outlines the revenue structure of a DC plan. On the demand side, within a firm, the employer (sponsor) provides the DC plan to employees. On the supply side, the integrated recordkeeper has two business departments: a recordkeeping department for recordkeeping services and an investment management department for asset management services. The investment management department of the recordkeeper and other investment providers can include their investment products on the investment menu of the DC plan. The firm (employer and employees) together pay two types of fees: recordkeeping fees for recordkeeping services and management fees for asset management services. Investment option providers other than the recordkeeper transfer part of the management fees to the recordkeeper under the revenue-sharing agreement.

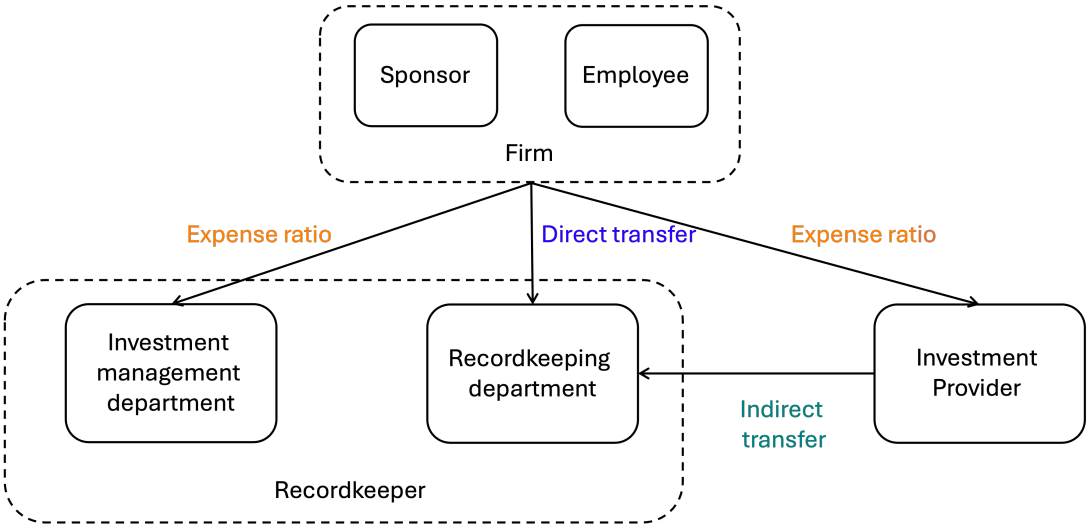

Figure 3: Proportion of plans with bundle sales

In this figure, the bar plot shows the proportion of bundle sales for the top 10 recordkeepers and the other integrated recordkeepers. The dashed line plots the average proportion of bundle sales among all integrated recordkeepers. A bundled sale is defined as a case in which the recordkeeper includes at least one affiliated investment option on the plan's investment menu. ADP, a non-integrated recordkeeper without proprietary investment products, is therefore left blank.

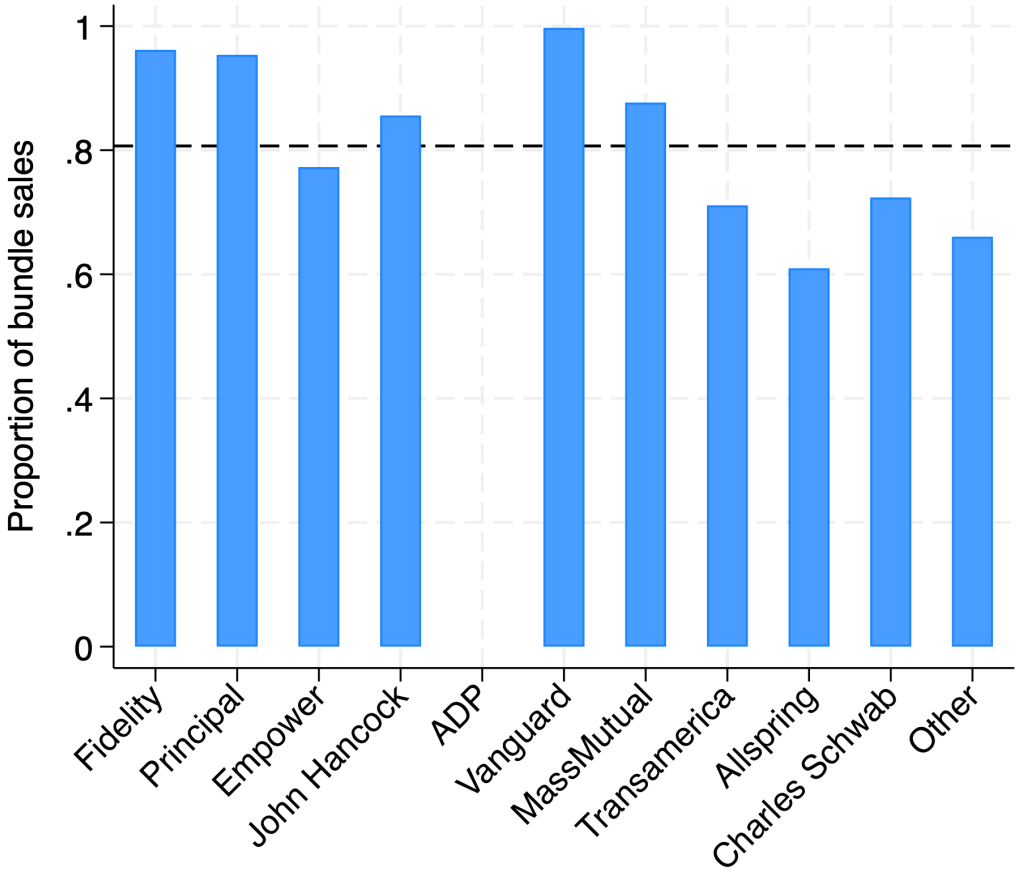

Figure 4: Market share and average client size

The figure plots the top 10 recordkeepers' market shares against their client size. Market share is defined as the number of DC plans each recordkeeper serves divided by the total number of plans in the market. Average client size is defined by the total number of participants served by the recordkeeper divided by the total number of plans served by the recordkeeper.

Figure 5: Direct and indirect transfers

The figure plots fitted relationships between a plan's average expense ratio and recordkeeping fees and the number of participants, estimated using fractional polynomial regressions. Panel (a) plots the expense ratio against the number of participants. A plan's expense ratio is defined as the asset-weighted average expense ratio. Panel (a) plots the recordkeeping fees against the number of participants. Recordkeeping fees are in dollars.

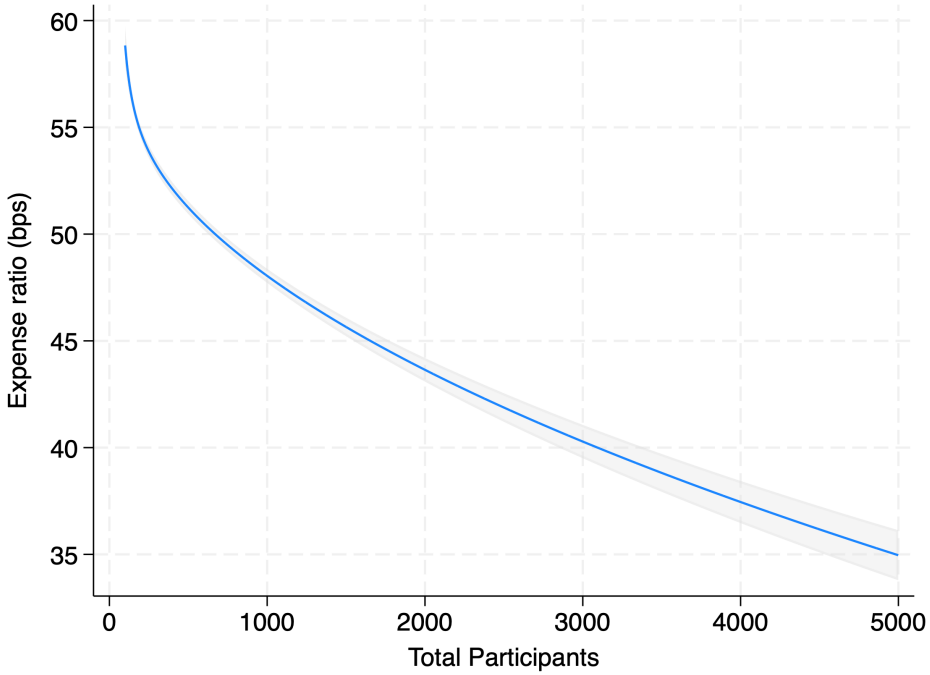

(a) Expense ratio

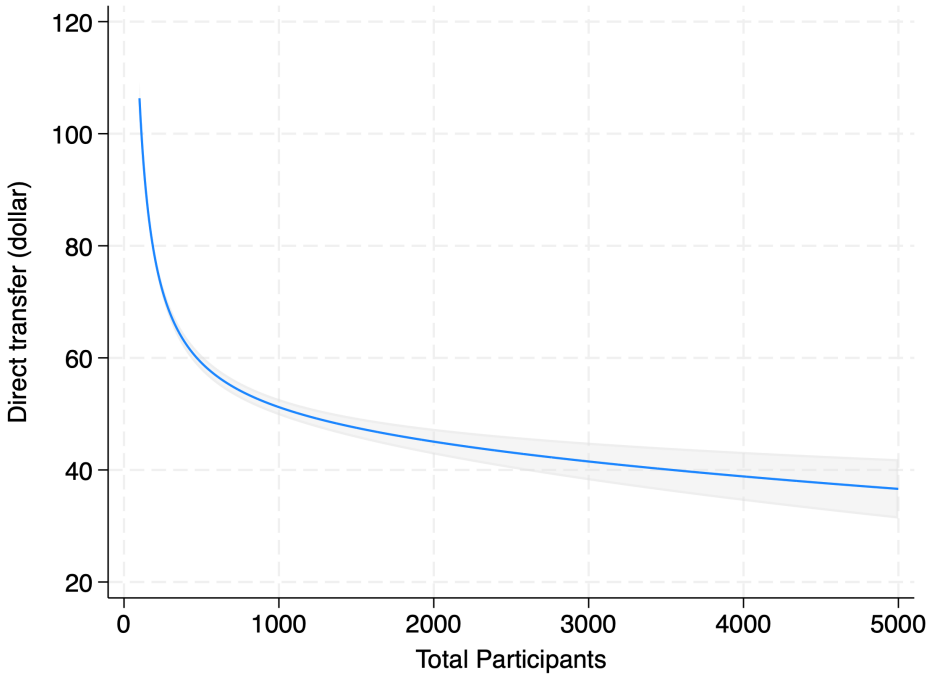

(b) Recordkeeping fees

Figure 6: Event study on plan size expansion

The figure conducts the event study based on the plan size's sudden expansion. Sudden expansion is defined as a plan's total number of participants increasing by at least 50% in a year. The event study regressions control for plan and year fixed effects. Panel (a) studies the change in average expense ratio around the event date. In Panel (a), the events include cases when the sponsor switches to a different recordkeeper at the event date. Panel (b) studies the probability of changing a recordkeeper around the event date. Panel (c) repeats the study as in Panel (a), but removes events where the sponsor switches to a different recordkeeper following the sudden expansion to isolate the effect of changing recordkeeper.

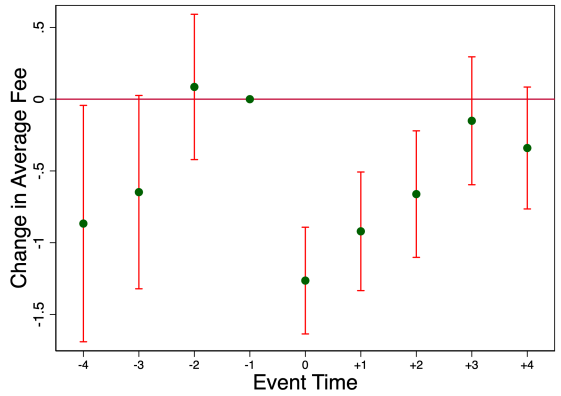

(a) Average expense ratio around plan size expansion

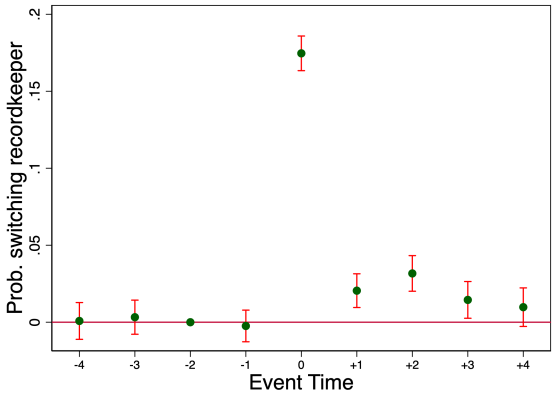

(b) Probability of switching a recordkeeper around plan size expansion

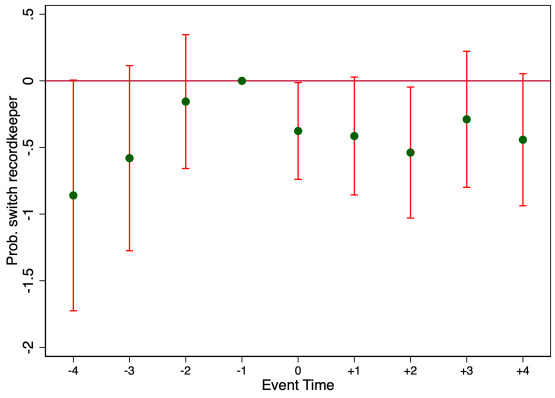

(c) Average expense ratio around plan size expansion excluding recordkeeper switching

Figure 7: Transfer, utility, and profit for different portfolios

The figure plots the transfer T, sponsor's utility u, and the recordkeeper's profit π as a function of portfolios with different management fees M. Panels differ in the sponsor's sensitivity k to management fees T. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$.

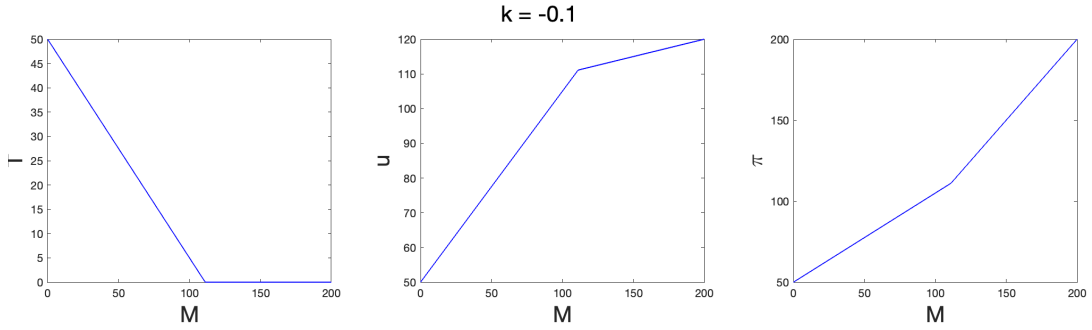

(a) $k = -0.1$

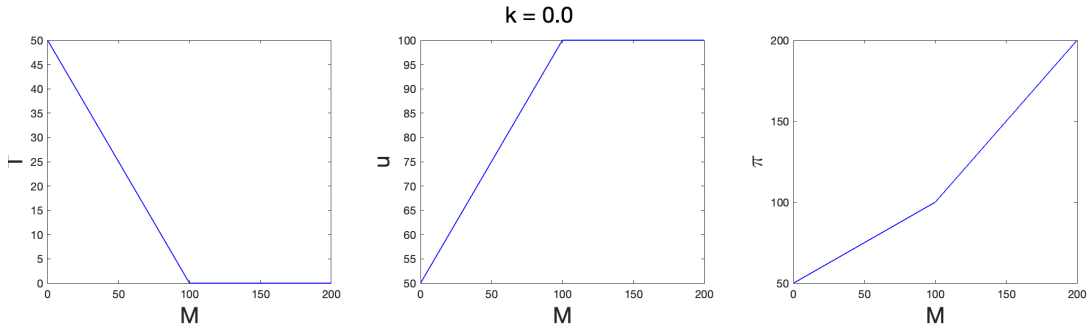

(b) $k = 0$

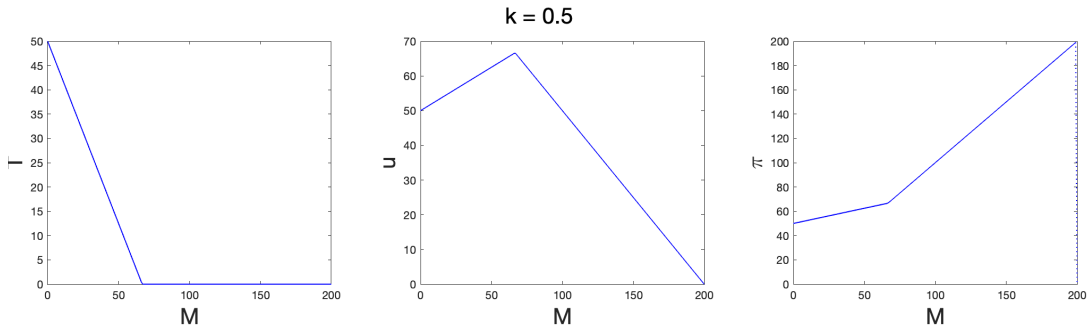

(c) $k = 0.5$

Figure 7: Transfer, utility, and profit for different portfolios (continued)

The figure plots the transfer T, sponsor's utility u, and the recordkeeper's profit π as a function of portfolios with different management fees M. Panels differ in the sponsor's sensitivity k to management fees T. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$.

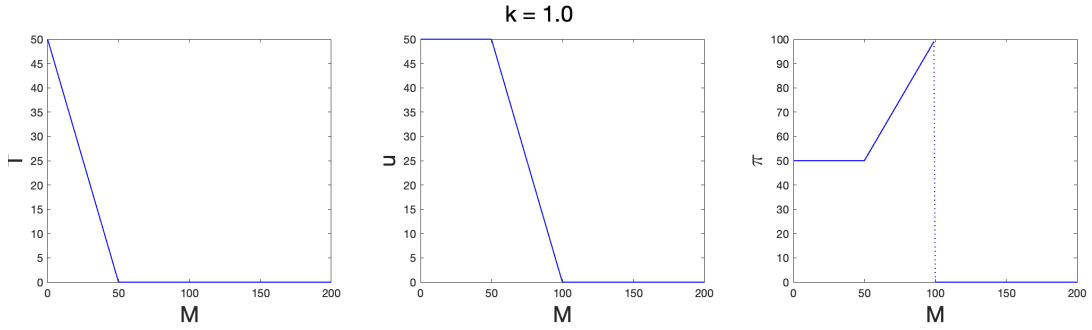

(d) $k = 1$

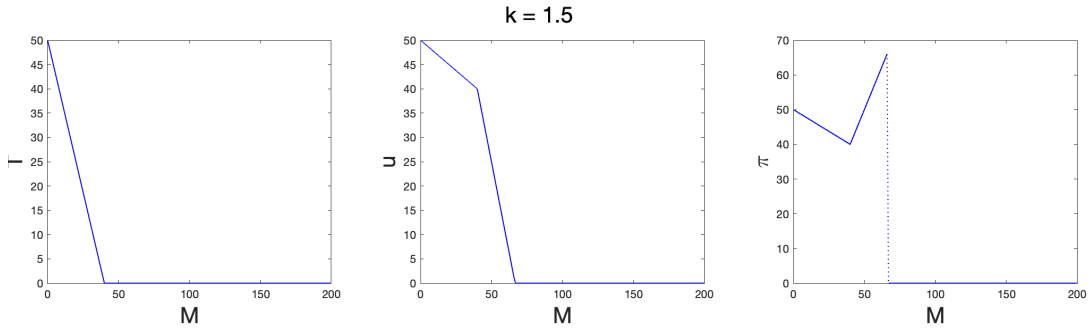

(e) $k = 1.5$

Figure 8: Transfer, utility, and profit with and without exclusions

The figure plots the transfer T, sponsor's utility u, and the recordkeeper's profit π as a function of portfolios with different management fees M with and without exclusions. Panels differ in the sponsor's sensitivity k to management fees T. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$. The sponsor is assumed to have 100 participants, and the recordkeeper has an exclusion of \$5,000. The solid blue lines represent the cases when there is no exclusion, while the dashed red lines represent cases when the recordkeeper has the exclusion.

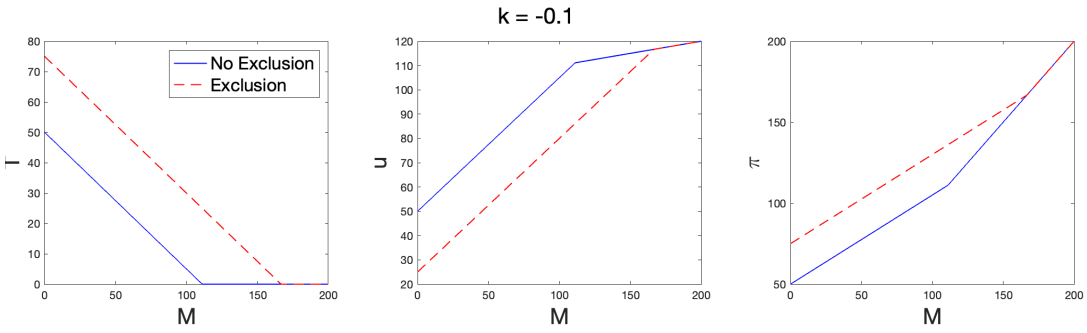

(a) $k = -0.1$

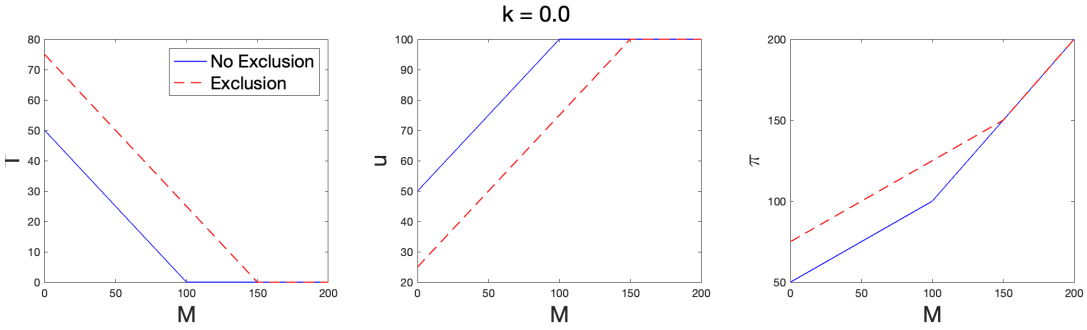

(b) $k = 0$

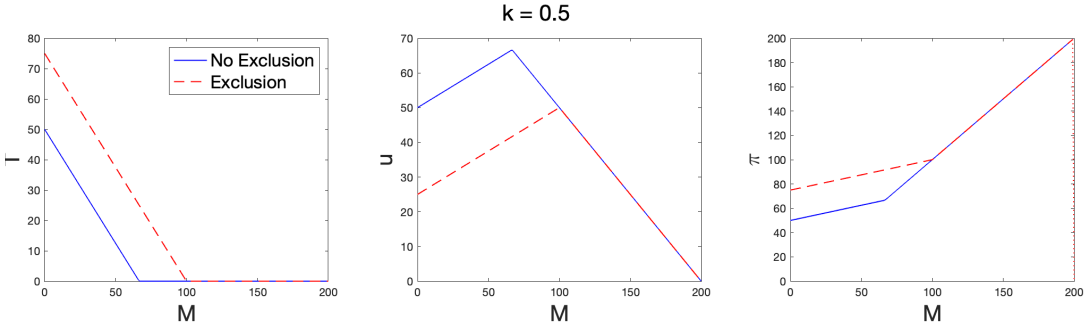

(c) $k = 0.5$

Figure 8: Transfer, utility, and profit with and without exclusions (continued)

The figure plots the transfer T, sponsor's utility u, and the recordkeeper's profit π as a function of portfolios with different management fees M with and without exclusions. Panels differ in the sponsor's sensitivity k to management fees T. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$. The sponsor is assumed to have 100 participants, and the recordkeeper has an exclusion of \$5,000. The solid blue lines represent the cases when there is no exclusion, while the dashed red lines represent cases when the recordkeeper has the exclusion.

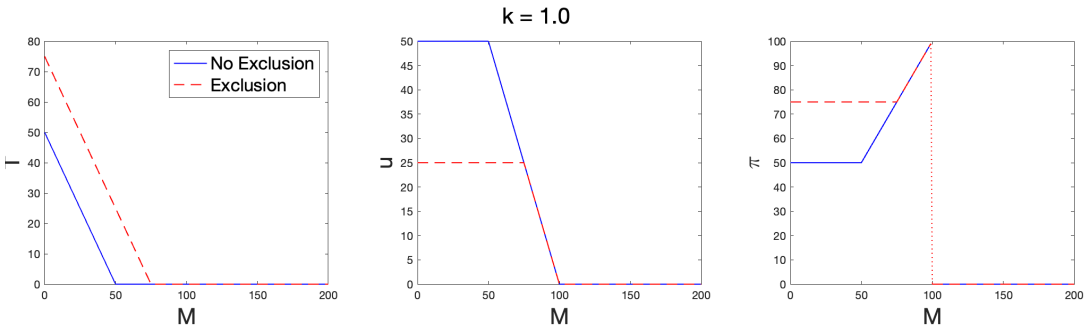

(a) $k = 1$

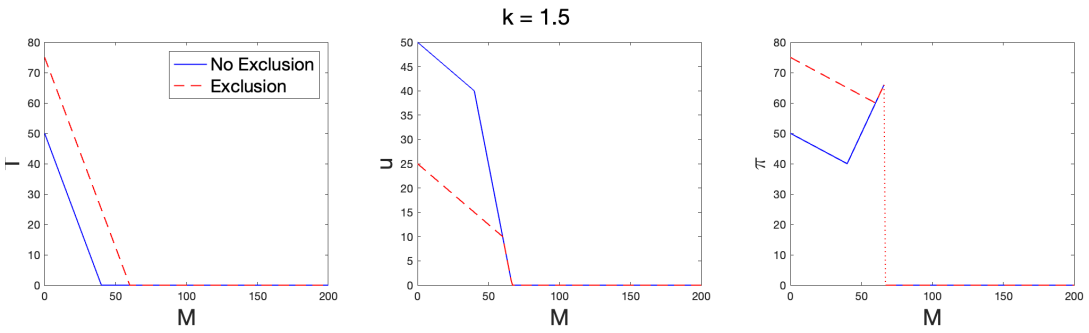

(b) $k = 1.5$

Figure 9: Transfer, utility, and profit for plans of different sizes

The figure plots the transfer T, sponsor's utility u, and the recordkeeper's profit π as a function of portfolios with different management fees M for plans of different sizes. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$. The sponsor has a sensitivity of $k = 0.5$, and the recordkeeper has an exclusion of \$5,000. The solid blue lines represent the cases when there is no exclusion, while the dashed red lines represent the cases when the recordkeeper has the exclusion.

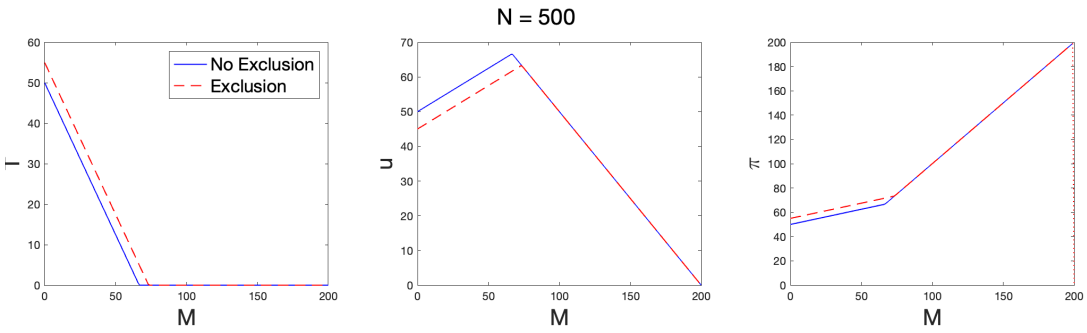

(a) $N = 500$

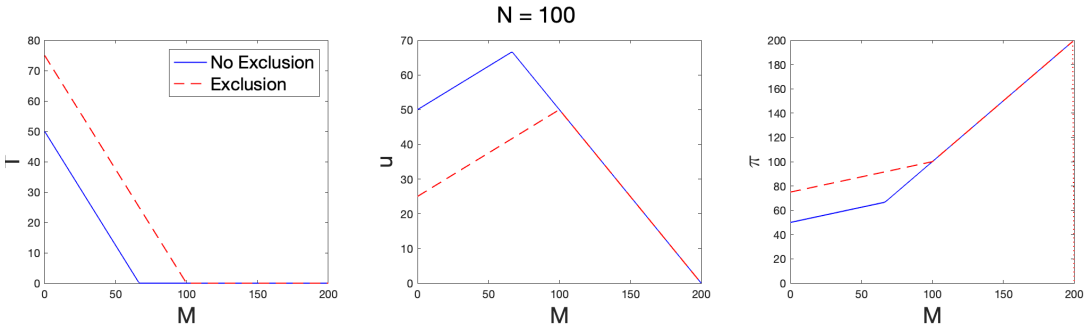

(b) $N = 100$

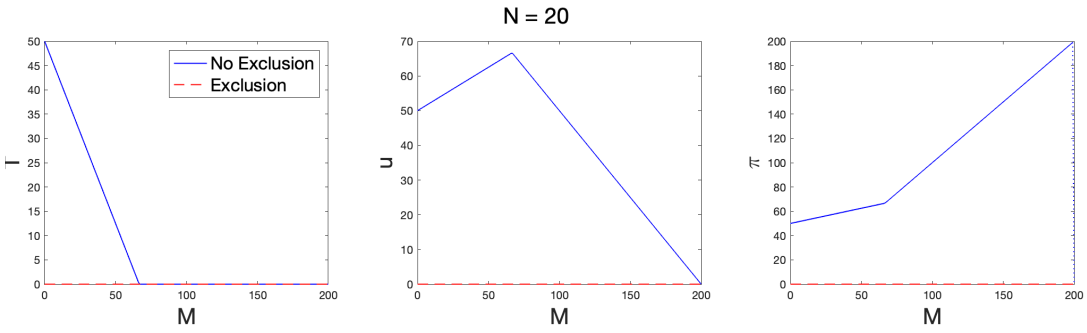

(c) $N = 20$

Figure 10: Portfolio choices for plans of different sizes.

The figure plots the transfer T, sponsor's utility u as a function of portfolios with different management fees M for plans of different sizes with and without the exclusion. The model assumes that $mc = 0$, $\eta = 0.5$, and $\delta(M) = 100 - kM$. The sponsor has a sensitivity of $k = 0.5$, and the recordkeeper has an exclusion of \$5,000. The solid blue lines represent the cases when there is no exclusion, while the dashed red lines represent the cases when the recordkeeper has the exclusion. The vertical lines represent two portfolios with $M = 50$ and $M = 100$. The intersections of the vertical lines of the utilities represent the portfolio choices.

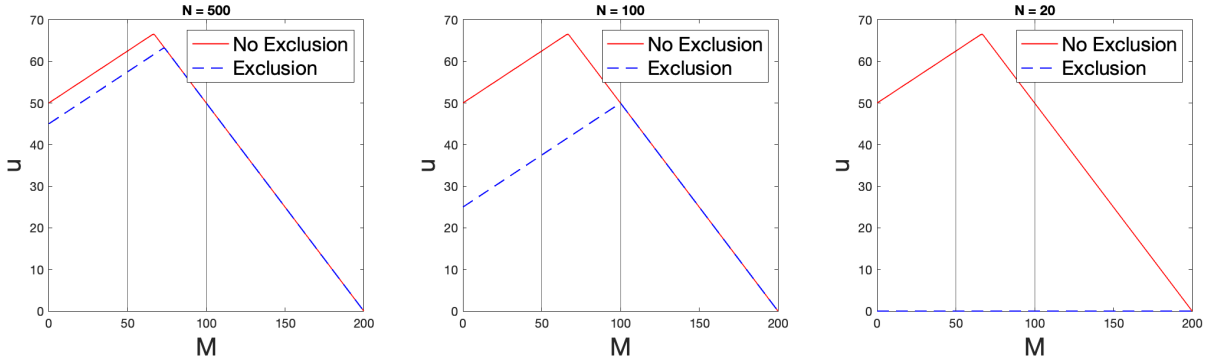

Figure 11: The recordkeeper's optimal exclusion

The figure plots the recordkeeper's optimal exclusion threshold against the mean of the sponsor size distribution, which is assumed to follow a Beta distribution. Panel (a) shows the optimal exclusion thresholds for sponsors with varying sensitivity k to management fees. Panel (b) displays the optimal exclusions for different portfolios offered by the recordkeeper.

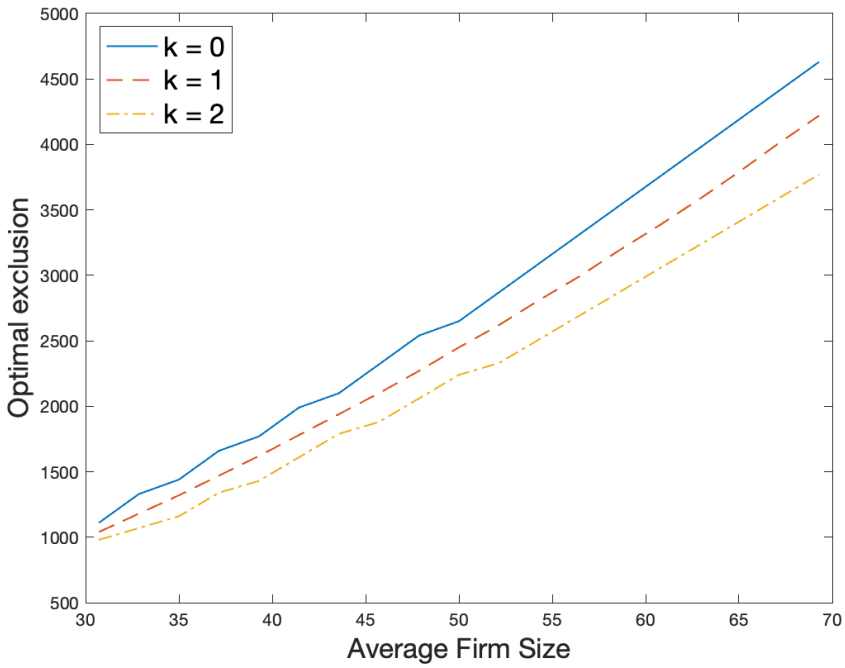

(a) Optimal exclusion for different k's

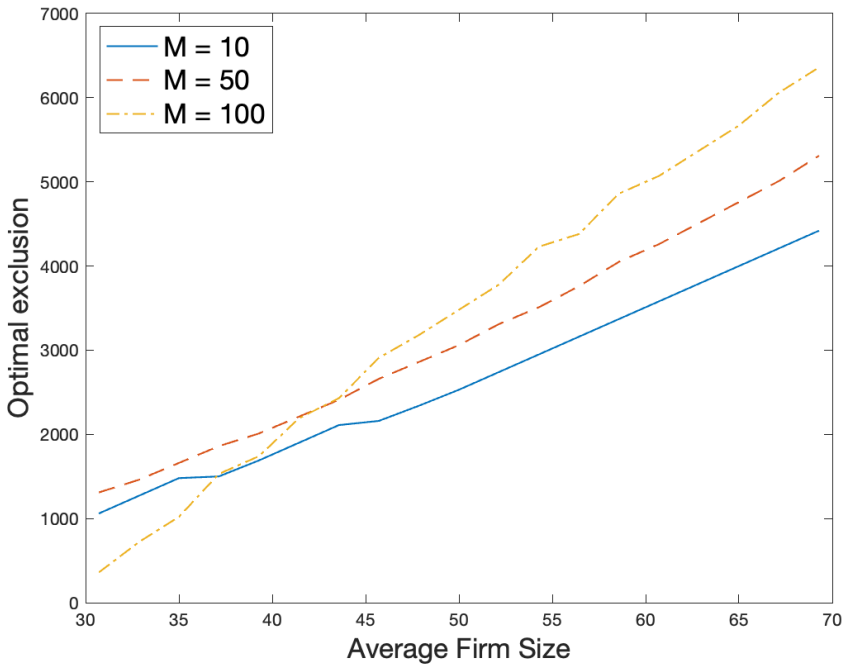

(b) Optimal exclusion for different M's

Figure 12: Effects of consideration set

The figure demonstrates the effects of consideration set in the model. Panel (a) plots the probability of being included in the consideration set against the normalized utility. Panel (b) plots the empirical probability of being included in the consideration set by the recordkeeper's rank in the prior set.

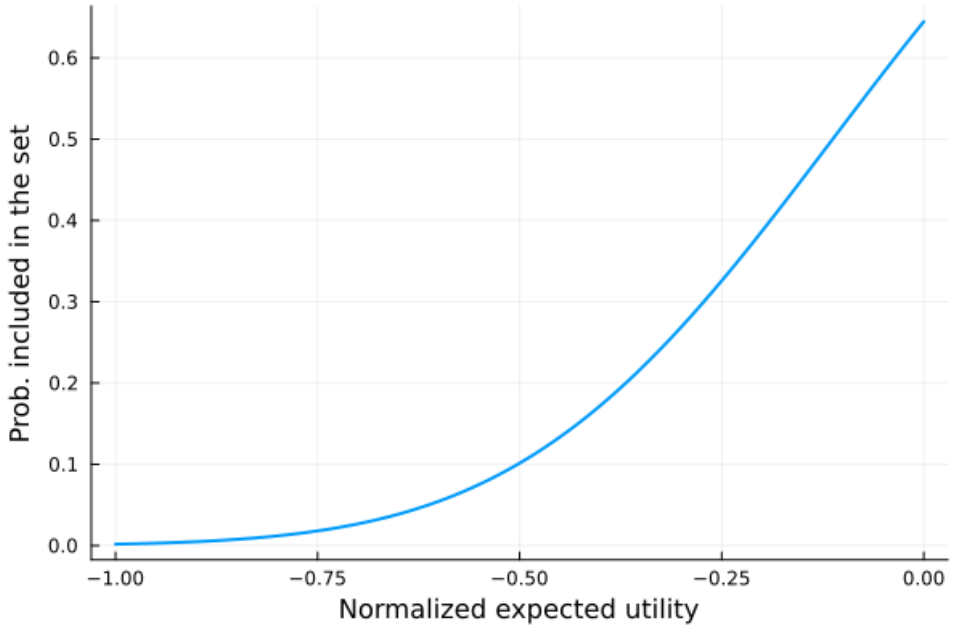

(a) Probability of inclusion

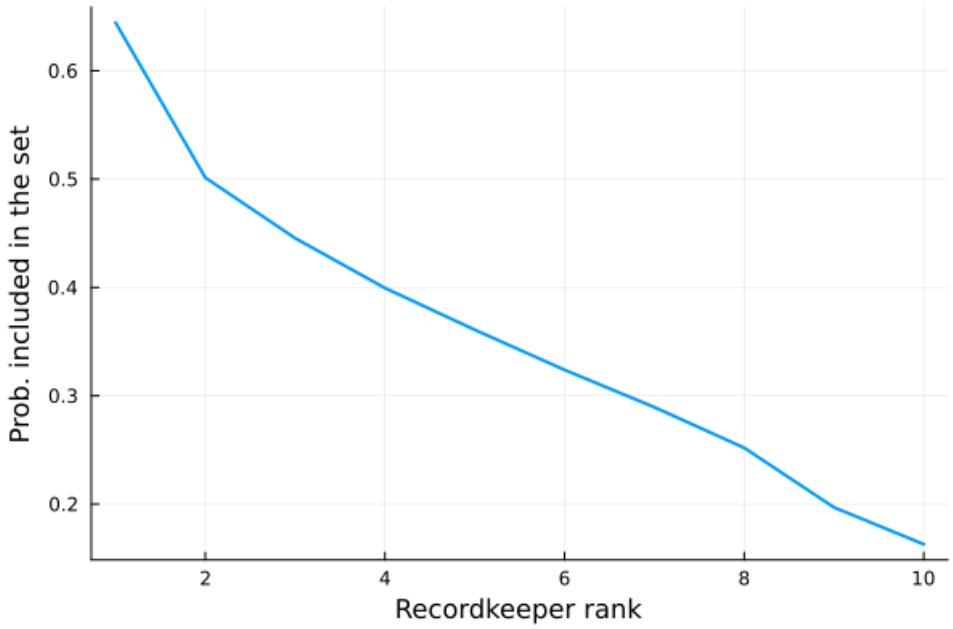

(b) Empirical probability of inclusion by rank

Figure 13: Mean utility and exclusion

The figure plots recordkeepers' mean estimated exclusions against their mean utilities. The top-10 recordkeepers each has its exclusion; other recordkeepers are grouped into four categories: large and small integrated recordkeepers, large and small non-integrated (pure) recordkeepers. Non-top-10 recordkeepers of the same category share the same exclusion. Mean utilities are the average of utilities over all sponsors served by the recordkeeper.

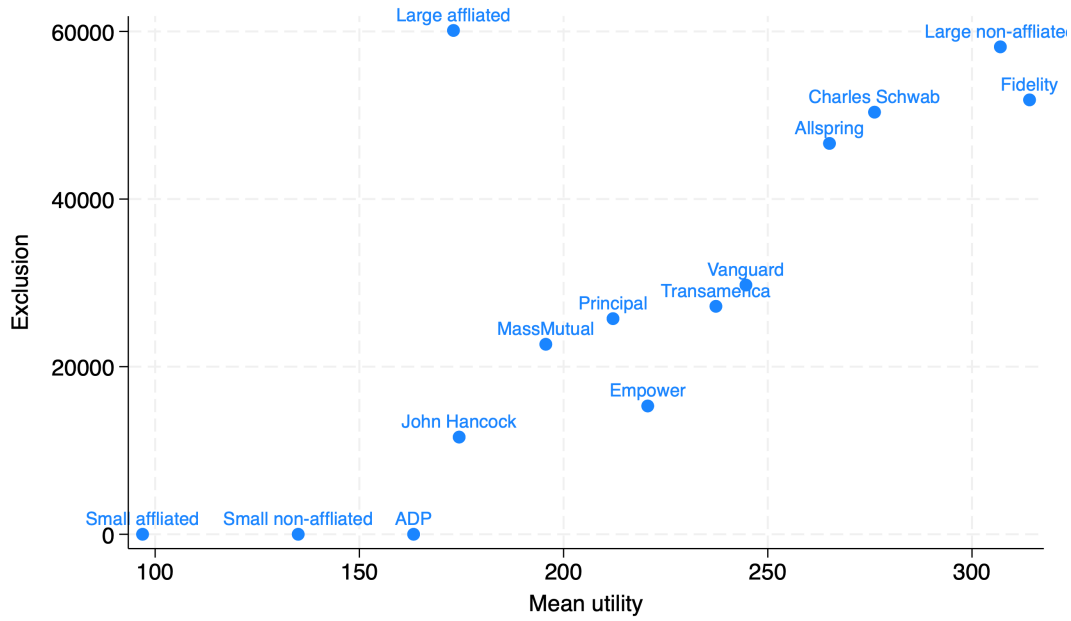

Figure 14: Model fit

This figure compares the market outcomes in the data to the model-fitted counterparts. The Panel (a) compares the market shares. Panel (b) compares the management fees measured by the average expense ratio in basis points. Panel (c) compares the recordkeeping fees in dollars.

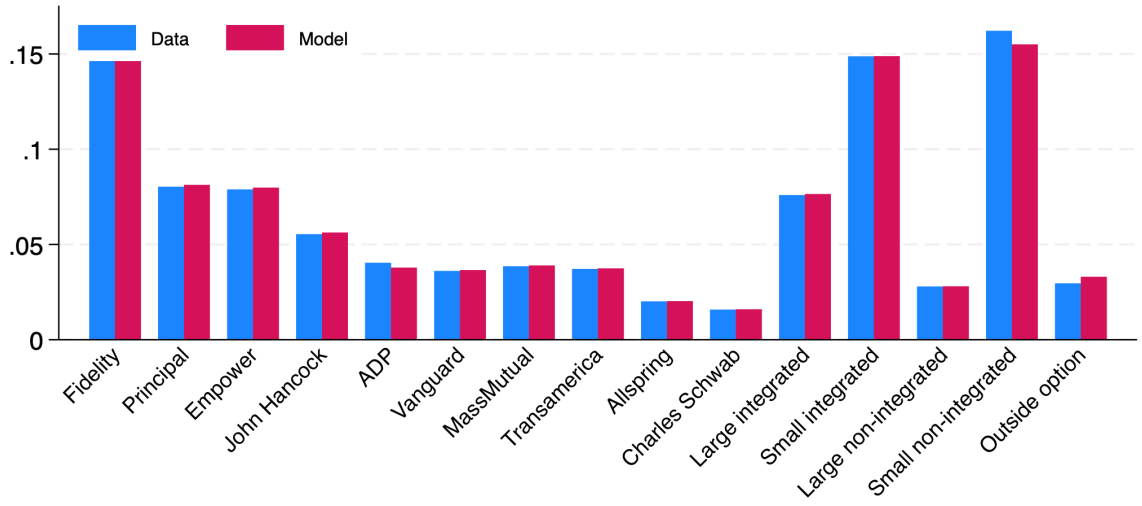

(a) Market share

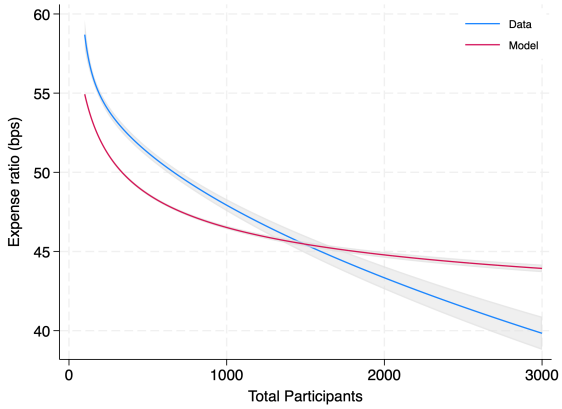

(b) Management fees

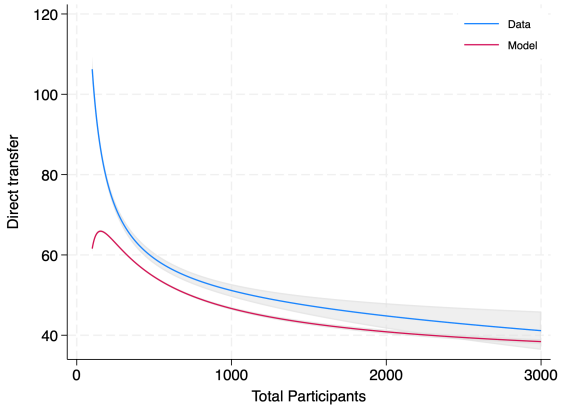

(c) Recordkeeping fees

Figure 15: Counterfactual market outcomes removing all exclusions

This figure compares the market outcomes in the counterfactual analysis where all recordkeepers' exclusions are removed to those in the base case. The base case is generated by the model with estimated parameters. Panel (a) reports market shares; Panel (b) compares management fees, measured by the average expense ratio in basis points; Panel (c) presents recordkeeping fees in dollars; and Panel (d) shows the per-participant expected welfare. Panel (e) plots the changes in each recordkeeper's total profit.

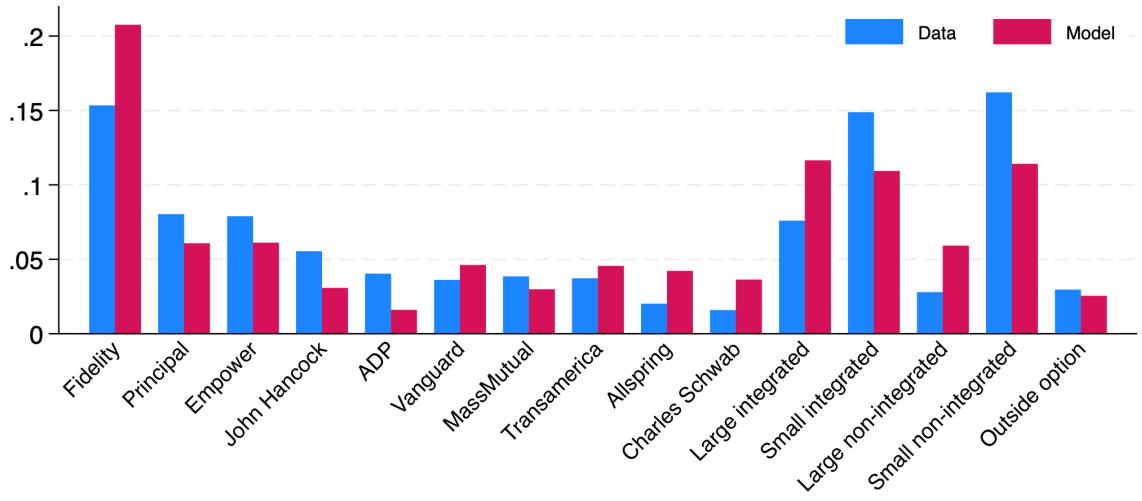

(a) Market share

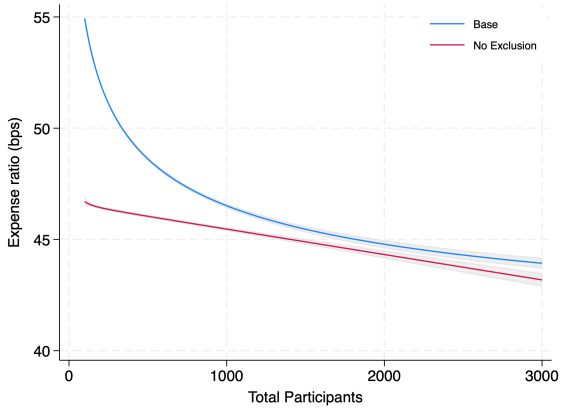

(b) Management fees

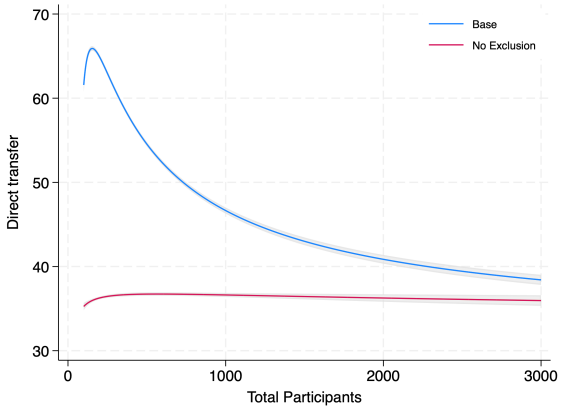

(c) Recordkeeping fees

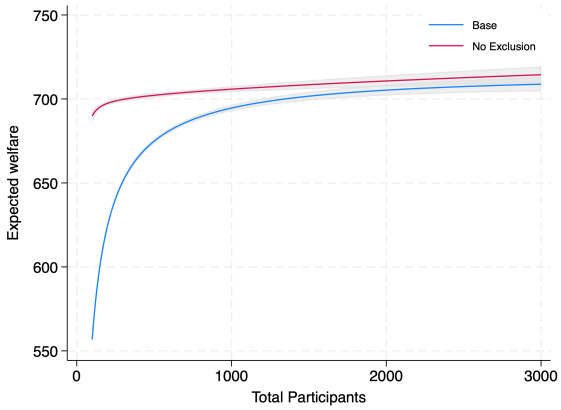

(d) Per-participant expected welfare

Figure 15: Counterfactual market outcomes removing all exclusions (continued)

This figure compares the market outcomes in the counterfactual analysis where all recordkeepers' exclusions are removed to those in the base case. The base case is generated by the model with estimated parameters. Panel (a) reports market shares; Panel (b) compares management fees, measured by the average expense ratio in basis points; Panel (c) presents recordkeeping fees in dollars; and Panel (d) shows the per-participant expected welfare. Panel (e) plots the changes in each recordkeeper's total profit.

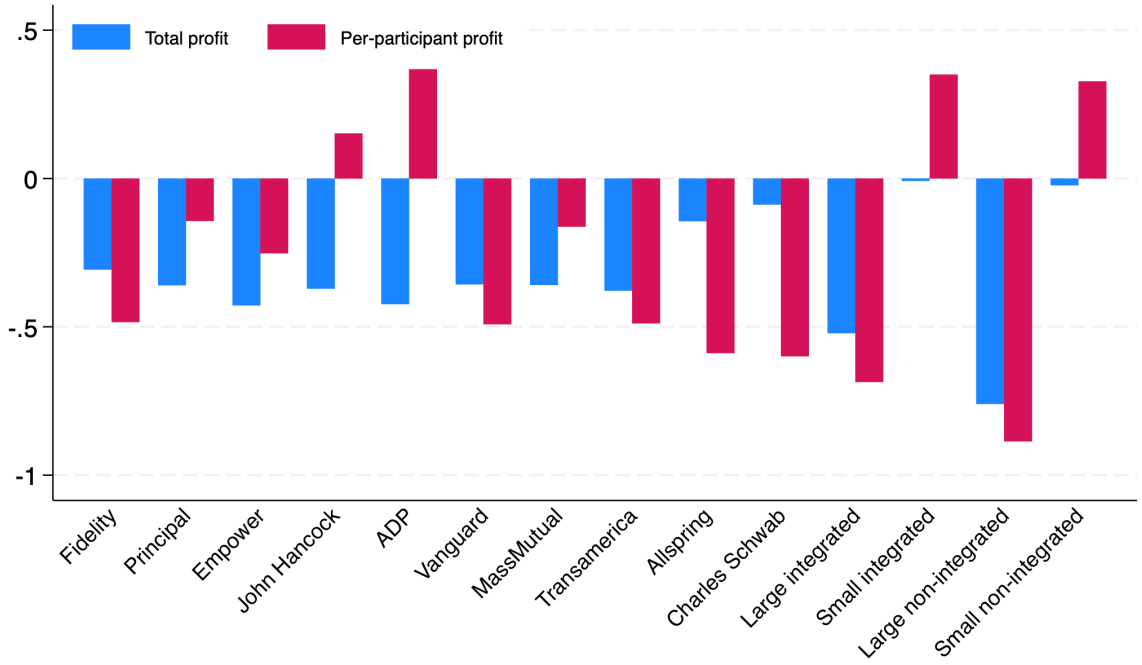

(e) Change in total profit

Figure 16: Counterfactual market outcomes with pooled employer plans

The figure compares market outcomes in counterfactual scenarios that introduce pooled employer plans. In this series of counterfactuals, the number of participants for all sponsors in the data is fixed at 10, 20, . . . , 100, and the model is simulated using these hypothetical plan sizes. Panel (a) shows the market shares of recordkeepers. Panel (b) plots the distribution of management fees across counterfactual plan sizes. Panel (c) reports the distribution of recordkeeping fees, and Panel (d) presents the distribution of per-participant expected welfare.

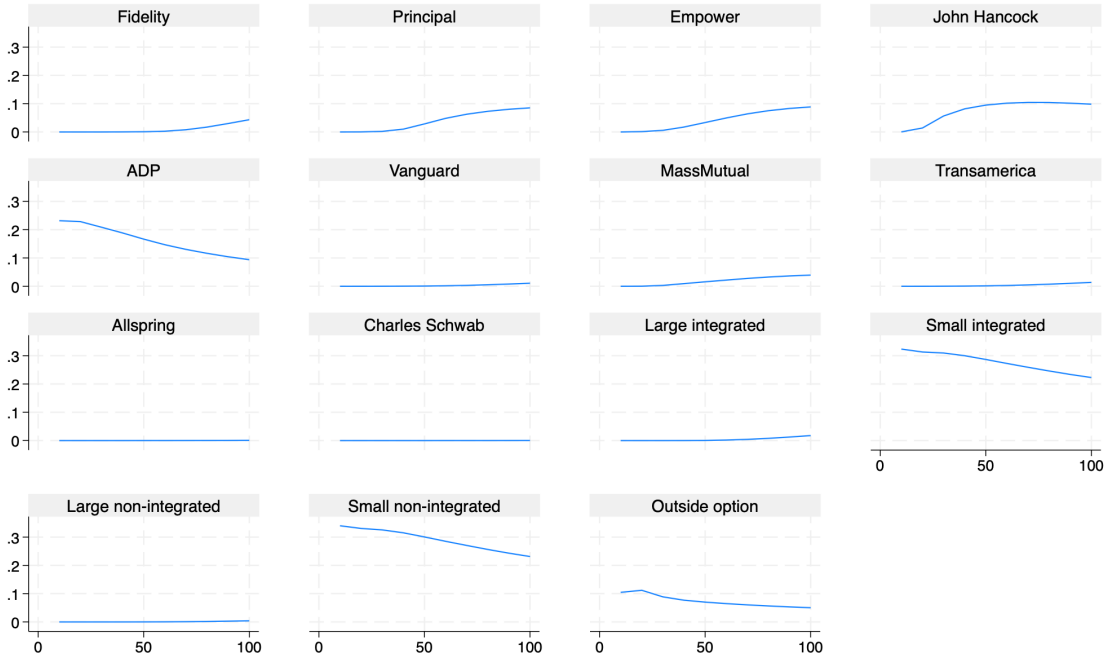

(a) Market share

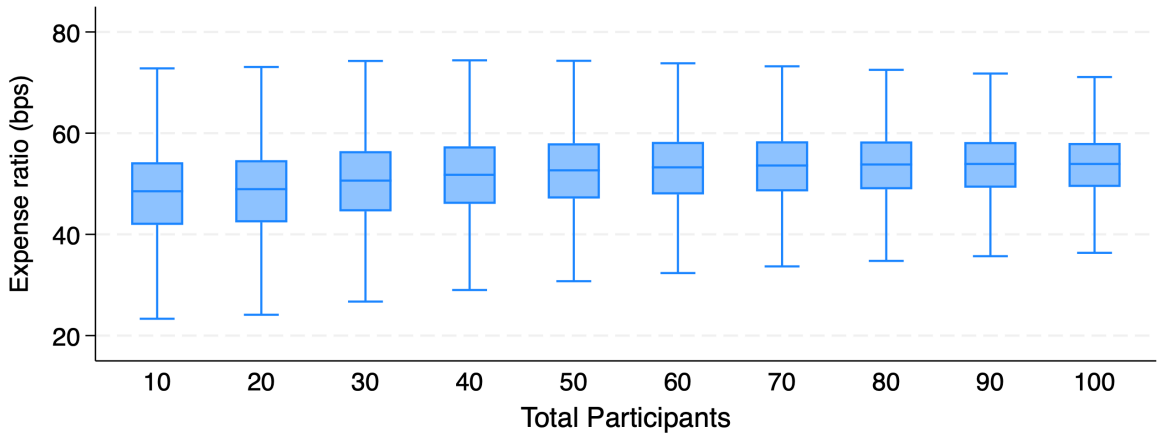

(c) Management fees

Figure 16: Counterfactual market outcomes with pooled employer plans (continued)

The figure compares market outcomes in counterfactual scenarios that introduce pooled employer plans. In this series of counterfactuals, the number of participants for all sponsors in the data is fixed at 10, 20, . . . , 100, and the model is simulated using these hypothetical plan sizes. Panel (a) shows the market shares of recordkeepers. Panel (b) plots the distribution of management fees across counterfactual plan sizes. Panel (c) reports the distribution of recordkeeping fees, and Panel (d) presents the distribution of per-participant expected welfare.

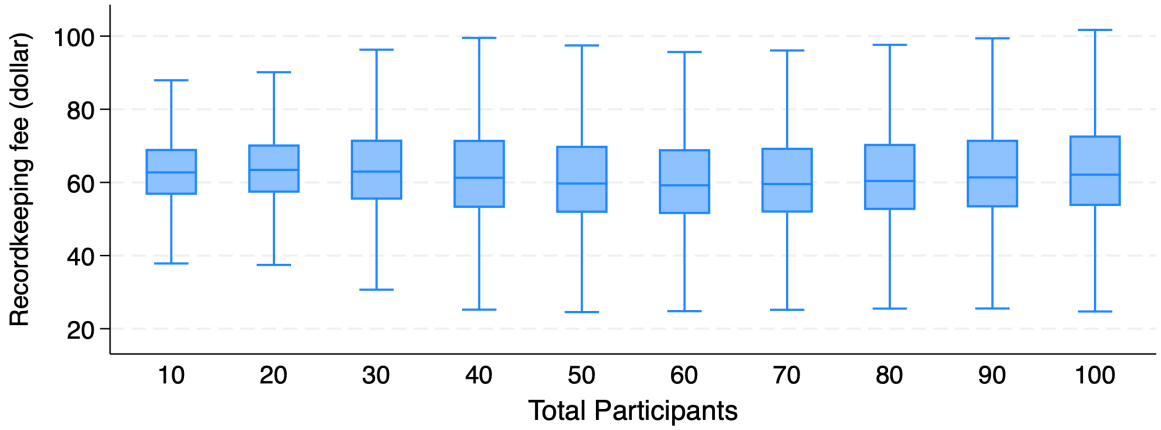

(c) Recordkeeping fees

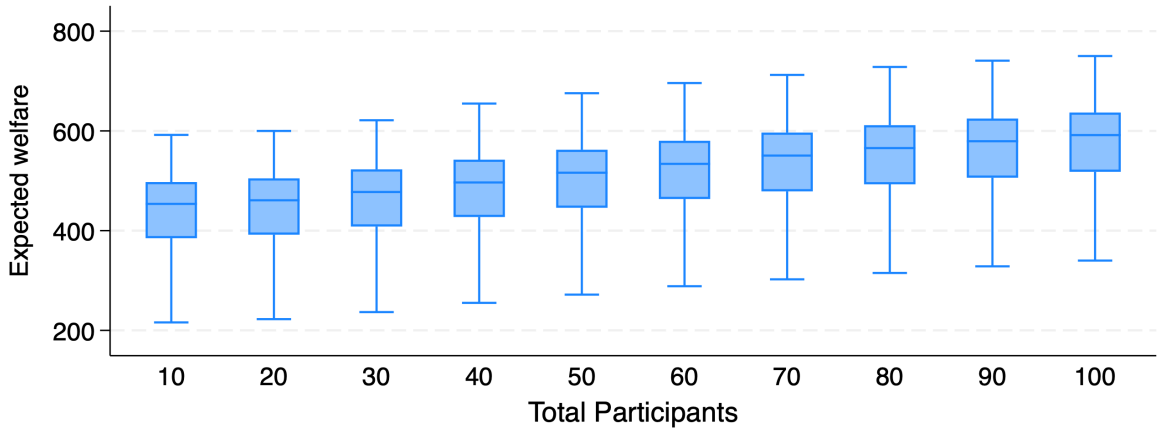

(d) Per-participant expected welfare

Figure 17: Counterfactual market outcomes with local pooled employer plans

This figure compares the market outcomes in the counterfactual analysis where sponsors with fewer than 300 participants in the same ZIP code area form a single pooled employer plan. Two cases are considered: one in which recordkeepers do not update their exclusions, and another in which recordkeepers reoptimize their exclusions. Panel (a) compares market shares. Panels (b), (c), and (d) show the differences in management fees, recordkeeping fees, and per-participant welfare relative to the base case for sponsors that participate in PEPs. Panels (e), (f), and (g) present the corresponding differences for sponsors that do not participate in PEPs.

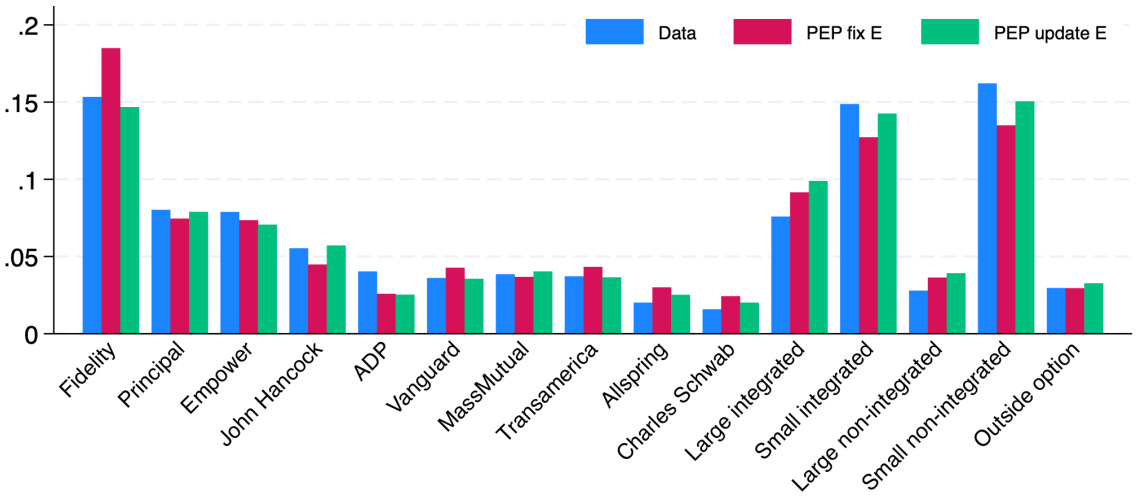

(a) Market share

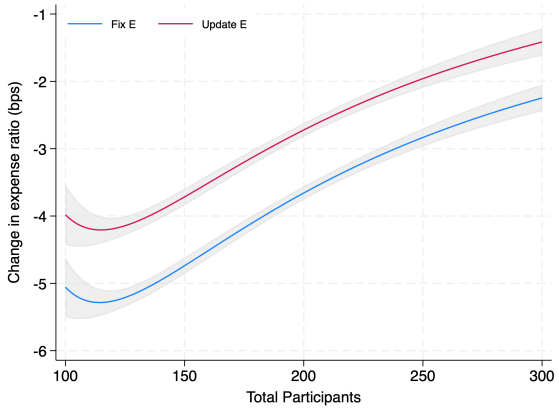

(b) Difference in management fees

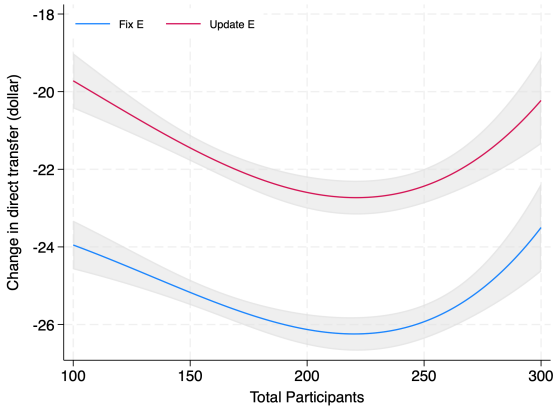

(c) Difference in recordkeeping fees

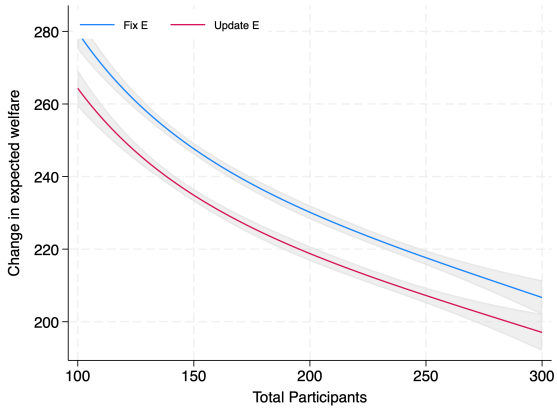

(d) Difference in per-participant welfare

Figure 17: Counterfactual market outcomes with local pooled employer plans (continued)

This figure compares the market outcomes in the counterfactual analysis where sponsors with fewer than 300 participants in the same ZIP code area form a single pooled employer plan. Two cases are considered: one in which recordkeepers do not update their exclusions, and another in which recordkeepers reoptimize their exclusions. Panel (a) compares market shares. Panels (b), (c), and (d) show the differences in management fees, recordkeeping fees, and per-participant welfare relative to the base case for sponsors that participate in PEPs. Panels (e), (f), and (g) present the corresponding differences for sponsors that do not participate in PEPs.

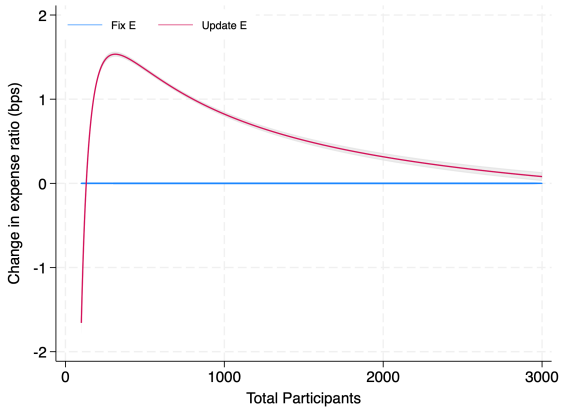

(e) Difference in management fees

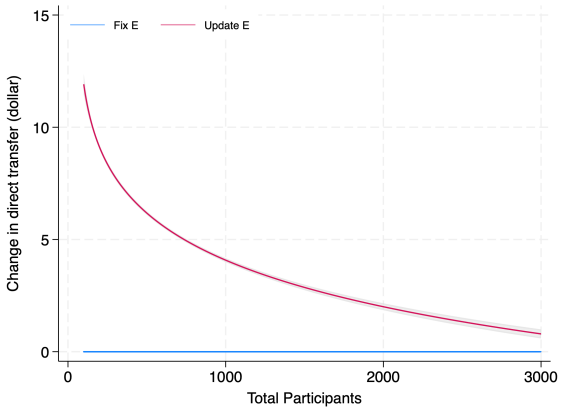

(f) Difference in recordkeeping fees

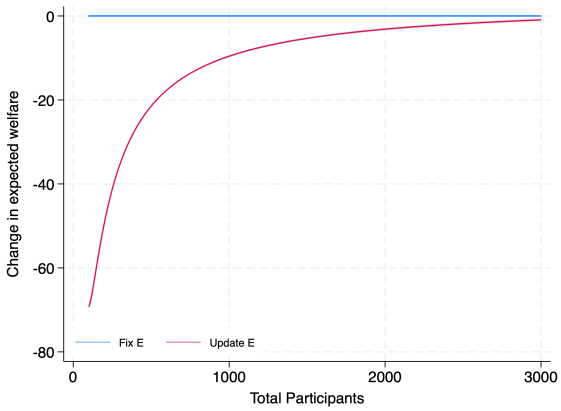

(g) Difference in per-participant welfare

Figure 18: Counterfactual market outcomes with Fidelity's PEP

This figure compares the market outcomes in the counterfactual analysis where Fidelity starts its own pooled employer plan with no exclusion to the base case. Panel (a) presents market shares; Panel (b) reports management fees, measured by the average expense ratio in basis points; Panel (c) shows recordkeeping fees in dollars; and Panel (d) displays the per-participant expected welfare.

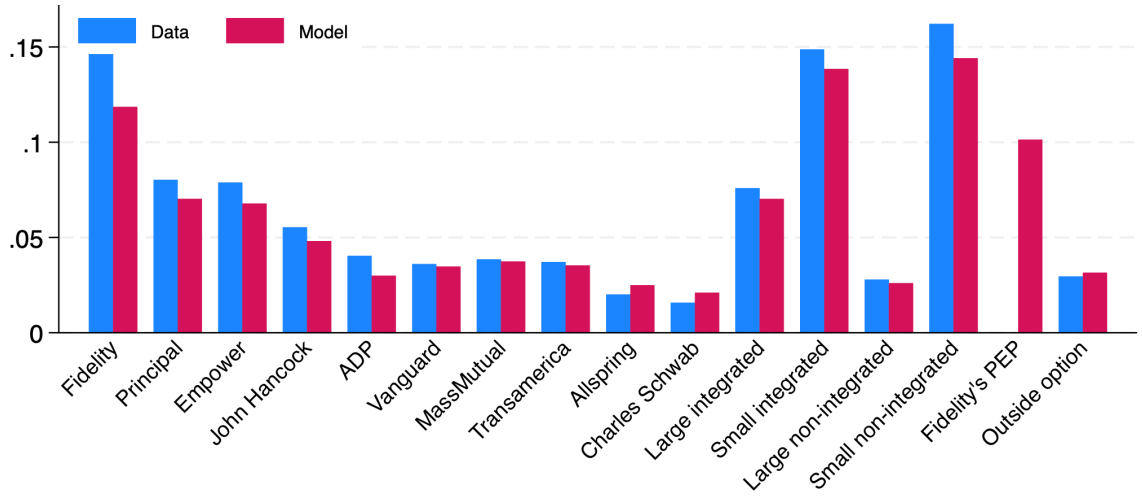

(a) Market share

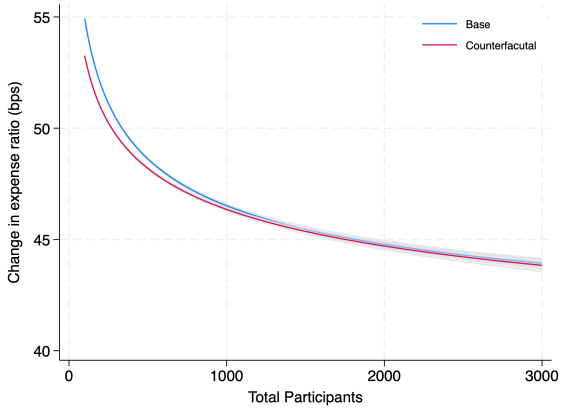

(b) Management fees

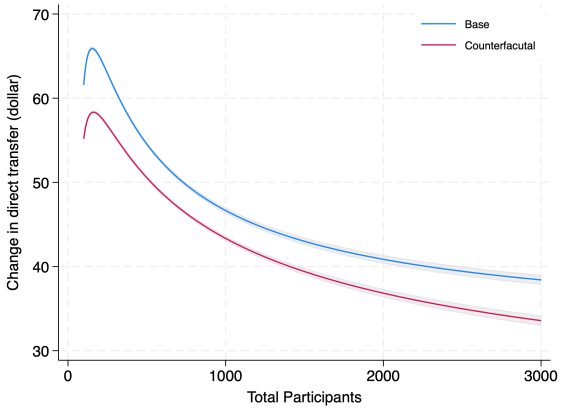

(c) Recordkeeping fees

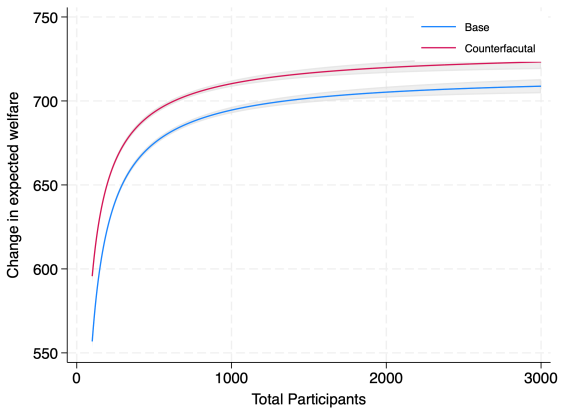

(d) Per-participant expected welfare

Figure 19: Counterfactual market outcomes with supply-side mergers

This figure compares the market outcomes in the counterfactual analysis where non-top-10 recordkeepers of the four types merge into four consolidated recordkeepers and optimally set their exclusions. Panel (a) presents market shares; Panel (b) reports management fees, measured by the average expense ratio in basis points; Panel (c) shows recordkeeping fees in dollars; and Panel (d) displays the per-participant expected welfare.

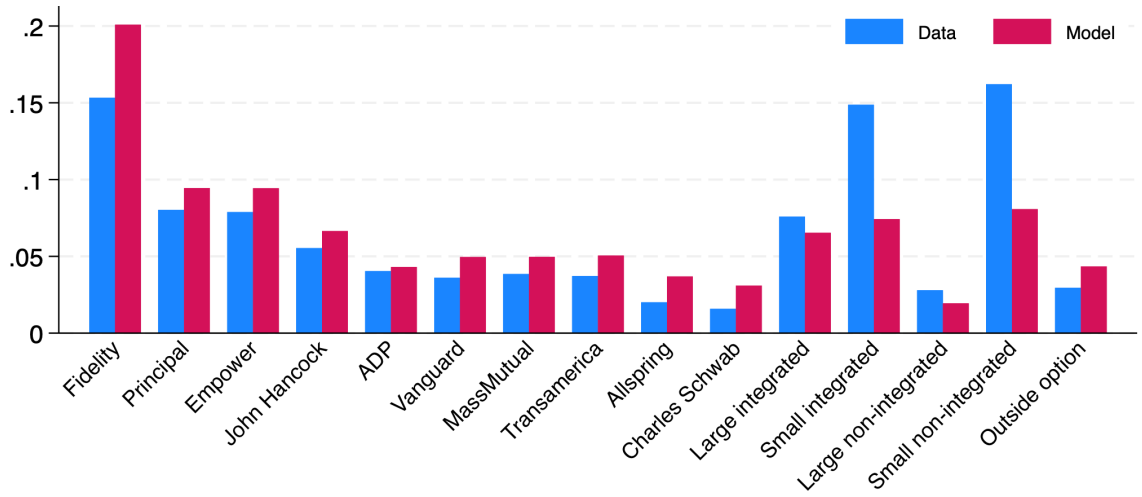

(a) Market share

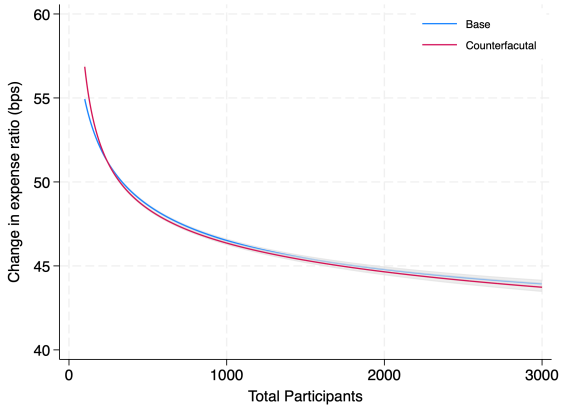

(b) Management fees

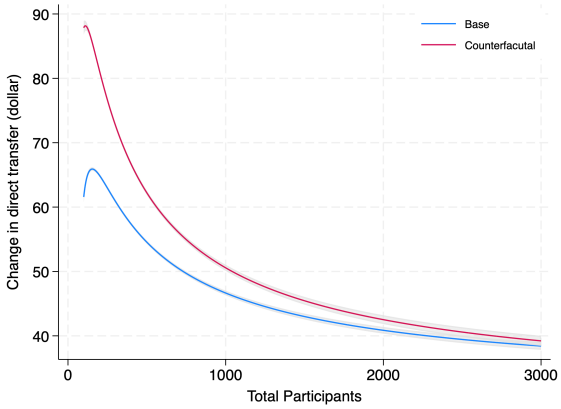

(c) Recordkeeping fees

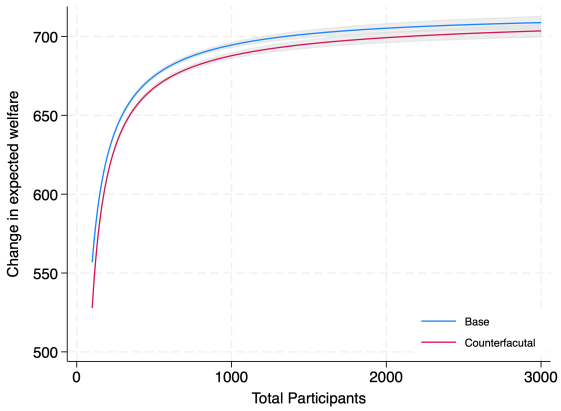

(d) Per-participant expected welfare

Appendix A: Data Construction Details

A.1 Match Form 5500 to BrightScope

Form 5500 and BrightScope have different identifiers for plans. The identifier for plan-year observation in Form 5500 is the *ack_id*, while the identifier for plan-year observation in BrightScope is *uniqueid*. I first match these two identifiers. Both Form 5500 and BrightScope provide the company identifier *companyein* of the sponsor. However, one company can have multiple plans. I match the plans in two steps: I first match plan names corresponding to the same *companyein* by fuzzy string matching using the FuzzyWuzzy package in Python. Then I check the number of participants and net assets of the matched plan to make sure they have exactly the same information of the two.

A.2 Extract revenue sharing information

Schedule C of Form 5500 provides information on revenue sharing. Item 2 of Schedule C includes information on the service providers of each plan. Item 3 of Schedule C gives information on the indirect income paid by other institutions to each service provider. The information on indirect income is given in two formats: a number or a text description. In the latter case, it's typically described as a percentage of plan assets. I need to extract the information from the texts.

I extract the revenue on revenue sharing in the following steps:

Step 1: Categorize the type of payors in Item 3. Following Bhattacharya and Illanes (2022), I assume that if the payor is a mutual fund, the service provider receives revenue sharing as a percentage of assets invested in that fund; if the payor is an asset management company, the service provider receives revenue sharing as a percentage of assets invested in all funds of the asset management company. Therefore, I need to categorize by name if a payor is a mutual fund or an asset management company.

To do this, I fine-tuned a Bert model. The training data comes from the CRSP mutual fund database. CRSP provides the full name of each mutual fund and the full name of the asset management company. I build a data set that maps a name to a label that indicates whether the name refers to a mutual fund or an asset management company, then fine-tune the Bert model for this classification task. With the trained model, I predict whether each payor is a mutual fund or an asset management company.

Step 2: Map service providers in Schedule C Item 2 to providers in BrightScope. In order to link match plan menu information from BrightScope, I need to map providers in Form 5500 to providers in BrightScope. Since I have mapped each plan in Form 5500

to each plan in BrightScope, I match the set of providers of each plan in Form 5500 to the set of providers in BrightScope. This is done by fuzzy string matching.

Step 3: Map provider names in Schedule C Item 2 to provider names in Schedule C Item 3. Item 2 and Item 3 do not have a consistent identifier for the same service provider of a plan, therefore, I need to match the providers by their name. This is also done by fuzzy string matching.

Step 4: Map payor names in Schedule C Item 3 to provider names in BrightScope. Payors are other institutions that transfer to the recordkeeper, and they correspond to each investment option provider on the menu provided by BrightScope. I match the payor names in Item 3 to investment option providers in BrightScope by fuzzy string match.

Step 5: Extract indirect transfer information from Schedule C Item 3. The text description of the indirect transfer has two main types: the majority gives a percentage, and the others give a per-participant dollar expense. I designed regular expressions to extract these numbers.

Unlike Bhattacharya and Illanes (2022), I don't assume that investment option providers that do not have revenue-sharing information in Form 5500 transfer the 12b-1 fees to the recordkeeper. Pool, Sialm, and Stefanescu (2022) find that revenue-sharing and 12b-1 fees are largely unrelated. Whenever I don't observe indirect transfer information in Form 5500, I assume there exists no revenue-sharing agreement between the recordkeeper and the investment option provider.

A.3 Predict revenue-sharing and management fees

Similar to Bhattacharya and Illanes (2022), I predict the per-participant revenue-sharing and per-participant management fees received by the recordkeeper using XGBoost. Instead of directly predicting the dollar revenue, I find that predicting revenue sharing and management fees as a percentage of plan balance has a higher accuracy. Therefore, I first predict the percentages, then multiply them by the plan balance, and finally divide the total number of participants.

The variables used in the predictions include the total number of options, total number of mutual funds, counts of investments in each BrightScope investment category, percentage of target-date funds, percentage of index funds, percentage of affiliated funds, HHI of investment option providers, per-participant balance, average fund expense ratio, maximum and minimum of expense ratios, 12b-1 fees, and management fees, percentages of funds in each fee decile, sponsor 2-digit NAICS2 codes, sponsor state codes, and recordkeeper rank indicators.

Appendix B: Model Details

Model setting. The utility function is

$$u(M) = \begin{cases} \delta(M) & \text{if } (1-\eta)\delta(M) - \eta(M - mc) \leq 0 \\ \eta(\delta(M) + M) & \text{if } (1-\eta)\delta(M) - \eta(M - mc) > 0 \end{cases}$$

Without loss of generality, assume $mc = 0$. Consider the cases where the utility is linear in M, that is

$$\delta(M) = a - bM$$

Assume that $M \in [0, \bar{M}]$.

Discussion of the shape of the utility function.

Define $T^*(M) = (1-\eta)\delta(M) - \eta M$. Plug in $\delta(M) = a - BM$ yields

$$T^*(M) = (1-\eta)a - \big((1-\eta)b + \eta\big)M$$

Denote the slope by

$$k = (1-\eta)b + \eta$$

When $k > 0$, the unique kink where $T^*(M) = 0$ is

$$M_s = \frac{(1-\eta)a}{k} > 0,$$

The corresponding utility at the kink is

$$u(M_s) = \frac{a\eta}{k}$$

The left and right slopes are

$$u'_L(M) = \eta(1-b) \qquad u'_R(M) = -b$$

Discussion of the maximizer M^*.

- **Case 1 ($k \leq 0$, i.e. $b \leq -\frac{\eta}{1-\eta}$):** ($T^*(M) > 0$ for $M \in [0, \infty)$) and u is strictly increasing. In this case,
$$M^* = \bar{M}$$
- **Case 2 ($k > 0$ and $b < 0$):** Both pieces increase ($u'_L > 0$ and $u'_R > 0$), so u is

strictly increasing. In this case,

$$M^* = \bar{M}$$

- **Case 3 ($k > 0$ and $b = 0$):** The function $u(M)$ increases up to the kink, then flat ($u'_L = \eta > 0$, $u'_R = 0$). In this case,

$$M^* = \begin{cases} \overline{M}, & \overline{M} \le M_s, \\ \text{any } M \in [M_s, \overline{M}], & \overline{M} \ge M_s, \end{cases} \qquad u^* = \begin{cases} \eta(a + \overline{M}), & \overline{M} \le M_s, \\ a, & \overline{M} \ge M_s. \end{cases}$$

- **Case 4 ($k > 0$ and $0 < b < 1$):** The function $u(M)$ increases on the left, decreases on the right ($u'_L > 0$, $u'_R < 0$). In this case,

$$M^* = \begin{cases} M_s = \dfrac{(1-\eta)a}{(1-\eta)b + \eta}, & M_s \le \overline{M}, \\ \overline{M}, & M_s > \overline{M}, \end{cases} \qquad u^* = \begin{cases} \dfrac{a\eta}{(1-\eta)b + \eta}, & M_s \le \overline{M}, \\ \eta\big(a + (1-b)\overline{M}\big), & M_s > \overline{M}. \end{cases}$$

- **Case 5 ($k > 0$ and $b = 1$):** The function $u(M)$ is flat on the left, decreases on the right ($u'_L = 0$, $u'_R = -1$). Here $M_s = (1-\eta)a$. In this case

$$M^* \in \big[0,\ \min\{\overline{M}, M_s\}\big], \qquad u^* = \eta a.$$

  (A canonical choice is $M^* = M_s$ when $M_s \le \overline{M}$, which also yields zero transfer at the kink.)

- **Case 6 ($k > 0$ and $b > 1$):** Both pieces decrease ($u'_L < 0$, $u'_R < 0$). In this case,

$$M^* = 0, \qquad u^* = \eta a.$$

Appendix C: Structural Model Details

The sponsor's utility function is

$$u_{ijp} = \bar{u}_{ijp} + \sigma\epsilon_{ijp} = \delta_{ijp} - T^{res}_{ijp} + \sigma\epsilon_{ijp},$$

The reservation transfer is

$$T^{res}_{ijp} = \max\left\{0, \frac{E_j}{N_i} + mc_j - R_{ijp} - M_{ijp}\right\}$$

Sponsor i chooses (j^*, p^*) by solving problem

$$(j^*, p^*) = \arg\max_{(j,p)}\{u_{ijp}\}$$

$$F_J\left(\epsilon_i\right) = \exp\left[-\sum_{j=1}^{J}\left(\sum_{p\in\mathcal{P}_j}\exp\left(-\frac{\epsilon_{ijp}}{\rho}\right)\right)^{\rho}\right],\ \epsilon_i = \{\epsilon_{ijp}\}_{jp}.$$

The choice probability is

$$s_{ijp} = \frac{\exp(\frac{\bar{u}_{ijp}}{\rho\sigma})\left[\sum_{p\in\mathcal{P}_j}\exp(\frac{\bar{u}_{ijp}}{\rho_j\sigma})\right]^{\rho-1}}{\sum_{j'}\left(\sum_{p'\in\mathcal{P}_{j'}}\left[\exp(\frac{\bar{u}_{ij'p'}}{\rho\sigma})\right]^{\rho}\right)+1}$$

For optimal choice (j^*, p^*), the recordkeeper and the sponsor bargain over the direct transfer

$$T = \arg\max_{T\geq 0}\left(T_{ij^*p^*} + R_{ij^*p^*} + M_{ij^*p^*} - mc_{j^*} - \frac{E_{j^*}}{N_i}\right)^{1-\eta}\left(\delta_{ij^*p^*} - T_{ij^*p^*} + \sigma\epsilon_{ij^*p^*} - \phi_{ij^*}(\epsilon_i)\right)^{\eta}$$

where

$$\phi_{ij^*}(\epsilon_i) = \max_{j\neq j^*, p\in\mathcal{P}_j}\{u_{ijp}\}$$

The solution is

$$T_{ijp} = \max\{0, T^*_{ijp}\}$$

where

$$
\begin{aligned}
T^*_{ijp} &= (1-\eta)\left(\delta_{ijp} + \sigma\epsilon_{ijp} - \phi_{ij}(\epsilon_i)\right) - \eta\left(R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}\right) \\
&= (1-\eta)\sigma\left(\epsilon_{ijp} - \phi_{ij}(\epsilon_i)/\sigma\right) + (1-\eta)\delta_{ijp} - \eta\left(R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}\right)
\end{aligned}
$$

By construction, the Nash bargaining is valid, since

$$
\begin{aligned}
&\left(T_{ijp} + R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}\right) + \left(\delta_{ijp} - T_{ijp} + \sigma\epsilon_{ijp} - \phi_{ij^*}(\epsilon)\right) \\
=&\left(T^{res}_{ijp} + R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}\right) + \left(\delta_{ijp} - T^{res}_{ijp} + \sigma\epsilon_{ijp} - \phi_{ij^*}(\epsilon)\right) \\
\geq& 0.
\end{aligned}
$$

The last inequality follows from (8) and $u_{ijp} \geq \phi_{ijp}(\epsilon)$.

Let C_{ijp} be the indicator that plan p by provider j is chosen by sponsor i. By construction,

$$
\{C_{ijp} = 1\} = \left\{u_{ijp} \geq \phi_{ijp}(\epsilon), \quad u_{ijp} \geq \max_{p' \in \mathcal{P}_j, p' \neq p} u_{ijp}\right\}.
$$

By the property of EVT1 distribution, define (e_1, e_2, e_3) with CDF

$$
F(e_1, e_2, e_3) = \exp\left(-\left(\exp(-e_1/\rho) + \exp(-e_2/\rho]\right)^\rho - \exp(-e_3)\right).
$$

Then

$$
\begin{aligned}
&\left(u_{ij^*p^*}/\sigma, \quad \max_{j \neq j^*, p \in \mathcal{P}_j}\{u_{ijp}/\sigma\}, \quad \phi_{ij^*}(\epsilon)/\sigma, \quad T^*_{ijp}\right) \\
&\sim (\mu_1 + e_1, \quad \mu_2 + e_2, \quad \mu_3 + e_3, \quad (1-\eta)\sigma(e_1 - e_3) + \bar{T})
\end{aligned}
$$

where

$$
\begin{aligned}
\mu_1 &= \frac{\bar{u}_{ij^*p^*}}{\sigma} \\
\mu_2 &= \log\left(\left[\sum_{p \in \mathcal{P}_{j^*} \setminus p^*} \exp\left(\frac{\bar{u}_{j^*,p}}{\rho_{j^*}\sigma}\right)\right]^{\rho_{j^*}}\right) \\
\mu_3 &= \log\left(\sum_{j \neq j^*}\left[\sum_{p \in \mathcal{P}_j} \exp\left(\frac{\bar{u}_{j,p}}{\rho_j\sigma}\right)\right]^{\rho_j} + 1\right) \\
\bar{T} &= (1-\eta)(\delta_{ijp} - \sigma\mu_3) - \eta\left(R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}\right)
\end{aligned}
$$

The conditional expectation of T_{ijp} is

$$\begin{aligned}
E\left[T_{ijp}|C_{ijp}=1\right] =& E\left[T_{ijp}|C_{ijp}=1,T_{ijp}>0\right]P\left(T_{ijp}>0|C_{ijp}=1\right) \\
&+E\left[T_{ijp}|C_{ijp}=1,T_{ijp}=0\right]P\left(T_{ijp}=0|C_{ijp}=1\right) \\
=& E(T_{ijp}|C_{ijp}=1,T_{ijp}>0)P(T_{ijp}>0|C_{ijp}=1) \\
=& E(T^*_{ijp}|C_{ijp}=1,T^*_{ijp}>0)P(T^*_{ijp}>0|C_{ijp}=1)
\end{aligned}$$

where

- $C_{ijp}=1$ if

$$\left\{\mu_1+e_1\geq\mu_3+e_3,\quad \mu_1+e_1\geq\mu_2+e_2\right\}=\left\{e_1\geq\mu_2-\mu_1+e_2,\quad e_1\geq\mu_3-\mu_1+e_3\right\}.$$

- $T^*_{ijp}>0$ if

$$e_1-e_3>-\frac{1}{(1-\eta)\sigma}\bar{T}$$

In sum,

$$\left\{C_{ijp}=1,T_{ijp}>0\right\}=\left\{e_1\geq A+e_2,\quad e_1\geq B+e_3\right\}$$

where

$$A=\mu_2-\mu_1,\qquad\qquad B=\max\left\{\mu_3-\mu_1,-\frac{1}{(1-\eta)\sigma}\bar{T}\right\}.$$

Then, let

$$\begin{aligned}
F_{13}(e_1,e_2,e_3) &= \frac{\partial^2 F(e_1,e_2,e_3)}{\partial e_1\partial e_3} \\
&= \left(e^{-\frac{e_1}{\rho}}+e^{-\frac{e_2}{\rho}}\right)^{\rho-1}\exp\left(-\frac{\rho\left(\left(e^{-\frac{e_1}{\rho}}+e^{-\frac{e_2}{\rho}}\right)^{\rho}+e_3\right)+e_1}{\rho}-e^{-e_3}\right).
\end{aligned}$$

The key is to calculate

$$\begin{aligned}
g(A,B;\rho) =& E\left[e_1 - e_3 | \, e_1 \geq A + e_2, \quad e_1 \geq B + e_3\right] \\
=& \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{e_1-B}\int_{-\infty}^{e_1-A}(e_1-e_3)f_J(e_1,e_2,e_3)\mathrm{d}e_2\mathrm{d}e_3\mathrm{d}e_1}{\int_{-\infty}^{\infty}\int_{-\infty}^{e_1-B}\int_{-\infty}^{e_1-A}f_J(e_1,e_2,e_3)\mathrm{d}e_2\mathrm{d}e_3\mathrm{d}e_1} \\
=& \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{e_1-B}(e_1-e_3)F_{13}(e_1,e_1-A,e_3)\mathrm{d}e_3\mathrm{d}e_1}{\int_{-\infty}^{\infty}F_1(e_1,e_1-B,e_1-A)\mathrm{d}e_1} \\
=& \frac{\left(e^{A/\rho}+1\right)^{\rho}+e^{B}}{\left(e^{A/\rho}+1\right)^{\rho-1}}\int_{-\infty}^{\infty}\int_{-\infty}^{e_1-B}(e_1-e_3)F_{13}(e_1,e_1-A,e_3)\mathrm{d}e_3\mathrm{d}e_1
\end{aligned}$$

Then,

$$\begin{aligned}
E\left[T_{ijp}^{*}|C_{ijp}=1,T_{ijp}^{*}>0\right] =& E\left[(1-\eta)\sigma(e_1-e_3)+\bar{T}|e_1\geq A+e_2, \quad e_1\geq B+e_3\right] \\
=& \bar{T}+(1-\eta)\sigma g(A,B;\rho).
\end{aligned}$$

$$\begin{aligned}
E\left[T_{ijp}|\, C_{ijp}=1\right] =& E\left[(1-\eta)\sigma(e_1-e_3)+\bar{T}|e_1\geq \mu_2-\mu_1+e_2, \quad e_1\geq \mu_3-\mu_1+e_3\right] \\
=& \bar{T}+(1-\eta)\sigma g(A,\mu_3-\mu_1;\rho).
\end{aligned}$$

Recordkeeper j's profit of providing plan p to sponsor i is

$$\pi_{ijp} = T_{ijp} + R_{ijp} + M_{ijp} - mc_j - \frac{E_j}{N_i}$$

The expected profit is thus

$$\begin{aligned}
\Pi_j =& \sum_i \sum_{p\in\mathcal{P}_{ij}} P\left(C_{ijp}=1\right) E\left[\pi_{ijp}|\, C_{ijp}=1\right] \\
=& \sum_i \sum_{p\in\mathcal{P}_{ij}} s_{ijp}\left(R_{ijp}+M_{ijp}-mc_j-\frac{E_j}{N_i}+E\left[T_{ijp}|\, C_{ijp}=1\right]\right) \\
=& \sum_i \sum_{p\in\mathcal{P}_{ij}} s_{ijp}\left(R_{ijp}+M_{ijp}-mc_j-\frac{E_j}{N_i}+\bar{T}+(1-\eta)\sigma\times g(A,\mu_3-\mu_1;\rho)\right)
\end{aligned}$$

Appendix D: Additional Tables and Figures

D.1 Supplementary Tables and Figures for Section 2

Figure D-1 supplements Figure 3 by reporting the proportion of affiliated options and the average expense ratio of affiliated options by plan size. Large plans tend to have fewer affiliated options, and those affiliated options tend to have a lower expense ratio.

Figure D-2 supplements Figure 1 by showing the ranks of recordkeepers by the total number of plans, total assets under management (AUM), and total number of participants.

[Insert Figure D-2 here]

Table D-1 and Table D-2 provide summary statistics of key plan features by each recordkeeper in addition to Table 1.

[Insert Table D-1 here]

Table D-1 reports the proportion of plans without revenue from management fees and revenue sharing. According to the tables, revenue sharing is not observed very frequently in the data. Therefore, when constructing the dataset, I deviate from Bhattacharya and Illanes (2022) where they predict revenue sharing for each plan using 12b-1 fees; instead, I assume that a plan has revenue sharing only if there exists revenue sharing agreements in the data. The amount of revenue sharing when the plan is provided to a different sponsor is predicted using machine learning methods as detailed in Appendix A.

[Insert Table D-2 here]

Table D-2 reports the summary statistics as reported in Figure 1 by recordkeeper. Figure D-3 visualizes the distributions of summary statistics by recordkeeper.

[Insert Figure D-3 here]

D.2 Supplementary Tables and Figures for Section 6

Table D-4 reports the geographic distribution of plans served by each recordkeeper. The plans are counted at the county level.

[Insert Figure D-4 here]

Table D-1: Revenue sources

This table reports the percentages of plans without management fees and revenue sharing served by each recordkeeper. The second column reports the percentages of plans served by each recordkeeper where there is no affiliated investment products of the recordkeeper. The third column reports the percentages of plans where there is no revenue sharing between the recordkeeper and other investment product providers.

	N	% Missing management fees	% Missing revenue sharing
Fidelity	8,398	0.036	0.151
Principal	4,170	0.041	0.950
Empower	4,137	0.217	0.995
John Hancock	2,868	0.136	0.969
ADP	2,073	1.000	0.993
Vanguard	1,970	0.006	0.831
MassMutual	2,022	0.124	0.413
Transamerica	1,969	0.266	0.845
Allspring	1,145	0.400	0.976
Charles Schwab	865	0.275	0.754
Integrated	11,929	0.389	0.750
Non-integrated	9,883	1.000	0.890

Table D-2: Summary statistics of plans by recordkeeper

This table reports the summary statistics of plans by recordkeeper. Panel A to Panel E report summary statistics of plan participants, per-participant recordkeeping fees, per-participant revenue sharing, per-participant management fees, and total revenue.

Panel A: Plan participants

	Mean	SD	p10	p25	p50	p75	p90
Fidelity	1,762	7,069	144	203	382	989	3,177
Principal	718	2,895	130	164	261	509	1,225
Empower	840	3,108	135	171	271	525	1,282
John Hancock	561	1,892	128	158	229	407	914
ADP	422	1,286	127	159	225	368	726
Vanguard	1,861	10,150	139	184	328	850	3,027
MassMutual	822	2,524	135	170	273	565	1,410
Transamerica	1,085	3,189	143	194	343	846	2,069
Allspring	3,162	14,181	157	244	561	1,635	5,124
Charles Schwab	1,974	6,340	169	267	615	1,717	3,808
Intergrated	1,272	7,420	135	172	279	597	1,631
Non-integrated	716	3,934	130	161	238	446	1,006

Panel B: Per-participant recordkeeping fee

	Mean	SD	p10	p25	p50	p75	p90
Fidelity	42.0	48.0	2.9	7.9	22.8	62.5	107.8
Principal	111.9	113.5	11.8	27.2	76.7	161.2	259.3
Empower	67.8	71.2	4.1	15.1	49.3	97.3	155.9
John Hancock	29.2	75.4	0.9	2.1	4.8	13.1	95.7
ADP	27.6	38.3	2.4	5.3	12.5	31.5	78.4
Vanguard	58.9	196.6	4.6	14.5	47.4	82.6	112.2
MassMutual	91.0	117.3	5.2	12.8	38.1	129.3	251.8
Transamerica	86.2	114.9	11.2	26.6	62.4	109.9	181.4
Allspring	111.9	107.9	13.7	39.0	82.3	148.8	249.5
Charles Schwab	30.0	59.4	0.0	0.8	6.9	30.3	82.7
Intergrated	76.7	107.7	2.4	10.2	40.3	105.2	192.1
Non-integrated	76.6	94.0	6.2	18.4	50.8	99.0	176.1

Table D-2: Summary statistics of plans by recordkeeper (continued)

This table reports the summary statistics of plans by recordkeeper. Panel A to Panel E report summary statistics of plan participants, per-participant recordkeeping fees, per-participant revenue sharing, per-participant management fees, and total revenue.

Panel C: Per-participant revenue sharing (conditional on positive)

	Mean	SD	p10	p25	p50	p75	p90
Fidelity	40.5	73.7	0.9	4.1	15.6	47.6	106.5
Principal	22.3	46.6	0.1	0.4	2.0	20.9	79.2
Empower	36.9	73.5	0.2	2.9	11.3	33.5	99.0
John Hancock	87.2	107.0	2.2	4.9	47.7	121.9	239.9
ADP	32.4	31.6	0.7	3.0	25.1	59.2	81.5
Vanguard	8.7	17.1	0.2	0.8	2.7	8.1	20.4
MassMutual	731.2	21,274.3	0.2	4.7	50.4	151.9	311.7
Transamerica	68.3	83.3	3.8	13.6	39.8	87.6	169.2
Allspring	42.9	62.9	1.0	4.0	11.0	44.0	144.5
Charles Schwab	24.2	75.8	0.0	0.1	6.3	19.3	48.9
Integrated	50.2	135.6	0.8	3.7	16.7	60.3	130.3
Non-integrated	28.8	56.7	0.1	0.4	4.4	30.3	89.4

Panel D: Per-participant management fee

	Mean	SD	p10	p25	p50	p75	p90
Fidelity	180.9	237.0	2.6	18.6	98.2	265.3	460.8
Principal	187.4	198.4	7.6	41.2	132.6	267.9	438.6
Empower	50.8	89.8	1.6	6.3	22.7	58.3	133.2
John Hancock	161.2	216.1	3.9	22.3	87.2	221.8	409.6
ADP	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Vanguard	97.7	132.9	2.7	13.7	44.5	132.4	270.2
MassMutual	83.5	136.6	0.7	7.2	33.7	94.4	231.4
Transamerica	25.4	53.0	0.2	1.9	8.4	25.6	67.4
Allspring	58.6	93.2	1.0	6.7	24.4	70.7	157.2
Charles Schwab	99.5	177.4	0.0	1.9	24.5	125.7	289.2
Integrated	111.7	200.1	1.1	6.7	32.0	128.3	327.5
Non-integrated	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table D-2: Summary statistics of plans by recordkeeper (continued)

This table reports the summary statistics of plans by recordkeeper. Panel A to Panel E report summary statistics of plan participants, per-participant recordkeeping fees, per-participant revenue sharing, per-participant management fees, and total revenue.

Panel E: Total revenue from a plan (thousand dollars)

	Mean	SD	p10	p25	p50	p75	p90
Fidelity	287.5	1,388.2	19.5	36.1	74.7	177.9	485.3
Principal	130.8	351.2	16.1	34.5	68.7	137.6	254.9
Empower	47.1	102.9	3.4	10.2	22.5	46.1	94.4
John Hancock	60.3	159.0	2.3	9.1	28.1	64.5	130.6
ADP	10.5	24.0	0.5	1.2	3.1	9.3	25.5
Vanguard	289.3	1,366.3	4.9	13.6	36.3	150.3	477.3
MassMutual	341.6	10,304.1	2.9	11.4	50.9	134.7	273.7
Transamerica	80.1	177.4	6.5	13.4	30.6	75.7	186.7
Allspring	148.0	273.1	10.9	28.6	65.4	156.8	320.1
Charles Schwab	95.6	275.2	3.2	9.3	28.6	95.1	205.0
Affiliated	104.6	401.6	1.4	10.0	28.7	69.9	173.5
Non-affiliated	29.6	99.8	1.1	5.8	14.1	29.8	57.3

Figure D-1: Bundling Patterns

This figure reports the proportion of affiliated investment options and the average expense ratio of affiliated investment options by plan size. Affiliated options refer to investment products provided by the asset management affiliates of the recordkeeper. Panel (a) reports the proportion of affiliated investment options on the plan menu; Panel (b) reports the average expense ratio of affiliated investment options.

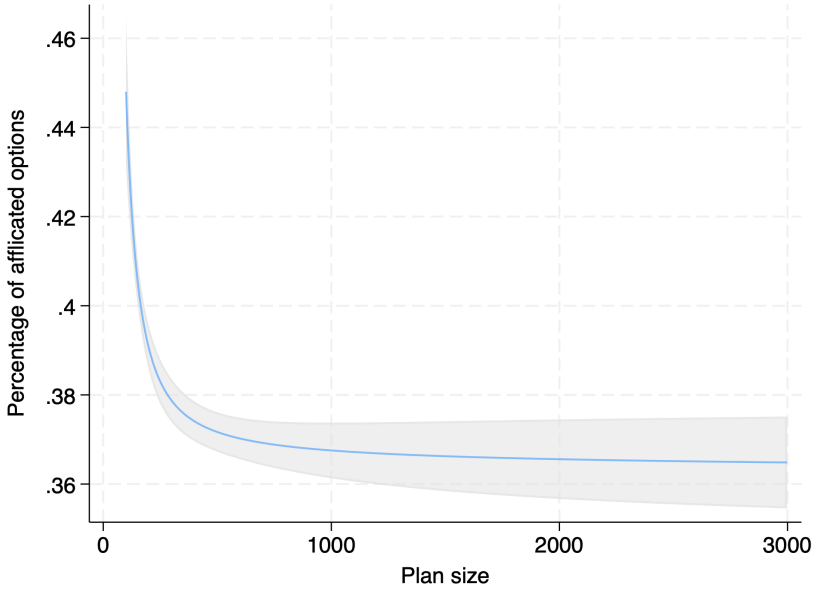

(a) Proportion of affiliated options

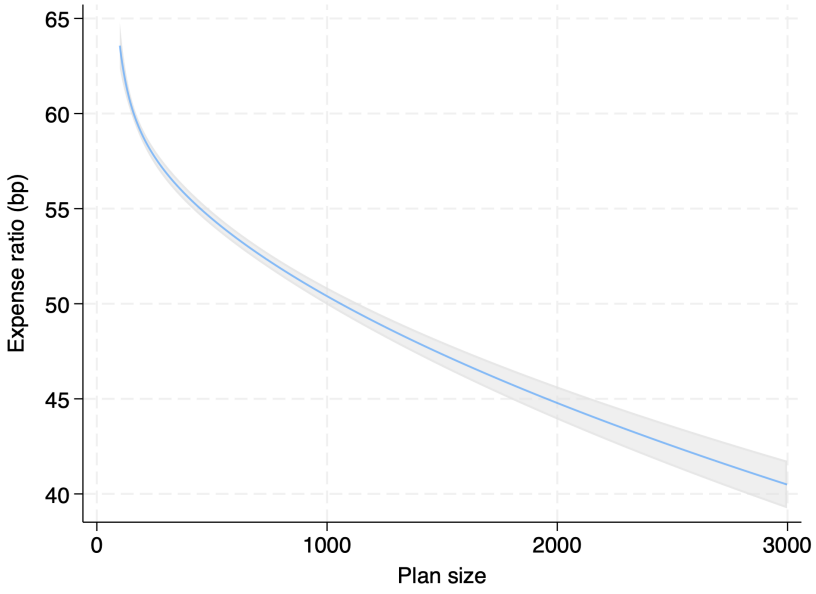

(b) Average expense ratio of affiliation options

Figure D-2: Recordkeeper ranks

This figure reports the ranks of the top 15 recordkeepers by different criteria. Panel (a) reports the ranks by the total number of plans served (i.e., market share). Panel (b) reports the ranks by the total assets under management (AUM). Panel (c) reports the ranks by the total number of plan participants.

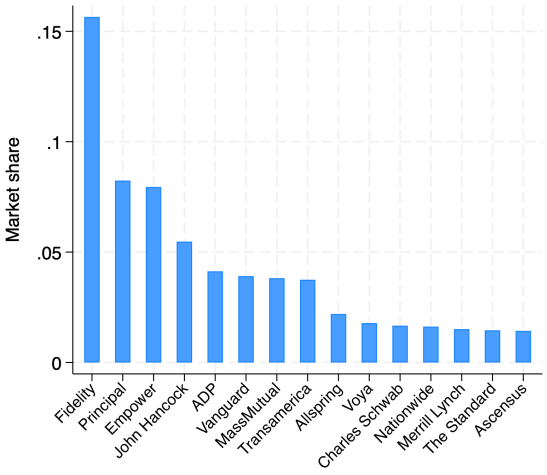

(a) By number of plans

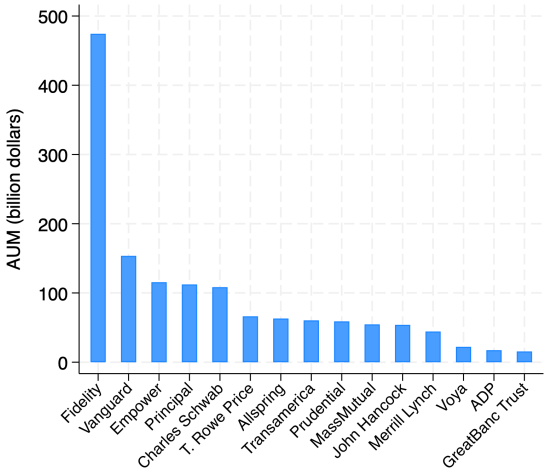

(b) By AUM

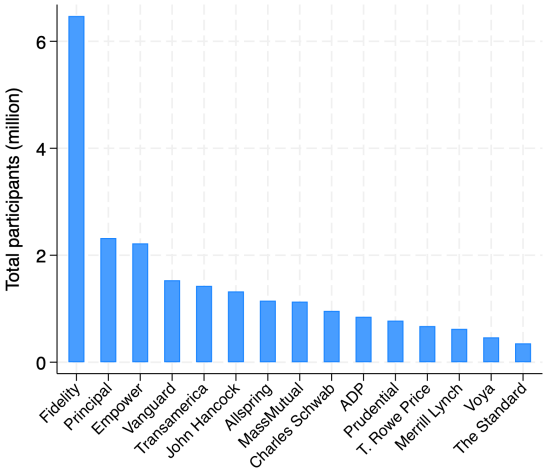

(c) By number of participants

Figure D-3: Plan summary statistics by recordkeeper

This figure reports the distributions of key plan statistics as reported in Table 2-1 by recordkeeper. Panels (a), (b), (d), and (e) use a logarithmic scale on the y-axis. Panel (a) reports the distribution of plan participants. Panel (b) reports the distribution of per-participant recordkeeping fees. Panel (c) reports the distribution of per-participant revenue sharing. Panel (d) reports the ratio of management fees to a plan's total revenue. Panel (e) reports the distribution of plan total revenues.

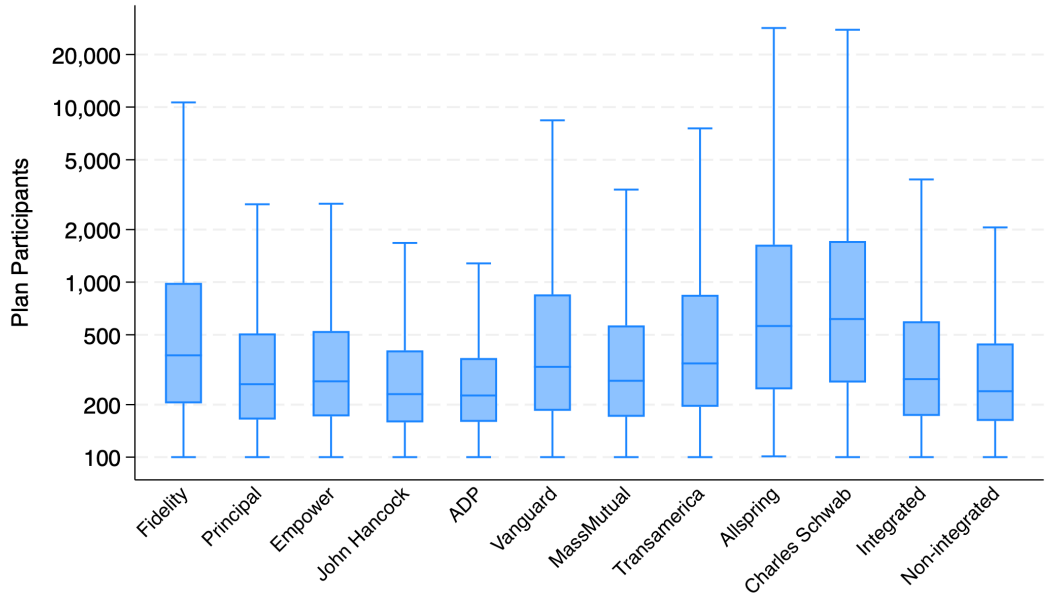

(a) Plan participants

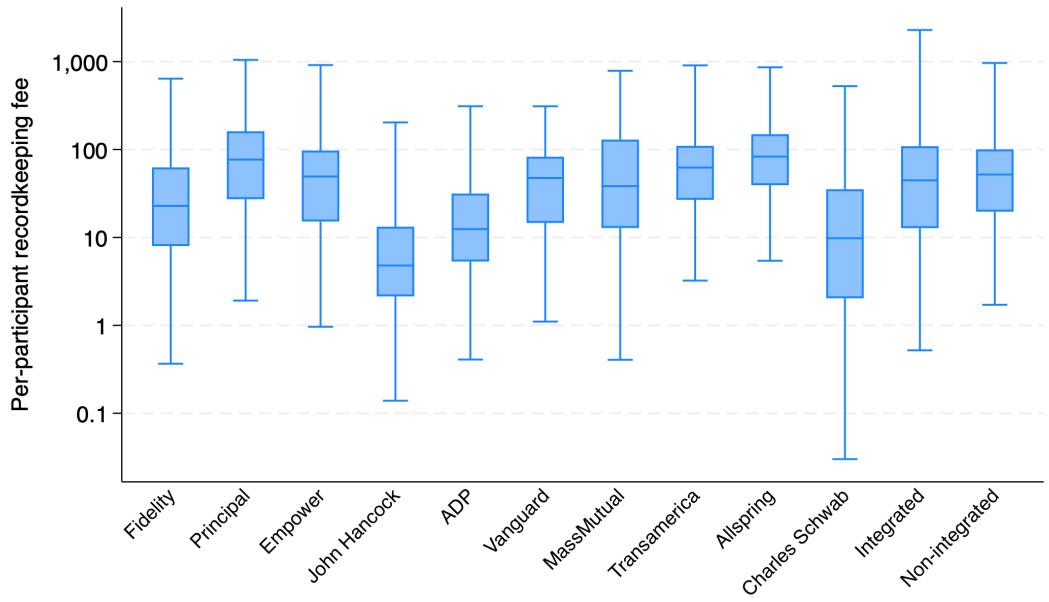

(b) Per-participant recordkeeping fee

Figure D-3: Plan summary statistics by recordkeeper (continued)

This figure reports the distributions of key plan statistics as reported in Table 2-1 by recordkeeper. Panels (a), (b), (d), and (e) use a logarithmic scale on the y-axis. Panel (a) reports the distribution of plan participants. Panel (b) reports the distribution of per-participant recordkeeping fees. Panel (c) reports the distribution of per-participant revenue sharing. Panel (d) reports the ratio of management fees to a plan's total revenue. Panel (e) reports the distribution of plan total revenues.

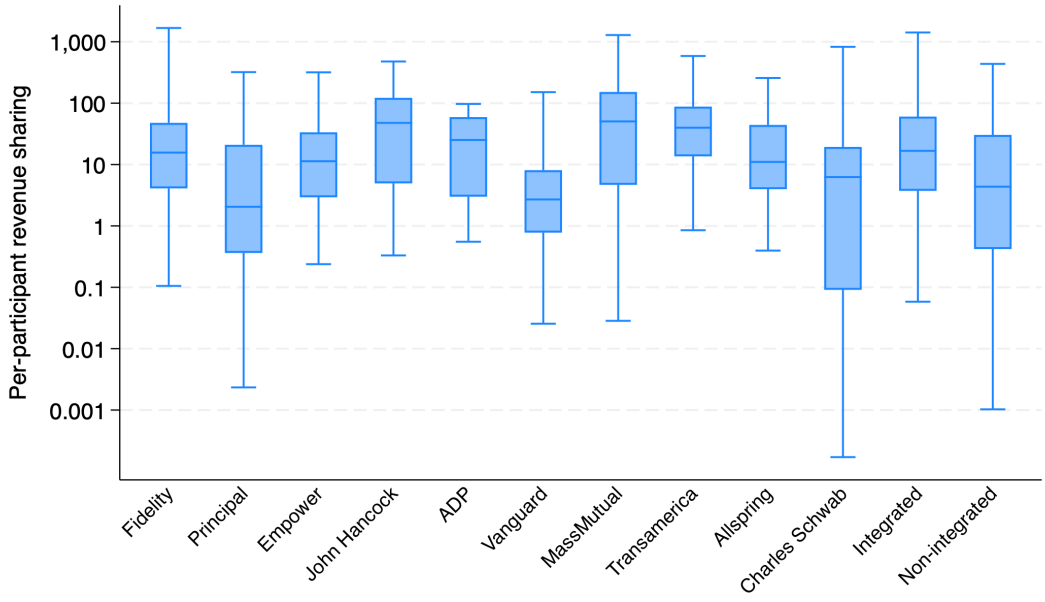

(c) Per-participant revenue sharing

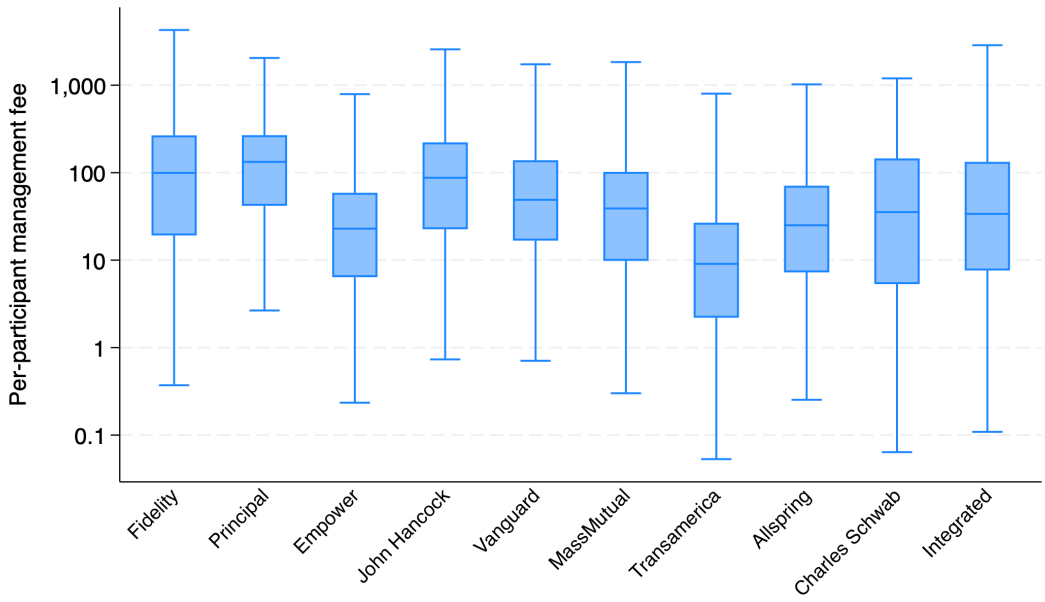

(d) Per-participant management fee

Figure D-3: Plan summary statistics by recordkeeper (continued)

This figure reports the distributions of key plan statistics as reported in Table 2-1 by recordkeeper. Panels (a), (b), (d), and (e) use a logarithmic scale on the y-axis. Panel (a) reports the distribution of plan participants. Panel (b) reports the distribution of per-participant recordkeeping fees. Panel (c) reports the distribution of per-participant revenue sharing. Panel (d) reports the distribution of per-participant management fees. Panel (e) reports the ratio of management fees to a plan's total revenue. Panel (f) reports the distribution of plan total revenues.

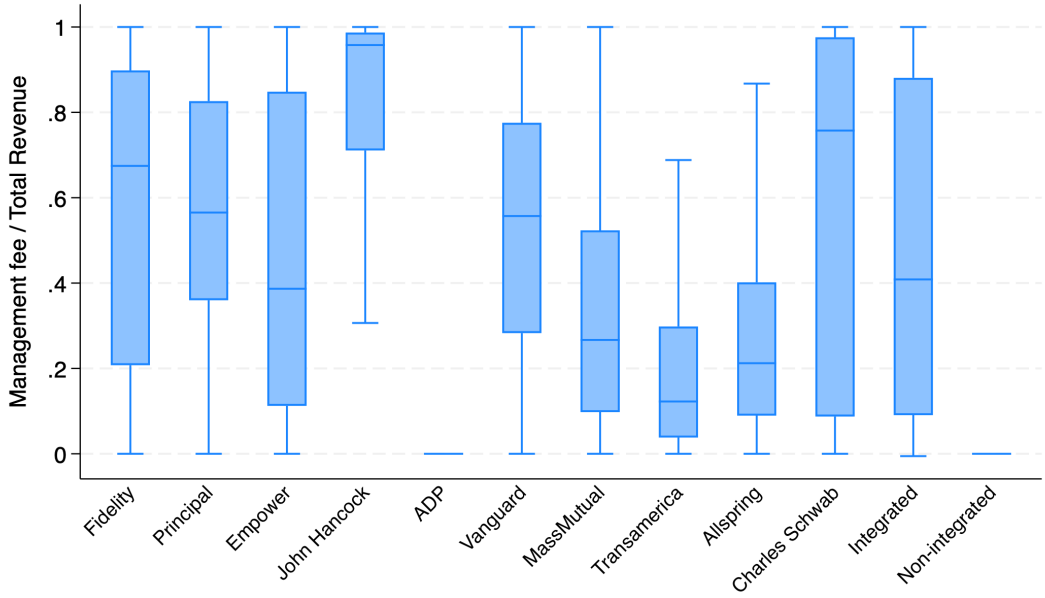

(c) Management fee to total revenue

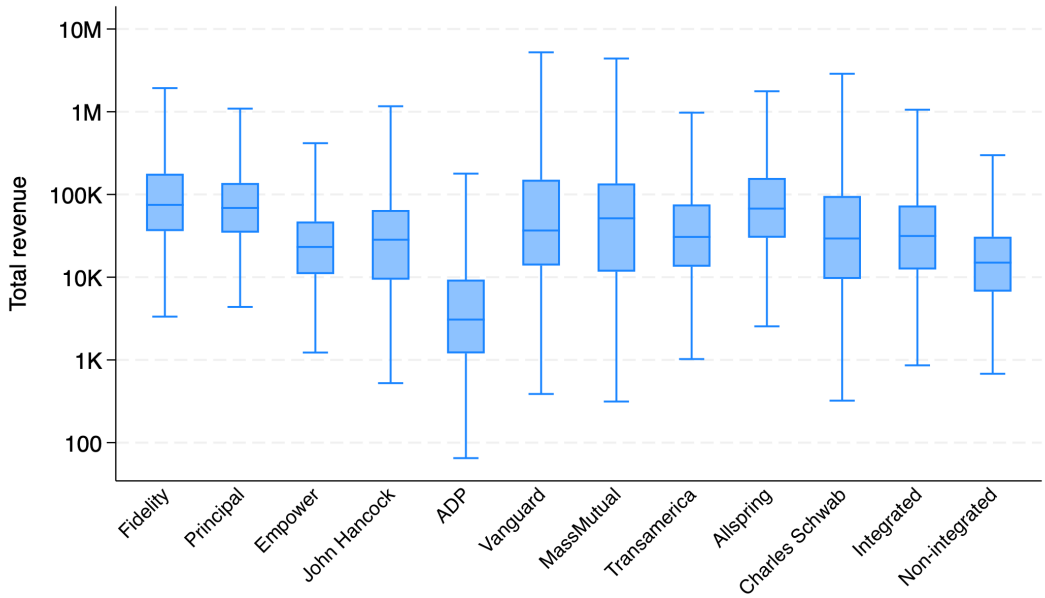

(d) Total revenue

Figure D-4: Geographic distribution of plans by recordkeeper

This figure records the geographic distributions of plans served by each recordkeeper. Plans are counted at the county level.

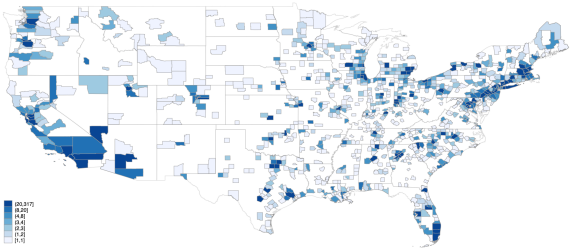

(a) Fidelity

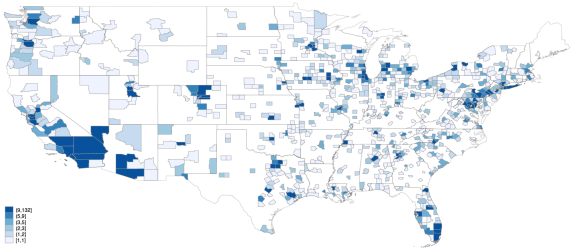

(b) Principal

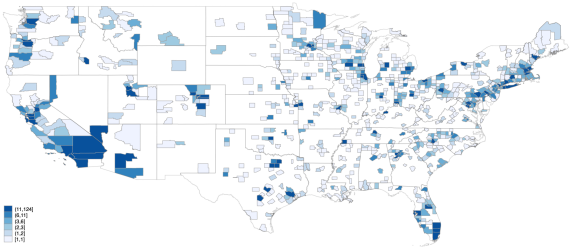

(c) Empower

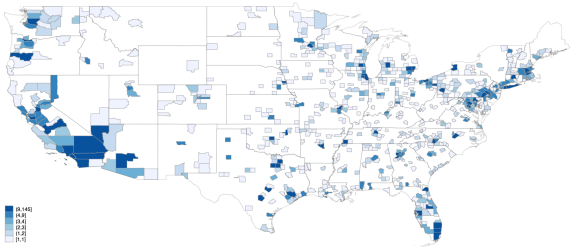

(d) John Hancock

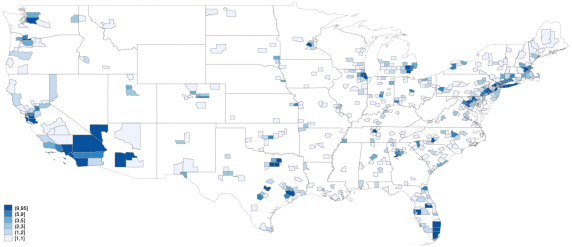

(e) ADP

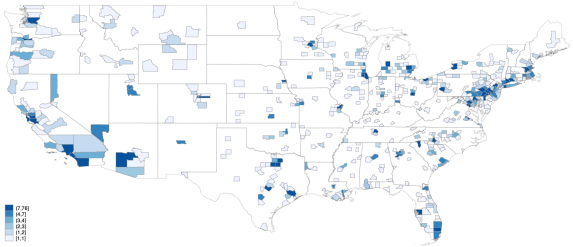

(f) Vanguard

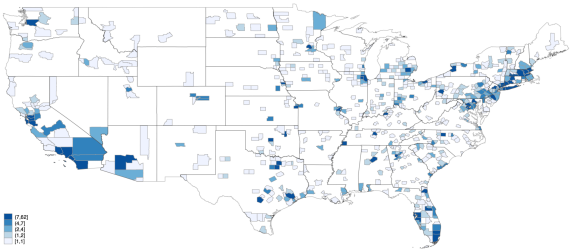

(g) MassMutual

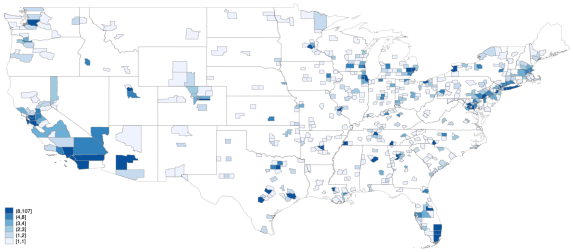

(h) Transamerica

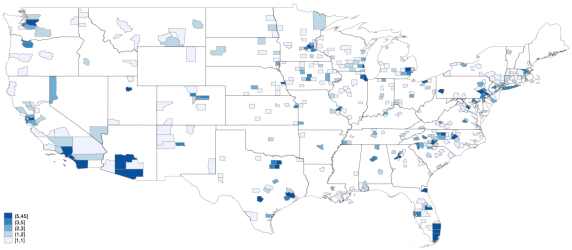

(i) Allspring

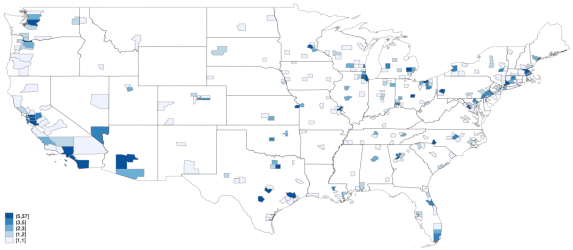

(j) Charles Schwab