NBER WORKING PAPER SERIES

THE DESIGN OF DEFINED CONTRIBUTION PLANS

Vivek Bhattacharya
Gastón Illanes

Working Paper 29981
http://www.nber.org/papers/w29981

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

We are grateful to Diego Cussen, Bernardo de Moura, and Dmitry Sedov for excellent research assistance. We would also like to thank Sebastien Betermier, Rob Clark, Ignacio Cuesta, Glenn Ellison, Hanming Fang, Kate Ho, JF Houde, Ali Hortac¸su, Panle Jia Barwick, Lynn Johnson, Jakub Kastl, Alex MacKay, Aviv Nevo, Rob Porter, Jim Poterba, Jimmy Roberts, Claudia Robles-Garcia, Paulo Somaini, Chad Syverson, and Pietro Tebaldi for helpful comments, as well as participants at the 2020 Winter Econometric Society Meetings, Carnegie Mellon, Chicago IO Fest, Cornell, the CSIO-TSE Conference, Cowles, the Duke-UNC-Fuqua IO Seminar, the Harvard-MIT IO Seminar, Microsoft Research, National University of Singapore, Princeton, Queen's University, Singapore Management University, Stanford, Stanford GSB, SITE, University of Chicago, University of Pennsylvania, and University of Wisconsin. We acknowledge financial support from the National Science Foundation under Grant No. SES 1824463 and the Center for the Study of Industrial Organization at Northwestern. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Design of Defined Contribution Plans
Vivek Bhattacharya and Gastón Illanes
NBER Working Paper No. 29981
April 2022
JEL No. C78, J32, L51

ABSTRACT

Defined contribution plans are a major vehicle for retirement savings in the US, holding almost $10 trillion in assets under management. In recent years, the quality and availability of these plans has been the subject of active policy attention and of several major lawsuits. This paper studies how employers and plan providers design these plans. We argue that low plan quality and limited provision can come from two sources. First, employer willingness to pay may be misaligned with that of workers or of regulators. Second, the market for plan provision may be imperfectly competitive. We propose a model of plan design and estimate that while both frictions are at play, significant changes to plan quality require modifying employer preferences. Accordingly, we evaluate proposed policies and conclude that only direct quality regulation can lead to significant quality improvements, although such regulation reduces plan provision.

Vivek Bhattacharya
Northwestern University
Department of Economics
and NBER
vivek.bhattacharya@northwestern.edu

Gastón Illanes
Northwestern University
Department of Economics
and NBER
gaston.illanes@northwestern.edu

A data appendix is available at http://www.nber.org/data-appendix/w29981

1. Introduction

Defined contribution (DC) plans are one of the most important vehicles for retirement savings in the United States. In 2019, over 100 million workers had DC plans, and total assets invested in them amounted to $9.1 trillion. In a DC plan, workers earn tax advantages for saving a portion of their salary into a pre-specified menu of investments, set in conjunction by the employer and a plan administrator, or "recordkeeper." Regulations place standards of care on employers to try to ensure that the plan contains high-quality investment options.

Nevertheless, many plans arguably do not meet this standard. Many limit workers' access to investment options that are typically considered higher-quality: for the firms in our sample, 42% of DC plans did not include an S&P 500 tracker, 26% did not include a target date fund, and 16% did not include any passively managed index fund at all. Fees for investment options are often high: the average asset-weighted expense ratio is 67 bp, and over 60% of all plans had an asset-weighted expense ratio higher than 57 bp, the average expense ratio among all mutual funds in the US (Simon, 2017). These patterns are even starker for small firms. Moreover, many employers do not even offer a plan.

These issues have sparked growing legal scrutiny. Over 300 lawsuits related to DC plans have been filed since 2015 (Schembari, 2022), and cases have even reached the Supreme Court (e.g., *Tibble v. Edison International* and *Hughes v. Northwestern University*). These lawsuits typically allege that large employers have violated fiduciary responsibilities, particularly by offering high-fee investment options.

This paper studies how employers design their defined contribution plans. We study how both market imperfections and misalignment in employers' willingness to pay for plan characteristics can distort the menu of investment options made available to workers, or cause employers to not offer a plan at all. We then analyze which policy levers can improve outcomes for workers, studying both the impact of policies designed to target plan provision and the impact of direct regulation on expenses. We place emphasis on evaluating policies that are under current debate or that are provisions of the 2019 SECURE Act.

Why would employers design a plan without high-quality investment options? There are two main incentives for the provision of high-quality investment options: (i) as with any other non-wage benefit, it might lead to wage savings, advantages in hiring and retention, or increases in worker morale; and (ii) offering a low-quality plan exposes the firm to litigation risk. Under the Employee Retirement Income Security Act of 1974 (ERISA), employers are typically considered fiduciaries to their workers and are expected to choose a diverse

menu and monitor investments prudently. In practice, however, there are several reasons that may lead a firm to choose lower-quality plans. The firm may simply lack the sophistication or the bandwidth to properly evaluate investment options. The quality of the plan may not be salient (at hiring or even later), limiting the incentives to provide higher quality. Even though ERISA-related lawsuits about fiduciary duty are growing, litigation risk remains low—especially for small firms, whom plaintiffs may not find worthwhile to take to court. Moreover, even if the firm internalized workers' objectives entirely, workers themselves may not value high-quality investment options, perhaps again due to lack of sophistication. Under these rationalizations, employers' willingness to pay for quality may be misaligned —either with workers' (an agency problem) or with the preferences a regulator would have.

Additionally, compensation structures in this setting can lead to agency frictions that further misalign employers' willingness to pay. Recordkeepers primarily earn revenue from two sources: direct payments from the employer and indirect payments from the investments on the menu. Indirect payments are typically a share of the expenses paid by workers (called "revenue sharing") and are thus larger for higher-fee investment options (Pool et al., 2021). Employers that do not fully internalize their workers' utility, or who face internal budget constraints for non-wage benefits, have an incentive to shift plan design towards higher-fee investments in order to save on direct payments.

Finally, the structure of the market for plan provision may also have imperfections. Recordkeepers may have market power, stemming from imperfect search by employers for potential plan providers, differentiation, or simply bargaining power in negotiations with employers. Since recordkeepers also participate in discussions on plan design, market power may skew the design of plans away from ones that are beneficial for workers—or even preferred by the employer.

In this paper, we develop a model of the market for DC plan provision that incorporates both these classes of frictions. We allow for a variety of sources of market power for recordkeepers: employers search for recordkeepers that may administer their plan, and they bargain with them over both the menu of funds that is made available to workers and the transfer that will be paid to the recordkeeper. We also account for the two sources of misalignment of willingness-to-pay. First, employers perceive a benefit from offering a plan that is a function of the plan's characteristics, including the number of investment options, expenses paid by workers, and whether the plan has certain investment vehicles. Second, employers may place differential disutility on spending money in direct transfers to recordkeepers, in employer contributions to workers, and in the money their workers spend

compensating recordkeepers through fees—which affects their willingness to substitute direct transfers paid by sponsors for expenses paid indirectly by workers.

We use data from Form 5500, collected by the Department of Labor, on all DC plans in 2016 with at least 100 participants. We observe the investment menu and the total allocations in each asset by both the employer and the employees. We also observe all sources of revenue for the recordkeeper, including the direct transfers and indirect revenues in the form of revenue sharing. Observing plan choice together with the payments between parties is useful in identifying the magnitude of the frictions outlined above: broadly, the characteristics of the chosen plan are informative of the preferences of the employer, and patterns in the transfer help us identify the extent of recordkeeper market power.

We find that both misalignment of willingness-to-pay and market imperfections are at play. In particular, small employers (fewer than 200 employees) value having index funds in their plan, even if that comes at the cost of adding other high expense investment options, while large employers (over 1,000 employees) are more directly focused on the expenses their employees pay. These gaps can come from several sources: different demands from workers, different competitive pressure from the labor market, or different exposure to fiduciary duty litigation would generate this pattern. Regardless of the source, we find that assigning small employers the preferences of large employers reduces expenses paid per worker by roughly $10 a year, more than half the difference in expenses between the two groups. Second, we also find that small employers are not especially averse to paying directly for plans: at a typical value of transfers, small employers treat a $1 increase in transfers as costing $1.09. Accordingly, alleviating this friction does not change plan characteristics appreciably.

We also find that there is scope for the exertion of market power. We estimate a bargaining parameter of 0.85 for large employers and parameters between 0.31 and 0.65 for smaller employers. We do not find much evidence for incomplete search being a source of market power for recordkeepers. On net, the incidence of market power lies primarily on transfer payments by employers and not on distorted plan characteristics.

We next analyze how policy can improve market outcomes. We first consider two provisions discussed in the SECURE Act. The first involves subsidies to sponsors for offering retirement plans, and the second allows small businesses to band together to establish multiemployer plans (MEPs). In principle, these proposals could improve plan characteristics. Subsidies could increase competition for an employer, leading them to be able to negotiate lower transfers or reduce plan expenses; improved bargaining power could

have similar effects. In line with the result from before, we find that while some of these policies could affect transfers, they would have limited effect on plan characteristics.

These results suggest that if the goal is to reduce expenses significantly, it is not effective to incentivize competition or subsidize provision. We thus consider a variety of policies that target plan characteristics. We first consider "must-carry provisions," which mandate that plans must include certain types of investments. Overall, we find this to be a blunt tool to influence plan expenses: mandating S&P 500 trackers has limited effect on plan expenses, and mandating target date funds can reduce plan expenses by up to $2/person. Mandating expense ratio caps on plans can have significantly larger effects on plan expenses: a cap of 40 bp on (unweighted) expenses in a plan reduces expenses by $9/person for small firms. Less stark approaches—such as penalizing firms for offering high-expense plans, or subsidizing them for offering low expense ones—can have similar impacts as expense caps.

Our results thus indicate that the policies that increase competition or lower market power of recordkeepers have an effect on direct transfers but do not reduce plan expenses appreciably. Instead, affecting expenses requires directly targeted policies. However, there is no free lunch. First, for all the policies we have discussed above, we also compute the extensive margin response on plan provision. Direct constraints on plan design improves expenses at the cost of a smaller proportion of firms offering plans at all. While we estimate this extensive margin to be small, and the policies under consideration typically do not affect participation in the defined contribution market by more than 1–2 pp, this trade-off may be relevant to policymakers considering regulation in this market. Second, since these policies often increase transfers paid by employers to recordkeepers, the degree to which these transfers are passed through to workers through lower compensation can lead to equilibria where workers are worse off even if plan expenses themselves decrease.

Related Literature. A large literature in household finance has focused on participant behavior in defined contribution plans, documenting a variety of behavioral frictions in participation, contributions, and asset allocations. Benartzi and Thaler (2007) and Choi (2015) provide survey of the behavioral literature, and we discuss relevant papers in Section 2.2; Egan et al. (2024) study these patterns through the lens of a demand model. This literature does not study the decision of which funds are offered in a plan in the first place: this is determined by the supply side of this industry, on which there is considerably less work. A set of recent papers documents potential conflicts of interest: plan providers give preference to their own funds (Pool et al., 2016), they trade off direct fees from the sponsor

with indirect fees from investments (Doellman and Sardarli, 2016; Badoer et al., 2020), and these indirect fees do affect whether funds are included in the plan (Pool et al., 2021). Our model of the market for plan provision takes this literature into account, respecting the possibility of behavioral frictions that may govern consumer decisions. Our contribution is to move beyond documenting that workers often invest suboptimally to studying why employers have incentives to provide them with the investment menus that we observe. We also shed light on the equilibrium effects of regulation, providing a comparison between approaches that have been debated in recent years.

More broadly, this paper fits into a literature on the industrial organization of financial markets, particularly using structural methods to study market structure and competition. Settings include car loans (Einav et al., 2012; Grunewald et al., 2019), credit cards (Nelson, 2020), insurance (Koijen and Yogo, 2016), mortgages (Allen et al., 2014, 2019; Robles-Garcia, 2020), municipal bonds (Brancaccio et al., 2020), pensions (Luco, 2019; Illanes, 2017; Illanes and Padi, 2021), advice (Bhattacharya et al., 2025; Egan, 2019; Guiso et al., 2022), and consumer and student loans (Bachas, 2019; Cuesta and Sepúlveda, 2019). Clark et al. (2021) provides an overview. Two recent papers that study the industrial organization of the market for defined contribution plans have followed our work. Loseto (2023) studies whether workers would design better menus than those that are set by employers, while Yang (2025) studies the role of transaction costs in updating plan menus.

This paper also contributes to the literature taking models of Nash bargaining, either between two parties or bilaterally between many interconnected pairs, to empirical work. Researchers have studied automobiles (Larsen, 2021), cable (Crawford and Yurukoglu, 2012; Crawford et al., 2018), medical devices (Grennan, 2013), financial markets (Robles-Garcia, 2020; Brancaccio and Kang, 2025), and hospital-insurer relationships (Gowrisankaran et al., 2015; Ho and Lee, 2017, 2019; Cuesta et al., 2019). Unlike most of these papers (Brancaccio and Kang (2025) is an exception), we study a setting in which quality is manipulable and a direct transfer between participants is a significant component of the payoffs. Understanding the interaction between these dimensions is a contribution of this paper.

2. Industry Background and Data

2.1. Setting and Data Sources

A defined contribution plan consists of a menu of investment options selected by the employer. Workers choose how much to contribute, and employers often match a portion of

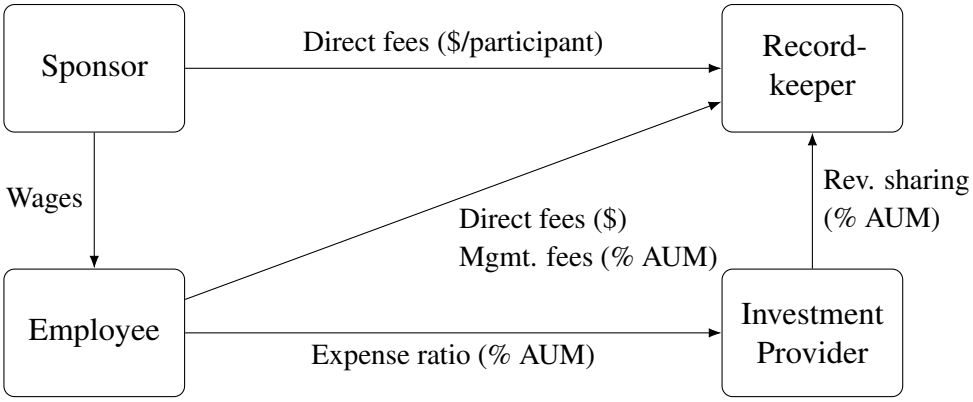

Figure 1: Illustration of DC plans, adapted from Collins et al. (2016).

workers' contributions; the level of matching varies across employers. Typically, workers keep these funds in their individual accounts until they are spent down in retirement. Importantly, employers are not required to provide DC plans or any other retirement benefit.

Figure 1 illustrates the structure behind the administration of defined contribution plans. Employers (often called a plan "sponsor") decide on a recordkeeper (called a "trustee" in other sources) who administers both employer's and employees' contributions into the plan.[1] For this service, recordkeepers charge sponsors direct fees, usually on a per-participant basis. In some cases, employees also pay direct and management fees to recordkeepers. Moreover, the recordkeeper and employer also decide on the set of investment options available through the plans, including index funds, actively managed funds, common stock of the employer, etc. These investments are offered by "investment providers." It is common for recordkeepers to be vertically integrated into investment provision and for them to offer their own funds. Investment providers pay recordkeepers fees for offering their funds, usually as a portion of assets under management, a practice called "revenue sharing." Finally, employees pay investment providers fees in the form of expense ratios.

Beyond these payments, both employers and recordkeepers face administrative costs when setting up a plan. These costs may include a fixed cost component, which would make it particularly expensive for small companies to set up a plan. Employers are also fiduciaries to their employees when designing a DC plan, and face the threat of legal action if plan quality is too low. For example, lawsuits often cite excessive expenses.

Currently, regulators have only instituted guidelines for certain features of the plan.

[1] In some cases, sponsors select multiple recordkeepers for employees to choose from. As this is rare, we describe the market ignoring this possibility.

For instance, plans must be designed to be "nondiscriminatory," i.e. benefit rank-and-file employees instead of just managers. Specific matching rules provide safe harbor provisions that exempt a plan from nondiscrimination testing, but no regulator mandates that plans follow such rules. Employers that automatically enroll their workers in the plan must choose a default investment, and they are liable for choosing a high-quality one. The Department of Labor issued a regulation in 2006 that designated certain types of funds—such as index or target date retirement funds—as "qualified" default investment alternatives (QDIAs), freeing employers from liability if they choose such funds but without mandating they do so.[2] Moreover, there are no specific mandates for the dimensions of plan quality that are often brought up in lawsuits—such as the types of funds included in the plan, the number of options on the plan, or the expense ratios charged by funds.

We leverage data from Form 5500, an annual filing with the Department of Labor that is required of all employers who offer benefits, to study this market. For all firms offering employee benefits, we observe the number of participants in each benefit plan, as well as their name, location, and NAICS code. Focusing on defined contribution plans, we can identify all the entities involved in the plan and almost all the transfers outlined above.[3] In particular, we observe revenue sharing outlays at the plan-asset level, contributions made by employees and employers, and matching and vesting rules. Finally, we observe the menu of investment options for all plans with 100 or more participants, including balances for each investment option. While all this information is made public by the DOL, some of it is not digitized, and we purchased a digitized version from Brightscope. We complement this information with data from CRSP and Morningstar regarding asset fees, star ratings, and investment styles. For the analysis that follows, we focus on plans sized between 100 and 5,000 participants, dropping sponsors that offer defined benefit (i.e., traditional pension) plans and a small handful that exhibit outsized fees. See Appendix A for further details.

2.2. Descriptive Facts

Figure 2 reports market shares for recordkeeping and investment provision. The market as a whole is not especially concentrated, but not all recordkeepers necessarily compete for the business of each sponsor. Comparing Panels (a) and (b) also highlights that vertical

[2]See https://www.irs.gov/retirement-plans/operating-a-401k-plan and https://www.betterment.com/work/resources/what-is-a-qualified-default-investment-alternative-qdia for more information.

[3]We do not have data on fees paid by the employer to the recordkeeper directly. However, our understanding is that most plans tend not to have such fees, while direct fees from the sponsor to the recordkeeper and revenue sharing are nearly universal.

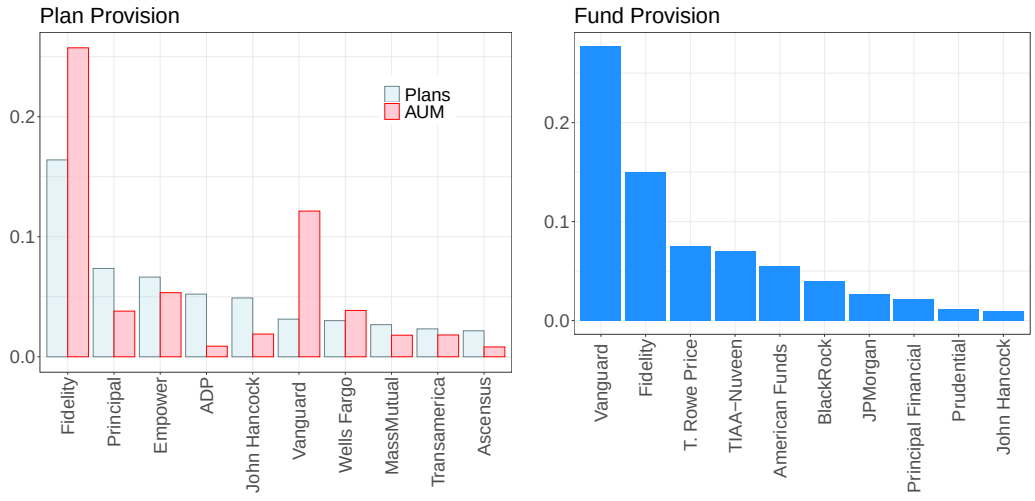

(a) Recordkeeping (b) Investment Provision

Figure 2: Market Shares

integration with investment providers is important, as Vanguard and Fidelity are market leaders in both plan provision (by AUM) and fund provision.

Table 1 reports summary statistics for our sample. Panel A illustrates the heterogeneity across sponsors and savings behavior in our sample. The average firm that offers a DC plan has 901 employees, with 165 and 556 as the first and third quartiles of the distribution of this variable. Mean per-person balances are $36,300. Mean employee and sponsor contributions are $2,560 and $1,340, respectively. Some sponsors, however, have considerably larger balances and contributions. The empirical model allows for this heterogeneity to lead to difference in potential revenues (through revenue sharing agreements) and costs (through economies of scale).

We next document a number of descriptive facts to motivate elements of the model. First, the extensive margin is sizable: among firms that do not offer defined benefit plans, about 23% also do not offer any DC plan. The extensive margin into plan provision is important, as such workers are only eligible for a much smaller tax-advantaged contribution through an IRA. Table D.1 in Appendix D.1 compares firms who offer and do not offer plans across other dimensions. Offering firms pay higher wages and also vary in industry composition.

Second, plans have heterogeneous quality, and a fair number of them are arguably of low quality as measured by various metrics. Panel B shows that 42% lack an S&P index fund, 26% lack a target date retirement fund, and 16% lack any sort of index fund—investments

Table 1: Summary Statistics

	N	Mean	SD	p5	p25	p50	p75	p95
A. Statistics for all firms								
Has a DC Plan	56,237	0.77	0.42	0	1	1	1	1
Total Participants	43,553	901	3,153	115	165	266	566	3,079
Balance / Person ($1K)	43,553	36.3	41.5	1.24	9.60	23.7	48.3	113
Employee Contribution / Person ($1K)	43,553	2.56	1.99	0.29	1.11	2.10	3.51	6.41
Employer Contribution / Person ($1K)	43,553	1.34	1.92	0	0.23	0.77	1.68	4.75
Weighted Expense Ratio (bp)	43,553	66.9	31.6	17.3	44.4	64.9	89.2	117
Expenses / Person ($)	43,553	211	232	7.27	53.8	141	290	641
Direct Transfer / Person ($)	43,553	60.2	256	0	2.31	15.4	62.4	226
Revenue Sharing / Person ($)	43,553	67.8	110	0.20	4.98	25.3	83.3	280
Direct Transfer / Total Revenue	43,553	0.42	0.39	0	0.03	0.29	0.85	0.99
B. Statistics for firms that report plan menus								
Has SP Tracker	43,553	0.58	0.49	0	0	1	1	1
Has Target Date Fund	43,553	0.74	0.44	0	0	1	1	1
Has Index Fund	43,553	0.84	0.37	0	1	1	1	1
Share in Index Funds	43,553	0.23	0.27	0	0.03	0.13	0.33	0.86
Average Expense Ratio (bp)	43,553	71.1	30.0	22.8	50.2	69.5	93.2	118
Weighted / Average Expenses	43,546	0.93	0.17	0.62	0.85	0.95	1.01	1.16

that are normally thought of as high-quality. Asset-weighted expense ratios are not low: they average 71 bp, but over 25% of firms have one that exceeds 93 bp. As mentioned earlier, the average across all mutual funds was 57 bp in 2016. In the model, high expenses and the lack of such funds can come from sponsors undervaluing their benefits to workers or recordkeepers exerting market power to avoid including lower-fee funds on the menu.

Third, workers tend to invest across funds in the menu instead of simply allocating the majority of their assets to index funds or other lower-fee options. Only about 23% of assets on average are invested in index funds. We also tabulate the ratio between the asset-weighted average expense ratio and the equal-weighted average expense ratio of funds in the menu; this ratio is only slightly smaller than 1 (0.93), again suggesting that workers are not directing funds towards lower-fee options. Figure D.1 in Appendix D.1 provides more details on these quantities.

These observations are in line with considerable evidence in this setting that decisions regarding contributions and asset allocations are made sub-optimally and subject to heuristics: examples include naive diversification by splitting assets roughly evenly across all options (Benartzi and Thaler, 2001; Huberman and Jiang, 2006; Iyengar and Kamenica, 2010), underdiversification by investing in only one fund or stock (Choi et al., 2005), ignoring assets in one account when deciding on the other (Choi et al., 2009), or limited sensitivity to fees when they are not salient (Kronlund et al., 2021). Moreover, contributions are also

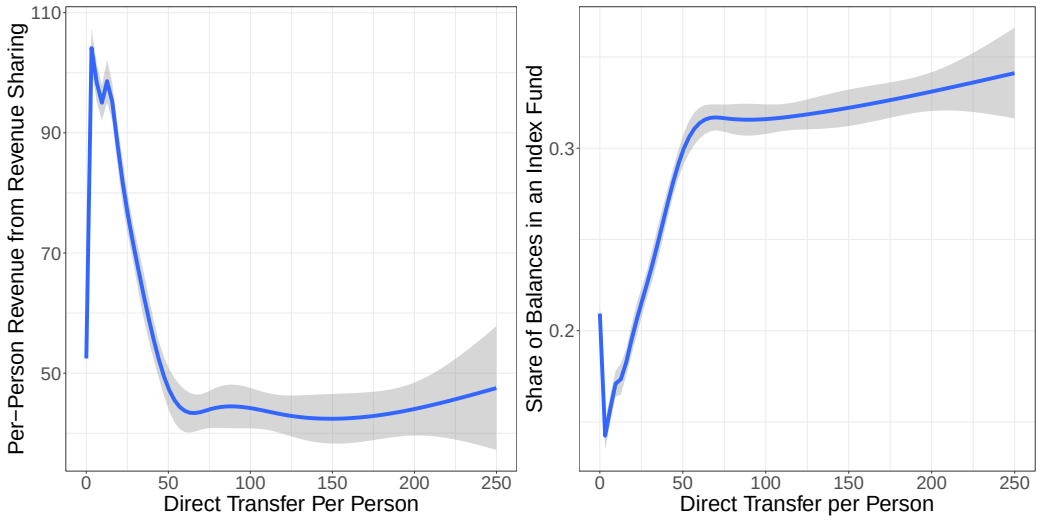

(a) Revenue Sharing (b) Share in Index Fund

Figure 3: Relationship between direct transfers and plan quality metrics

subject to behavioral frictions, which would affect the map between the plan menu and total expenses or revenue generated.[4] From a modeling perspective, we avoid conceptualizing asset allocations as the outcome of an optimal portfolio allocation problem, instead relying on flexible predictions of the outcomes necessary for estimation and counterfactuals.

Fourth, the two main sources of revenue mentioned above—revenue sharing agreements and direct transfers—are both important to the recordkeeper. The per participant means for both these objects ($68 and $60, respectively) are of similar magnitudes, as are the quantiles. In fact, the share of total revenues to the recordkeeper due to direct transfers is on average 42%, but there is heterogeneity on this dimension as well, with significant mass near both 0 and 1. Moreover, Panel (a) in Figure 3 shows a trade-off between the two sources of revenue, with the exception of the case where transfers are especially small. The implication of this observation for modeling purposes is that, unlike other settings with bargaining, we cannot abstract away from this direct transfer; how we model the trade-off between these two revenue sources has implications for the nature of competition in this market. Additionally, the fact that revenue sharing is low when direct transfers are zero is an

[4]Participants are often reluctant to invest money even when deposits so would lead to employer contributions and could be immediately withdrawn (Choi et al., 2011). Conditional on contributing, participants often set contribution rates at "round" portions of income (5% or 10%), as documented by Benartzi and Thaler (2007). Increasing choice is not always always associated with greater takeup (Papke, 2004; Sethi-Iyengar et al., 2004; Huberman et al., 2007; Beshears et al., 2013).

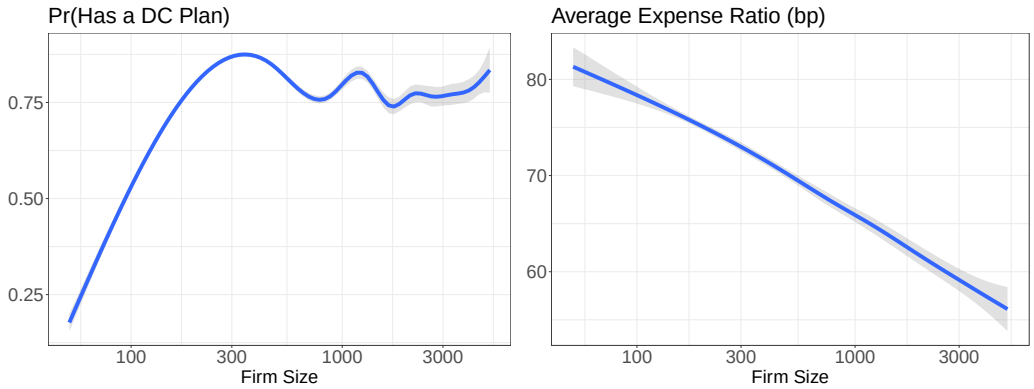

(a) Has a DC plan

(b) Mean expenses (bp)

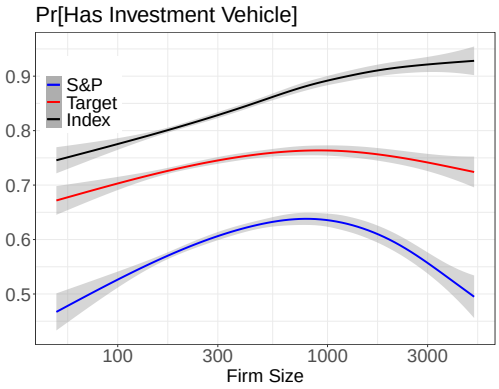

(c) Has investment options

Figure 4: Relationship of plan provision and quality metrics with firm size

implication of our model; we return to this issue in Section 3. Panel (b) corroborates that plans with higher direct transfers have higher quality investment outcomes, consistent with recordkeepers being compensated either through direct transfers or through higher expenses.

Finally, plan size is an important driver of the heterogeneity documented above: both plan availability and quality are correlated with it.[5] Figure 4(a) shows that plan availability increases about 30 pp when moving from sponsors with 100 employees to those with 300. Figure 4(b) shows that average asset-weighted expenses move from 80 bp to 55 bp when moving from firms with 100 to 2,000 employees. The availability of various investments also generally increases with firm size, as show in Figure 4(c). In the empirical implementation of the model, we allow the magnitude of the frictions to depend on firm size to disentangle

[5]For firms that offer a plan, we use the maximum between the number of plan participants and number of participants in the health insurance plan as the measure of firm size. For firms that do not, we use the number of participants in the health insurance plan.

the sources of the gradients documented in Figure 4.

3. Model

In this section, we develop a model of the interactions between plan sponsors and recordkeepers that lead to observed plan choice.

3.1. Payoffs to the Parties

A set $\mathcal{R}_s$ of recordkeepers is capable of providing a DC plan to sponsor s. The sponsor s chooses a subset $\tilde{\mathcal{R}}_s \subseteq \mathcal{R}_s$ with whom to negotiate the design of the plan. The sponsor s and a recordkeeper r can decide on a plan $p \in \mathcal{P}_{rs}$ and a dollar transfer T from s to r. If the parties decide on a plan and transfer (p, T), the per-participant payoff to the sponsor s is

$$\pi^S_{rs}(p, T) \equiv B_{rs}(p) - C^S - T - \kappa(T). \tag{1}$$

In (1), $B_{rs}(p)$ is the per-participant net benefit to the sponsor from offering a plan p. We conceptualize this benefit as coming from four main sources. First, there is a compensating differential between wage and non-wage benefits, so that offering a higher quality plan allows the employer to also offer lower wages. Second, there may be additional benefits to employers to offering a better plan beyond pure wage savings. Third, plans with certain characteristics—lower expenses, more index funds, etc.—reduce the probability that the sponsor is sued for breach of fiduciary duty, and sponsors may be differ in the likelihood of being sued and in legal costs. Costs of fiduciary insurance would fall into this category. Finally, sponsors often offer matching contributions, and these are deducted from benefits. Next, C^S is the per-person administrative cost faced by the sponsor.

The final two terms of (1) relate to the direct transfer T from the sponsor to the recordkeeper. First, this transfer is a cost that directly impacts utility. Second, the function $\kappa(T)$ denotes an additional cost of paying for a defined contribution plan. We assume $\kappa(\cdot) \geq 0$, $\kappa'(\cdot) \geq 0$, and $\kappa''(\cdot) \geq 0$. A convex $\kappa(\cdot)$ corresponds to the case where the firm's marginal willingness to pay decreases with the transfer: the extreme case of $\kappa(\cdot)$ being especially large after some $\bar{T}$ would correspond to a firm unwilling to pay more than $\bar{T}$ for a DC plan. This cost can capture liquidity constraints, internal budgeting, managerial incentives, or any other frictions that would prevent a firm from wanting to pay especially large transfers or guardrails that would prevent the managers in charge from doing so.

The per-participant payoff to the recordkeeper r is

$$\pi^R_{rs}(p,T) \equiv \Pi_{sr}(p) - C^R_r + T. \tag{2}$$

The term $\Pi_{sr}(p)$ in (2) constitutes the per-participant revenues the recordkeeper earns. The primary component of this comes from revenue sharing agreements: the sum of the expected dollar amount invested by each of the sponsor's workers in fund f times q_{fr}, is the revenue sharing that r earns from fund f. If f is part of r's fund family, then q_{fr} is the expense ratio and otherwise it is the revenue sharing rate. Next, C^R_r is the per-participant administrative cost of the recordkeeper; in estimation, we let it depend on the number of participants to capture economies of scale. Finally, in the empirical implementation we allow for an idiosyncratic shock at the sponsor-recordkeeper level.

3.2. Plan Choice

As discussed in Section 2, sponsors decide on plans jointly with recordkeepers, and it is a priori unclear which party sets the terms of the plan. We thus assume that plans and transfers are set through Nash bargaining between the two parties. In particular, we posit that if a sponsor and recordkeeper are negotiating over a plan design, they solve

$$\begin{aligned} &\arg\max_{p\in\mathcal{P}_{rs},T} \left(\pi^S_{rs}(p,T) - \pi^{Sd}_{rs}\right)^{\eta} \cdot \left(\pi^R_{rs}(p,T)\right)^{1-\eta} \\ \text{subject to}\quad &\pi^R_{rs}(p,T) \geq 0, \quad \pi^S_{rs}(p,T) \geq \pi^{Sd}_{rs} \\ &T \geq 0. \end{aligned} \tag{3}$$

We discuss the terms in (3) one-by-one. First, η is the bargaining power of the sponsor. The first set of constraints are the IR constraints. The second constraint ($T \geq 0$) reflects the restriction that the recordkeeper cannot pay the sponsor for the right to provide the plan.[6]

Writing the Nash program requires us to take a stance on the "disagreement payoffs" of the two parties. Serving sponsor s does not prevent recordkeeper r from serving any other sponsor, and if the bargain between recordkeeper r and s were to fall apart, r would simply earn zero profits. Thus, the disagreement payoff for r is zero, and the recordkeeper's utility enters the Nash program as $\left(\pi^R_{rs}(p,T)\right)^{1-\eta}$. The disagreement payoff to the sponsor is

[6] A natural concern is that the recordkeeper may offer "free" services in lieu of such payments that make the plan seem more attractive. In the empirical implementation, we let the mean utility for the recordkeeper depend on firm size, which would capture systematic differences in quality provision by size. However, our model does not allow us to introduce costly unobserved quality.

non-zero: the sponsor can contract with a different recordkeeper if this bargain were to fall apart. We posit that the disagreement payoff π_{rs}^{Sd} to the sponsor corresponds to extracting the maximum surplus possible from the second-best recordkeeper, subject to meeting its individual rationality constraint. That is, we let

$$\bar{\pi}_{rs} \equiv \max_{p \in \mathcal{P}_{rs}, T} \pi_{rs}^{S}(p, T) \text{ s.t. } T \geq 0 \text{ and } \pi_{rs}^{R}(p, T) \geq 0, \tag{4}$$

and $\pi_{rs}^{Sd} \equiv \max\{0, \max_{r' \in \tilde{\mathcal{R}}_s, r' \neq r} \bar{\pi}_{r's}^{S}\}$. Thus, if the bargain were to break down between r and s, s could go to its next best option r', who would then be willing to serve s at zero profits, or to the outside option of not having a plan.

One can conceptualize the process that leads to this bargaining outcome as consisting of two steps. First, the sponsor runs a back-and-forth negotiation between recordkeepers in its choice set. During this negotiation, the sponsor solicits offers from recordkeepers and asks for counteroffers from others. Since a recordkeeper would never be willing to let a competitor win when it could instead offer a profitable contract that the sponsor prefers, all but one competitor would be willing to offer a contract (p, T) up to the point where they earn zero profits. The best utility that the sponsor earns from these contracts, if r is the remaining recordkeeper, is thus π_{rs}^{Sd}. When only one firm r is remaining, the sponsor engages in a Nash bargain with this firm, with its threat point being going back to the second-best firm and accepting the previously offered contract that provided a utility π_{rs}^{Sd}. The extensive form of this final bargain follows the traditional framework from Rubinstein (1982), with a small exogenous probability of breakdown in negotiations. This approach is what the literature has referred to as "split-the-difference" bargaining, following Binmore et al. (1989): a stronger second-best option for the recordkeeper always improves its bargaining outcome. This extensive form mirrors the "second-price auction followed by bargaining" framework of Larsen (2021). It has also been used by Bhattacharya (2021) in a setting of a single buyer choosing between sellers.

An alternative approach is used in the Nash-in-Nash with Threat of Replacement (NNTR) framework (Ho and Lee, 2019). The similarity between these approaches is that in both models, the second-best option for the sponsor is held to its zero-profit constraint. The main difference is that NNTR uses the so-called "deal-me-out" framework from Binmore et al. (1989); here, bargaining breakdown is strategic, so weak second-best options do not improve the sponsor's bargaining position since it cannot credibly threaten to accept them.[7] Which

[7] A second difference between NNTR and our setting is that NNTR is designed to model bargaining over an

approach is appropriate is in principle an empirical question: Binmore et al. (1989) provide experimental evidence in favor of the deal-me-out approach, although results from other experimental work (Ochs and Roth, 1989; Kahn and Murnighan, 1993) are more in line with "split-the-difference." In practice, previous literature assumes the format of bargaining rather than letting the data adjudicate between them, and we follow this approach.

Fixing a recordkepeer, (3) and (4) specify how the plan is chosen. The sponsor first selects the consideration set $\tilde{\mathcal{R}}_s$: how this is chosen is not relevant for the remainder of the procedure, and we discuss our empirical implementation in Section 4.1. Given this choice set, the sponsor selects the recordkeeper that maximizes its own utility $\pi^S_{rs}(p^*, T^*)$ at the optimum (p^*, T^*), and the plan and transfer are given by this bargaining procedure. Inspecting (3) and (4) also shows that they implicitly specify how the recordkeeper itself is chosen by the sponsor: the chosen recordkeeper is the one in $\tilde{\mathcal{R}}_s$ that provides the highest $\bar{\pi}_{rs}$. Otherwise, there would be no (p, T) that would satisfy the IR constraint of the sponsor: a sponsor would never contract with a recordkeeper from which it could never possibly extract as much surplus as it could from another recordkeeper.

3.3. Properties of the Model

Broadly, the model captures market imperfection in two ways. First, the model captures search frictions naturally: sponsors only consider a limited subset $\tilde{\mathcal{R}}_s$ when deciding on the plan menu, and this could be a strict subset of all potential recordkeepers $\mathcal{R}_s$. Second, bargaining power η, together with the surplus generated by the second-best option, captures recordkeeper market power. The model also captures misalignment in willingness to pay through two channels. First, since we do not take a stance on the determinants of the benefit function $B_{rs}(p)$, the sponsor may not place the same emphasis on plan characteristics—e.g., having low-fee funds, or incorporating certain types of investment options—that informed workers would if they were selecting plans themselves. Second, $\kappa(\cdot)$ is an explicit agency friction: a sponsor has an incentive to reduce expenses (like the direct transfer) that are incident on itself, even if doing so comes at the cost of increasing expenses incident on workers.

These frictions interact with each other in the model: the bargaining procedure allows for market power to affect plan characteristics, and how much it does so depends on the

entire network: there are multiple buyers and multiple sellers. In our setting, there is only one buyer and multiple sellers, although the constraint that the buyer only contracts with one seller introduces a role for a "threat of replacement."

misalignment in willingness to pay. To understand the economic forces embedded in this model, consider the case where we relax the $T \geq 0$ constraint and set $\kappa(\cdot) = 0$. Then, the model reduces to Nash bargaining with transferable utility, and the solution to (3) is to set the plan p to maximize joint sponsor-recordkeeper surplus $\pi_{rs}(p) \equiv \pi^S_{rs}(p, T) + \pi^R_{rs}(p, T)$, which does not depend on T by design. The transfer T is then set so that the sponsor captures a share η of the incremental surplus generated by the bargain:

$$T = \tilde{\pi}^S_{rs}(p^*) - \eta \left(\pi_{rs}(p^*) - \pi^{Sd}_{rs} \right), \tag{5}$$

where $\tilde{\pi}^S_{rs}(\cdot)$ denotes the utility of the sponsor ignoring the transfer. This is an application of the so-called bilateral contracting principle (Lee et al., 2021). Here, the choice of the plan (and recordkeeper) incorporates agency frictions—captured in the benefit function. However, reducing the market power of the recordkeeper does not affect plan choice: increasing η or increasing π^{Sd}_{rs} (which captures an increase in the level of competition in the market) only reduces the direct transfer.

The constraint $T \geq 0$ and the cost function $\kappa(\cdot)$ allow bargaining and competition to directly affect plan characteristics. Consider an extreme case where $\kappa(T) = 0$ for $T \leq \bar{T}$ and prohibitively large otherwise, perhaps if the HR department of a firm has a budget to spend on retirement plans. If (5) is such that T would be negative, then the solution sets $T = 0$ and the plan to maximize $\left(\pi^S_{rs}(p, 0) - \pi^{Sd}_{rs}\right)^{\eta} \cdot \left(\pi^R_{rs}(p, 0)\right)^{1-\eta}$: since the recordkeeper cannot transfer surplus to the sponsor by lowering the transfer further—which would be efficient from the perspective of joint surplus—it does so by changing the plan characteristics to ones the sponsor prefers. Conversely, if the sponsor's bargaining power is so low (or competition is so weak that π^{Sd}_{rs} is low) that T would exceed $\bar{T}$, the sponsor would instead limit the transfer to $\bar{T}$ but be willing to accept plans the recordkeeper preferes more. Smoother functional forms for $\kappa(\cdot)$ embed analogous forces, and thus $\kappa(\cdot)$ is an important parameter that governs how competition and market power translates to plan characteristics.

The elements can rationalize the descriptive facts outlined in Section 2.2. The model incorporates an extensive margin: both the sponsor and at least one recordkeeper in the consideration set have to meet their individual rationality constraints for there to be a plan. Direct inspection of (3) shows that recordkeepers view direct fees and revenue sharing as substitutes, setting up a trade-off between them. High expenses or the lack of index funds could be explained by a low preference for them on the part of the sponsor ($B_{rs}(\cdot)$), that such funds bring especially high revenues to the recordkeeper and thus increase the Nash product

in (3), or that $\kappa(\cdot)$ is large and the sponsor is averse to high transfers (and thus providing the recordkeeper revenues through revenue sharing increases the Nash product). That surplus must be provided to the sponsor through plan characteristics when $T = 0$ explains the low average revenue sharing documented in Figure 3 for plans with transfers at this bound.

Finally, we model plan design as a static game. Since in practice employers do not change their plan provider on a yearly basis, a natural concern is the possibility of inertia. Not accounting for inertia would cause us to miss potential dynamic pricing incentives due to lock-in or cause use to underestimate sensitivity to quality. We do not consider this concern to be first-order in our setting. First, survey evidence suggests that sponsors do actively evaluate their DC plans: a 2021 survey by Fidelity of more than 1,000 sponsors showed that 61% evaluate their providers and advisors annually, almost half of whom evaluate it more than once a year.[8] Moreover, these evaluations are not merely passive: 74% of sponsors made changes to plan menus in the past two years. Overall, these observations suggest that sponsors do not simply set a menu once and leave it be for many years; thus, we find the approximation of active choice to be sensible.

4. Identification and Estimation

4.1. Empirical Model

To flexibly estimate how model primitives change with firm size, bin together firms with 200 or fewer employees, between 201 and 500 employees, between 501 and 1,000 employees, and more than 1,000 employees (Groups 1–4, respectively), and estimate the model separately by bin. All parameters noted below are different for each group. We assume that there is no further dependence on firm size. We then parameterize

$$\begin{aligned} B_{sr}(p) &= \boldsymbol{\theta}_s \cdot \boldsymbol{Q}_s(p) + \delta_r + \epsilon^S_{rs}, \\ \Pi_{sr}(p) &= R_s(p) - C^R_r + \epsilon^R_{rs}. \end{aligned} \tag{6}$$

In the specification of the benefit function, we include a dependence on sponsor-plan properties $\boldsymbol{Q}_s(p)$ capturing plan characteristics: whether the plan has a large number of investment options (more than the median of 24); indicator variables for whether the plan has an index fund, an S&P 500 tracker, or a suite of target retirement date funds;

[8] See https://institutional.fidelity.com/app/item/RD_13569_26306/plan-sponsor-attitudes.html for the 2021 Plan Sponsor Attitudes Survey, conducted by Fidelity.

average expenses paid by employees; and per-person employer contributions. We allow the coefficient multiplying characteristics to be sponsor-specific and, for each characteristic q, parameterize $\theta_{qs} \sim N(\mu_q, \sigma_q^2)$. Per-person costs to recordkeepers C_r^R are constant within bins of firm size: if there are economies of scale, this value will decrease over firm size groups. We also assume unobserved (net) preferences for recordkeepers δ_r are fixed within bins of firm size. We can conceptualize $\delta_r = \xi_r - C^S$, where ξ_r captures differences in quality or services provided by recordkeepers which may make them differentially attractive to sponsors, as well as the mean utility difference between having a plan and not having one. It also encapsulates the sponsor cost (C^S in the notation from Section 3.1); in the empirical model, disentangling the preferences for recordkeepers from costs to the sponsor will require an instrument that shifts one or the other, and doing so is not necessary for our counterfactuals of interest. Finally, we include an unobserved shock $\epsilon_{rs}^S(p) \sim \text{T1EV}(\sigma^S)$ to the benefit function, which we interpret as a "match value": a sponsor may prefer a particular recordkeeper because of a pre-existing relationship between the companies, or because of the possibility of getting more firm-specific services from them. The recordkeeper's revenue $\Pi(\cdot)$ consists of revenue from revenue-sharing agreements $R_s(p)$, per-person costs C_r^R, and a shock $\epsilon_{rs}^R \sim \text{T1EV}(\sigma^R)$. This shock also captures any unobserved prior relationship or synergies that may affect the revenues (or the service cost) for a particular pair.

The remaining elements of the model are the bargaining parameter η, the cost of transfers $\kappa(\cdot)$, and the search process that leads to $\tilde{\mathcal{R}}_s$. We parameterize $\kappa(T) = \kappa \cdot T^2$: the quadratic function allows for an increasing marginal cost of transfers, and firms with high κ are more willing to accept lower benefit plans in exchange for lower transfers. The quadratic form also allows for closed-form computation of the transfer conditional on a candidate plan, which is useful in estimation and counterfactuals.

Finally, we specify the search process as follows. We assume that $\mathcal{R}_s$, the set of all potential recordkeepers sponsor s might select, consists of all the recordkeepers that operate in the sponsor's state: Table D.3 reports geographic coverage for these recordkeepers and shows some evidence of geographic specialization.[9] From this set, a sponsor s constructs the consideration set $\tilde{\mathcal{R}}_s$ by picking each of the top 19 (by number of plans nationwide) recordkeepers who operate in the state with probability ρ. The sponsor can also select, with probability ρ, one recordkeeper at random from the fringe of remaining firms that are present

[9]We define operating in a state as serving at least 5 sponsors in that state in our dataset in 2015. While we are unaware of heterogenous licensing requirements by state, this is a low enough number that any sponsor that does not meet this bar likely does not invest much in marketing to firms in that state.

in the state. Additionally, if the sponsor offers a 403(b) plan rather than a 401(k) plan, we eliminate any recordkeeper that does not offer any 403(b) plans in the country from $\mathcal{R}_s$. For the top 19 recordkeepers, this restriction only affects ADP.

This is a reduced-form approach to consideration set formation, similar in spirit to Goeree (2008). We adopt this approach in lieu of a model of optimal search partly for its parsimony but also because we find it appropriate for the setting: a model of optimal search requires a level of sophistication on the sponsor side that is at odds with them choosing plan characteristics in a potentially sub-optimal way. Since we think that allowing for behavioral frictions is of first-order importance, we are not comfortable imposing a high degree of sophistication in the search process. Moreover, we do not find evidence of limited search, so we do not expect results to be sensitive to alternative models of search.

We specify P_{rs}, the set of plans that a recordkeeper r could offer sponsor s, to be the set of all plans that r offers any sponsor in any state. That is, r can offer any combination of funds (and associated expense ratios and revenue sharing agreements) to s, although the characteristics of sponsor s will change the realization of quantities such as expenses paid per person or revenue sharing given this set of funds. The rationale behind this assumption is that there are no technical or legal restrictions to plan design. One may be concerned that small sponsors do not have access to the plans of large sponsors due to minimum investment restrictions. We explore this issue at length in Appendix E, and we find that investment minima for retirement share classes are often low and that they can also be waived by discretion. Therefore, we assume that it is theoretically possible for small sponsors to select a plan with the features of a large sponsor's: that many do not is an equilibrium feature of bargaining that our model rationalizes.

To construct the elements of the plans of P_{rs} that would enter (6), we thus first start with all plans that r offers any sponsor. Some of the characteristics of a plan p enter directly into (6): we know whether the plan would contain an index fund, whether it has a high number of options S&P tracker, etc. However, per-person employer contributions, expenses, and revenue sharing are a function of asset allocations and are thus not observed for these plans p for sponsor s (unless the plan is the one that s selected). We treat the computation of these quantities for counterfactual plans as a prediction problem. In particular, we use stochastic gradient boosting (Friedman, 2002) to predict these quantities from a large vector of plan and sponsor characteristics; Appendix B provides a description of the procedure. This procedure leads us to being able to construct the vector of plan-sponsor-specific covariates that would enter into (6) for each plan p, and the set of all such vectors forms $\mathcal{P}_{rs}$.

We take this approach since it allows us to be agnostic about how workers decide on asset allocations, thus respecting the literature (and facts shown in Section 2.2) about behavioral biases in contribution and allocation decisions. Of course, the main assumption needed for this approach to be valid is that plans are not selected on the basis of unobserved sponsor-level variables that affect asset allocations.[10] This approach limits counterfactuals to ones in which we expect the map from plans to asset allocations to be unchanged. This could be an issue for, for instance, "information" counterfactuals that educate workers or sponsors about fees, but we do not consider this to be a central concern in our policy counterfactuals.

4.2. Identification

For intuition about identification, it is informative to consider the benchmark with $\kappa = 0$ and no constraint that transfers must be positive, and with $\rho = 1$. As discussed in Section 3.3, the parties would maximize joint surplus in this model. Like in standard discrete choice settings, plan choice would identify the sponsor's preferences for characteristics as well as the mean utility for a particular recordkeeper net of total costs ($\delta_r = \xi_r - C^S$). Ignoring (ϵ^S, ϵ^R) shocks, we could then compute the joint surplus $u^{(1)}$ of the chosen plan and that of the second-best option $u^{(2)}$. Since the recordkeeper earns $(1 - \eta)$ of the surplus, the equation that $R^{(1)} + T - C^R = (1 - \eta) \cdot (u^{(1)} - u^{(2)})$, where $R^{(1)}$ is the revenue-sharing by the chosen plan, would identify C^R and $(1 - \eta)$ given variation in the incremental surplus generated.[11]

Using variation in the incremental surplus to identify a bargaining parameter mirrors the argument of Grennan (2013). In our setting, this variation could be induced by differences in choice sets across sponsors. Firms in different industries and geographies have different sets of plans from which to choose, which is embedded into how we construct choice sets. Moreover, due to variation in firm characteristics (including policies such as match rates),

[10] An example would be a sponsor that chooses to contract with Fidelity since it knows its workers have an affinity for Fidelity funds; we find this concern to be unlikely. In general, threats to this strategy seem to require some sophistication in asset allocation decisions by workers (a preference for certain types of funds or a propensity for certain investment decisions, which must be known to the sponsor), which we find to be broadly at odds with the behavioral literature on worker investments.

[11] We make two comments about this discussion. First, identifying the scale of u from the discrete choice requires a numeraire; in our case, this is revenue sharing (which is observed and enters joint surplus with a coefficient of 1). Second, the logic in this paragraph can be extended to the case with ϵ shocks as well: if the observed component of the incremental surplus generated by the bargain is sufficiently large, one can show that the expected value of the terms due to ϵ in the equation is zero. Thus, a regression of transfers plus revenue sharing on incremental surplus—for sufficiently large incremental surplus—would return $1 - \eta$ as the slope and C^R as the constant.

firms would even face different revenue sharing levels and contributions for the same plan.

Simple departures from this baseline model help connect this intuition back to the setup in this paper. Consider adding back incomplete consideration, so that ρ could be less than 1 and must be identified. We make three observations. First, a sponsor still considers all plans offered by a recordkeeper; thus, variation in plan characteristics within the chosen recordkeeper can still identify preferences for characteristics. Second, the extensive margin of offering plans is informative of ρ. If $\rho = 1$, then sponsors who have choice sets or firm characteristics such that the gain to offering some plan is very large are almost surely going to offer one. If such sponsors do not, then the model would attribute this to limited search: the model cannot explain this through the bargaining parameter, since as long as there is a mutually beneficial plan in the choice set, we would see a plan offered regardless of the bargaining parameter. Our third point about identification of ρ is more subtle: we have two sources of variation in incremental surplus—the joint surplus of the chosen plan and that of the second-best option. With $\rho = 1$, having two different sources of surplus is immaterial for identification: the slope of total revenues of the recordkeeper with respect to either is $1 - \eta$ in absolute value. However, when $\rho < 1$, the candidate for the second-best option among all potential options may not be in the actual (unobserved) consideration set. Thus, finding that total revenues to the recordkeeper are less sensitive to the utility of the second-best choice among all potential options than they are to the utility of the chosen option provides evidence of limited search.

A third observation in the baseline model is that plan choice within a recordkeeper is not related to the utility of the second-best option, as the joint surplus is not affected by it. However, if variation in the outside option induces variation in the characteristics of a chosen plan conditional on the choice of recordkeeper, it must be the case that $\kappa > 0$. Moreover, the fact that the cost of transfers is convex means that the sensitivity of plan choice to the outside option changes with the level of the transfer. Loosely, if transfers are near 0, variation in the outside option primarily translates to variation in transfers in a manner dictated by η, much like the baseline model. Additionally, once the constraint on transfers is imposed again, the probability that the transfer is 0 is informative of η (and the cost of the recordkeeper) as well. As transfers increase, variation in the outside option also induces variation in the characteristics of the chosen plan, in a manner governed by a combination of both η and κ. With knowledge of η and κ, the choice of plan within recordkeeper is still informative of the sponsor's sensitivity to plan quality; instead of maximizing joint surplus, however, the parties maximize the Nash product.

The intuition in the above paragraph is incorporated into estimation through the choice of moments to match. All moments we target take the form $\mathbb{E}[X \cdot \mathbb{1}(s \text{ chooses } r)]$: the expectation of the product of some variable X and an indicator for firm s choosing recordkeeper r. The first set of moments are simply properties of the chosen plans and transfers: (i) an indicator for the transfer equaling 0, (ii) the value of the transfer, (iii) each quality dimension $Q_s(p)$ of the chosen plan, and (iv) revenue sharing.

The second set of moments captures correlations between transfers and quantities that would affect surplus of the chosen plan. These moments are (i) interactions between each quality dimension and T, (ii) interactions between revenue sharing and T, and (iii) the same interactions, replacing T with an indicator that $T = 0$. These moments capture the intuition that the correlation between elements of $u^{(1)}$ and transfers is informative of η. The final set of moments mirrors this second set but replaces characteristics of the chosen plan with characteristics of non-chosen plans in the choice set. These moments mirror the classic "rivals' characteristics" instruments of Berry et al. (1995) as shifters of price; as discussed above, here they shift plan choice and transfers, and how they do so informs ρ and κ.

In Appendix D.2, we provide some descriptive evidence to illustrate the patterns that inform our estimates. A difficulty with looking for such patterns is that while the structural procedure has an internally coherent model of the second-best option for any candidate set of parameters, the identity of the second-best option in is unclear directly from the data. Thus, we use proxies for the second-best option: the recordkeeper that offers the minimum revenue sharing, minimum expenses, or minimum employer contributions. While none of these proxies are perfect, they together show that improvements in the outside option translate to reductions in transfers for the plan. To look for evidence in favor of search, we document that the average expenses of the plans in a sponsor's choice set shift the probability that that sponsor offers a plan. For sponsors who have especially attractive choice sets, the probability of offering a plan is high, which feeds into our estimates of large search probabilities.

4.3. Estimation

Estimation of this model can be thought of as joint estimation of a demand-side as in Berry et al. (2004) ("micro-BLP"), and a supply-side that allows for the optimal setting of transfers as the outcome of a bargaining game between employers and recordkeepers. Our procedure involves inverting a share equation for the choice of recordkeepers to recover mean utilities for recordkeepers, and then using these recovered parameters to compute the micro-moments described in Section 4.2. This section provides an outline; see Appendix C for more details.

First, we reduce the dimensionality of the choice set. As discussed above, we allow firms to potentially select all plans offered by the top 19 recordkeepers, plus all plans offered by a fringe provider. This creates choices over tens of thousands of plans for each sponsor. To estimate the model, we reduce this dimensionality by running a k-means clustering algorithm on the set of characteristics that enter the benefit function plus revenue sharing, and divide the set of plans offered by each recordkeeper to each sponsor into 100 groups. We then randomly sample one plan from each of these groups to build $\mathcal{P}_{rs}$, the set of plans sponsor s and recordkeeper r bargain over if the recordkeeper is considered.

We begin with a guess of random coefficient parameters for θ_s, of the variance of recordkeeper and sponsor errors σ^R and σ^S, of disutilities from transfers κ, of recordkeeper costs C_r^R, and of search cost probabilities ρ. Call this vector of parameters Ω for notational convenience. Fix these values and consider a particular consideration set. The plan with the highest extractable surplus for each recordkeeper in the consideration set satisfies

$$V_{rs}^*(\epsilon_{rs}^R;\theta,\sigma^R,\kappa,C_r^R) \equiv \max_{p\in\mathcal{P}_{rs}} \boldsymbol{\theta}_s \cdot \boldsymbol{Q}_s(p) - \max\left[0, C_r^R - R_r(p) - \sigma^R\epsilon_{rs}^R\right] - \kappa \cdot \max\left[0, C_r^R - R_r(p) - \sigma^R\epsilon_{rs}^R\right]^2. \quad (7)$$

Note that (7) reflects the impact of two constraints: the recordkeeper's individual rationality constraint and the restriction that transfers cannot be negative. Note that within-recordkeeper plan choice does not depend on sponsor shocks ϵ^S, on unobserved preferences for recordkeepers ξ_r, or on sponsor-level administrative costs, as these do not vary within recordkeeper. Because of the lack of a plan-sponsor-recordkeeper shock, our model for within-recordkeeper plan choice is akin to a pure characteristics demand system (Berry and Pakes, 2007). We make this modelling assumption to ensure that the expected utility obtained from a particular recordkeeper does not increase mechanically when the number of plans offered by the recordkeeper increases. This is particularly important in counterfactuals where we restrict the set of admissible plans. Finally, in (7) we impose the scale normalization that the utility of a dollar of transfer payments is equal to -1.

Given a consideration set, the probability sponsor s chooses recordkeeper r is given by

$$\Pr\left[V_{rs}^*(\epsilon_{rs}^R;\theta,\sigma^R,\kappa,C_r^R) + \delta_r + \epsilon_{rs}^S \geq V_{r's}^*(\epsilon_{r's}^R;\theta,\sigma^R,\kappa,C_{r'}^R) + \delta_{r'} + \epsilon_{r's}^S \text{ for all } r'\right] = \iiint \frac{\exp(\frac{1}{\sigma^S}(V_{rs}^*(\epsilon_r;\theta,\sigma^R,\kappa,C_r^R)+\delta_r))\cdot\mathbb{1}[r\in\tilde{\mathcal{R}}_s]}{1+\sum_{r'\in\tilde{\mathcal{R}}_s}\exp(\frac{1}{\sigma^S}(V_{r's}^*(\epsilon_{r'};\theta_s,\sigma^R,\kappa,C_{r'}^R)+\delta_{r'}'))} dF(\theta)dF(e)dF(\tilde{\mathcal{R}}_s). \quad (8)$$

In (8), we impose the location normalization that the mean utility of not having a plan is equal to 0. We compute the integral over the distribution of random coefficients $F(\theta)$ using monomial rules for Gaussian quadrature, the integral over the vector of recordkeeper shocks ϵ_r by simulation, and the integral over consideration sets using importance sampling.

Denote the left-hand side of (8) by $s_{rs}(\Omega, \delta)$. Averaging over firms yields the market-level shares $s_r(\Omega, \delta) = \frac{1}{S}\sum_s s_{rs}(\Omega, \delta)$. Since this equation satisfies the conditions outlined in Berry et al. (2013), given a guess of the parameters in Ω, we can invert the share equation and recover the vector of mean utilities δ. Having obtained mean utilities, we can compute the aforementioned moments. We estimate the parameters using GMM.

5. Parameter Estimates

Table 2 reports estimates from the model outlined in Section 4. The model is estimated separately for each group of firms, binned by the number of participants. Each column provides a separate set of estimates.

Panel A shows estimates relating to the market imperfections embedded in the model. We find that search probabilities ρ are large. Small firms have a lower search probability than the other groups do, but even these are sufficiently high that limitations on choice set size are not a primary driver of the difference in outcomes between groups. However, it is important to remember that differentiation (embodied by σ_S) could still prevent certain recordkeepers from being viable alternatives to a particular sponsor.

We find differences in bargaining parameters across groups. The largest firms (Group 4) have a significantly larger bargaining parameter than the firms in the other three groups do, and their parameter (0.85) suggests that they capture most of the surplus generated by the plans. Bargaining parameters for other firms are between 0.31 and 0.65, significantly below 1: recordkeepers capture a substantial portion of the incremental surplus generated by these plans. We find that small and large firms (Groups 1 and 4) have a larger bargaining parameter than firms in Groups 2 and 3. For large employers, this could arise from the fact that they are more likely to employ dedicated benefits teams, which experience and resources to negotiate deals with providers. While small employers are less likely to have such teams, our result is consistent with recordkeepers finding the back and forth with these employers more onerous than with middle-sized ones.[12] Note that a larger bargaining parameter does not means that

[12]For instance, small firms may be less likely to run formal requests for proposals and instead engage in informal discussions with potential recordkeepers with whom they already have a relationship, which could lead them capturing more of the surplus than a firm would if they ran arms-length RFPs. Small firms may

	Group 1 ≤ 200	Group 2 201–500	Group 3 501–1000	Group 4 > 1000
A. Frictions				
ρ	0.84	1.00	0.98	1.00
	(0.05)	(0.05)	(0.15)	(0.03)
η	0.65	0.40	0.31	0.85
	(0.03)	(0.04)	(0.08)	(0.01)
Marginal Effect of κ (at \$50)	1.09	1.00	1.08	1.66
	(0.05)	(0.02)	(0.07)	(0.03)
B. Sensitivity to Characteristics (θ)				
Means				
Expenses per Person (\$1K)	-0.52	-0.64	-0.89	-1.52
	(0.03)	(0.01)	(0.03)	(0.04)
Employer Contributions (\$1K)	-1.08	-1.62	-1.65	-0.53
	(0.21)	(0.06)	(0.12)	(0.03)
Has High # of Options	0.00	0.02	0.02	0.02
	(0.00)	(0.00)	(0.00)	(0.00)
Has Target Date Fund	0.06	3.26	2.20	0.02
	(0.02)	(0.14)	(0.36)	(0.00)
Has Index Fund	3.40	3.94	2.73	0.02
	(3.09)	(0.90)	(0.51)	(0.01)
Has S&P 500 Tracker	0.04	0.01	0.10	-0.03
	(0.01)	(0.00)	(0.02)	(0.00)
Standard Deviations				
Constant	4.48	3.81	2.65	3.07
	(3.43)	(0.34)	(0.47)	(1.85)
Expenses per Person (\$1K)	0.00	0.00	0.02	0.02
	(0.91)	(0.50)	(0.15)	(0.77)
Employer Contributions (\$1K)	0.00	0.06	0.01	1.59
	(1.58)	(0.38)	(0.47)	(0.08)
Has High # of Options	0.00	0.04	0.03	0.06
	(0.47)	(0.01)	(0.01)	(0.00)
Has Target Date Fund	1.02	4.12	2.90	0.22
	(0.55)	(0.27)	(0.33)	(0.07)
Has Index Fund	6.43	3.82	2.82	0.00
	(3.14)	(0.46)	(0.37)	(0.65)
Has S&P 500 Tracker	1.03	0.00	0.62	0.06
	(0.30)	(0.09)	(0.27)	(0.02)
σ_S	0.23	0.06	0.06	1.13
	(0.03)	(0.00)	(0.00)	(0.09)
C. Recordkeeper Costs				
σ_R	0.00	0.00	0.00	0.02
	(0.23)	(0.03)	(0.02)	(0.01)
Share-Weighted Average Cost	0.00	0.05	0.04	-0.02
	(0.12)	(0.01)	(0.01)	(0.01)

Table 2: Parameter estimates, in thousands of dollars. Standard errors are in parentheses.

small firms have larger surplus in equilibrium, as their choice sets (and thus their second-best options) could be weaker than those of larger firms; it is solely a reflection of how much of the incremental surplus is captured by each party. Moreover, the heterogeneity in bargaining parameters is not reflective of the probability of offering a plan, as in our framework a deal will be struck whenever it is efficient, regardless of η.

We find that κ plays a small role in this setting, except for Group 4. At a transfer of \$50 dollars, employers in Groups 1–3 employers penalize paying an extra dollar of transfers by no more than 9 cents. Group 4 employers, however, penalize it by an additional 66 cents. As mentioned above, these are the employers who are most likely to have a dedicated benefits team, and it is plausible that they would enact the most guardrails against paying transfers. As suggestive corroboration of this fact, we report in Section 6 that their preferences imply that this group would have a drastic increase in transfers if κ were set to 0.

Panel B shows estimates of sensitivity to plan characteristics. We find that employers are sensitive to the expenses their employees pay, and the sensitivity increases with firm size. For Groups 1–3 the sensitivity is less than dollar-for-dollar: Group 1 sponsors, for example, exhibit a disutility of 52 cents for a dollar of expenses, suggesting underinternalization of costs to employees. Group 4, however, has a larger disutility than one dollar. The increasing sensitivity to expenses with size could capture increasing litigation risk, and such legal recourse could justify sensitivity to expenses that exceed one-for-one: excessive fee litigation has mostly targeted large employers. At the same time, these differences could also be reflective of differences in sophistication by those in charge of making benefits decisions, or of the sensitivity of workers to the fees charged by the plan. Overall, we do not see much heterogeneity within group along this dimension.

We see a similar pattern for employer contributions: Group 4 also has a lower sensitivity than Groups 1–3. When $\kappa = 0$, a coefficient of -1 would indicate that the employer trades off a marginal dollar of transfers equally to a marginal dollar of contributions. It is not a priori clear that this number should be -1: contributions are usually tax-advantaged (a force towards lowering the magnitude of the coefficient), but they may also be less salient to workers than expenses in other benefits (a force towards increasing the magnitude). They may also be harder to target to particular groups of workers than wage increases. Overall, we find that firms in Groups 1–3 are weakly more sensitive to contributions than dollar-for-dollar. Group 4 firms view the cost of a dollar of contributions as 53 cents. Large

also be more "patient" than medium-sized firms in these negotiations, as they may not be in a rush to start 401(k) plans.

employers may recognize that paying benefits is an unavoidable part of attracting certain types of workers, may better internalize that non-wage benefits trade-off with wages, or may make this form of compensation more salient. We also observe that these firms have more heterogeneity along this dimension on preferences than other groups do, reflecting that some firms penalize employer contributions much more, while others see it as more of a benefit.

We next turn to the indicators of having certain plan features. It is a priori unclear whether a large number of options would be viewed as a benefit: while there are diversification benefits, too many options may induce choice overload and lead to suboptimal investment choice, exposing firms to liability. We find a limited role for this on average, with valuations no more than $20 for having more than the median number of options in the data (24). We find that Groups 2 and 3 value target date funds between $2,000 and $3,000 per person, while Groups 1–3 value having index funds in the same range. These values are large: given that these groups internalize expenses less than one-for-one, one explanation is that they place emphasis on investment options that are viewed as reasonable (e.g., as qualified defaults) as a shield for liability conditional on offering a plan. They may also simply view them as higher-quality independent of the fees. By contrast, Group 4 firms place much larger emphasis on fees directly rather than on other features of funds per se. They still value low fee index funds, but they do not exhibit an outright preference for the fund category. It is worth noting that for some of these preferences, we estimate sizable heterogeneity in the random coefficient, suggesting that substitution patterns are highly correlated with these characteristics. Finally, we see that having an S&P 500 tracker does not provide much value: these market trackers are already index funds, so this should be interpreted as the additional value beyond an index fund. Overall, we see these results as indicating that small to medium employers value having index funds and target date funds in the plan, even if other funds lead to large expenses, while large employers are much more focused on the overall expenses of the plan.

The last row of Panel B reports the standard deviation of the sponsor preference shocks, σ_S. This variable dictates the degree of differentiation between recordkeepers. We find that Groups 1 and 4 exhibit the most differentiation, while Groups 2 and 3 exhibit the least. Group 1 may be purchasing the service of defined contribution plan management from a company that already offers it other services, creating differentiation. For example, ADP and Empower also offer payroll services, and firms in this group may have a higher match value from managing their defined contribution plans along with their payroll. Group 4, on the other hand, may place a higher value on particular recordkeepers due to services like

financial advice or consulting to avoid exposure to litigation risk.

Panel C reports results for recordkeeper costs. We find that sponsor-side cost shocks (σ_R) play a small role, so that heterogeneity in revenue sharing is the main driver of differences across sponsors from the perspective of the recordkeeper. Share-weighted average costs are up to about $50/person. Table D.6 in Appendix D.3 reports recordkeeper-level cost estimates. Costs are heterogeneous across recordkeepers. We see some evidence of economies of scale within-recordkeeper, as well as of specialization at different sizes.

We consider the possibility that sponsors also value returns of the funds in the plan rather than simply expenses. In Appendix D.4, we report results from a specification where we replace the expenses with "adjusted expenses," by adding back in a measure of risk-adjusted returns at the plan level. Results are largely unchanged, although we do find that Group 4 firms are more sensitive to both expenses and employer contributions in this specification. Given the salience of expenses in litigation surrounding defined contribution plans, and concerns about whether risk-adjusted returns are really persistent, we use the results in Table 2 as our preferred specification for counterfactual policy analysis.

6. Quantifying the Frictions in Plan Design

The parameter estimates identify both market imperfections and misaligned willingness to pay as present in this market. How much do these forces contribute to limited participation and to observed plan characteristics?

Figure 5 shows how outcomes would change if each friction were alleviated. To simulate these effects, we change model parameters to alleviate frictions one by one and solve for counterfactual equilibria. In doing so, we keep all choice sets $\mathcal{R}_s$ equal to the ones used in estimation and re-draw consideration sets $\tilde{\mathcal{R}}_s$. We then recompute the maximum extractible surplus to the sponsor s from all recordkeepers in this consideration set, let the sponsor pick the recordkeeper that offers the largest surplus, and recompute the solution to the Nash bargaining problem by letting the sponsor and the chosen recordkeeper select a new plan from $\mathcal{P}_{rs}$ and an associated transfer. This process allows sponsors to change both the recordkeeper and the plan they choose (and thus the characteristics of the chosen plan), allowing us to simulate the effects of counterfactual policies on both the extensive and intensive margins of plan offerings.

Our counterfactual analysis captures what we believe to be a salient implication of vertical integration in this industry: vertically integrated recordkeepers still receive the full

expense ratio and thus will have an incentive to steer sponsors towards plans that offer their own funds, and they will be willing to give sponsors lower transfers to do so. However, we do not allow recordkeepers to renegotiate the revenue sharing agreements that they have with fund providers: if sponsor s and recordkeeper p decide on a new plan $p \in \mathcal{P}_{rs}$, the revenue sharing that r earns will be a function of the previously negotiated fund-level revenue sharing agreements.[13]

Figures 5–7 show results for firms in each of the four size groups. Each group's bar has as a base value the average of the outcome in that group in the baseline, and the size and direction of the bar corresponds to the effect of the counterfactual.

Alleviating Market Imperfections. The first set of columns in Figure 5 shows the impact of alleviating market imperfections: we evaluate setting $\rho = 1$, $\eta = 1$, and then both. The first column in Figure 5 sets $\rho = 1$. Since ρ is already high for all groups, the effects of this change are negligible: employers who do not offer a plan choose not to because they earn limited benefits from doing so, not because of limited search over recordkeepers.

The second column removes recordkeeper market power by setting $\eta = 1$. In this model, market power does not lead to quantity constraints: as long as there is some transfer and plan that is profitable for both the recordkeeper and the sponsor, an agreement will be reached. Accordingly, we see in Panel (a) that the probability of offering a plan is not affected. Panel (d) shows that eliminating market power has a large impact on transfers. However, since κ is small for Groups 1–3, the effects on plan characteristics are limited if the choice set is not improved. Group 4 has a large κ, so effects on expenses could have been large. However, since this group already penalizes expenses heavily, the reduced transfer does not map to further reductions in this outcome.

The third column shows the effects of setting both $\rho = 1$ and $\eta = 1$. In our model, this represents the limit of what competition can do to market outcomes. We find that while transfers and revenue to the recordkeeper fall, this does not translate to significantly cheaper plans for workers; expenses per person are nowhere close to those of plans offered by large employers, for example. Panel (f) provides a metric to summarize these observations. We report the share of total expenses (transfer payments plus dollar payments through expense ratios) that are due to expense ratios; this metric provides the relative incidence of policies on sponsors versus workers. Alleviating all market imperfections increases the relative

[13] Moreover, many funds offered in defined contribution plans are ones that could be bought outside retirement plans, and thus their revenue sharing agreements would depend on more than just the market for DC plans.

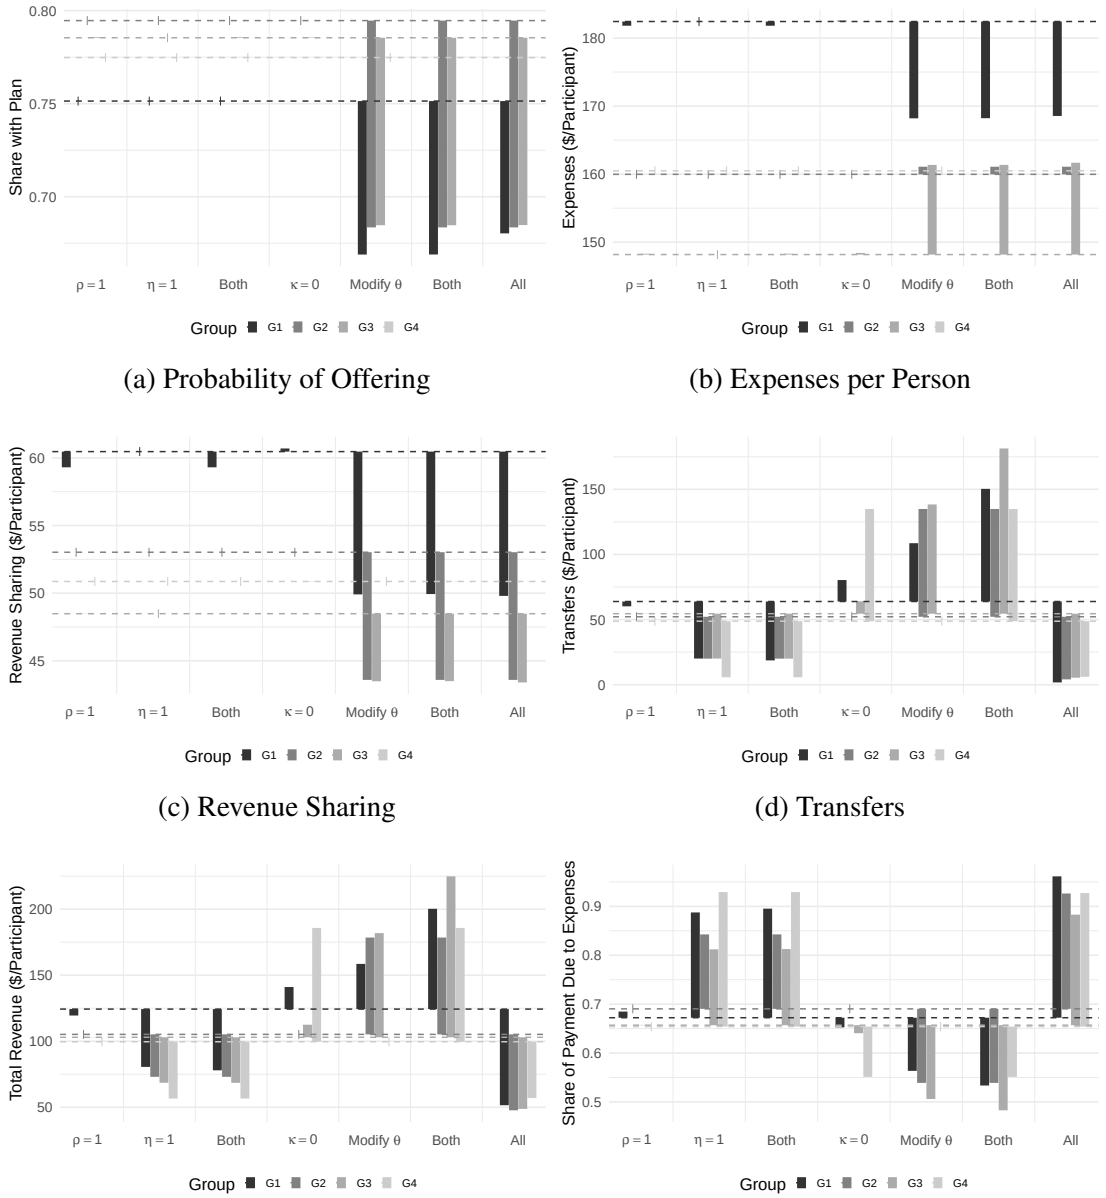

(a) Probability of Offering

(b) Expenses per Person

(c) Revenue Sharing

(d) Transfers

(e) Total Revenue

(f) Share of Expenses

Figure 5: Effects of alleviating various market frictions. The bars originate from dashed horizontal lines that denotes the mean value of the metric for each group.

burden on workers, from about 0.65–0.7 in the baseline to 0.80–0.93. This suggests that the incidence of market power lies primarily on the sponsor: given sponsor preferences (θ and κ), recordkeepers exercise market power by increasing transfers rather than by affecting plan characteristics. Moreover, we find limited effects of market power on the extensive

margin of plan provision too, as search is extensive enough that sponsors who would want to offer a plan are able to find an acceptable recordkeeper already.

Aligning Willingness to Pay. To what extent do sponsor preferences distort plan characteristics and provision? The second set of columns in Figure 5 answers this question alleviating the misalignment in willingness-to-pay for plans. One direct source of misalignment is that $\kappa > 0$, which induces two distortions. The first is on the intensive margin: a positive κ implies that the sponsor is willing to trade off plan features that go into its benefit function if doing so reduces the transfer, especially if the transfer is already high. The second is on the extensive margin. Each recordkeeper has a minimum transfer that would make it willing to serve a particular sponsor; reducing κ makes it more likely that this minimum transfer would be feasible for the sponsor as well and thus increases plan provision. However, if recordkeeper costs are such that recordkeepers can make positive profits even without transfers from the sponsor, κ will have limited effect on plan provision itself.

In line with this intuition and the fact that κ is typically close to 0, we see that the share of firms with a plan does not change if $\kappa = 0$. Moreover, in line with κ being small for Groups 1–3, we see almost no effect on plan characteristics such as expenses and revenue sharing; Group 4 already penalizes expenses heavily, so reductions in κ will not lead to further reductions in expenses. Panel (d) shows that the main effect of decreasing κ is an increase in transfers—especially for Group 4, and Panel (e) puts these results together to show that recordkeepers greatly benefit. In summary, the fact that employers have a disutility of a dollar of transfers that is greater than a dollar has limited effects on plan characteristics and coverage, but due to market power, recordkeepers capture a large share of the additional surplus generated when this friction is removed.

The second source of misalignment is that θ is such that sponsors may value plan characteristics differently from either what a worker would or what an expert would choose. Unlike for ρ, η, or κ, there is no natural benchmark for alleviating any frictions induced by a misaligned θ, as it is a priori unclear what the "right" preferences should be. Thus, we consider the experiment of moving θ to that for large firms (and move δ to this level as well to account for changes in the mean benefit of offering plans). Of course, aligning preferences of small employers with preferences of large employers may not be feasible, so in Section 7.2 we discuss direct quality regulation.

The main difference between preferences for Group 1 and Group 4 is that the former penalizes expenses much less. If Group 1 penalized expenses as much as Group 4, it would

choose plans with significantly lower expenses, as shown in Panel (b). This reduces revenue sharing (Panel (c)), which leads to two main effects: a shutdown of plans, as shown in Panel (a), and an increase in transfers, as shown in Panel (d). Total revenue conditional on offering a plan increases (Panel (e)), but this is a compositional effect: plans with low levels of revenue for recordkeepers are shut down. Thus, we find that giving small employers the preferences of large employers would reduce expenses, shrinking the gap between small and large employers by around 2/3, at the cost of a reduction in plan provision of around 7.5 pp. Results are more subtle for Groups 2 and 3, where expenses increase. This is due to the fact that these groups value index and target date funds more than Group 4, and have less of a gap in expense sensitivities. Nevertheless, revenue sharing falls for these groups, so that we also observe a contraction in plan offering and an increase in transfers.

The final column in the second set of columns in Figure 5 aligns θ and sets $\kappa = 0$. The effect on expenses is overall similar to just aligning θ, as changing κ had little effect on this outcome. The effects on transfers compound. Overall, adjusting willingness to pay and sensitivity to characteristics without addressing market power would benefit recordkepeers greatly as well. Our main finding is that while eliminating market power could affect transfers, adjusting θ or κ is necessary to have a noticable effect on plan characteristics like expenses—especially for small firms in Group 1.

Interaction Between Market Power and Willingness to Pay. The final column in Figure 5 shows the effect of both alleviating market imperfections and aligning willingness to pay. The main effect of alleviating market power after aligning incentives is on transfers themselves: setting $\eta = 1$ reduces transfers considerably, to below the baseline levels—and to below the levels of just eliminating market power. Recall that eliminating market power by increasing bargaining power η does not undo the extensive margin impact of aligning preferences, as efficient bargains will always be struck in this setting, regardless of η.

7. How Does Regulation Affect Plan Design?

The message of Section 6 is that material changes to plan characteristics require directly changing preferences. This observation guides our analysis of counterfactual regulation. As discussed below, some of them can be viewed as attempts to address market imperfections indirectly while others attempt to target quality directly. We analyze such policies in Sections 7.1 and 7.2, respectively, and discuss the results together in Section 7.3. The

takeaways from this section are in line with the lessons from Section 6: policies that directly promote or constrain plan characteristics are necessary to appreciably change them, and indirect measures are much less effective.

7.1. Targeting Market Imperfections

In this section, we study two approaches included in the SECURE Act, signed into law in December 2019 (Commito, 2020). One of its provisions provides tax credits of up to $250/participant-year to subsidize the provision of DC plans—albeit the focus is on especially small plans (effectively fewer than 20 employees). Here, we consider the impact of expanding subsidies to larger firms, which would both directly incentivize provision and encourage competition. Another provision allows small employers to band together to form "multiple-employer plans" (MEPs) or pooled-employer plans, changing an earlier Department of Labor rule that drastically limited the ability of firms to form such plans.[14] One rationale for MEPs is to allow recordkeepers and small firms to realize economies of scale in plan provision; if these are large, it would directly increase plan offerings and potentially also competition by increasing the set of recordkeepers willing to serve the sponsor. Another rationale is that these MEPs would give small firms more buying power to negotiate with plan providers, providing one mechanism to increase η.[15]

Subsidizing Plan Provision. We investigate the effects of subsidies around the levels proposed by the SECURE Act. The first set of columns for Figure 6 shows the effects of subsidies of $100 and $500 per participant.[16] They have a direct impact on provision, with $500 subsidies increasing participation of Group 1 firms by about 1.1 pp, or roughly halfway to that of Group 4 firms. The increase for other groups is larger and is about 4–5

[14] Previously, firms were only allowed to form them if they were tightly linked, e.g., in the same industry and geography ("nexus"). Moreover, such plans were subject to the "one bad apple" rule, which made all firms liable for a fiduciary violation from one of them. Industry participants claimed this sharply increased costs.

[15] For commentary from industry participants discussing this rationale, see "SECURE Act Alters 401(k) Compliance Landscape" (SHRM, 2020) at https://www.shrm.org/resourcesandtools/hr-topics/benefits/pages/secure-act-alters-401k-compliance-landscape.aspx or "How the SECURE Act is Changing Retirement" (QRPS, 2020) at http://www.qrps.com/Resources/Articles/ArticleID/1371/How-the-SECURE-Act-is-Changing-Retirement.

[16] We make these subsidies nominally incident on the sponsor to capture the setup of the SECURE Act. In a world with efficient bargaining ($\kappa = 0$ and no positive transfer constraint), nominal incidence of the subsidy would not matter for real outcomes. However, in our model, nominal incidence does matter; a subsidy incident on the recordkeeper will not pass through to lower transfers after the recordkeeper hits the zero lower bound on transfers.

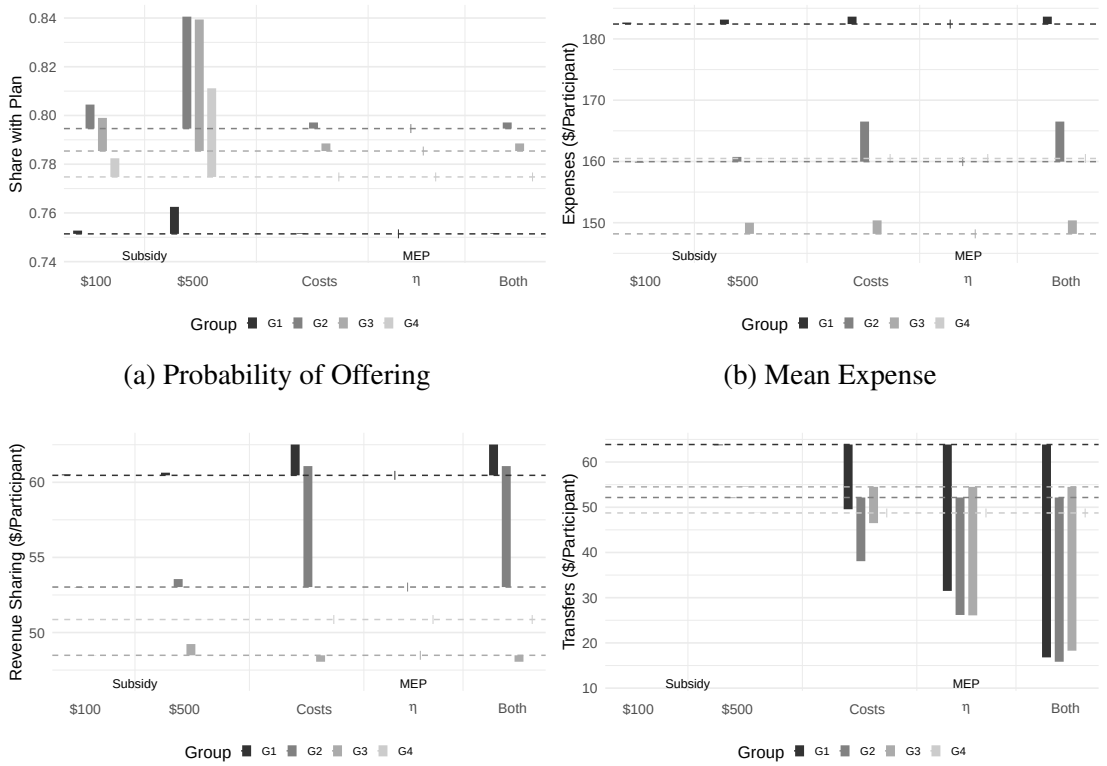

(a) Probability of Offering

(b) Mean Expense

(c) Revenue Sharing

(d) Transfers

Figure 6: Effects of subsidies and multiemployer plans. The figure shows two values of untargeted subsidies ($100 and $500 per person) and three versions of MEPs (adjusting costs, adjusting bargaining parameter η, and adjusting both). The bars originate from dashed horizontal lines that denotes the mean value of the metric for each group.

pp. However, subsidies do not have a material impact on other metrics. To understand why, note that a sponsor would collect a flat subsidy for plan provision regardless of the choice of the recordkeeper; thus, in the bargaining problem, if the second-best option were another recordkeeper (rather than not offering a plan at all), the incremental surplus to the sponsor would be unaffected by the level of the subsidy. Therefore, transfers and plan characteristics are also unaffected. Subsidies affect these outcomes only insofar as they change the outside option for a sponsor from not offering a plan to offering a plan from a different recordkeeper, or through a compositional change in the firms who are offering plans. The small effects we observe on expenses and revenue sharing are due to a combination of these channels. These results again underscore that without direct changes to the benefit function, it is difficult to change plan characteristics significantly.

Effects of MEPs. Our general approach to modeling MEPs is to change certain parameters of small firms to those of large firms while keeping choice sets and other parameters fixed. Since the policy intends to exploit potential economies of scale, we start with the recordkeeper's cost parameters. The third column of Figure 6 shows that in our setting, however, the effects of such cost changes on plan provision are small. While (point estimates for) share-weighted average costs are lower for Group 1 than for Group 4, the heterogeneity in costs across recordkeepers matters as well: this result arises because costs fall for some vertically integrated recordkeepers who offer plans with higher levels of revenue sharing. That is, the policy shifts the identity of the chosen recordkeeper, increasing revenue sharing and expenses while decreasing transfers, without large overall effects on the intensive margin. A natural question is what happens if the sponsors' costs change as well. Our empirical strategy allows us to estimate δ, the mean net benefit of offering a plan with each recordkeeper, but does not separate costs and benefits within δ. However, sponsor subsidies are isomorphic to cost reductions, so even an extremely large cost reduction of \$100 per worker will not lead to changes in plan characteristics, and they would also have limited effects on plan participation for Group 1 firms.

What if the effect of MEPs is to improve buyer power? We can model this directly by increasing the bargaining power of small firms to those of large ones. Accordingly, this counterfactual is a less extreme version of the one presented in Section 6 to alleviate market power. We thus find similar, but smaller, results: transfers decrease considerably, but we find little impact on expenses: they decrease by 1.6 bp for small firms and are mostly unchanged for large firms. Putting them together, we see that the net effect of MEPs would be a sizeable decrease in transfers, as the effects of the reallocation from costs and the increase in bargaining power compound. However, we do not predict an improvement in expenses and rather expect a modest increase in this quantity.

Of course, if MEPs were somehow able to cause small firms to adopt the willingness to pay of large firms, we would recover the results from Section 6 and affect plan characteristics directly. One might envision policies that could be coupled with MEPs that could directly target θ. For instance, small firms might not value low expense ratios because of limited sophistication by their decision-makers, or because their workers do not value them and thus the firm does not accrue any labor market benefits from putting them on the plan; in such cases, informational policies provided when setting up MEPs may help. Alternatively, small firms may not face much litigation risk for lower-quality menus; an MEP may be a more effective target for a lawsuit, and workers in MEPs may thus have more options for legal

recourse. On net, however, we find it implausible that these policies could completely close the gap in willingness to pay between small and large employers, as significant labor market differences will always remain. Thus, we conclude that MEPs are unlikely to effectively close the outcome differences between small and large employers.

7.2. Direct Regulation of Plan Design

Targeting market imperfections provides an indirect method to regulating characteristics. A more direct approach would constrain the set of plans a sponsor can offer, or provide subsidies or penalties as a function of plan characteristics. As discussed in Section 2.2, no regulation currently imposes such restrictions. Nevertheless, this approach could be envisioned as a natural strengthening of the guidelines provided by the IRS or DOL in certain situations. For instance, regulators could impose "must-carry provisions," which mandate that all plans contain investments that would be appropriate as a QDIA, rather than simply offering safe harbor provisions for offering a QDIA. Regulation could also introduce "expense caps," prohibiting recordkeepers from plans with high average expense ratios.

In this subsection, we evaluate policies that directly regulate plan design. The fundamental trade-off here is between the intensive and extensive margin. On the one hand, such policies mechanically affect plan characteristics conditional on selecting a plan—at least on the dimension that is regulated. On the other, imposing constraints on acceptable plans makes it less likely a sponsor offers a plan: a sponsor may have to compensate a recordkeeper more when plan characteristics are changed, and it may be unwilling to do so—both because it perceives a lower benefit from it and because it prefers to avoid high transfers. A less drastic measure is to target subsidies to sponsors specifically for plans that meet certain criteria.

Figure 7 shows results from these policy experiments. To implement must-carry provisions, we modify the choice set of plans by adding an S&P 500 tracker or a suite of target date funds to any plan that does not meet the must-carry provision. Further, we update predicted values of expenses per person, revenue sharing, and employer contributions per person by adding the most common S&P 500 tracker / target date suite offered by each recordkeeper to the plan. See Appendix B.3 for further details. We also consider targeted subsidy policies for expense caps; for this counterfactual, we retain the same choice set as in the data but add a fixed dollar amount (per participant) to the sponsor's payoff if the plan has average expenses less than 50 bp.

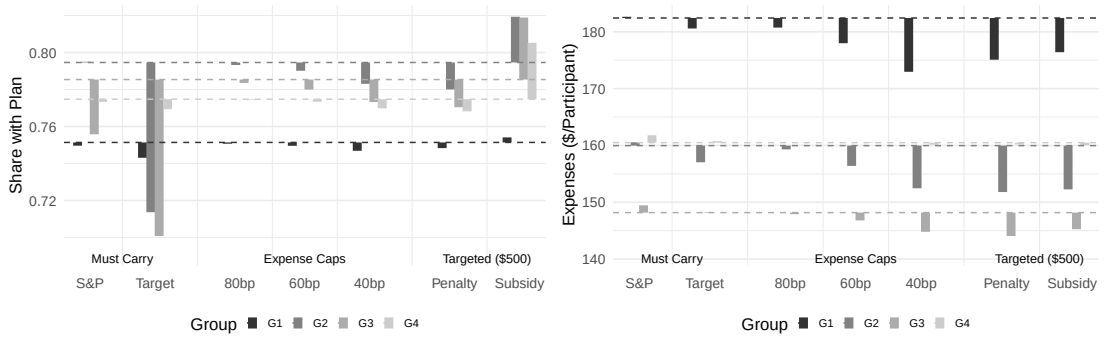

(a) Probability of Offering

(b) Mean Expense

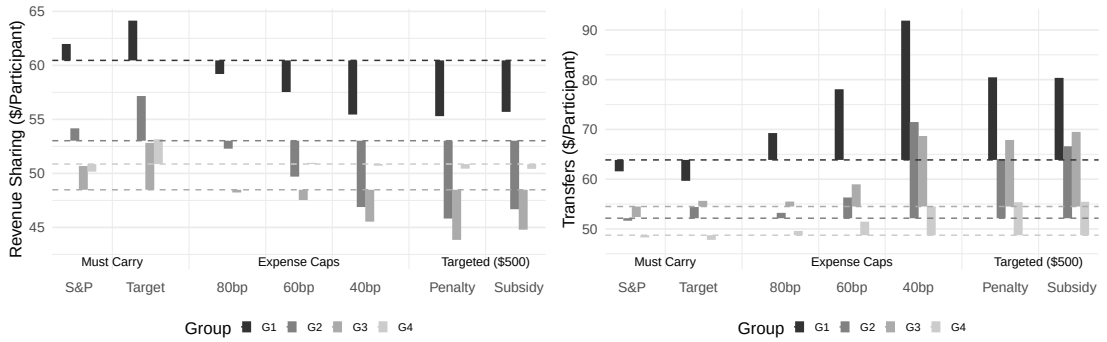

(c) Revenue Sharing

(d) Transfers

Figure 7: Effects of must-carry provisions, expense caps, and targeted subsidies or penalties (for plans with average expenses of 50 bp). The bars originate from dashed horizontal lines that denotes the mean value of the metric for each group.

Must-Carry Provisions and Expense Caps. We first consider two simple and easily implementable must-carry provisions—mandating either an S&P index fund or a target date fund—modeled after strengthening QDIAs.[17] The effect on plan provision is minimal for S&P mandates, but mandating target date funds can affect plan provision by 8–9 pp for certain sponsors. This is driven by the fact that preference heterogeneity is larger for target date funds than for S&P 500 trackers. We find negligible effects on expenses, and an increase in revenue sharing. This latter effect is a combination of a compositional shift—plans that drop out have low levels of revenue sharing—and the fact that target date funds for vertically integrated recordkeepers accrue the entirety of expenses as revenue sharing.

These must-carry provisions have generally small effects on expenses because they

[17] We do not consider a policy that forces plans to carry an index fund. Given there is large heterogeneity in the types and goals of index funds, this seems to us like an unrealistic policy to implement.

target the wrong metric. Simply including an S&P index or a target date fund does not ensure that workers will invest heavily in these funds, nor does it ensure that higher-fee funds cannot be on the menu. Thus, we find that sponsors choose similar types of plans on dimensions such as expenses even though the plan now contains these mandated funds.

These results suggest that if the goal is to lower expenses, a policy should directly target them. Thus, we evaluate a expense ratio cap, which forces sponsors to only offer plans whose mean expense ratios are lower than a specified amount. In Figure 7, we consider successively stringent caps of 80 bp, 60 bp, and 40 bp. These caps directly lower expenses by considerably more than must-carry provisions. Since expense ratios are correlated with revenue sharing, we also see corresponding drops in this outcome. Transfers increase significantly: not only do recordkeepers now need to be compensated for the loss of revenue sharing, but the sponsor's second-best option during the bargaining process also weakens, as it is able to extract less from any recordkeeper given the legal constraints on plan design. The increase in transfers is especially large for Groups 1–3 since they have a relatively small value of κ and thus do not face a penalty for paying them. Our results indicate that despite these large changes in transfers and plan characteristics, the extensive margin effects are somewhat small and do not exceed 2 pp for any group: the firms that offer plans gain enough benefit from them (and have a low enough aversion to transfers) that they are able to find a mutually beneficial plan even when they are constrained.

Targeted Penalties and Subsidies. A perhaps better-targeted version of a must-carry provision is to subsidize firms for offering certain kinds of plans (e.g., those with low expenses)—or to penalize them for not doing so. We consider policies that subsidize plans with expense ratios lower than 50 bp (or penalize the opposite) and call these reforms targeted penalties and subsidies. The final set of columns in Figure 7 shows the effects of $500/participant targeted penalties and subsidies. A penalty generally has a qualitatively similar effect on outcomes as an outright cap. Quantitatively, the effect of the penalty is similar to that of an expense cap of 40 bp. Of course, the penalty need not be this large, and Figure 8 (discussed in Section 7.3 below) shows that certain values of targeted penalties can dominate must-carry provisions on plan provision and expenses. A $500 subsidy has similar effects on mean expenses and revenue sharing, but it leads to a small increase in plan provision (of up to 4 pp, depending on the group). The government must balance these benefits with the cost of the transfer itself. In both targeted policies, transfers to the recordkeeper increase. The rationale is the same as the one discussed for expense caps: recordkeepers must be compensated

for lower revenue sharing from these policies, and constrained plan choice also lowers the sponsor's outside option when bargaining with the recordkeeper.

7.3. Discussion

Figure 8(a) summarizes the results of the policy counterfactuals for small firms (Group 1). For each policy, it plots the mean expense ratio of chosen plans against the share of firms who carry a DC plan. The dashed vertical and horizontal lines indicate the baseline values for these firms. Since mean expense ratios are plotted in descending order, any policy in the top right quadrant is preferred on these metrics, while those in the top left or bottom right are not unambiguously comparable. A variety of counterfactual policies considered in the previous two subsections are plotted with crossmarks, expense caps are plotted as circles connected by solid lines, and outcomes with subsidies (targeted or not) are plotted as patterned lines connecting various values of subsidies.

Our results indicate that in the market for DC plan provision, there is no free lunch—at least if a policymaker has as his objective function plan provision and expenses paid by participants. Almost all policies considered lead to a reduction in expenses at the cost of reduced plan provision, or an increase in plan provision at the cost of increased expenses. Must-carry provisions improve expenses (or leave them unchanged) but reduce plan provision slightly. MEPs have limited effects on outcomes. Untargeted subsidies increase plan provision but increase expenses slightly. Targeted penalties, plotted as dashed black lines, can deliver higher coverage for any given level of plan quality than must-carry provisions, as do expense ratio caps. The model provides quantitative comparisons between these policies: for instance, targeted penalties and expense caps are quite similar in their effects on expenses and plan provision, and both strictly dominate must-carry provisions. The only policy that improves both expenses and plan provision are targeted subsidies—but this comes at a cost to the government, and even a $500/person subsidy fails to change expenses and provision to the level seen in large firms. Note, however, that extensive margin changes are small, and policymakers may find it worthwhile to bear a small loss in plan provision to significantly lower expenses paid by workers.

As discussed, the force underlying these results is that the main barrier to small firms offering low-expense plans is not market power or aversion to transfers but rather limited sensitivity to certain plan characteristics. In the absence of direct methods to increase this sensitivity, the policy options to change these plan characteristics must either constrain plan choice and thus reduce provision, or subsidize it directly, which would be expensive.

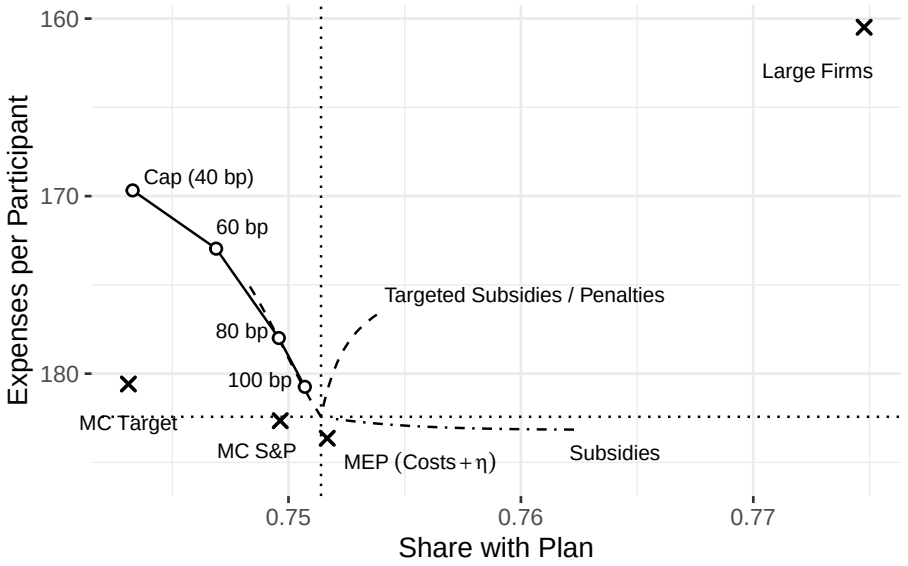

(a) Expenses against probability of offering a plan

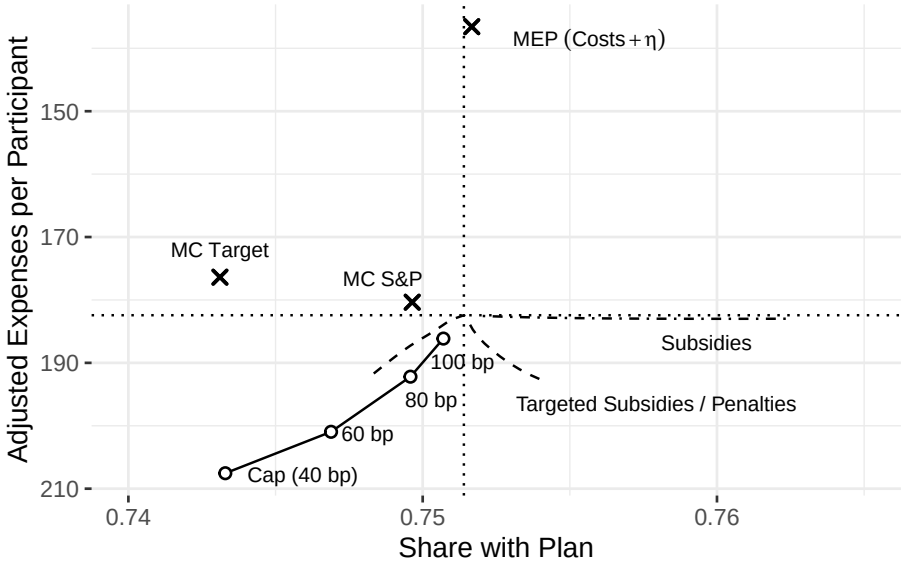

(b) Adjusted expenses against probability of offering a plan

Figure 8: Summary of policy counterfactuals. The dotted axes indicate the values for Group 1 firms. The black crossmarks show values for must-carry policies (MC) and multi-employer plans (MEP). The bold dotted line shows targeted penalties or subsidies (from a $500 penalty to a $500 subsidy) for plan with mean expenses under 50 bp. The dash-dotted lines show the effects of untargeted subsidies (from $0 to $500). Panel (b) shows the same policies but assumes that changes in transfers are fully passed through to workers.

The results in the preceding sections also indicate that these policies can have large effects on transfers. One might expect that changes in the characteristics of DC plans, or the transfers paid by the employer for these plans, will pass on to wages or other non-wage compensation. Without data on other compensation, we are unable to estimate this elasticity, and despite the fact that the substitution between wages and non-wage benefits is a classic question in labor economics (Woodbury, 1983), we are unaware of any such estimates in the literature in the context of DC plans. We thus compute an extreme benchmark: Figure 8(b) reports "adjusted" expenses, where we assume that the change in transfers induced by the policy are fully passed through to employees as wages or other methods of non-wage compensation. This benchmark changes the relative rankings of the counterfactuals. The quantitative impacts on must-carry provisions and untargeted subsidies is small. However, the large transfer changes in expense caps and targeted policies—coming from both the reduction in revenue sharing and the harm to the bargaining position coming from constraints on the second-best plan—overshadow the improvement in expense ratios. On the positive side, the analysis highlights one potential benefit of MEPs. If MEPs greatly increase the bargaining power of small firms, then small firms are able to negotiate lower transfers with employers; if these transfers are then passed through to workers, then MEPs can improve effective worker expenses considerably. However, we reiterate that the plan characteristics themselves do not change appreciably according to our model, and thus the benefit of MEPs is predicated on the passthrough of transfers to other forms of compensation rather than on a direct improvement in plan characteristics.

8. Conclusion

This paper studies the equilibrium effects of regulation in the market for defined contribution plans, through a model of the bargaining between employers and recordkeepers. While employers are charged with keeping their workers' best interests in mind when designing these plans, there is significant policy concern over the prevalence of what many consider to be low-quality plans: those that have high fees or lack particular investment options. Using data on plan menus and payments, we show that both market imperfections and a misalignment in willingness to pay (which incorporates agency frictions between the employer and the worker) play a role here. However, alleviating market power itself does not have a significant impact on plan characteristics that directly impact workers. Since plan characteristics are largely governed by employers' willingness to pay for them, which may

be difficult to impact directly, we conclude that regulation just target plan characteristics directly, either by constraining the design of plans or by subsidizing or penalizing certain plan features, if it is to be effective at improving them.

We view this paper as broadly contributing to our understanding of the market for non-wage benefits, which are often provided by other companies. Research on regulating the supply-side of these markets has largely focused on health insurance so far, but this paper highlights that market frictions are present in other settings as well. We view this area as a fruitful avenue for future research.

References

ALLEN, J., R. CLARK, AND J.-F. HOUDE (2014): "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry," *American Economic Review*, 104, 3365–96.

——— (2019): "Search Frictions and Market Power in Negotiated-Price Markets," *Journal of Political Economy*, 127, 1550–1598.

BACHAS, N. (2019): "The Impact of Risk-Based Pricing and Refinancing on the Student Loan Market," Working Paper.

BADOER, D. C., C. P. COSTELLO, AND C. M. JAMES (2020): "I Can See Clearly Now: The Impact of Disclosure Requirements on 401(k) Fees," *Journal of Financial Economics*, 136, 471–489.

BENARTZI, S. AND R. THALER (2007): "Heuristics and Biases in Retirement Savings Behavior," *Journal of Economic Perspectives*, 21, 81–104.

BENARTZI, S. AND R. H. THALER (2001): "Naive Diversification Strategies in Defined Contribution Saving Plans," *American Economic Review*, 91, 79–98.

BERRY, S., A. GANDHI, AND P. HAILE (2013): "Connected Substitutes and Invertibility of Demand," *Econometrica*, 81, 2087–2111.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 841–890.

——— (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112, 68–105.

BERRY, S. AND A. PAKES (2007): "The Pure Characteristics Demand Model," *International Economic Review*, 48, 1193–1225.

BESHEARS, J., J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2013): "Simplification and Saving," *Journal of Economic Behavior & Organization*, 95, 130–145.

BHATTACHARYA, V. (2021): "An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program," *Econometrica*, 89, 2189–2224.

BHATTACHARYA, V., G. ILLANES, AND M. PADI (2025): "Fiduciary Duty and the Market for Financial Advice," *Econometrica*, Forthcoming.

BINMORE, K., A. SHARED, AND J. SUTTON (1989): "An Outside Option Experiment," *Quarterly Journal of Economics*, 104, 753–770.

BRANCACCIO, G. AND K. KANG (2025): "Search Frictions and Product Design in the Municipal Bond Market," Tech. rep., NYU.

BRANCACCIO, G., D. LI, AND N. SCHÜROFF (2020): "Learning by Trading: The Case of the US Market for Municipal Bonds," Tech. rep., NYU Stern.

CHOI, J. J. (2015): "Contributions to Defined Contribution Pension Plans," *Annual Review of Financial Economics*, 7, 161–178.

CHOI, J. J., D. LAIBSON, AND B. C. MADRIAN (2009): "Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect," *American Economic Review*, 99, 2085–95.

——— (2011): "$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans," *Review of Economics and Statistics*, 93, 748–763.

CHOI, J. J., D. I. LAIBSON, AND B. C. MADRIAN (2005): "Are Empowerment and Education Enough? Underdiversification in 401(k) Plans," *Brookings Papers on Economic Activity*, 2005, 151–213.

CLARK, R., J.-F. HOUDE, AND J. KASTL (2021): "The Industrial Organization of Financial Markets," in *Handbook of Industrial Organization*, Elsevier, vol. 5, 427–520.

COLLINS, S., S. HOLDEN, J. DUVALL, AND E. B. CHISM (2016): "The Economics of Providing 401(k) Plans: Services, Fees, and Expenses, 2015," *ICI Research Perspectives*, 22.

COMMITO, T. F. (2020): "SECURE Act Enhances Small Employer 401(k) Plans," *Journal of Financial Service Professionals*, 74, 12–15.

CRAWFORD, G. S., R. S. LEE, M. D. WHINSTON, AND A. YURUKOGLU (2018): "The Welfare Effects of Vertical Integration in Multichannel Television Markets," *Econometrica*, 86, 891–954.

CRAWFORD, G. S. AND A. YURUKOGLU (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102, 643–85.

CUESTA, I. AND A. SEPÚLVEDA (2019): "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access," Tech. rep., Stanford University.

CUESTA, J. I., C. NOTON, AND B. VATTER (2019): "Vertical Integration Between Hospitals and Insurers," Tech. rep., Stanford University.

DOELLMAN, T. W. AND S. H. SARDARLI (2016): "Investment Fees, Net Returns, and Conflict of Interest in 401(k) Plans," *Journal of Financial Research*, 39, 5–33.

EGAN, M. (2019): "Brokers vs. Retail Investors: Conflicting Interests and Dominated Products," *Journal of Finance*, 74, 1217–60.

EGAN, M. L., A. MACKAY, AND H. YANG (2024): "What Drives Variation in Investor Portfolios? Evidence from Retirement Plans," Tech. rep., National Bureau of Economic Research.

EINAV, L., M. JENKINS, AND J. LEVIN (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80, 1387–1432.

FRIEDMAN, J. H. (2002): "Stochastic Gradient Boosting," *Computational Statistics & Data Analysis*, 4, 367–78.

GOEREE, M. S. (2008): "Limited Information and Advertising in the U.S. Personal Computer Industry," *Econometrica*, 76, 1017–1074.

GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 105, 172–203.

GRENNAN, M. (2013): "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, 103, 145–77.

GRUNEWALD, A., J. LANNING, D. LOW, AND T. SALZ (2019): "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects," Tech. rep., MIT.

GUISO, L., A. POZZI, A. TSOY, L. GAMBACORTA, AND P. E. MISTRULLI (2022): "The Cost of Steering in Financial Markets: Evidence from the Mortgage Market," *Journal of Financial Economics*, 143, 1209–1226.

H2O.AI. (2020): "h2o: R Interface for H2O. R package version 3.30.0.6." .

HO, K. AND R. S. LEE (2017): "Insurer Competition in Health Care Markets," *Econometrica*, 85, 379–417.

——— (2019): "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *American Economic Review*, 109, 473–522.

HUBERMAN, G., S. S. IYENGAR, AND W. JIANG (2007): "Defined Contribution Pension

Plans: Determinants of Participation and Contributions Rates," *Journal of Financial Services Research*, 31, 1–32.

HUBERMAN, G. AND W. JIANG (2006): "Offering Versus Choice in 401(k) Plans: Equity Exposure and Number of Funds," *Journal of Finance*, 61, 763–801.

ILLANES, G. (2017): "Switching Costs in Pension Plan Choice," Tech. rep., Northwestern University.

ILLANES, G. AND M. PADI (2021): "Retirement Policy and Annuity Market Equilibria: Evidence from Chile," Tech. rep., National Bureau of Economic Research.

IYENGAR, S. S. AND E. KAMENICA (2010): "Choice Proliferation, Simplicity Seeking, and Asset Allocation," *Journal of Public Economics*, 94, 530–539.

KAHN, L. M. AND J. K. MURNIGHAN (1993): "A General Experiment on Bargaining in Demand Games with Outside Options," *American Economic Review*, 1260–1280.

KOIJEN, R. S. J. AND M. YOGO (2016): "Shadow Insurance," *Econometrica*, 84, 1265–1287.

KRONLUND, M., V. K. POOL, C. SIALM, AND I. STEFANESCU (2021): "Out of Sight No More? The Effect of Fee Disclosures on 401(k) Investment Allocations," *Journal of Financial Economics*.

LARSEN, B. J. (2021): "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*, 88, 851–882.

LEE, R. S., M. D. WHINSTON, AND A. YURUKOGLU (2021): "Structural Empirical Analysis of Contracting in Vertical Markets," in *Handbook of Industrial Organization*, Elsevier, vol. 4, 673–742.

LOSETO, M. (2023): "Plan Menus, Retirement Portfolios, and Investors' Welfare," Tech. rep., Bocconi University.

LUCO, F. (2019): "Switching Costs and Competition in Retirement Investment," *American Economic Journal: Microeconomics*, 11, 26–54.

NELSON, S. (2020): "Private Information and Price Regulation in the US Credit Card Market," Working Paper.

OCHS, J. AND A. E. ROTH (1989): "An Experimental Study of Sequential Bargaining," *American Economic Review*, 355–384.

PAPKE, L. E. (2004): "Individual Financial Decisions in Retirement Saving Plans: The Role of Participant-Direction," *Journal of Public Economics*, 88, 39–61.

POOL, V. K., C. SIALM, AND I. STEFANESCU (2016): "It Pays To Set the Menu: Mutual Fund Investment Options in 401(k) Plans," *Journal of Finance*, 71, 1779–1812.

——— (2021): "Mutual Fund Revenue Sharing in 401(k) Plans," Tech. rep., Vanderbilt.

ROBLES-GARCIA, C. (2020): "Competition and Incentives in Mortgage Markets: The Role of Brokers," Tech. rep., Stanford University.

RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

SCHEMBARI, J. (2022): "The Supreme Court Rules for Participants in Hughes v. Northwestern," https://www.kutakrock.com/newspublications/publications/2022/02/the-supreme-court-rules-hughes-v-northwestern.

SETHI-IYENGAR, S., G. HUBERMAN, AND W. JIANG (2004): "How Much Choice is Too Much? Contributions to 401(k) Retirement Plans," in *Pension Design and Structure: New Lessons from Behavioral Finance*, ed. by O. S. Mitchell and S. P. Utkus, Oxford University Press, chap. 5, 83–95.

SIMON, J. (2017): "Morningstar reports record low in fund fees," https://www.plansponsor.com/morningstar-reports-record-low-in-fund-fees/.

STROUD, A. (1971): *Approximate Calculation of Multiple Integrals*, London: Prentice Hall.

WOODBURY, S. A. (1983): "Substitution Between Wage and Nonwage Benefits," *American Economic Review*, 73, 166–182.

YANG, H. (2025): "What Determines 401(k) Plan Fees? A Dynamic Model of Transaction Costs and Markups," Tech. rep., London Business School.

SUPPLEMENTAL APPENDIX

A. Data Appendix

Appendix A.1 describes the data sources, and Appendices A.2 and A.3 describes the cleaning procedure use to extract information from the sources and arrive at the final dataset for estimation.

A.1. Data Sources

All firms that offer non-wage benefits are required to file a version of Form 5500 annually with the Department of Labor to comply with the requirements of ERISA. Plans with fewer than 100 participants are required to file a shorter form, Form 5500-SF, which contains a subset of the information in Form 5500. Our dataset comes from these public filings, but we import it from two sources. First, the DOL provides machine-readable data for a large subset of attachments in Form 5500 for direct download.[18] However, some information is submitted to the DOL by the firms in the form of PDF files without any common structure; these pdf files are also public and can be downloaded from the DOL's EFAST filing system.[19] Our second data source consists of these pdf files; we purchase data from BrightScope that extracts and cleans plan menus and asset allocations for plans from these pdf files.

The body of Form 5500 contains the employer's identification number (EIN), plan number, industry codes, main location, and participant counts. It also includes pension and welfare benefit codes, which allow us to identify firms that offer health insurance without offering a retirement plan or ones that offer defined benefit plans. Form 5500-SF also contains all this information.

Schedule C of Form 5500 (but not Form 5500-SF) contains information about service providers and the fees that they collect. Items 1 and 3 of Part 1 of the Schedule C contain information about which service providers received "indirect" compensation from the fund. The exact amount of the compensation is included in a format that is not easily machine-readable, and in Appendix A.3 we outline how we extract it. Item 2 lists service providers that received direct compensation for services, including payment amounts. This is the data source for what we refer to as transfers in the model. We use the payment to the

[18] This data can be accessed at https://www.dol.gov/agencies/ebsa/employers-and-advisers/plan-administration-and-compliance/reporting-and-filing/form-5500.

[19] See https://www.efast.dol.gov/welcome.html.

party designated as the main recordkeeper (see Appendix A.2) as our measure of T, but descriptive results are similar if we use total payment to all service providers as the measure.

Finally, Schedule H of Form 5500 contains information the plan's general financial positions, including total balances, contributions (by employers and employees), and expenses. The machine-readable data from the DOL contains this information in aggregate at the plan level. In addition, firms are also required to report information about the plan in Line 4(i) of this schedule. However, responses to Line 4(i) are in pdf files, and this is the information we obtain from BrightScope (who processes these pdf files into a machine-readable format). BrightScope records investment menus, asset allocations, and contributions by investment. It also adds tickers for publicly traded mutual funds and includes expense ratios and other fees of these funds, along with a description of investment styles.

Our sample consists of all firms who file a version of Form 5500 to declare non-wage benefits. We exclude firms that offer defined benefit plans (even if they offer defined contribution plans), as we do not wish to model the interaction between the two types of plans. We focus on plans sized between 100 and 5,000 participants for the main analysis in the paper; for firms that do not offer a defined contribution plan, we take the number of participants in the benefit that they do offer (usually health insurance) as the plan size. We drop plans that report matching outlays higher than $35,000 per participant, revenue sharing outlays higher than $1,000 per participant, expense ratio outlays higher than $3,000 per participant, or direct transfers higher than $10 million. This corresponds to 102 observations. We also drop 55 plans reporting negative values for one of these outlays, as these are likely data entry mistakes.

A.2. Identifying the Recordkeeper

In the model, we consider bargaining between the sponsor and the recordkeeper—treating the recordkeeper as the "main" entity involved in the negotiation. To identify the main provider, we use BrightScope data on service providers; this is based on Schedule C of Form 5500, which lists the services provided by each firm for the plan (based on a set of pre-specified service codes). If a firm is listed as providing a trustee service, we designate it the main provider. If no such firm is listed, we choose the firm listed as the recordkeeper. We then go down a pre-specified list of roles (e.g., "investment management services" and then "investment advisory services") until we find a firm that fulfills one of these roles. We have checked that variations of this order do not appreciably change the designated main providers, since the vast majority of plans have one firm designed as either a trustee or a

recordkeeper. This choice is relevant for the estimation procedure since (i) we allow costs and mean utilities to depend on the identity of the main provider, and (ii) we treat a firm as capturing the full expense ratio from funds with which it is vertically integrated.

A.3. Extracting Revenue Sharing Agreements

The primary source of revenue sharing agreements is Line 3 of Schedule C, which includes information on indirect compensation, including the type of payment, the payer and the payee, and some information about the formula for payment. There is no consistent reporting format for the formula, so we developed a series of regular expressions to extract as much information from the explanations as possible. We are left with either (i) a single number corresponding to the revenue sharing rate for that particular payer, (ii) a set of rates corresponding to multiple payers, or (iii) some indication that the revenue sharing agreement cannot be extracted from the data in Line 3.

Since revenue sharing is given in rates, we have to merge these back to the menu to compute the total amount of revenue sharing. The payer listed in Line 3 can be a fund provider or a specific fund. In the case in which the payer is a fund provider, we assume that all funds provided by that provider have the same revenue sharing rate. In the case where the payer is a fund, we attempt to match back to the plan menu. Since the fund name is listed differently in Line 3 than in BrightScope, we first match the fund name to a ticker.[20] We then match the ticker to a row in the plan menu by merging on ticker if possible, and using minimum distance matching if not. If we are unable to match on tickers, we match funds in the plan menu to Line 3 based purely on the fund provider, averaging revenue sharing rates across funds in Line 3 if needed.

If we are unable to match a fund to a revenue sharing agreement from Line 3 using the procedure above, we use the listed 12b-1 fee as the revenue sharing rate.

B. Predictive Model

B.1. Details of the Baseline Predictive Model

This section discusses implementation details for predicting revenue from expense ratios, revenue from revenue sharing agreements, and expenses due to matching outlays. Our

[20] To do so at scale, we set up three Google Custom Search Engines to search the Morningstar, Financial Times, and US News websites. We programmatically input the strings specified in the payer field into the CSEs and accept a match if at least one returns a match and the different CSEs do not disagree.

objective is to construct a model for each of these outcomes that predicts their counterfactual values under alternative plan designs. This entails two steps: first, distilling features from each plan menu into a set of variables that can serve as inputs into a model; and second, selecting a suitable model. We discuss each of these issues in turn.

Recall that a firm's plan consists of a series of investment vehicles, each with their own features. We convert these heterogenous investment menus into a series of variables that describe the menu appropriately. The variables can be characterized as either counts, fractions, and minima/average/maxima.

The count variables included in the model are number of investment options, number of mutual funds, number of target retirement date funds, number of index funds, number of S&P 500 tracker funds, number of funds offered by each fund provider, number of bond alternatives by bond investment style, number of bond alternatives by bond credit quality, number of bond alternatives by bond interest rate sensitivity and number of funds by decile of the expense ratio distribution (unweighted) within Brightscope and Morningstar fund categories. We calculate minima, unweighted averages, and maxima for the following variables: expense ratio, 12b-1 fee or revenue sharing agreements, management fee, front end load, and deferred load. Finally, we also the fraction of investment vehicles by value of Morningstar rating (for example, fraction of options that are rated as 5 stars by Morningstar). Note that by including fees and loads, along with Morningstar ratings (which are in part based on historical risk-adjusted returns), we are implicitly including the possibility of dependence of the predicted quantities on historical returns themselves—which could happen through allocations, say.

Beyond features of the investment menu, we also incorporate features of the plan itself and firm characteristics. Features of the plan include matching rates and amounts, vesting rules, total balances, total participants, active participants, number of participants with a positive balance, retired participants, participant contributions, identity of the main provider, whether the plan is collectively bargained, the date it was first opened, and participation rates. For firm characteristics, we incorporate NAICS code, state, region, and FIPS code. We also construct interactions between region and 2-digit NAICS code, and augment the dataset with marginal tax, wage, and rent data at the county level from NBER. Including these variables allows us to control for variation in observables between firms that offer and do not offer plans when predicting the latter's outcomes in the counterfactual where they do offer a plan.

We build a predictive model for revenues from expense ratios per plan participant,

Table B.1: In-Sample and Out-of-Sample Fit for Predictive Models

	RMSE		MAE	
Outcome	Training	Testing	Training	Testing
Expenses per Person ($)	17.53	60.22	8.00	24.06
Revenue Sharing per Person ($)	46.08	49.82	20.63	23.53
Employer Contribution per Person ($)	759.74	1172.11	441.20	602.15

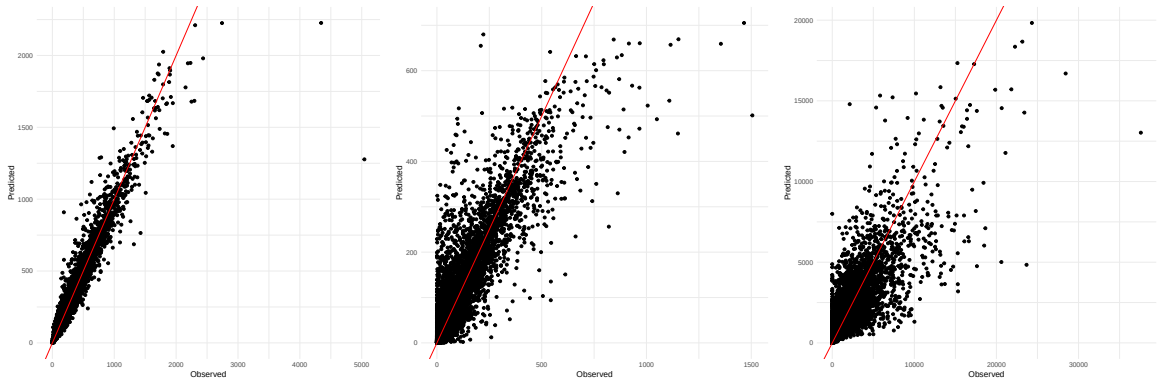

(a) Expenses (b) Revenue Sharing (c) Matching

Figure B.1: Model Fit, Full Sample

revenues from revenue sharing agreements per plan participant, and outlays from matching per plan participant using stochastic gradient boosting (Friedman, 2002).[21] We use 70% of the data for training the model, and the remainder for testing.

Table **??** provides measures of root mean squared error (RMSE) and mean absolute error (MAE) for both samples. These objects allow us to assess the quality of the predictions. The models deliver low error values in the testing dataset, regardless of the error metric. For expense ratio revenues, RMSE is $60 per person in the testing dataset, while MAE is $24. Revenues from revenue sharing are $50 and $24 per person, respectively. Finally, errors for matching expenditures are higher, but still reasonable. RMSE is $1,172 in the testing dataset, and MAE is $602. Overall, these values reassure us that the procedure is working well. Additionally, Figure B.1 reports fit for each outcome using the testing sample.

To dig deeper into the determinants of these outcomes, Figure B.2 shows variable importance plots. For all three outcomes, the main predictor is balance per person. For expenses, the next predictors in terms of importance are mean expenses (an unweighted average of expense ratios) and state. This is consistent with expense ratios mechanically

[21] This model was implemented in R via the h2o package (H2O.ai., 2020).

falling with balances, due to share classes, and with asset allocations that on aggregate are similar to a $\frac{1}{N}$ rule. For revenue sharing, the next most important predictors are the number of investment vehicles without a star rating from Morningstar, the identity of the main provider, the minimum expense ratio in the plan, and the number of funds offered by provider 50 (who is particularly likely to use revenue sharing agreements). We believe that these determinants are also reasonable. Since main providers differ in their propensity to use revenue sharing, particularly because some of them are vertically integrated, the identity of the main provider mattering makes sense. Unrated funds tend to be newer, smaller, and more expensive, so higher revenue sharing makes sense. Overall, the main predictors in these models seem sensible.

B.2. Default Options

One potential missing component in these predictions is the identity of the plan's default option. Unfortunately, we do not observe the identity of the default option directly. However, our predictive model includes a large degree of information about the funds in the plan, and it is possible that this information is already controlling for the features of the default. By law, default options must be Qualified Default Investment Alternatives (QDIAs), i.e. they must be a "life-cycle or targeted-retirement-date fund, or a balanced fund, or a professionally managed account."[22] Thus, while we cannot capture variation in the characteristics of the specific default fund in our data, but we do capture variation in characteristics of the types of funds that are likely to be default funds. As a result, if fees or the fund provider of the "potential" default options change, we expect our predictive models to take this into account, and for predicted revenue sharing, expense ratios, and matching contributions to adjust.

To test this notion, we have guessed the identity of the default for each plan by taking the QDIA with the largest balances (aggregating target date funds offered by the same provider into a single value, in case the default is the age-appropriate member of a suite of target date funds). The assumption is that defaults are important, so they are likely to have the highest asset holdings among the set of eligible funds. We can then revisit the prediction exercise augmenting the set of plan observables to include characteristics of the predicted default. In line with the arguments discussed above, we find that the correlation in predictions between the models that include this information on defaults and our baseline models to be very high: 99.4% for expenses, 93.1% for revenue sharing, and 94.7% for employer contributions.

[22] See https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/resource-center/fact-sheets/default-investment-alternatives-under-participant-directed-individual-account-plans

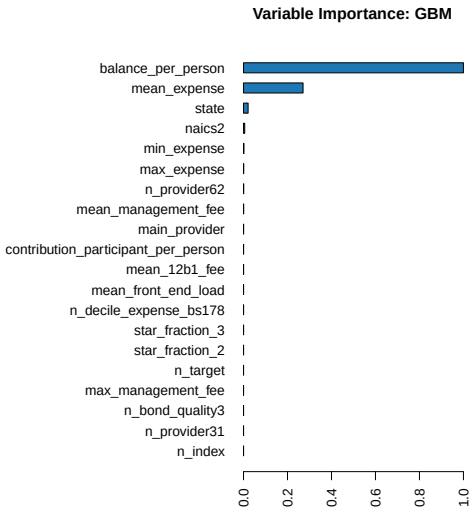

(a) Expenses

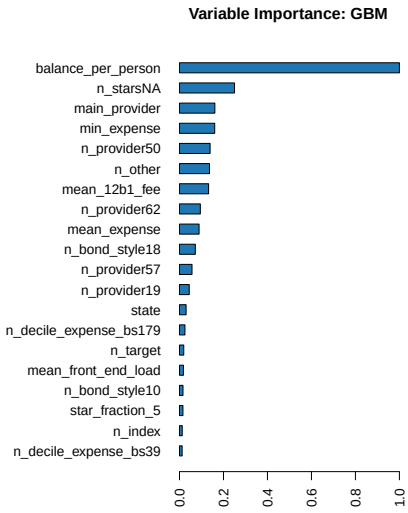

(b) Revenue Sharing

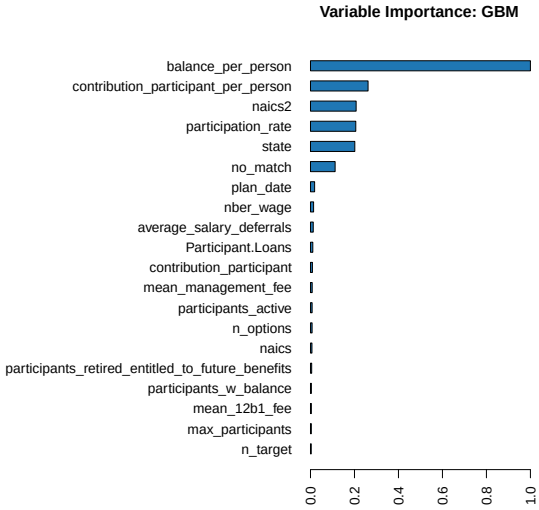

(c) Matching

Figure B.2: Relative importance of inputs into the predictive model

B.3. Predictions for Must Carry Provision Counterfactuals

For must carry provision counterfactuals, we add an a S&P 500 tracker or a suite of target date funds to all plans that do not offer them. Since adding these funds may change expenses per person, revenue sharing per person, and employer contributions per person,

we are predicting these outcomes once again after updating plan features to reflect the characteristics of the added funds.

More precisely, we take all plans that do not offer a S&P 500 tracker or a target date fund, identify the fund provider with the highest balances, and assume that this provider offers the mandated fund. We then find the average features of S&P 500 trackers or target date funds offered by this provider, and assume that the newly added fund has these features. With these assumptions, we update the variables that enter into the predictive exercise. For example, we increase the number of funds offered by the assumed provider by one, the number of funds of the decile of the expense ratio distribution where the added fund belongs to by one, and re-compute average expense ratios, 12b-1 fees, and fractions of investment vehicles by value of Morningstar rating. We then recompute our predictions of expenses per person, revenue sharing per person and employer contributions per person using these updated characteristics.

C. Estimation

C.1. Integration

We solve the integrals in (8) as follows. The integral over θ is approximated using monomial rules for Gaussian quadrature, calculated by Stroud (1971).[23] We integrate over e and $\tilde{\mathcal{R}}_s$ by simulation. In particular, we draw $D = 1808$ consideration set probabilities ρ uniformly from the grid $\{0.1, 0.2, ..., 0.9, 0.99\}$,[24] and for each draw ρ_d, draw whether each possible recordkeeper from $\mathcal{R}_s$ is in the consideration set from a Bernoulli distribution with probability ρ_d. This gives us D draws of consideration sets, $\tilde{\mathcal{R}}_{s,d}$, and a value for the probability of observing the draw, ω_d. We hold these sets fixed throughout the estimation procedure, and for each guess of parameter values, solve for choice probabilities using importance sampling. The probability sponsor s chooses recordkeeper r is

$$\sum_{d=1}^{D} \frac{\omega_d(\rho)}{\omega_d} 1[r \in \tilde{\mathcal{R}}_{s,d}] \iint \frac{\exp(\frac{1}{\sigma^S}(V^*_{rs}(\theta, \sigma^R, \kappa, C^R_r) + \delta_r))}{1 + \sum_{r' \in \tilde{\mathcal{R}}_{s,d}} \exp(\frac{1}{\sigma^S}(V^*_{r's}(\theta_s, \sigma^R, \kappa, C^R_{r'}) + \delta'_r))} dF(\theta) dF(e)$$

where $\omega_d(\rho)$ is the probability that consideration set $\tilde{\mathcal{R}}_{s,d}$ is drawn if the search cost is ρ.

[23] We use rule en_r2_07_2 in https://people.math.sc.edu/Burkardt/m_src/stroud/stroud.html.
[24] We use a multiple of the number of quadrature nodes (452).

C.2. Moments

Recall that all moments take the form $\mathbb{E}[X \cdot \mathbb{1}[s \text{ chooses } r]]$, where X are different features of the chosen plan. This expectation is equal to

$$E[X_p^*] = \frac{1}{S}\sum_{s,d}\frac{\omega_d(\rho)}{\omega_d}\sum_{r\in\tilde{\mathcal{R}}_{sd}}\sum_{p\in\mathcal{P}_{rs}} X_{rsp}\cdot \Pr[s \text{ chooses } r \text{ and } p \text{ in draw } d]. \tag{9}$$

Let $\pi_s^d \equiv \max_{r'\in\mathcal{R}_{rs}\setminus r} V^*_{r's}(\theta,\sigma^R,\kappa,C^R_{r'}) + \delta_{r'} + \sigma^S\epsilon^S_{r's}$ denote sponsor s's disagreement payoff, and define $l_{rs} \equiv \sigma^S\epsilon^S_{rs} - \pi^d_s$. When sponsor s chooses recordkeeper r and plan p, it must be the case that $V^*_{rs}(\theta,\sigma^R,\kappa,C^R_r) + \delta_r + l_{rs} \geq 0$ and

$$\begin{aligned} p &= \arg\max_{p\in\mathcal{P}_{rs}}\max_{T_p\geq 0}(\boldsymbol{\theta}_s\cdot\boldsymbol{Q}_s(p) + \delta_r + l_{rs} - T_p - \kappa\cdot T_p^2)^\eta\cdot(R_r(p) - C^R_r + \sigma^R\epsilon^R_{rs} + T_p)^{(1-\eta)} \\ &\equiv \arg\max_{p\in\mathcal{P}_{rs}} NP(l_{rs},\epsilon^R_{rs},\theta_s;\delta,\kappa,C^R_r,\eta). \end{aligned}$$

With this notation, the probability in (9) can be expressed as

$$\begin{aligned} &\Pr[s \text{ chooses } r \text{ and } p \text{ in draw } d] \\ &\quad= \Pr[s \text{ chooses } p \text{ in draw } d | s \text{ chooses } r \text{ in draw } d]\Pr[s \text{ chooses } r \text{ in draw } d] \\ &\quad= \Pr[p = \arg\max_{p\in\mathcal{P}_{rs}} NP(l_{rs},\epsilon^R_{rs},\theta_s;\delta,\kappa,C^R_r,\eta) | l_{rs} \geq -V^*_{rs}(\theta_s,\sigma^R,\kappa,C^R_r) - \delta_r]] \\ &\qquad\qquad\times s_{rs}(\epsilon^R_{rs},\theta_s;\sigma^R,\kappa,C^R,\sigma^S,p,\delta) \\ &\quad= \int\int\int_{l_{rs}\geq -V^*_{rs}(\theta,\sigma^R,\kappa,C^R_r)-\delta_r} \mathbb{1}[p = \arg\max_{p\in\mathcal{P}_{rs}} NP(l,e,\theta;\delta,\kappa,C^R_r,\eta)] \\ &\qquad\qquad\times s_{rs}(e,\theta;\sigma^R,\kappa,C^R,\sigma^S,p,\delta)dF(l)dF(\theta)dF(e). \end{aligned} \tag{10}$$

To solve the integral in (10), we need the distribution of $l_{rs} \equiv \sigma^S\epsilon^S_{rs} - \pi^d_s$ conditional on being greater than $-V^*_{rs}(\theta,\sigma^R,\kappa,C^R_r) - \delta_r$. We begin by deriving the distribution of the disagreement payoff $\pi^d_s \equiv \max_{r'\in\mathcal{R}_{rs}\setminus r} V^*_{r's}(\theta,\sigma^R,\kappa,C^R_{r'}) + \delta_{r'} + \sigma^S\epsilon^S_{r's}$. Note $V^*_{r's}(\theta,\sigma^R,\kappa,C^R_{r'}) + \delta_r + \sigma^S\epsilon^S_{rs}$ distributes Gumbel with location $V^*_{r's}(\theta,\sigma^R,\kappa,C^R_{r'}) + \delta_r$ and scale σ^S. Since the maximum of independent Gumbel draws with the same scale parameter is also Gumbel, π_d distributes Gumbel with location

$$\sigma^S\log\left(1 + \sum_{r'\in\mathcal{R}_{rs}\setminus r}\exp\left(\frac{1}{\sigma^S}(V^*_{r's}(\theta,\sigma^R,\kappa,C^R_{r'}) + \delta_{r'})\right)\right)$$

and scale σ^S. Since the difference between independent Gumbels with the same scale parameter is Logistic, $\sigma^S \epsilon^S_{rs} - \pi^d_s$ distributes Logistic with location

$$\mu \equiv -\sigma^S \log\left(1 + \sum_{r' \in \mathcal{R}_{rs} \setminus r} \exp\left(\frac{1}{\sigma^S}(V^*_{r's}(\theta, \sigma^R, \kappa, C^R_{r'}) + \delta_{r'})\right)\right)$$

and scale σ^S. This is the unconditional distribution of $\sigma^S \epsilon^S_{rs} - \pi^d_s$. To derive the conditional distribution, note that for $s > a$,

$$\Pr[\sigma^S \epsilon^S_{rs} - \pi^d_s \leq s | \sigma^S \epsilon^S_{rs} - \pi^d_s \geq a] = \frac{\Pr[a \leq \sigma^S \epsilon^S_{rs} - \pi^d_s \leq s]}{\Pr[\sigma^S \epsilon^S_{rs} - \pi^d_s \geq a]}$$

$$= \left[\frac{1}{1+\exp(-\frac{s-\mu}{\sigma^S})} - \frac{1}{1+\exp(-\frac{a-\mu}{\sigma^S})}\right] \cdot \left[1 - \frac{1}{1+\exp(-\frac{a-\mu}{\sigma^S})}\right]^{-1} = \frac{1-\exp(-\frac{a-\mu}{\sigma^S})}{1+\exp(-\frac{s-\mu}{\sigma^S})}.$$

so if $y = \Pr[\sigma^S \epsilon^S_{rs} - \pi^d_s \leq s | \sigma^S \epsilon^S_{rs} - \pi^d_s \geq a]$, $s = \mu + \sigma^S \left[\log(y + \exp(\frac{a-\mu}{\sigma^S})) + \log(1-y)\right]$. To integrate numerically over this conditional distribution, we obtain D uniform (0,1) draws as y and transform them following this last expression, with $a = -V^*_{rs}(\theta, \sigma^R, \kappa, C^R_r) - \delta_r$.

Finally, we estimate the model on a random subsample of 2,000 firms for each of the four groups.

D. Additional Tables and Figures

D.1. Additional Descriptive Statistics

Table D.1 compares observable characteristics between employers that offer and do not offer a defined contribution plan. Beyond the size differences documented earlier, we find that employers who offer DC plans pay higher wages and are located in counties with slightly higher marginal tax rates. Moreover, the pool of employers that offers a DC plan has fewer mining and transportation companies, and more utilities and manufacturing firms. These characteristics are used as inputs into the predictive models discussed in Section B.

Table D.2 presents correlations of plan characteristics of offered plans. For example, there is a positive correlation between having a high number of options and offering at least one target date fund, index fund, or S&P 500 tracker.

Table D.3 reports geographic coverage and market shares, by number of plans, for the

Table D.1: Comparison Across Firms that Offer and Do Not Offer DC Plans

Variable	No Plan	Has Plan	t-Statistic
Number of Employees	489.85	552.42	-8.49
	(6.42)	(3.61)	
NAICS-County Average Wages	63,035.92	65,821.79	-5.67
	(368.12)	(208.28)	
NBER County Tax	26.12	26.61	-5.54
	(0.07)	(0.04)	
Pr[Unclassified]	0.002	0.001	1.72
	(0.000)	(0.000)	
Pr[Agriculture, Forestry, Fishing and Hunting]	0.014	0.011	2.3
	(0.001)	(0.001)	
Pr[Mining, Quarrying, and Oil and Gas Extraction]	0.116	0.068	15.63
	(0.003)	(0.001)	
Pr[Utilities]	0.167	0.201	-8.78
	(0.003)	(0.002)	
Pr[Construction]	0.157	0.172	-3.93
	(0.003)	(0.002)	
Pr[Manufacturing]	0.265	0.309	-9.73
	(0.004)	(0.002)	
Pr[Wholesale Trade]	0.163	0.155	2.18
	(0.003)	(0.002)	
Pr[Retail Trade]	0.068	0.047	8.77
	(0.002)	(0.001)	
Pr[Transportation and Warehousing]	0.045	0.036	4.19
	(0.002)	(0.001)	
Pr[Information]	0.003	0.001	4.37
	(0.001)	(0.000)	

Comparison of means across firms that offer and do not offer a DC plan. NAICS-County Average Wages are the average wages in the county where the firm is located for firms sharing the 6-digit NAICS code of the firm, according to the BLS. NBER County Tax is the marginal tax rate for a 45 year old married worker with two dependents and no other sources of income but the average wage referenced above, according to NBER's TAXSIM. The remaining columns denote the share of each group that belongs to each industry, as determined by the first digit of the employer's NAICS code.

Table D.2: Correlations among plan characteristics

Feature	High Options	Has Target Date Fund	Has Index Fund	Has S&P 500 Tracker	Expenses per Person	Employer Contribution per Person
High Options	1	0.42	0.21	0.20	-0.02	-0.01
Has Target Date Fund		1	0.14	0.06	-0.03	-0.02
Has Index Fund			1	0.49	-0.12	0.02
Has S&P 500 Tracker				1	-0.06	0.00
Expenses per Person					1	0.67
Employer Contribution per Person						1

Recordkeeper	Number of States	Market Share (by Number of Plans)
Fidelity	50	16.4%
Principal	46	7.4%
Empower	47	6.6%
ADP	40	5.2%
John Hancock	48	4.9%
Vanguard Group	38	3.1%
Wells Fargo	39	3.0%
MassMutual	37	2.7%
Transamerica	40	2.3%
Ascensus	33	2.2%
Charles Schwab	30	2.2%
American Funds (Capital Group)	35	2.0%
Nationwide	37	2.0%
Bank of America Merrill Lynch	31	1.8%
Voya Financial (ING)	29	1.8%
T. Rowe Price	26	1.4%
Newport Group Securities, Inc.	29	1.4%
TIAA-Nuveen	25	1.1%
The Standard	20	1.1%

Table D.3: State coverage and market share of top recordkeepers

top 19 recordkeepers in our estimation sample. Table D.4 shows average characteristics of plans by recordkeeper. We find that there is substantial cross-recordkeeper heterogeneity in the types of plans they offer. Moreover, coefficients of variation within a recordkeeper are substantial, suggesting that each recordkeeper offers many types of plans. This is relevant for our analysis for two reasons. First, it provides auxiliary evidence that there is scope for negotiation with a recordkeeper when forming a plan. Second, it shows that the geographic variation identified in Table D.3 can be a useful part of the variation in choice sets. Of course, sponsor-specific covariates provides another source of heterogeneity in choice sets.

Figure D.1 provides more details about the asset allocation. Panel (a) shows the share of

Recordkeeper	High Options (%)	Target Date Fund (%)	Index Fund (%)	S&P 500 Tracker (%)	Expenses ($)	Expenses (CV, $)	Employer Contribution ($)	Contribution (CV, $)	Revenue Sharing ($)	Revenue Sharing (CV, $)
Fidelity	73.8%	93.1%	93.5%	72.7%	177.10	0.37	1230.13	0.40	67.14	0.61
Principal	42.3%	74.6%	89.6%	78.8%	234.05	0.39	1205.83	0.41	135.71	0.62
Empower	50.5%	75.8%	90.7%	65.5%	181.51	0.41	1238.08	0.41	32.12	1.20
ADP	41.2%	80.4%	46.4%	32.9%	256.40	0.33	1079.90	0.58	75.05	0.73
John Hancock	46.2%	72.5%	86.4%	57.5%	265.57	0.26	1219.54	0.40	162.30	0.44
Vanguard Group	63.0%	89.6%	99.0%	66.1%	73.77	0.67	1249.60	0.37	34.36	0.64
Wells Fargo	24.2%	71.3%	88.6%	31.9%	166.42	0.45	1315.50	0.34	32.17	1.17
MassMutual	39.9%	70.6%	80.6%	60.1%	210.45	0.42	1269.54	0.36	53.65	0.92
Transamerica	40.7%	61.2%	81.2%	37.7%	176.85	0.49	1251.29	0.39	40.93	0.95
Ascensus	33.0%	54.5%	66.9%	51.1%	211.78	0.45	1222.89	0.43	53.88	0.89
Charles Schwab	27.8%	54.0%	95.8%	66.3%	141.95	0.49	1354.02	0.33	18.01	1.22
American Funds (Capital Group)	50.1%	92.6%	34.2%	18.3%	225.01	0.42	1183.29	0.41	118.82	0.51
Nationwide	72.0%	74.8%	95.2%	67.0%	181.96	0.47	1205.23	0.42	39.79	1.06
Bank of America Merrill Lynch	25.9%	35.9%	82.9%	60.1%	213.31	0.38	1310.94	0.34	46.11	0.88
Voya Financial (ING)	36.6%	78.8%	84.1%	32.2%	185.09	0.51	1228.41	0.41	67.52	0.84
T. Rowe Price	62.3%	88.9%	92.6%	54.0%	168.90	0.23	1273.34	0.34	116.25	0.40
Newport Group Securities, Inc.	37.4%	60.6%	87.8%	64.6%	148.22	0.56	1242.21	0.39	27.68	1.32
TIAA-Nuveen	85.6%	92.4%	97.2%	44.2%	159.92	0.34	1289.63	0.26	98.19	0.41
The Standard	27.9%	48.2%	95.1%	78.7%	132.34	0.42	1282.61	0.36	22.14	1.22
Fringe	36.0%	65.5%	82.8%	54.8%	174.36	0.40	1237.52	0.37	39.09	1.02

Table D.4: Choice Set Variation by Recordkeepers

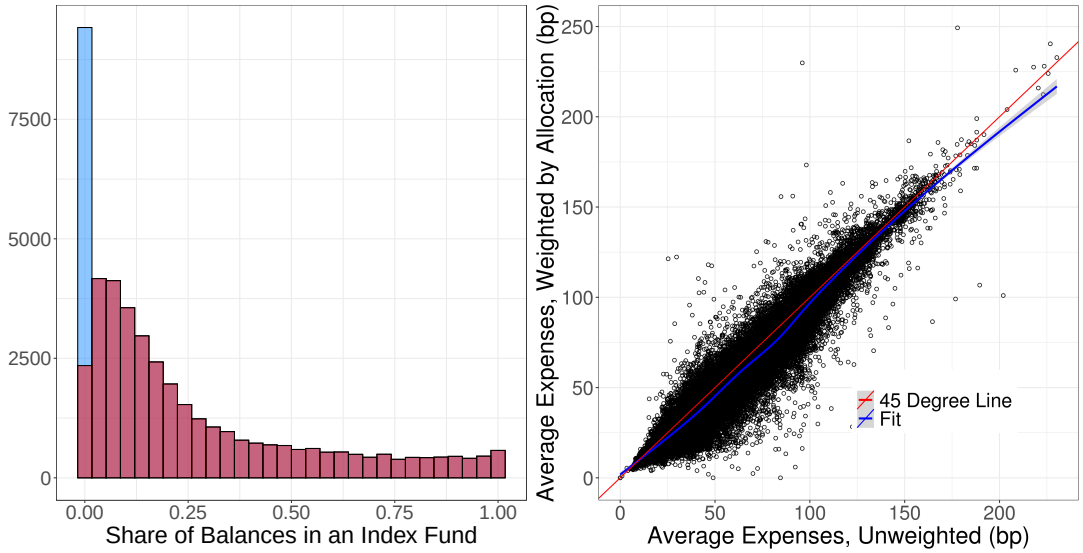

(a) Share in index fund

(b) Weighted vs. unweighed expense ratios

Figure D.1: Information about asset allocation

assets in an index fund across all plans. The section of the histogram in blue illustrates plans that do not have an index fund at all (in which case the share is mechanically zero), and the remainder of the histogram shows the distribution for plans with an index fund. Panel (b) shows a scatter of the weighted expense ratio against the unweighted expense ratio for all plans, along with a locally linear best fit line. The best fit line is very close to the 45° line (but slightly lower), suggesting that workers do not seem to be weighing low expense ratio funds more in their allocation.

D.2. Descriptive Results Related to Identification

Table D.5 shows evidence that properties of the plan respond to proxies of the outside options. We construct proxies of the outside option by looking at the set of plans offered by non-chosen recordkeepers, and selecting plans with extremum values of different characteristics. In Columns (1)–(3), we use the plan with the minimum revenue sharing as our proxy for the outside option. In Columns (4)–(6), we use the plan with the minimum expenses. In Columns (7)–(9), we use the plan that induces the minimum employer contribution, under the assumption that firms prefer to avoid employer contributions. We then regress the transfer of the chosen plan, the total revenue of the chosen recordkeeper (i.e., the transfer plus the revenue sharing), and the expenses of the chosen plan on characteristics of the

Table D.5: Correlation between proxies of the second-best option and values of the choice set.

	Minimum Revenue Sharing Plan			Minimum Expenses Plan			Min. Employer Contribution Plan		
	Transfer	Total RK Comp	Expenses	Transfer	Total RK Comp	Expenses	Transfer	Total RK Comp	Expenses
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Revenue Sharing	-0.55	-1.5	-1.4	0.10	0.56	1.5	-0.01	-0.12	-0.26
	(0.15)	(0.23)	(0.33)	(0.05)	(0.07)	(0.08)	(0.01)	(0.02)	(0.02)
Expenses	0.04	0.36	1.0	0.02	0.54	2.1	0.04	0.26	0.68
	(0.02)	(0.02)	(0.01)	(0.04)	(0.07)	(0.10)	(0.007)	(0.010)	(0.01)
Employer Contribution	0.008	0.009	0.02	0.008	0.01	0.02	0.008	0.01	0.07
	(0.002)	(0.002)	(0.002)	(0.0009)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)
Observations	43,547	43,547	43,547	43,547	43,547	43,547	43,547	43,547	43,547

Columns (1)–(3) use the minimum revenue sharing plan as the proxy for the outside option. Columns (4)–(6) use the minimum expense plan, and Columns (7)–(9) use the plan with minimum employer contributions. All units are in thousands of dollars. All regressions include size bucket fixed effects.

this purported second-best option. So, for example, columns (1) regresses transfers on the revenue sharing payments, expense payments, and employer contribution payments of the minimum revenue sharing plan.

For concreteness, consider Columns (1)–(3), where we take the stance that the minimum revenue sharing plan from a rival recordkeeper is the second-best option. We would expect that all of the following would improve the second-best option of the sponsor: an increase in revenue sharing, since the sponsor could extract a larger surplus from this recordkeeper; a decrease in expenses; or a decrease in employer contributions induced by this plan (under the assumption that these contributions are a cost to the sponsor). We generally see that improvements in the outside option are correlated with reductions in total recordkeeper compensation and expenses. The effect on transfers themselves would be ambiguous in the model, but we see that an increase in revenue sharing is also correlated with a reduction in transfers. Note that the effects of employer contributions from the posited outside option are muted: that this only corresponds to a weakening of the outside option to the extent that sponsors want to avoid paying out contributions, and this could be a worse proxy for the payoff from the second-best option.

Columns (4)–(6) repeat this exercise with the minimum expense plan as the proxy for the outside option. Here, the correlation in the outcomes with expenses and employer contributions are similar, but the correlations with revenue sharing have the opposite sign as in Column (1)–(3). We do not have a good explanation for this opposite correlation, other than to note that changes in revenue sharing of the minimum expense plan from a rival might not be a good proxy for changes in the outside option. The results in Columns (7)–(9) are similar in sign to those in Columns (1)–(3), although they are smaller in magnitude. Again,

the minimum employer contribution plan might not be a good proxy for the outside option. Overall, though, these correlations paint a picture that the transfers (and characteristics) of the chosen plan are responses to potential measure of the outside option. The structure of the model provides a better measure of the outside option and a quantification of how correlations like these translate to η and κ.

We also provide some descriptive evidence in favor of limited search. Geographic variation in choice sets, induced by which recordkeepers operate in which states, induces variation in the average expenses that a firm faces for its potential plans. This variation is related to the probability of offering a plan: a simple regression of offering a plan on the average expenses shows that a 1 bp increase in this mean is correlated with a 1.1 pp reduction in the probability that a firm offers a plan (s.e. 0.08). Taking this as a proxy for the benefit of offering a plan, we consider what happens when the average expenses in the choice set are especially low (this benefit is especially high). For sponsors with a choice set in the bottom 10%, 5%, 2%, and 1% of expenses, the probability of offering a plan is 0.86, 0.89, 0.92, and 0.96, respectively. This suggests that there is no evidence that sponsors with especially high benefits of offering a plan fail to do so, and the model interprets this as limited evidence for search.

D.3. Estimated Recordkeeper Costs

Table D.6 reports estimated costs by recordkeeper. We make a few comments about these results. First, there is some evidence of heterogeneity in costs across recordkeepers. Second, while the cost estimates for Group 1 are noisily estimates, there is some specialization: the recordkeepers that are cheapest for one group are not necessarily the cheapest ones for another. Third, while economies of scale within-recordkeeper are not universal, it is often the case that costs for Group 4 within-recordkeeper are lower than those for Groups 1 or 2, which provides some evidence in favor of claims by the industry. (Note that market shares are determined by preferences for recordkeepers too, which are group-specific, so the highest share recordkeeper is not necessarily the cheapest one in each group.) We finally note that we sometimes estimate costs to be negative for certain groups. We interpret these costs as representing other sources of revenue a recordkeeper might earn: workers in large firms might be more likely to start other relationships with the recordkeeper outside the defined contribution plan, say (e.g., through an general brokerage account).

	Group 1 ≤ 200	Group 2 201–500	Group 3 501–1000	Group 4 > 1000
Fidelity	0.02	0.03	0.03	-0.04
	(0.13)	(0.02)	(0.01)	(0.01)
Principal	0.30	0.28	0.16	0.08
	(0.13)	(0.01)	(0.02)	(0.01)
Empower	0.02	0.03	0.02	-0.01
	(0.13)	(0.01)	(0.01)	(0.00)
ADP	-0.01	0.02	-0.05	-0.05
	(0.12)	(0.01)	(0.01)	(0.00)
John Hancock	-0.09	0.01	0.03	-0.04
	(0.10)	(0.02)	(0.02)	(0.04)
Vanguard Group	0.04	0.04	0.05	0.02
	(0.11)	(0.02)	(0.02)	(0.01)
Wells Fargo	0.12	0.06	0.07	0.03
	(0.11)	(0.02)	(0.01)	(0.01)
MassMutual	0.03	0.04	0.05	-0.03
	(0.13)	(0.02)	(0.01)	(0.01)
Transamerica	0.08	0.18	0.22	0.07
	(0.11)	(0.01)	(0.01)	(0.01)
Ascensus	0.02	0.06	0.05	0.01
	(0.14)	(0.02)	(0.01)	(0.01)
Charles Schwab	0.00	0.02	0.02	-0.00
	(0.13)	(0.02)	(0.01)	(0.01)
American Funds (Capital Group)	0.03	0.05	-0.02	0.00
	(0.13)	(0.01)	(0.02)	(0.01)
Nationwide	0.19	0.06	0.03	0.04
	(0.13)	(0.01)	(0.01)	(0.01)
Bank of America Merrill Lynch	-0.10	0.03	0.01	-0.04
	(0.12)	(0.02)	(0.01)	(0.01)
Voya Financial (ING)	-0.09	0.02	0.00	-0.10
	(0.13)	(0.02)	(0.00)	(0.00)
T. Rowe Price	0.02	0.02	0.08	0.01
	(0.13)	(0.01)	(0.04)	(0.03)
Newport Group Securities, Inc.	0.12	0.06	0.11	0.04
	(0.13)	(0.02)	(0.01)	(0.01)
TIAA-Nuveen	0.01	0.00	0.02	-0.10
	(0.12)	(0.03)	(0.03)	(0.04)
The Standard	0.15	0.16	0.06	0.04
	(0.12)	(0.02)	(0.01)	(0.02)
Fringe	-0.10	0.00	0.01	-0.07
	(0.13)	(0.02)	(0.01)	(0.01)

Table D.6: Recordkeeper cost coefficients, in thousands of dollars per worker. Standard errors are in parentheses.

	Group 1		Group 2		Group 3		Group 4	
	Baseline	Adjusted	Baseline	Adjusted	Baseline	Adjusted	Baseline	Adjusted
A. Frictions								
ρ	0.84	0.82	1.00	1.00	0.98	0.98	1.00	1.00
	(0.05)	(0.10)	(0.05)	(0.02)	(0.15)	(0.16)	(0.03)	(0.03)
η	0.65	0.67	0.40	0.43	0.31	0.33	0.85	0.85
	(0.03)	(0.02)	(0.04)	(0.04)	(0.08)	(0.08)	(0.01)	(0.04)
Marginal Effect of κ (at \$50)	1.09	1.10	1.00	1.00	1.08	1.07	1.66	1.87
	(0.05)	(0.05)	(0.02)	(0.01)	(0.07)	(0.08)	(0.03)	(0.06)
B. Sensitivity to Characteristics (θ)								
Means								
Expenses per Person	-0.52	-0.46	-0.64	-0.57	-0.89	-0.85	-1.52	-2.12
	(0.03)	(0.03)	(0.01)	(0.01)	(0.03)	(0.03)	(0.04)	(0.08)
Employer Contributions	-1.08	-0.90	-1.62	-1.73	-1.65	-1.66	-0.53	-3.15
	(0.21)	(0.20)	(0.06)	(0.04)	(0.12)	(0.08)	(0.03)	(0.50)
Has High # of Options	0.00	0.00	0.02	0.02	0.02	0.02	0.02	0.02
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Has Target Date Fund	0.06	0.06	3.26	3.38	2.20	2.25	0.02	0.06
	(0.02)	(0.01)	(0.14)	(0.33)	(0.36)	(0.20)	(0.00)	(0.01)
Has Index Fund	3.40	2.39	3.94	4.05	2.73	2.79	0.02	0.73
	(3.09)	(2.86)	(0.90)	(0.47)	(0.51)	(0.28)	(0.01)	(0.16)
Has S&P 500 Tracker	0.04	0.04	0.01	0.01	0.10	0.05	-0.03	-0.06
	(0.01)	(0.00)	(0.00)	(0.00)	(0.02)	(0.02)	(0.00)	(0.00)
Standard Deviations								
Constant	4.48	3.93	3.81	4.00	2.65	2.64	3.07	20.00
	(3.43)	(3.59)	(0.34)	(0.35)	(0.47)	(0.48)	(1.85)	(7.20)
Expenses per Person	0.00	0.00	0.00	0.00	0.02	0.01	0.02	0.02
	(0.91)	(0.87)	(0.50)	(0.11)	(0.15)	(0.39)	(0.77)	(2.25)
Employer Contributions	0.00	0.00	0.06	0.11	0.01	0.00	1.59	5.44
	(1.58)	(1.98)	(0.38)	(0.13)	(0.47)	(0.24)	(0.08)	(0.88)
Has High # of Options	0.00	0.00	0.04	0.04	0.03	0.02	0.06	0.19
	(0.47)	(0.29)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.04)
Has Target Date Fund	1.02	0.89	4.12	4.28	2.90	2.99	0.22	0.58
	(0.55)	(0.25)	(0.27)	(0.54)	(0.33)	(0.26)	(0.07)	(0.18)
Has Index Fund	6.43	5.18	3.82	3.99	2.82	2.81	0.00	0.77
	(3.14)	(3.62)	(0.46)	(0.32)	(0.37)	(0.47)	(0.65)	(0.32)
Has S&P 500 Tracker	1.03	0.97	0.00	0.00	0.62	0.65	0.06	0.07
	(0.30)	(0.38)	(0.09)	(0.11)	(0.27)	(0.16)	(0.02)	(0.01)
σ_S	0.23	0.25	0.06	0.07	0.06	0.07	1.13	1.25
	(0.03)	(0.04)	(0.00)	(0.00)	(0.00)	(0.01)	(0.09)	(0.18)
C. Recordkeeper Costs								
σ_R	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.02
	(0.23)	(0.12)	(0.03)	(0.01)	(0.02)	(0.03)	(0.01)	(0.01)
Share-Weighted Average Cost	0.00	-0.01	0.05	0.05	0.04	0.04	-0.02	-0.03
	(0.12)	(0.06)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)

Table D.7: Parameter estimates: baseline vs. adjusted expense specification. Parameters in thousands of dollars. Standard errors in parentheses.

D.4. Using Adjusted Returns as a Plan Characteristic

Table D.7 combines the baseline results with results obtained by replacing the expenses per person variable with a adjusted expenses metric. This metric adjusts for the fact that different funds have exhibited different returns in the past, and sponsors may make decisions as if this performance is persistent. More precisely, to form adjusted expenses we estimate α for each

fund using a one-factor CAPM model, compute the average α for the plan, multiply it by plan balances, and subtract this from expenses per person. Thus, if a plan has positive alpha on average, we would reduce expenses to compensate for this. The correlation between expenses and adjusted expenses is 0.81, which is high and suggests that higher expenses are not systematically compensation for excess returns.

Nevertheless, since this correlation is not 1, we still re-estimate the model replacing expenses per person with adjusted expenses per person. Table D.7 reports results for both the baseline and the adjusted amounts. In general, the results are similar. We find economically similar values of the frictions. Plan characteristics are generally valued similarly to the baseline too. The two exceptions are that we find that Group 4 firms have a somewhat larger sensitivity to adjusted expenses than they do to expenses, and we also estimate a much larger sensitivity to employer contributions. Of additional note is that estimates of some of the standard deviations of random coefficients are larger for Group 4 than in the baseline. Overall, though, we still see that the main message applies: firms in Group 4 are significantly more sensitive to expenses than ones in Groups 1–3, at the expense of direct preferences for other plan characteristics.

E. Minimum Investments and Plan Restrictions

An important modeling assumption is that recordkeepers and sponsors bargain over the entirety of the set of plans offered by the recordkeeper. Thus, there are no restrictions to the set of plans a recordkeeper can offer a small employer. The only such potential constraint we are aware of comes from share class minimums: less expensive share classes are advertised to be restricted to larger investment sizes. However, we are aware of anecdotal evidence that investment providers are willing to waive share class minimums for employer-sponsored retirement plans. This section studies this phenomenon systematically.

For each of the ten investment providers with the highest balances in our dataset, we identified their highest balance fund. We then found the different share class versions of the same fund: funds that correspond to the same underlying investment but simply charge different fees. For example, one top ten investment provider is Vanguard, and their highest balance fund is VINIX, the institutional share class version of their S&P 500 tracker. In our data we also observe VOO, VFIAX, VIIIX and VFFSX, which also track the S&P 500 but charge different expense ratios. We then manually went through the prospectus for each fund to extract the investment minimum. Note that this is not necessarily the same value

one would see from searching the ticker on Morningstar, for instance: the prospectus often lists different minima for different types of investment vehicles. Table E.1 lists these share classes, the minimum investment amounts, some properties of the distribution of actual investment amounts, and the proportion of plans that have investment amounts below the listed minimum. For this latter outcome, we report both the frequency that the invested amount in the fund is below the minimum and the frequency that invested amounts summed across all funds owned by the provider are below the fund's minimum.[25]

There are three takeaways from this analysis. First, investment minimums are often very low. Only a handful of share classes have a minimum above even a nominal amount ($10K), and many are identically zero. This is partly because minimums are lower for employer-sponsored retirement plans. Thus, in theory any plan could include these funds.

Second, when minimums are large, they do not seem to bind. For every fund in the list with an investment minimum of $1M, over 25% of plans with that fund have less than the minimum in the plan. Moreover, it is often substantially less: the first quartile is usually much lower than the minimum. Even if the minimum investment amount were defined by adding all the investments in a provider's funds within a plan, we observe frequent discretionary violations of the minimum (see final column).

Third, there is substantial overlap in investment amounts for different share classes. The Institutional share classes for Vanguard are a good example of this. Plans with Institutional Plus have larger investment amounts in the fund than those with Institutional, but the first quartile of investment amounts for Instituational Plus is lower than the mean and third quartiles for those with the Institutional class. Plans with Institutional Select classes (the cheapest one with the largest minimum) are actually smaller. In general, we have not found any evidence that plans are binned into the "appropriate" share class as a function of investment amounts.

Our reading of the prospectuses is that this is consistent with institutional and contractual details of the setting. Almost all include a degree of discretion on the part of the fund provider for waiving the investment minimum. For instance, Vanguard notes that the "minimum investment amount required to open and maintain a Fund account for Institutional Select Shares is generally $5 billion" and that "Investment minimums may differ for certain categories of investors,[26] without providing further specifics. State Street Global Advisers

[25] One fund, from Dodge & Cox, only had a single share class, so we omit this from the table. We also drop all share classes present in fewer than 25 plans.

[26] See Page 5 of https://personal.vanguard.com/pub/Pdf/spi1785.pdf?2210120649 and Page 22 of https://personal.vanguard.com/pub/Pdf/i854.pdf?2210120672.

Ticker	Description	Expense Ratio (bp)	N	Minimum Investment	Balance ($M)			Share Below Cutoff	
					Mean	Q1	Q3	Fund	Provider
BlackRock Global Allocation Fund									
MCLOX	Investor C	186	39	$0	0.12	0.02	0.11	-	-
MRLOX	Class R	144	317	$0	0.16	0.02	0.20	-	-
MDLOX	Investor A	112	1,496	$0	0.45	0.04	0.48	-	-
MALOX	Institutional	84	610	$0	4.33	0.09	1.77	-	-
MKLOX	Class K	76	310	$0	0.31	0.03	0.36	-	-
Fidelity Contrafund Fund									
FCNTX	Contrafund	74	3,765	$0	5.25	0.26	2.64	-	-
FCNKX	Class K	65	846	$0	19.10	3.03	19.77	-	-
JP Morgan SmartRetirement 2030									
JSMAX	Fund-A	86	1,352	$0	0.67	0.07	0.63	-	-
JSMSX	Fund-I	70	818	$1M	4.58	0.68	3.81	31.05%	5.75%
Principal Lifetime 2030 Fund									
PXASX	R-1 Class	159	37	$0	0.26	0.09	0.29	-	-
PTCMX	R-3 Class	128	174	$0	0.49	0.10	0.59	-	-
PTCAX	Class A	109	66	$1K	0.44	0.15	0.49	3.03%	0.00%
PTCSX	R-4 Class	109	65	$0	1.22	0.34	1.68	-	-
PTCPX	R-5 Class	97	136	$0	1.86	0.29	1.98	-	-
PMTIX	Institutional	72	124	$0	5.75	1.16	5.14	-	-
T. Rowe Price Retirement 2030 Fund									
RRTCX	R Class	117	1,187	$0	0.61	0.05	0.50	-	-
PARCX	Advisor	92	860	$0	1.57	0.15	1.34	-	-
TRRCX	Investor	67	4,106	$1K	3.63	0.21	2.81	0.63%	0.00%
TRPCX	I Class	53	62	$500K	2.27	0.28	2.13	32.26%	3.23%
TIAA-CREF Lifecycle 2040 Fund									
TCLOX	Retirement	69	615	$0	1.03	0.12	1.29	-	-
TCZPX	Premier	59	76	$5M	3.88	0.82	3.70	88.16%	2.63%
TCOIX	Institutional	44	145	$2M	6.13	0.38	6.12	41.38%	4.14%
Vanguard Institutional Index Fund									
VFINX	Investor	14	2,855	$0	1.51	0.13	0.78	-	-
VOO	ETF	4	140	$0	1.66	0.07	0.66	-	-
VFIAX	Admiral	4	8,993	$0	2.45	0.27	2.45	-	-
VINIX	Institutional	3.5	2,488	$5M	32.33	4.65	33.41	25.64%	16.12%
VIIIX	Institutional Plus	2	141	$100M	220.32	15.91	286.30	46.81%	32.62%
American Funds EuroPacific Growth Fund									
RERAX	Class A	159	114	$250	0.16	0.01	0.13	3.51%	0.88%
RERBX	Class R-4	157	289	$250	0.16	0.01	0.17	1.73%	0.00%
RERCX	Class F-1	113	2,389	$250	0.53	0.05	0.50	0.84%	0.08%
AEGFX	Class R-6	85	25	$250	0.25	0.03	0.12	0.00%	0.00%
REREX	Class R-3	83	2,146	$250	1.83	0.15	1.48	0.51%	0.19%
AEPGX	Class R-5	82	2,977	$250	0.58	0.05	0.47	0.84%	0.20%
RERFX	Class R-2	53	1,086	$250	2.79	0.16	2.22	0.09%	0.09%
RERGX	Class R-1	49	7,147	$250	4.10	0.15	2.41	0.64%	0.18%
PIMCO Total Return Fund									
PTTCX	Class C	164	52	$10K	0.06	0.01	0.08	32.69%	25.00%
PTRRX	Class R	114	686	$0	0.22	0.02	0.23	-	-
PTTAX	Class A	89	2,480	$10K	0.38	0.04	0.43	12.94%	9.92%
PTRAX	Administrative Class	80	1,274	$1M	1.35	0.11	0.99	75.43%	73.08%
PTTPX	Class I-2	65	59	$1M	3.47	0.07	1.37	69.49%	64.41%
PTTRX	Institutional Class	55	3,934	$1M	5.66	0.16	2.94	53.76%	51.19%

Table E.1: Alternative share classes for most frequent funds in the dataset.

regularly notes in their propectuses that "The investment minimum may be modified, waived or reduced for certain types of investors (e.g., 401(k) or 403(b) plans)."[27] T. Rowe Price notes that "the I Class required a $500,000 minimum investment per fund per account registration, although the minimum is generally waived or reduced for financial intermediaries, eligible retirement plans..."[28] Dodge & Cox note that "The Funds, in their sole discretion, reserve the right to modify or waive minimum investment amounts for certain financial intermediaries [...] For example, the Funds may waive or lower the minimum investment amount for certain financial intermediaries that use the Funds as part of an asset allocation program, certain retirement plans, and accounts that hold the Funds in omnibus name."[29]

To summarize, we find that minimums are often irrelevant empirically. Moreover, it is not just technically feasible but also common practice codified in statutory documents to ignore these minima. We believe this provides justification for our assumption that all firms have access to any plan. Fidelity could offer the cheap share class to smaller plans by waiving investment minima, but the employer and Fidelity together choose not to: this is the role of the model.

[27] See Page 33 of https://www.ssga.com/us/en/institutional/ic/resources/doc-viewer\#svspx&prospectus.

[28] See https://www.troweprice.com/literature/public/country/us/language/en/literature-type/prospectus/sub-type/mf?productCode=RPC&currency=USD.

[29] See https://www.dodgeandcox.com/content/dam/dc/us/en/pdf/prospectuses/dc_statutory_prospectus.pdf.