The Design of Retirement Plan Match Schedules*

Guillermo Carranza Taha Choukhmane Fiona Greig
Cormac O'Dea Lawrence Schmidt

September 2025

Please do not cite or circulate

Abstract

We use survey responses to hypothetical scenarios linked with administrative 401(k) data to investigate the efficiency and equity implications of employer 401(k) match formulas. We find that (i) survey responses can be used to accurately predict saving responses to both observed and hypothetical plans, (ii) saving is inelastic to the match rate, (iii) non-elective contributions do not crowd out saving. These patterns imply that (iv) plans combining lower match rates with non-elective contributions generate higher savings and more equitable match distributions, and (v) many existing plans—including safe-harbor formulas—are strictly dominated along both dimensions.

*Carranza: Yale University,guillermo.carranzajordan@yale.edu; Choukhmane: MIT Sloan School of Management, tahac@mit.edu; Greig: Vanguard fiona_greig@vanguard.com; O'Dea: Yale University, cormac.odea@yale.edu; Schmidt: MIT Sloan School of Management, ldws@mit.edu. We are very grateful to Kelly Hahn and Anna Madambda for illuminating comments and suggestions. We are very grateful to the Yale University Tobin Center for Economic Policy for co-funding this work. Any errors are our own.

1 Introduction

Every year, U.S. private employers contribute over $200 billion to their employees' retirement savings plans, mostly in the form of matching contributions. These incentives tie employee compensation to voluntary saving behavior and vary widely across employers, with more than 1,000 different matching formulas featuring different match rates, caps, and often multiple tiers.[1] The stated goal of these programs is to increase employee contributions for more than 100 million eligible workers while complying with federal regulations that limit the share of benefits accruing to highly paid workers. Given these twin goals of increasing employee savings and mitigating inequality, this paper asks a simple question: which matching formula should an employer choose to increase savings, reduce inequality, while containing costs? The answer to this question matters not only for employers seeking to align benefit design with these goals, but also for public policy. Over half of firms adopt one of four standardized "safe harbor" formulas, which are exempt from non-discrimination testing. Yet these widely adopted designs have never been evaluated against cost-equivalent alternatives.

We provide the first comprehensive evaluation of employer matching formulas by combining administrative data from 1.8 million Vanguard participants with a novel survey methodology that elicits responses to counterfactual matching designs. A distinctive feature of the survey is that it asks respondents about their likely behavior under match formulas that are rare or non-existent in the current landscape. We integrate these survey responses with a predictive model of saving behavior and show that, when calibrated with our survey responses, the model closely replicates observed saving choices. This provides validation for using the model, augmented with survey input, to predict responses under a wide range of counterfactual policies. This approach allows us to overcome the limitations of existing variation and consider a wide range of potential match formulas, including many that are never observed in practice.

There is good reason to believe that current matching formulas could be improved on both equity and efficiency grounds. Using administrative data from nearly 1,000 retirement plans managed by Vanguard, we first show that most low-income workers do not contribute enough to receive the full match. Even among high-income workers, around 20% fall short of the match threshold. This pattern implies significant inequality in the distribution of matching resources, one that aligns with the pattern observed in nationally-representative tax data by Choukhmane et al. (2024). Second, we find that the majority of employer match dollars accrue to *inframarginal* savers—employees who already contribute beyond

[1]For instance, a single-tier plan might match 100% of employee contributions up to 5% of pay. A two-tier plan could match 100% of contributions up to 3% of pay, and then 50% of contributions on the next 2% of pay.

the maximum matched amount. This implies that most matching dollars are not efficiently targeted at creating incentives for workers to contribute more.

Next, we use a new survey of retirement contributions under hypothetical match formulas to study the potential of alternative formulas to improve outcomes along both dimensions. We have four main results. First, employee self-reported saving is inelastic with respect to the match rate: dollar-for-dollar match rates are much more costly than 25-cents-on-the-dollar match, and yet only generate slightly more additional saving. Second, non-elective contributions—employer contributions made regardless of employee saving—do not appear to crowd out employee contributions. This holds in both survey responses (where hypothetical non-elective contributions do not reduce stated preferences) and administrative data (where actual increases in non-elective contributions do not reduce employee contributions). Third, bringing together the survey evidence and the predictive model, we characterize the frontier of the equity-efficiency trade off. In other words, we identify the set of contribution formulas where, for a given employer cost, reducing inequality necessarily requires lowering contributions, and vice versa. Formulas on the frontier tend to pair lower match rates with non-elective contributions. Fourth, we find that many existing retirement plans, including nearly all safe harbor plans, are strictly dominated in the sense that cost-neutral alternatives could increase saving without increasing inequality. We estimate that moving all plans in our sample from their current formulas to the dominating alternatives we identify would increase average employee contributions by 50%, without raising employer costs or exacerbating inequality. Alternatively, moving all plans to the efficiency frontier could reduce inequality in the distribution of employer contributions by 40%, again without reducing voluntary saving or increasing plan costs.

Our work relates to three strands of the literature. The first concerns the causal effects of saving incentives. Researchers have addressed this question using a variety of empirical strategies, including cross-sectional variation in match schedules (e.g., Engelhardt and Kumar (2007); Mitchell et al. (2007)), firm-level changes in match rates (e.g., Papke (1995); Kusko et al. (1996); Choi et al. (2002)), variation in tax incentives (e.g., Ramnath (2013); Chetty et al. (2014)) and experimental designs (e.g., Duflo et al. (2006)). A review by Choi (2015) concludes that while matching has a small positive effect on participation, its effect on average contribution rates is ambiguous.

We advance this literature by addressing a fundamental challenge: the difficulty of identifying responses to multidimensional and heterogeneous saving incentives. Our approach estimates stated contribution preferences across a wide range of match formulas that extend well beyond those observed empirically. By combining a model with purpose-built survey data designed specifically for model calibration, we can predict contribution behavior under

a large set of match schedules, thereby overcoming the limitations of existing variation in incentive structures.

Second, we contribute to the literature on the design of retirement saving institutions. Going back to Diamond (1977), the literature has identified two primary rationales for intervention in retirement savings: paternalism (addressing undersaving) and redistribution (both across people and states of the world). Beginning with Feldstein (1985), studies have used theoretical models to characterize optimal paternalistic retirement saving policies. These models typically assume that individuals undersave either because of behavioral biases (e.g., Beshears et al. (2025)) or because of the disincentives created by means-tested social safety nets (e.g., Hubbard et al. (1995)). In parallel, redistribution and insurance have been studied as an independent source of welfare gains in the design of these policies (e.g., Hosseini and Shourideh (2019), O'Dea (2018)). Recent work increasingly recognizes that these motives cannot be considered in isolation. For instance, Moser and Olea de Souza e Silva (2019) analyze optimal savings policies under joint concerns for retirement undersaving and income inequality, while Choukhmane and Palmer (2025) estimate a structural lifecycle model to quantify the welfare implications of alternative policy designs when policymakers value both objectives. While this literature provides important insights for the design of saving policies, it often relies on stylized settings that abstract from the complexities of actual retirement saving incentives. Real-world matching formulas involve multiple design parameters (i.e., match rates, caps, and often multiple tiers) and must account for wide heterogeneity in employee responses and employer financial resources, making analytical characterization often intractable. We contribute to this literature by providing the first systematic answer to which matching formulas actually achieve the two objectives identified in the theoretical literature, increasing savings and reducing inequality, while accounting for real-world constraints and both employee and employer heterogeneity.

Third, we contribute to the growing literature that combines administrative data, models of economic behavior, and custom surveys (Fuster and Zafar (2023); Stantcheva (2023); Caplin (2025)). These surveys bring rich variation to identify preferences, beliefs, and constraints that are generally missing from most administrative datasets but critical for model calibration (see Koşar and O'Dea (2023) for a review). Recent work has applied this approach to household consumption and saving behavior: Giglio et al. (2021) and Carranza and Goodman (2024) show that survey responses predict actual behavior, while Ameriks et al. (2020) and Indarte et al. (2025) use surveys to calibrate key parameters in lifecycle consumption models. We contribute to this literature by showing how survey responses not only predict actual behavior but also enable quantitative policy evaluation.

The remainder of the paper proceeds as follows. Section 2 describes the institutional set-

ting and reports descriptive analysis of contributions using the administrative data. Section 3 introduces a simple model of retirement saving in the presence of employer matches. Section 4 presents our survey design and validation methodology. Section 5 brings together the model and survey data. Section 6 characterizes the efficiency frontier and evaluates existing plans, and Section 7 concludes.

2 Matching Incentives: Institutional Context and Key Empirical Patterns

2.1 Employer Matches in the US

Defined contribution (DC) plans are the dominant vehicle for private retirement saving in the United States. Two-thirds of workers of US workers have access to employer-sponsored account (U.S. Bureau of Labor Statistics (2024)), and Americans hold $12 trillion of assets within them. A central feature of these plans is the employer match, which provides a direct financial incentive for employees to contribute.

Match formulas vary widely in their structure. Key elements include:

- the **match rate** (e.g., 50% or 100%) applied to employee contributions,
- the **cap** on contributions eligible for a match (e.g., up to 6% of pay),
- the **number of tiers**, with some plans offering a higher match rate on initial contributions and a lower rate thereafter,
- **non-elective contributions**, where employers contribute regardless of employee action.

Other plan design features also shape participation and savings behavior, including automatic enrollment, vesting schedules, eligibility rules, and withdrawal options. Accordingly these plan design features may also influence whether the worker takes full advantage of the match and associated employer costs.

2.2 Non-Discrimination Testing and Safe Harbor Plans

While employers decide whether to offer a match and choose its specific formula, legislation and regulation play a pivotal role in guiding these choices. To retain tax-advantaged status, retirement plans must pass annual non-discrimination tests that ensure benefits,

and in particular employer contributions, do not disproportionately accrue to highly compensated employees. Failure to pass these tests requires corrective action and may trigger excise taxes. Employers can bypass testing altogether by adopting standardized "safe harbor" plan designs, which are automatically deemed compliant. Through the structure of non-discrimination tests and the design of safe harbor provisions, public policy can effectively steer employers toward specific matching formulas. Reflecting this influence, Vanguard (2025) show that a third of Vanguard plans have adopted a safe-harbor design, while Arnoud et al. (2021) find that nearly half of plans in a nationally-representative sample adopt a match formula that satisfies the safe harbor employer matching requirements.

2.3 Administrative Data

We use data from administrative records for Vanguard plan participants. We use information on current and historical employer and employee retirement, as well as information on plan characteristics including the match formula, auto/enrollment status and loan availability. Additionally, we have information on hardship withdrawals and loan origination inside the plan which we use as measures of liquidity demand.

For the majority of participants, we have data on the income received from the plan sponsor. Income data comes from two different sources. First, for a subset of plan for which the plan manager conducts non-discrimination testing, we have W-2 income data directly from the plan sponsor. For the remaining plans, we estimate income using the elected contribution rates and observed contribution amounts. We use W-2 income data from non-discrimination testing when available and use the second imputed income data otherwise.

Finally, we have some additional administrative information on basic demographics including age, gender and geographical location.

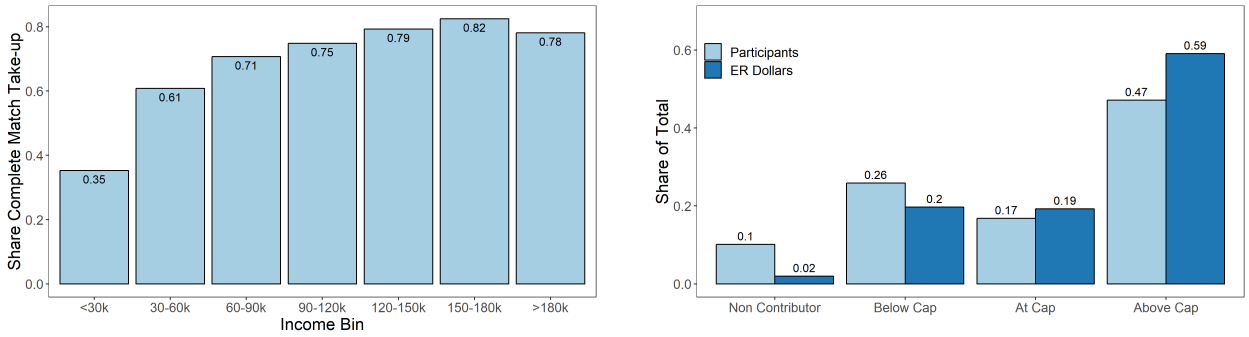
2.4 Two facts about DC retirement saving behavior

Before turning to our analysis of counterfactual policies, we summarize two empirical facts about defined contribution (DC) saving behavior in our data.

1. Take-up is incomplete, especially among low-income workers. Despite strong positive gradients in saving, many high earners do not contribute enough to receive the full employer match. Figure 1a shows that nearly 20% of workers earning $150,000 or more fall below the match threshold.

2. Most match dollars go to inframarginal savers. Close to half of of employees contribute more than the match ceiling, and this group receives almost 60% of total match dollars. If the goal of matching is to increase saving at the margin, the heavy allocation of match dollars to inframarginal savers suggests that existing incentives may be poorly targeted.

Figure 1: Patterns of Match Take-Up and Dollar Flows



(a) Match Take Up by Income

(b) Distribution by Marginality

Notes: Panel (a) shows the proportion of workers in the administrative data who do not full take up the match they are offered. Panel (b) divides the distribution of workers in groups of those who do not participate, those who participate but save less than match cap, those who save at the match cap, and save above the match cap. The figure also shows the proportion of match dollars going to employees in each of those four groups.

3 Predicting Behavior Under Counterfactual Plans

Our aim in this paper is to evaluate both the positive and normative features of match formulas, including those not currently observed in practice. To do so, we present a simple model of retirement saving in the presence of employer matches. We show that this model, combined with a small set of variables collected through our survey, can be used to predict saving behavior across a wide range of match formulas.

3.1 An Elementary Model of Retirement Saving Choices

An individual i has income Y to allocate between current consumption (C) and saving (S), such that $Y = C + S$. Tomorrow's wealth (W) comprises saving S plus an employer match $M(S)$, which is a function of saving. Individuals have quasilinear utility:

$$\max_{S} \quad \alpha_i \ln(Y - S) + \big(S + M(S)\big)$$

where α_i captures how much agent i values current consumption relative to future wealth. The quasilinear specification assumes that there are no wealth effects associated with the receipt of matches. We return to the assumption of no wealth effects below.

We consider match formulas that are weakly increasing and piecewise linear. We consider concave formulas (which are much more common). For illustration, we use a simple one-tier match: the employer matches saving at rate $\bar{m}$ up to a cap S^{cap}, and provides no match beyond that point.[2]

$$M(S) = m \cdot \min(S, S_{\text{cap}}) \tag{1}$$

We will refer to this below as the 'Partial Match', and refer to the saving that i does under this as $S_i^{\text{Partial Match}}$. Our aim is to obtain an estimate ($\hat{S}_i^{\text{Partial Match}}$) of this.

Suppose we observe two pieces of data for individual i (these will be the the objects that we will collect in our survey). These are i) the saving that they would do under no match ($S_i^{\text{No Match}}$) and ii) the saving saving that they would do under an *unlimited* match at rate m ($S_i^{\text{Full Match}}$).

There are three distinct cases to consider in search of of our estimate ($\hat{S}_i^{\text{Partial Match}}$) . These are (suppressing the i subscript):[3]

1. **Low Saver**: $S^{\text{Full Match}} < S_{\text{cap}}$. That is, even with an unlimited match, agent i would save less than the cap. Because the cap is not binding, their behavior under the partial match is the same as that under the full match:

$$\hat{S}^{\text{Partial Match}} = S^{\text{Full Match}} \tag{2}$$

   Figure 2a illustrates this case for an individual with 100 units of income and a 50% match up to 50 units.

2. **High Saver**: $S^{\text{No Match}} > S_{\text{cap}}$. That is, the agent would exceed the cap even without a match. Since the marginal incentive from matching under the partial match is irrelevant, the agent saves as they would without any match:

$$\hat{S}^{\text{Partial Match}} = S^{\text{No Match}} \tag{3}$$

   Figure 2b illustrates this cases.

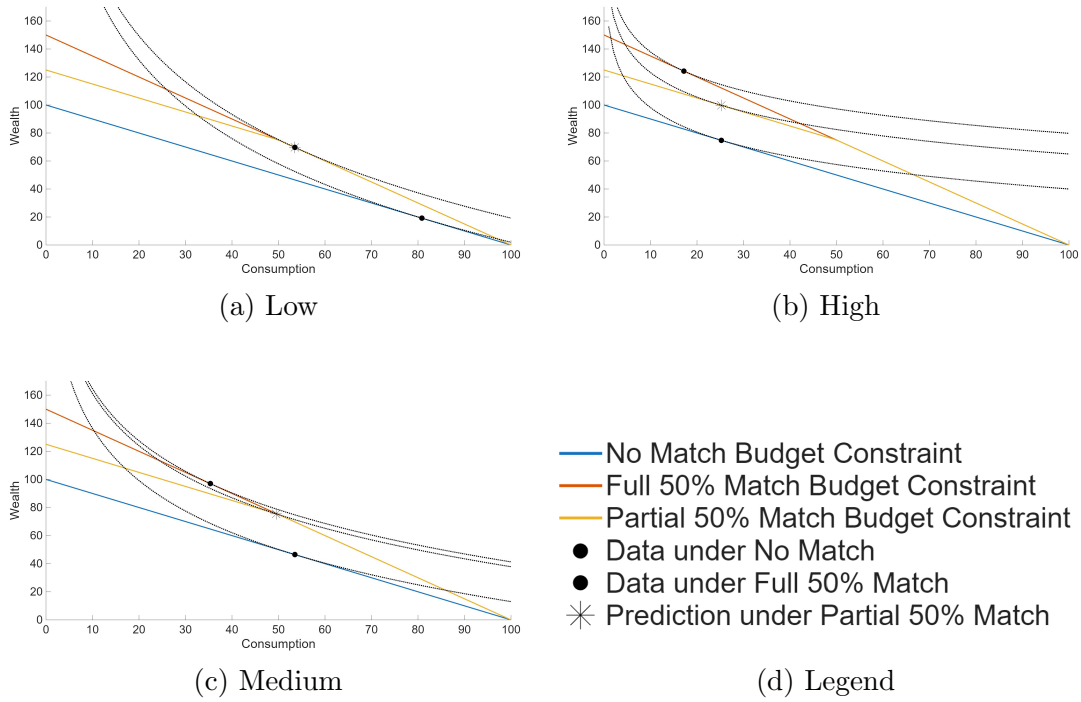
[2]The approach extends straightforwardly to multi-tier plans, as described in the appendix.
[3]A more formal exposition of the model predictions in each case is given in Appendix A.2.

3. **Medium Saver**: $S^{\text{Full Match}} > S_{\text{cap}}$ but $S^{\text{No Match}} < S_{\text{cap}}$. That is, the agent would exceed the cap under a full match, but fall short of it without any match. In this case the agent optimally bunches at the cap:

$$\hat{S}^{\text{Partial Match}} = S_{\text{cap}} \tag{4}$$

Figure 2c illustrates this case.

Figure 2: Model Predictions for one-tier concave match formulas



(a) Low (b) High

(c) Medium (d) Legend

Notes: This figure illustrates the optimal period 1 consumption and period 2 wealth choices for three different agents under three budget constraints: a case where there is no matching, a case where there is unlimited matching at 50% and a case where there is matching at 50% up to a cap (S_{cap}). Panel (a) (a 'low' saver) shows the optimal consumption and wealth allocation for an agent who would save below (S_{cap}) under both no match and the unlimited match. Panel (b) (a 'high' saver) shows the optimal allocation for an agent who would save above (S_{cap}) under both no match and the unlimited match. Panel (c) (a 'medium' saver) shows the optimal allocation for an agent who would save below (S_{cap}) under no match and above (S_{cap}) the unlimited match.

This model, together with data on the saving individuals would do under each of the two rates of return that feature in the partial match (0% and m%), can be used to predict saving under any cap. We show in the appendix that this approach extends naturally to match formulas with multiple tiers. For example, to evaluate saving under a counterfactual formula that matches at 100% up to one level and 25% up to another, we would need data

on saving under no match, an unlimited 25% match, and an unlimited 100% match. The logic holds because, with quasilinear utility, the marginal return to saving fully determines the agent's optimal choice — and at each segment, that marginal return is known.

4 Survey Data Collection and Validation

4.1 Survey Design and Implementation

The survey was designed to serve two main purposes: (i) to evaluate saving responses under counterfactual match formulas, and (ii) to generate responses that could be used to calibrate our model.

Survey Design The core of the survey consisted of a series of hypothetical contribution vignettes. Respondents were asked to evaluate several retirement plan designs. Before viewing the vignettes, they were presented with a set of assumptions that would apply to all scenarios. Figure 3a displays the introductory instructions. Participants were told to assume that employer contributions would be immediately vested, that employee contributions would receive the same tax treatment as in their current plan, and that all other aspects of their jobs would remain unchanged.

Following this, respondents were asked how much they would contribute under each hypothetical match schedule. Each vignette followed a standard format. Figure 3b shows the interface for the "No Match" vignette. Respondents used a slider to choose their contribution rate as a percentage of income. As the slider moved, the display updated in real time to show the employee contribution, the employer contribution, and the total combined amount.

The match schedules tested in the survey included:

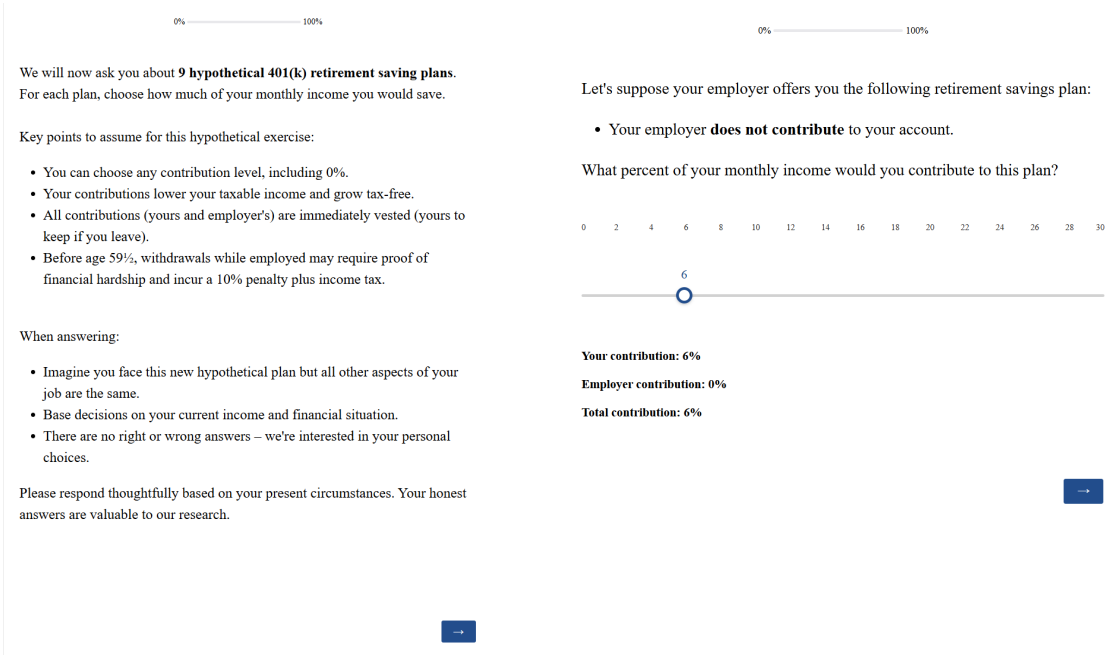
1. One vignette with no match.
2. Three vignettes with unlimited matching at rates of 25
3. One vignette with a match rate below 6
4. One vignette with no match but a universal employer contribution of 10

The vignettes in items 1 and 2 provide the main inputs for predicting matching contributions under alternative plan designs. The vignette in item 3 is used to test the consistency of responses and assess the predictive accuracy of our model. The vignette in item 4 allows us to evaluate the size of wealth effects.

The survey also included questions about current contribution behavior, as well as measures of behavioral traits such as present bias, risk aversion, self-control, and liquidity constraints.

Survey Implementation The survey was fielded in February 2025. The sample consisted of current plan participants for whom we had administrative data on contributions and income in 2023. We oversampled participants who had contributed less than the plan match cap in that year.

Table 1 compares the respondent sample to the full administrative population across age, income, and tenure. The final sample closely matches the population in terms of income and tenure but includes a somewhat higher proportion of older participants. In our empirical analysis, we re-weight responses to match the observable characteristics of the full participant population.



(a) Introductory Instructions (b) No Match Vignette

Figure 3: Survey Interface for Hypothetical Contribution Scenarios

Notes: Panel (a) displays the introductory instructions presented before the nine hypothetical contribution vignettes. These instructions explain key assumptions that apply across all scenarios, including contribution flexibility, tax treatment, vesting, and withdrawal rules. Panel (b) shows the interface for the "No Match" vignette. Respondents use a slider to choose the share of monthly income they would contribute, with real-time updates displaying their own contribution, the employer contribution, and the total.

Table 1: Demographic Characteristics of Survey Sample

Sample	Administrative Population	Survey Population	Final Sample
Count	1,862,579	348,372	1,754
Age			
25-35	23%	27%	6%
35-45	28%	28%	14%
45-55	25%	25%	29%
55-65	20%	16%	38%
65-75	4%	4%	14%
Income			
0-30k	10%	10%	9%
30-60k	24%	25%	23%
60-90k	24%	25%	24%
90-120k	15%	15%	16%
120-150k	10%	10%	13%
150-180k	6%	6%	7%
180k+	10%	10%	9%
Tenure			
0-3	21%	23%	17%
3-6	21%	22%	16%
6-9	14%	15%	11%
9-12	11%	11%	10%
12-15	7%	7%	8%
15+	27%	22%	39%

Notes: This table compares demographic characteristics across three groups: (i) the full administrative population of active plan participants, (ii) the subset of participants selected to receive the survey, and (iii) the final set of respondents who completed the survey. Percentages are calculated within each group for age, annual income, and job tenure categories. The final survey sample over-represents older participants and those with longer tenure compared to the administrative population, while the income distribution is broadly similar across groups. To account for these differences, all estimates in the empirical analysis are re-weighted to match the observable characteristics of the administrative population.

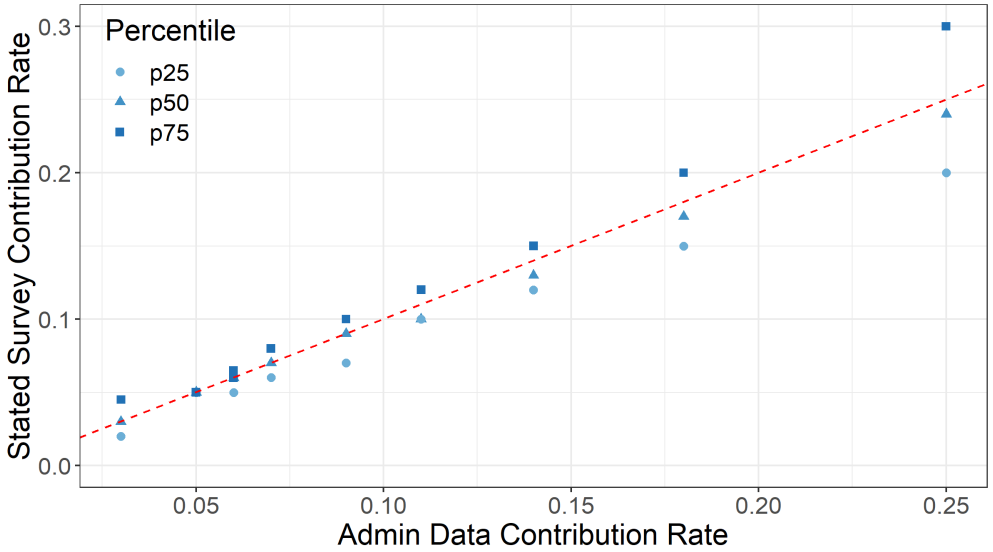
4.2 Survey Internal Validity

To test the internal validity of the survey responses, we compare survey responses to administrative data on measures for which both sources are available. First, we compare self-reported contribution rates to employer-sponsored retirement accounts with the actual contribution rates recorded in administrative data. Figure 4 shows the comparison. For most respondents, reported values align closely with observed contributions, with the distribution tightly clustered around the 45-degree line. The median response falls directly on this line, indicating that participants generally report their current contributions accurately.

Second, we examine responses to the hypothetical "50 on 6" and "100 on 6" match scenarios for the subset of respondents who actually face those match formulas in reality. Figure 5 plots each respondent's hypothetical response against their actual contribution rate. The correlation is high: 0.82 for "50 on 6" and 0.80 for "100 on 6". While the dispersion around the 45-degree line is greater (partly due to smaller sample sizes of around 200 participants per vignette). These exercises provide evidence that respondents understand and respond meaningfully to the hypothetical incentives presented in the survey.

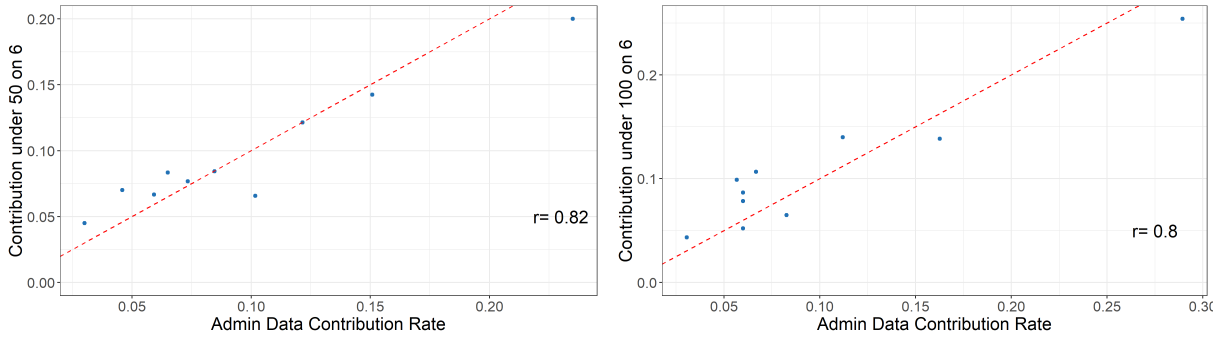
As a final piece of internal validation, Appendix Figure 18 examines internal consistency by comparing responses to a hypothetical "no match" scenario elicited at the start of the survey with responses to the same scenario elicited at the end. The stability of responses over time suggests that answers are not sensitive to survey fatigue or framing effects.

Figure 4: Comparing contributions in administrative data with those stated in the survey



Notes: This figure compares the contribution rates that are observed in the administrative data (horizontal axis) with the distributions reported by respondents in the survey (vertical axis).

Figure 5: Hypothetical stated contributions compared to admin for those who have formula



(a) 50% match up to 6%

(b) 100% match up to 6%

Notes: This figure uses samples of individuals who, in reality, have match formulas of 50% up to 6% of their pay (panel (a)), and have a match of 50% up to 6% of their pay (panel (b)). The figures compare their actual saving, measured in administrative data, with their hypothetical saving under their actual plan, as measured in the survey.

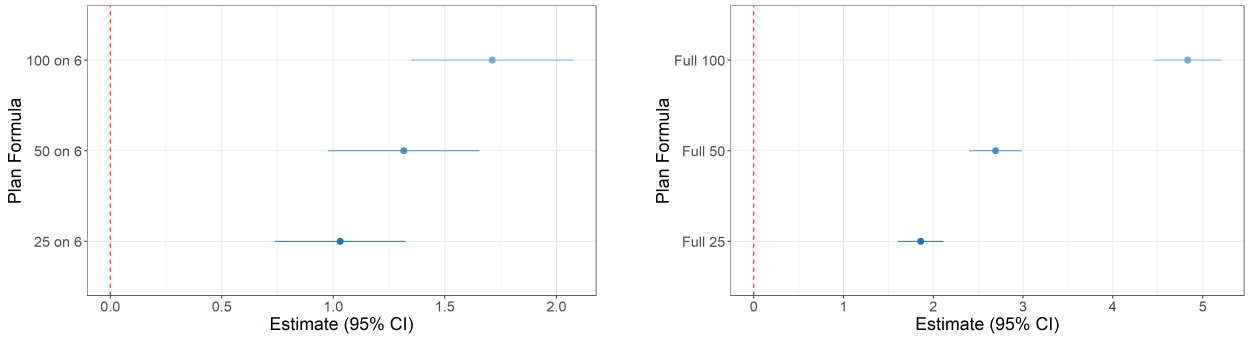
4.3 Survey Evidence on Saving Responses

Before using the survey responses to calibrate our model, we present initial evidence on the economic mechanisms reflected in the survey data.

Match *incentives* have a modest impact on saving Figure 6 illustrates how saving responds to different match rates, either capped at 6% of pay (panel (a)) or offered without a cap (panel (b)). In each case, we show the average increase in saving relative to the no-match baseline. Under the capped plans, a 25% match up to 6% of pay increases saving by about 1 percentage point; a 50% match yields 1.3pp in extra saving, and a 100% match increases saving by approximately 1.7 percentage points. The unlimited match plans generate substantially larger responses, with treatment effects of 1.9 and 4.8 percentage points for 25% and 100% match rates, respectively.

Low match rates provide better bang per buck than high match rates These patterns suggest that while saving is responsive to the presence of a match, the elasticity with respect to the match rate is modest. Table 2 illustrates this point. For both capped and uncapped matches, it reports the treatment effects (from Figure 6), the implied employer cost, and the cost per percentage point increase in employee saving. If the goal is to encourage saving, lower match rates deliver substantially greater value per dollar spent. Under the unlimited match, for example, it costs 1.5% of payroll to generate a one percentage point increase in saving using a 25% match rate. In contrast, it costs 2.9% of payroll to achieve the same effect using a 100% match.

Figure 6: Survey Treatment-Effect Predictions Under Two Incentive Structures



(a) Matching Up to 6%

(b) Unlimited Matching

Notes: Panel (a) shows the estimated average treatment effect of the introduction of matching up to 6% of *Treatment Effect, Cost* , at different rates (given on the vertical axis). This is calculated as the average difference between saving under the match and saving under no match. Panel (b) shows the estimated average treatment effect of the introduction of unlimited matching, at different rates (given on the vertical axis). This is calculated as the average difference between saving under the match and saving under no match.

Table 2: Treatment Effect, Cost, and Cost Per Percentage Point of Treatment Effect

	(a) Match up to 6%			(b) Unlimited Match		
	Employee savings increase (p.p.)	Cost	Cost per p.p. increase	Employee savings increase (p.p.)	Cost	Cost per p.p. increase
25% Match	1.0	1.4	1.3	1.9	2.8	1.5
50% Match	1.3	2.8	2.1	2.7	6.0	2.2
100% Match	1.7	5.6	3.3	4.8	14.1	2.9

Notes: This table shows the estimated effect on employee saving of the introduction of a match formula. Panel (a) gives the effect of matching, at a rate given in the first column, up to 6% of pay. Panel (b) gives the effect of matching, at a rate given in the first column, of any saving the employee does, up to 100% of their pay. The first column in each panel gives the mean treatment effect, calculated as the average difference between saving under the match and saving under no match. The second column in each panel gives the cost of the introduction as a proportion of payroll. The third column is the cost per percentage point of additional saving, calculated by dividing column (2) by column (1).

Wealth effects from non-elective contributions are small Many firms offer both matching and non-elective contributions to employee 401(k) accounts. Non-elective contributions do not affect the marginal return to saving, but they do increase wealth, which could, in principle, cause employees to reduce their own contributions. Our model assumes no such crowd-out. To test this assumption empirically, we present three pieces of evidence: (i) responses to one of our survey counterfactuals, (ii) panel regressions with participant fixed effects in the administrative data, and (iii) an event study exploiting large changes in non-elective contributions in selected plans.

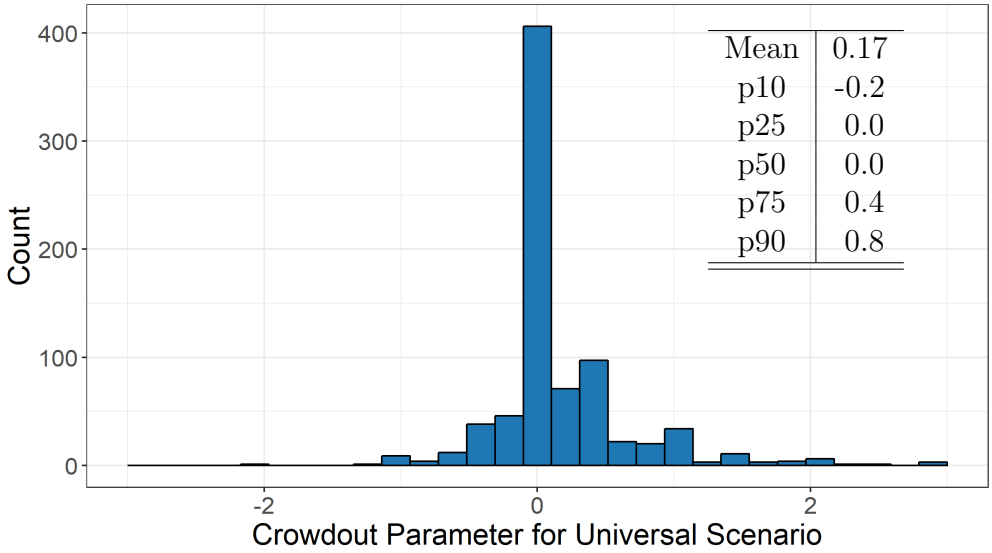
First, we examine how much additional self-reported contributions change (relative to

the no-match case) when the employer contributes 10% of pay as a non-elective contribution in one of our survey vignette. Figure 7 shows the distribution of changes in employee contributions, expressed as a share of the employer contribution that is crowded out. For the vast majority of respondents, the non-elective contribution has little to no effect on employee self-reported contribution preferences. The median crowd-out is zero, and even at the 10th percentile, only 20% (or 2 percentage points) of the employer contribution is offset. The mean effect is slightly positive, driven by a few respondents who report increasing their employee saving in response to the non-elective contributions. We interpret these results as evidence that wealth effects are small and that non-elective contributions have minimal impact on stated saving behavior.

We next examine evidence in our administrative data. Using a longitudinal panel of plan participants from 2014–2023, we select a 1% random sample of plan participants. Table 3 presents regression estimates indicating that any wealth effects from non-elective contributions are economically negligible and statistically indistinguishable from zero. Our specification includes participant fixed effects to absorb time-invariant heterogeneity in saving behavior, along with year fixed effects and standard demographic controls. The resulting coefficient is –0.001 (SE = 0.003), meaning that even a sizable 10 p.p. increase in non-elective contributions predicts only a 0.01 p.p. reduction in employees' voluntary contribution rates. While the analysis is correlational, the administrative evidence is consistent with our survey results and suggests that non-elective employer contributions do not crowd out participant saving in practice.

Finally, we leverage plan-level events in which non-elective contribution rates change sharply. Figure 8a plots the time series of non-elective and employee contribution rates for the six largest affected experiencing such changes. These plans exhibit large, discrete jumps in non-elective contribution rates. For instance, Plan C raised its average non-elective contribution rate from 0 to 5 p.p. of salary in 2012. Despite these sharp increases, employee contributions remain on a stable trend, showing no discernible break and no evidence of crowd-out. Because these non-elective increases occur in different calendar years across plans, we exploit this staggered timing using a difference-in-differences event-study design. Figure 8b reports the estimated coefficients from a specification with participant fixed effects and demographic controls (age, income, and tenure). The estimates show no pre-trend and only very small, if anything positive, changes in saving behavior following an increase in non-elective contributions. Across these three analyses, we find no evidence that employer non-elective contributions crowd out participants' voluntary saving.

Figure 7: Change in contribution with introduction of 10% non-elective contribution



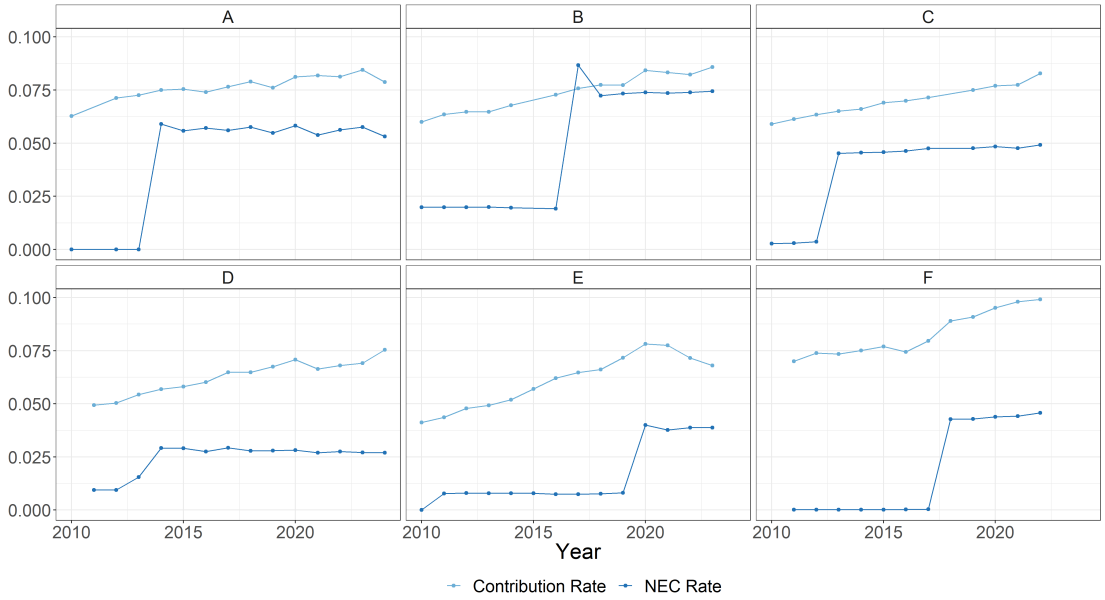
Notes: This figure illustrates the estimated treatment effect arising from the introduction of a 10% non-elective contribution. We estimate a crowd-out parameter that is change in employee elections, divided by 10. The distribution is summarized in the table.

Table 3: Non-elective Contribution Rate FE Regression

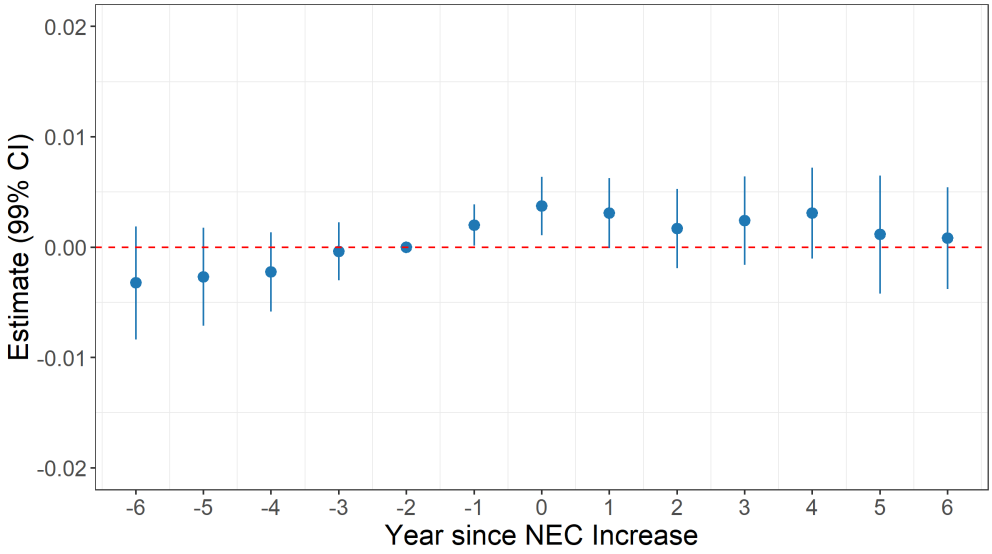
	Dependent variable:		
	Administrative Contribution Rate		
	(1)	(2)	(3)
Non-elective Rate	−0.001	−0.0004	−0.001
	(0.003)	(0.003)	(0.003)
Participant FE	No	Yes	Yes
Year FE	No	Yes	Yes
Demographic Controls	No	No	Yes
Observations	493,991	493,991	493,991
Note:			*p<0.1; **p<0.05; ***p<0.01

Notes: This table reports estimates from linear regressions of participants' voluntary contribution rates on employer non-elective contribution rates, using a 1% random sample from the 2014–2023 administrative panel. Column (1) reports the baseline specification with only the non-elective rate. Column (2) adds participant fixed effects and year fixed effects. Column (3) further includes demographic controls (log income, age, and tenure quadratic) and contribution-formula fixed effects.

Figure 8: Changes in non-elective contribution rates



(a) Evolution of non-elective and contribution rates for six large plans



(b) Staggered D-i-D event study coefficients

Notes: Panel (a) displays the time series of non-elective contribution rates and employee contribution rates for six large plans that experienced substantial changes in non-elective contributions at different points in time. Panel (b) presents the corresponding event-study coefficients from a staggered difference-in-differences model that incorporates participant fixed effects and demographic controls (age, income, and tenure). The estimates indicate no detectable pre-trend and only small positive post-event responses, suggesting that increases in employer non-elective contributions crowd out participants' voluntary saving.

5 Putting the Model to Work

5.1 Implementation

Our survey data provide two key reference points for each respondent: how much they would save if the employer offered no match, and how much they would save if the employer offered an unlimited match at different rates. We denote saving in the no-match case as $S_i^{\text{No Match}}$, and saving under an unlimited 50% match as $S_i^{\text{Full Match}}$. While we focus here on the 50% case for concreteness, the same reasoning applies to the 25% and 100% match rates.

The idea is straightforward. Knowing how someone saves with no incentive and how they save with a very strong incentive allows us to predict their behavior under any intermediate incentive. Using $S_i^{\text{No Match}}$ and $S_i^{\text{Full Match}}$, we apply equations (2)–(4) to estimate saving under a 50% match with any cap on matched contributions.[4] This approach allows us to map from the responses to a small number of hypothetical scenarios to predicted behavior under the much larger space of possible plan designs.

In Table 4, we illustrate the prediction process with three hypothetical respondents, labelled A, B, and C. For each, we show their stated saving under a no-match plan and under an unlimited 50% match. We then use the model to predict how much they would save under a plan offering a 50% match up to 6% of pay.

- **Respondent A** saves nothing with no match and 3% with an unlimited 50% match. Because their saving under the unlimited match is below the 6% cap, the capped match is equivalent to the unlimited one, and the model predicts they would also save 3% under the capped plan.

- **Respondent B** saves 6% without a match and 9% with an unlimited match. Under the 50% match up to 6%, they would already be at the cap in the no-match case, so the model predicts they would bunch at 6%.

- **Respondent C** saves below the cap with no match (3%) and above it with an unlimited match (9%). In this case, the capped match increases their contributions, but only up to the 6% limit, so the model again predicts bunching at 6%.

5.2 Model Internal validation

We validate the model's predictions using internal survey data for the one-tier concave match vignettes (25%, 50%, and 100% matches up to 6%).

[4] Appendix A.3 shows how the same approach extends to plans with two-tier match structures.

Table 4: Prediction Examples under 50% on 6%

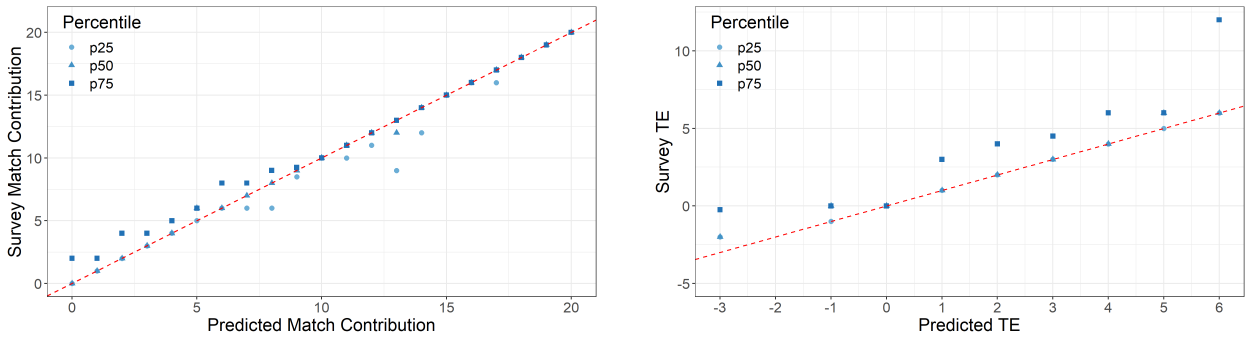
Person	Survey reported saving under...		Prediction
	...no match	...unlimited match	
A	0	3	3
B	6	9	6
C	3	9	6

Notes: The table shows predicted contributions for three survey responses, comparing "No Match" and "Full Match 50%" scenarios, with the prediction under a 50% match on 6%.

Panel (a) of Figure 9 shows that predicted and actual saving levels align closely, with median responses tracking the 45-degree line. Panel (b) evaluates predictions of treatment effects—saving under the match formula minus saving in the no-match vignette—which will be central to our optimality analysis. Predicted and actual treatment effects match well at the median, though dispersion is greater at the upper end of the distribution.

These patterns suggest that the model predicts both levels and treatment effects with sufficient accuracy to support the counterfactual exercises that follow.

Figure 9: Comparison of hypothetical contribution and model-predicted outcomes



(a) Concave Prediction

(b) Prediction of Survey Treatment Effect

Notes: Figure (a) compares respondents' response to a questions which asks them their hypothethical contribution under one of three formulas (25% up to 6%, 50% up to 6%, and 100% up to 6%), to the model predicted amount that they would save under that match. Figure (b) uses the same sample, but compares the treatment effect estimated by compare answers to saving under hypothetical match formulas to treatment effects estimated using the model.

5.3 Validation against Administrative Data

We next compare the model's predictions to administrative records of actual contribution rates.

The analysis focuses on the 1,085 respondents (65% of the sample) whose real-world plans are simple one-tier or two-tier matches with caps below 8%. Figure 10 plots the

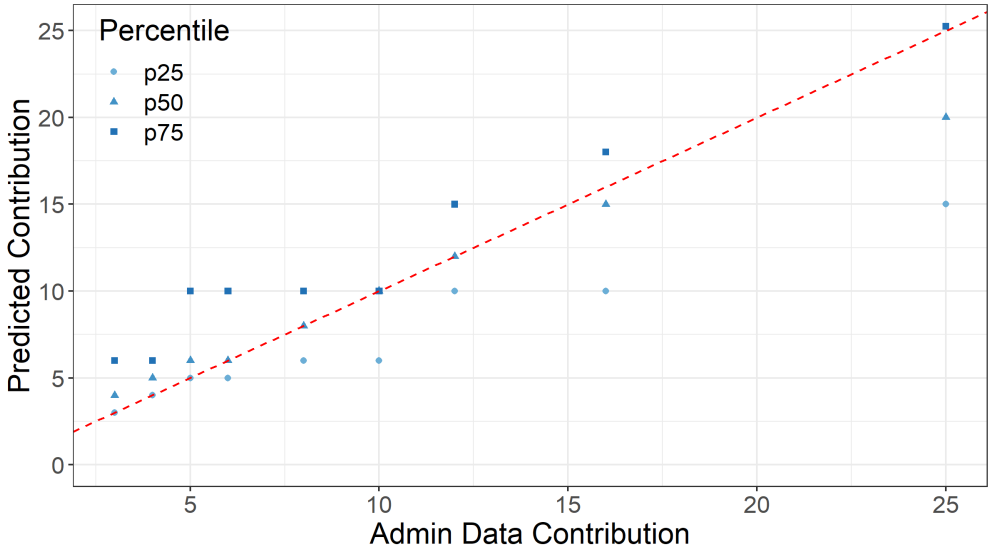
distribution of model predictions against the observed administrative contribution rates. Median predictions track the 45-degree line closely, indicating good alignment in central tendency. Dispersion is greater than in the survey validation, with noise increasing at very high contribution levels. Nonetheless, both the 25th and 75th prediction percentiles rise systematically with observed contribution rates.

These results suggest that the model captures broad patterns in real-world saving behavior, while being less precise at the extremes.

As a more formal test, Table 5 reports regressions of administrative contribution rates on our predicted rates, with alternative sets of controls. The results show that administrative rates rise nearly one-for-one with predicted rates, and the R^2 values are high—between 0.46 and 0.48. Most of the explained variance comes from the prediction algorithm itself. The coefficient on the predicted rate remains large and statistically significant even after adding the two survey measures used in the prediction, as well as a range of demographic variables known to correlate with saving.

We conclude that the model's predictions are strongly correlated with observed contribution rates, explaining roughly half of their variance. The predictive power is not simply due to the survey inputs or demographics alone, but reflects additional structure captured by the model.

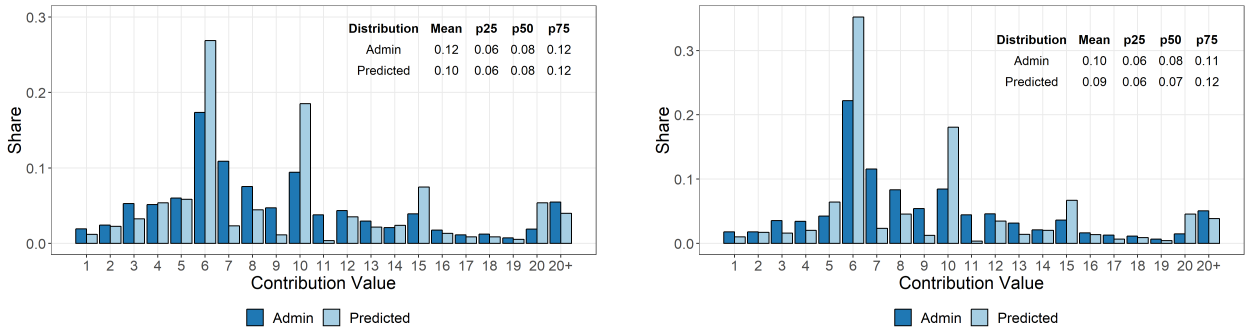
Figure 10: Model Prediction of Current Administrative Plan Saving Decision



Notes: This figure compares respondents' saving as measured int he administrative data (on the horizontal axis), with the model-predicted contribution that they would do under their observed plan (on the vertical axis).

Finally, we assess the extent to which our predicted algorithm used on our sample of plan participants can replicate plan-level distributions of contributions. Our ability to match real-

Figure 11: Predicted Distributions under Two Match Formulas



(a) Prediction for 50% up to 6% Plan

(b) Prediction for 100% up to 6% Plan

Notes: These figures compares, for two specific plans, distributions of saving in the administrative data, with model predicted distributions under such a plan. We weight survey responses to match demographics in the administrative data.

world contribution distributions depends on both the accuracy of the prediction algorithm and the representativeness of our sample data. We demonstrate that by appropriately re-weighting for observable demographics, we can capture the key features of contribution distributions under alternative formulas.

To achieve this, we use a post-stratification weighting procedure that adjusts survey responses to align with known characteristics of the target plan population. We divide both the survey sample and the administrative target population into mutually exclusive strata based on age and income bins, creating a two-dimensional stratification matrix that reflects the joint distribution of these demographic characteristics. For each age-income stratum, we calculate the proportion of observations in both the survey sample and the administrative data, then derive stratification weights as the ratio of the administrative share to the sample share. Each survey observation is assigned the weight corresponding to its age-income combination.

Figure 11 shows the administrative distribution of contributions alongside our reweighted distribution of predicted rates for two plan formulas: 50% on 6% and 100% on 6%. We capture the main features of the administrative distribution, including the peak at the 6% cap. The predicted distribution exhibits higher levels of bunching at the peak and at multiples of 10% (e.g., 10%, 15%), which is a known feature of survey respondent data. Despite this bunching, we accurately match the mean, median, and the 25th and 75th percentiles of the distribution, as shown in the tables within each plot. From this exercise, we conclude that our prediction algorithm, combined with proper reweighting of our survey data, is able to reproduce the distribution of real-world contributions under alternative matching formulas.

Table 5: Model Predictive Power

	Dependent variable:		
	Administrative Contribution Rate		
	(1)	(2)	(3)
Predicted Rate	0.906***	1.033***	1.001***
	(0.030)	(0.161)	(0.160)
Full Match		0.127***	0.126***
		(0.030)	(0.030)
No Match		−0.216	−0.215
		(0.142)	(0.142)
Log Income			0.616**
			(0.256)
Age			−0.304**
			(0.133)
Age^2			0.003***
			(0.001)
Tenure			−0.018
			(0.057)
$Tenure^2$			0.001
			(0.001)
Observations	1,085	1,085	1,085
R^2	0.464	0.475	0.485
Note:			*p<0.1; **p<0.05; ***p<0.01

Notes: This table reports the results of linear regressions of the contribution rate observed in the administrative data on the model-predicted values. Column (1) includes on a constant and the predicted rate. Column (2) includes the answers to the questions which are used to estimate the model. Column (3) additionally includes demographics.

6 Are plan formulas on the equity efficiency frontier?

Our normative analysis begins with the assumption, commonly adopted in theoretical analyses of optimal retirement systems (starting with Diamond (1977)), that the social planner values both increased retirement saving and a more equal distribution of employer contributions. These two goals reflect the dual objectives of efficiency (encouraging greater saving) and equity (reducing inequality).

Using our survey data and model, we can predict employee saving behavior under any match formula. Using these predictions, we can evaluate the equity and efficiency properties of all possible match formulas. We focus on one measure of efficiency: the increase in employee saving relative to a no-match scenario. For equity, our primary measure is the coefficient of variation in total employer contributions. Motivated by the fact that those with lower incomes are less likely to take up the match under the status quo, we also show the covariance between income and employer contributions.

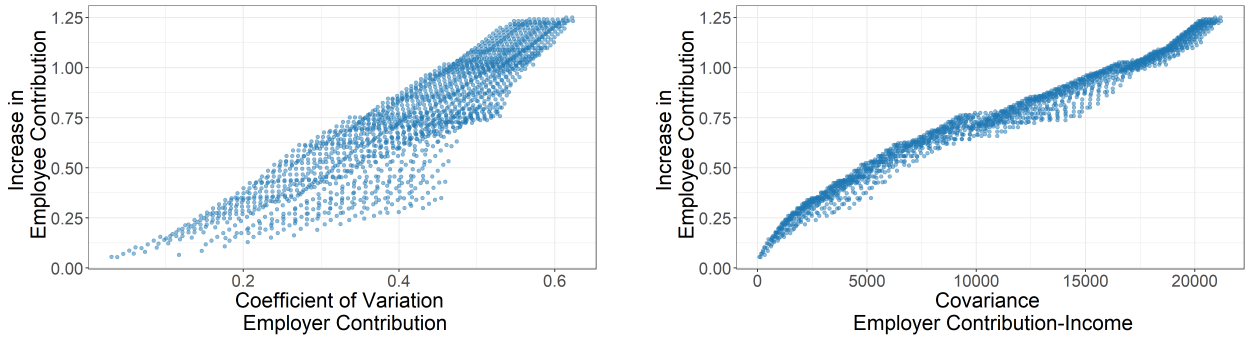
Figure 12 shows the equity and efficiency outcomes for the set of plans that meet three criteria: they cost the employer 6 percent of payroll, include either one or two tiers, and have a match rate that does not increase with employee contributions.

Two key patterns emerge from the figure. First, there is a clear equity-efficiency tradeoff in match design: along the frontier, holding employer cost constant, increasing employee saving requires accepting greater inequality in employer contributions. Second, many existing plans are strictly dominated in the sense that there are alternative plans that (i) cost the same, (ii) generate higher employee saving, and (iii) yield more equitable distributions of employer contributions.

6.1 Characterizing the Efficient Frontier

To characterize the properties of plans on the frontier, we begin with single tier plans, defined as those with a constant match rate up to a cap, and restrict attention to plans that cost 6% of payroll (i.e., the maximum cost of the most common formula in our sample: 100% match up to 6%). We consider a wide range of caps for each of three match rates (25%, 50%, and 100%). For each plan, we set the non-elective contribution so that, given the match rate and the saving behavior it induces, total employer cost equals 6% of payroll. We group the resulting plans into six categories, based on the match rate and whether the cap is above or below 5% of pay. Figure 13 plots the equity and efficiency of all plans, with colors indicating group membership. Average non-elective contributions for each group are shown in the legend and illustrate the fact that, holding cost constant, the higher the match rate, and the higher the cap, the lower is the non-elective contribution.

Figure 12: Equity-Efficiency Frontier under Two Equity Measures



(a) Coefficient of Variation Equity

(b) ER-Income Covariance Equity

Notes: These figures illustrate the equity and efficiency characteristics of the set of plans which cost 6% of payroll, are either single tier or two tier, and have a match rate that is non-increasing in employee contribution. In both panels, efficiency is measured using the increase in employee saving. In panel (a) the the equity dimension is represented using the coefficient of employer contributions. In panel (b) the equity dimension is represented using the covariance of employer contribution rate and income.

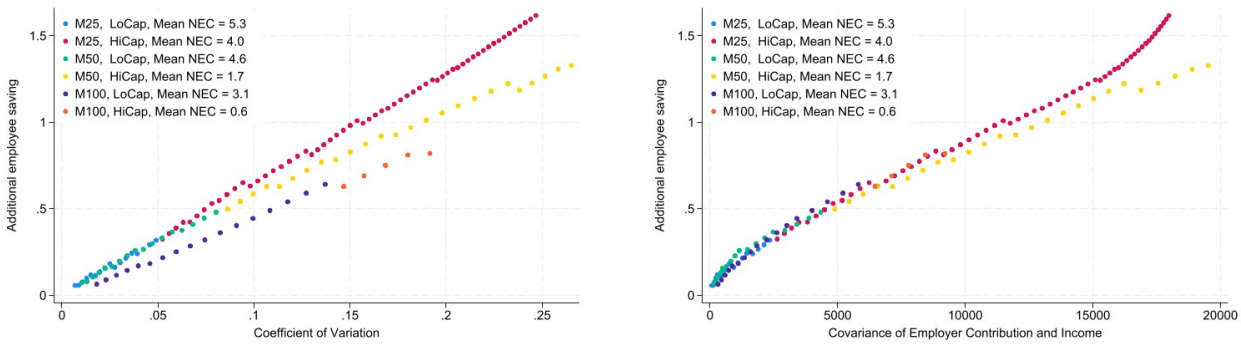
The figure shows that plans with lower match rates tend to lie on the frontier, while those with higher match rates are often dominated. This normative result follows from the positive finding that saving is inelastic with respect to the match rate. Holding total cost fixed, higher match rates allocate more dollars to infra-marginal savers (individuals who would have contributed even without the incentive) and therefore generate little additional saving. These plans are also more expensive, leaving less room in the budget for non-elective contributions that help reduce inequality. In contrast, lower match rates can be applied over a broader contribution range, directing fewer dollars to infra-marginal savers. As a result, they are more cost-effective and leave more budget available for non-elective contributions that mitigate inequality.

The level of the cap determines where a plan falls along the frontier, or, more generally, along the ray from the origin defined by its match rate. Holding the match rate fixed, increasing the cap allows the match to apply over a broader contribution range, which generates more marginal saving. However, it also increases inequality in employer contributions.

Figure 14 compares single tier plans, which apply a constant match rate up to a cap, with two tier plans, which apply a higher match rate on initial contributions and a lower rate on contributions beyond a threshold. Most single tier plans on the frontier are dominated by a two tier alternative. This is unsurprising, as two tier plans introduce two additional parameters (a second match rate and a second cap) which provide greater flexibility. However, as we document below, the efficiency and equity gains are often modest, and two tier plans come at the cost of added complexity. The two tier plans on the frontier exhibit the same patterns as the single tier plans: those with lower match rates tend to dominate those with

higher match rates. For example, a plan that matches 50% up to a first cap and 25% beyond that will often dominate a plan that matches 100% up to a first cap and 50% beyond. As before, the level of the caps determines where the plan lies along the frontier.

Figure 13: Frontier Characterization under Two Equity Metrics for Single-Tier Plans

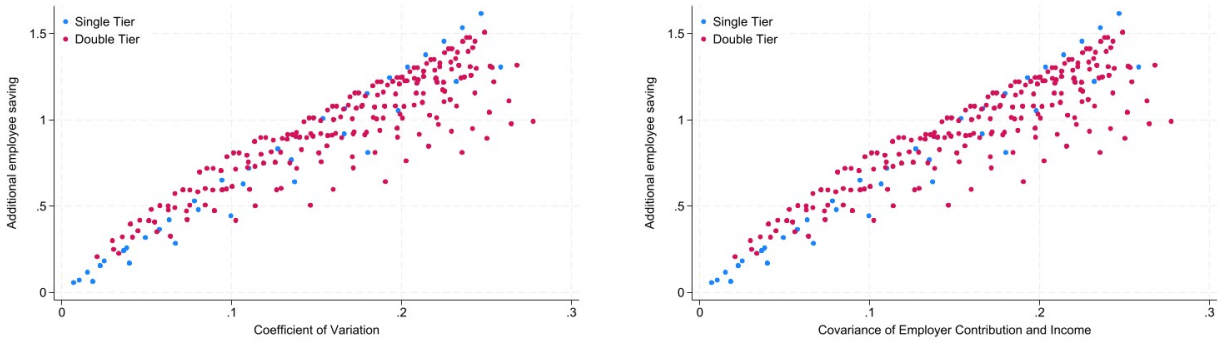


(a) Coefficient-of-Variation Equity

(b) ER–Income Covariance Equity

Notes: These figures illustrate the equity and efficiency characteristics of the set of single tier plans which cost 6% of payroll. In both panels, efficiency is measured using the increase in employee saving. In panel (a) the the equity dimension is represented using the coefficient of employer contributions. In panel (b) the equity dimension is represented using the covariance of employer contribution rate and income. Plans are distinguished their match rate (25%, 50%, or 100%), and whether the cap is high or low (resp. above or below 5% of pay); giving six groups in total. The payroll budget is balanced using non-elective contributions. Average non-elective contributions in each group are given in the legend.

Figure 14: Comparing Single and Two Tier Plans



(a) Coefficient of Variation

(b) Cov. of Employer Contribution and Income

Notes: These figures illustrate the equity and efficiency characteristics of the set of plans which cost 6% of payroll, are either single tier or two tier, and have a match rate that is non-increasing in employee contribution. In both panels, efficiency is measured using the increase in employee saving. In panel (a) the the equity dimension is represented using the coefficient of employer contributions. In panel (b) the equity dimension is represented using the covariance of employer contribution rate and income. Plans are distinguished by whether they are single tier or have two tier.

6.2 Moving to the frontier for the most common plans

We have shown that many plan formulas are dominated, in the sense that cost-neutral reforms could generate higher saving without increasing inequality in employer contributions. We now ask whether such dominated plans are actually observed in practice.

The most common plan in our data offers a 100% match on contributions up to 5% of pay. Using our model, we identify all single and two tier plans that are equally costly to the firm and compute, for each, the resulting increase in saving (relative to no match) and the coefficient of variation in employer contributions. Figure 15 plots these outcomes. The blue triangle represents the 100% match up to 5% plan, which lies well inside the frontier.

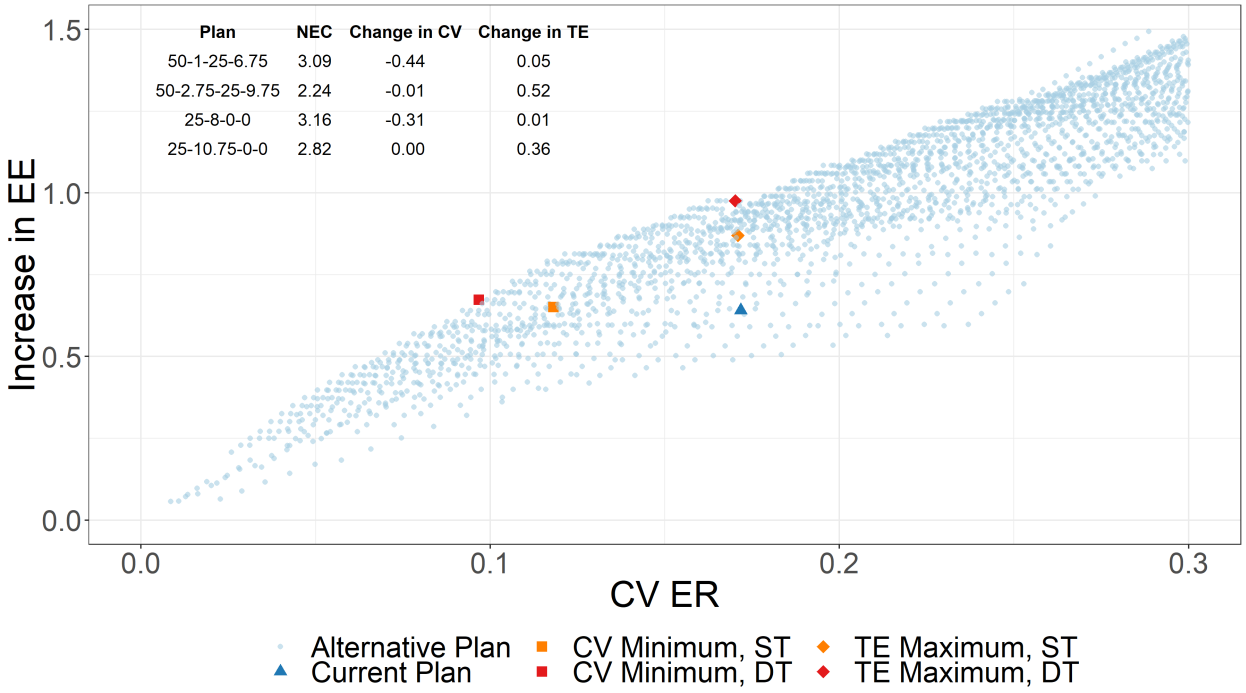
We highlight four alternative plans. The red square shows the plan that minimizes inequality, subject to achieving at least as much saving as the current plan. The red diamond shows the plan that maximizes saving, subject to no increase in the coefficient of variation in employer contributions. The orange markers show the analogous optima when we restrict attention to single tier plans.

The optimal plans have lower match rates and higher caps, and they balance the budget with non-elective contributions. The saving-maximizing double (single) tier increases employee contributions by 52% (36%) from a baseline of about 0.6 percentage points. The inequality-minimizing double (single) tier reduces the coefficient of variation in employer contributions by 46% (34%). Figure 16 shows a similar exercise for the covariance equity measure: the optimal plan in that case is also characterized by lower match rates, higher caps, and non-elective contributions.

We can extend this analysis beyond the single most common plan. Table 6 reports the gains from moving the five most common plans in our data to their cost-equivalent, dominating alternatives, showing both the increase in employee saving and the reduction in inequality in employer contributions. Appendix Table 7 presents the corresponding results for the 25 most common plans,[5] revealing a similar pattern. Most of these widely used plans lie well inside the efficiency–equity frontier. Averaging across all plans, we estimate that replacing each with its identified dominating alternative would increase average employee contributions by more than 50%, without raising employer costs or increasing inequality. Alternatively, moving all plans to the frontier could reduce inequality in employer contributions, measured by the coefficient of variation, by nearly 40%, again without reducing voluntary saving or increasing costs.

[5] Appendix Table 8 shows the results when counterfactuals are restricted to single tier plans.

Figure 15: Moving to the frontier for plans which have a 100% up to 5%



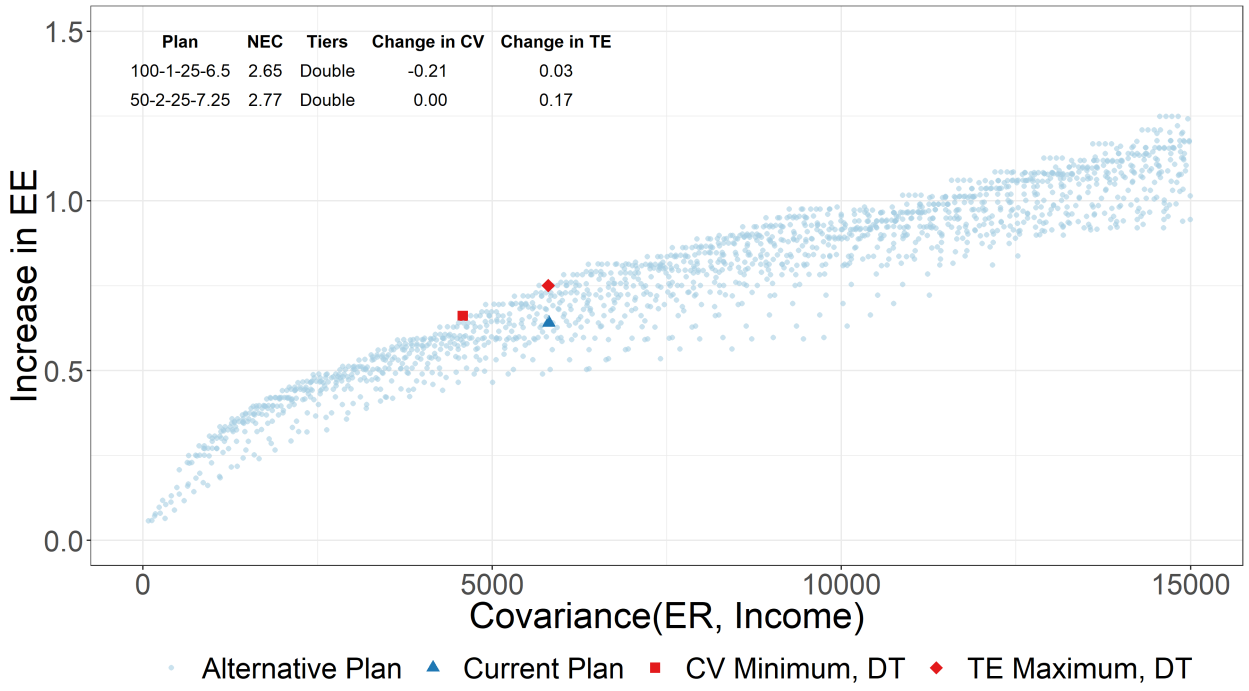
Notes: This figure shows the set of equity and efficiency characteristics for the set of plans which are as costly as the most common plan: 100% up to 5% of pay. We show four alternative plans. We show the double tier and single tier plans which maximizes employee saving, subject to no increase in inequality. These are the red and blue diamonds. We also show the double tier and single tier plans which minimize inequality, subject to no decrease in employee saving. These are the red and blue squares. The characteristics of plans which achieve these frontier points, as well as the changes in equity and efficiency they generate are given in the legend.

Table 6: The effect of moving to the frontier two tiers: five most common plans

	Minimize Inequality (CV)			Maximize Employee Savings (TE)		
Current Plan	Optimal Plan	Optimal NEC	Change CV	Optimal Plan	Optimal NEC	Change TE
100-5-0-0	50-1-25-6.75	3.1	-0.57	50-2-25-9.75	2.4	0.42
100-3-50-5	50-1-25-6.25	2.2	-0.46	50-1-25-8.75	1.8	0.37
100-6-0-0	50-2-25-7.75	3.5	-0.62	50-2-25-13.5	3	0.39
50-6-0-0	50-1-25-6.5	1.1	-0.37	50-1-25-8.5	0.8	0.29
100-1-50-6	50-1-25-7.5	1.4	-0.21	50-1-25-8.75	1.2	0.14

Notes: This table shows, for the five most common plans, the counterfactual plan (wether it is it single tier or double tier) which are equally costly and either minimize inequality (subject to no fall in employee saving), or maximize saving (subject to no increase in inequality). See Appendix Table 7 for further plans.

Figure 16: Moving to the frontier for plans which have a 100% up to 5%



7 Conclusion

This paper develops a framework for evaluating 401(k) match formulas that combine a simple economic model with purpose-built survey data. We show that survey responses to hypothetical scenarios can accurately predict saving behavior under both observed and counterfactual plan designs. Using this approach, we estimate saving responses to a wide range of match schedules, including those not observed in reality, and trace out their implications for both efficiency and equity. Our findings highlight a central trade-off in plan design: increasing saving often requires accepting greater inequality in employer contributions.

Notwithstanding the fact that there is a trade-off in plan design between equity and efficiency, we find that many plans, including the most common plans in our administrative data, and several which have safe-harbor status lie well inside the frontier. More saving, and lower inequality in employer contribution rates are attainable within existing budget constraints. This points to the value of incorporating both normative and predictive tools into retirement plan evaluation.

References

Ameriks, J., J. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy 128*(6), 2375–2451.

Arnoud, A., T. Choukhmane, J. Colmenares, C. O'Dea, and A. Parvathaneni (2021). The Evolution of U.S. Firms' Retirement Plan Offerings. Evidence from a New Panel Data Set. Technical report, NBER.

Beshears, J., J. J. Choi, C. Clayton, C. Harris, D. Laibson, and B. C. Madrian (2025). Optimal illiquidity. *Journal of Financial Economics 165*, 103996.

Caplin, A. (2025). Data engineering for cognitive economics. *Journal of Economic Literature 63*(1), 164–196.

Carranza, G. and A. Goodman (2024). Retention or regressivity? the empirical effects of 401 (k) vesting schedules. *The Empirical Effects of 401*.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

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

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax policy and the economy 16*, 67–113.

Choukhmane, T., J. Colmenares, C. O'Dea, J. L. Rothbaum, and L. D. Schmidt (2024, August). Who benefits from retirement saving incentives in the u.s.? evidence on gaps in retirement wealth accumulation by race and parental income. Working Paper 32843, National Bureau of Economic Research.

Choukhmane, T. and C. Palmer (2025). How Do Consumers Finance Increased Retirement Savings? Technical report.

Diamond, P. A. (1977). A framework for social security analysis. *Journal of public economics 8*(3), 275–298.

Duflo, E., W. Gale, J. Liebman, P. Orszag, and E. Saez (2006). Saving incentives for low-and middle-income families: Evidence from a field experiment with H&R Block. *The Quarterly Journal of Economics 121*(4), 1311–1346.

Engelhardt, G. V. and A. Kumar (2007). Employer matching and 401 (k) saving: Evidence from the health and retirement study. *Journal of Public Economics 91*(10), 1920–1943.

Feldstein, M. (1985). The optimal level of social security benefits. *The Quarterly Journal of Economics 100*(2), 303–320.

Fuster, A. and B. Zafar (2023). Survey experiments on economic expectations. In *Handbook of Economic Expectations*, pp. 107–130. Elsevier.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

Hosseini, R. and A. Shourideh (2019). Retirement financing: An optimal reform approach. *Econometrica 87*(4), 1205–1265.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary saving and social insurance. *Journal of political Economy 103*(2), 360–399.

Indarte, S., R. Kluender, U. Malmendier, and M. Stepner (2025). Consumption wedges: Measuring and diagnosing distortions.

Koşar, G. and C. O'Dea (2023). Expectations data in structural microeconomic models. In *Handbook of Economic Expectations*, pp. 647–675. Elsevier.

Kusko, A. L., J. M. Poterba, and D. W. Wilcox (1996). Employee decisions with respect to 401 (k) plans.

Mitchell, O. S., S. P. Utkus, and T. Yang (2007). Turning workers into savers? Incentives, liquidity, and choice in 401 (k) plan design. *National Tax Journal 60*(3), 469–489.

Moser, C. and P. Olea de Souza e Silva (2019). Optimal paternalistic savings policies. *Columbia Business School Research Paper* (17-51).

O'Dea, C. (2018). Insurance, efficiency and the design of public pensions. In *2018 meeting papers*, Volume 1037, pp. 373–383.

Papke, L. E. (1995). Participation in and contributions to 401 (k) pension plans: evidence from plan data. *Journal of Human Resources 30*(2), 311–326.

Ramnath, S. (2013). Taxpayers' responses to tax-based incentives for retirement savings: Evidence from the Saver's Credit notch. *Journal of Public Economics 101*, 77–93.

Stantcheva, S. (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics 15*(1), 205–234.

U.S. Bureau of Labor Statistics (2024, Apr). Access to and participation in defined contribution plans, march 2023. The economics daily, U.S. Department of Labor. 67

Vanguard (2025, June). How america saves 2025. Technical Report 24, Vanguard Workplace Solutions. 24th edition; data reflect savings behavior and plan trends through 2024.

A Model Appendix

A.1 Saving under convex formulas

A similar approach can be taken with convex formulas, in which the matching rate *increases* with saving. Consider the following match formula in which provide no employer match up to a floor and then matches saving at rate $\bar{m}$ on all saving above the floor.

$$M(S) = m \cdot \max(S - S^{\text{floor}}, 0) \tag{5}$$

We can again distinguish between three cases.

1. **Low Saver**: $S^{\text{Full Match}} < S_{\text{floor}}$. That is, even with an unlimited match, agent i would save less than the the point at which the match kicks in. Because the introduction of the match doesn't change their marginal incentive to save, their behavior under the partial convex match is that which would hold in the No Match case:

   $$\hat{S}^{\text{Partial Match}} = S^{\text{No Match}}$$

   Figure 17a illustrates this case.

2. **High Saver**: $S^{\text{No Match}} > S_{\text{floor}}$. That is, the agent would exceed the floor even without a match. Since the marginal incentive from matching is the same here as in the Full Match case, the agent saves as they would in that case:

   $$\hat{S}^{\text{Partial Match}} = S^{\text{Full Match}}$$

   See Figure 17b.

3. **Medium Saver**: $S^{\text{Full Match}} > S_{\text{floor}}$ but $S^{\text{No Match}} < S_{\text{floor}}$. That is, under a full match the agent would save more than the match floor, but without any match they would save less than it.

   The agent's optimal saving depends on whether the additional utility from crossing the floor and accessing the match outweighs the utility loss from lower current consumption. There are two possible outcomes:

   - If the gain from the match is sufficiently strong, the agent will increase saving above the floor. Once they have done this, their optimal level of saving will be pinned down by their choice under the full match:

     $$\hat{S}^{\text{Partial Match}} > S_{\text{Full Match}}$$

     See Figure 17c.

   - If the match is not attractive enough to offset the required cut in consumption, the agent will ignore the match and save as in the No Match case:

     $$\hat{S}^{\text{Partial Match}} = S^{\text{No Match}}$$

See Figure 17d.

There will be some threshold of saving propensities (α_i in the utility function, or equivalently some saving level in the 'No Match cases') which will separate the two cases. Unlike the concave case, where bunching at a cap is optimal, we predict no bunching at the convex kink. The prediction depends on how far $S^{\text{No Match}}$ lies below S_{floor}.

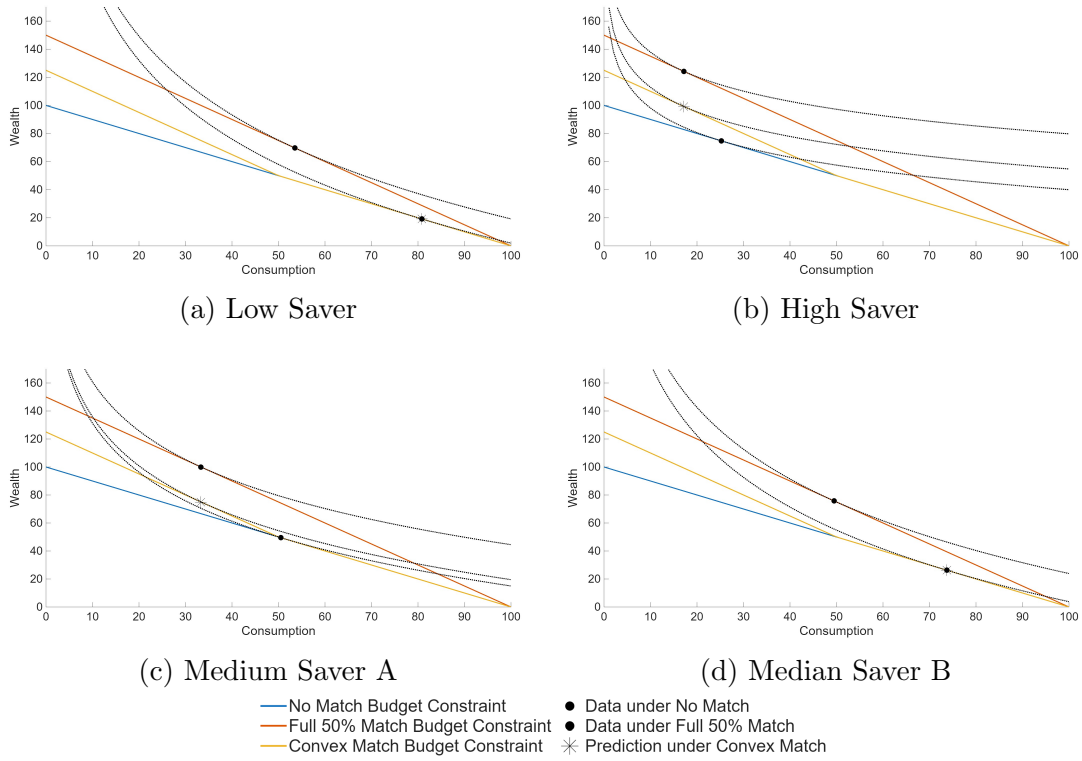


Figure 17: Model Predictions for one-tier convex match formulas

A.2 Proof cases for concave one-tier match formula

A.2.1 Case 1: Low Saver

Assumption: $S^{\text{Full Match}} < S^{\text{cap}}$
Claim: $\hat{S}^{\text{Partial Match}} = S^{\text{Full Match}}$
Proof: Under a full match ($M(S) = mS$), the problem becomes:

$$\max_S \alpha_i \ln(Y - S) + (1 + m)S$$

The first-order condition (FOC) is:

$$-\frac{\alpha_i}{Y-S}+(1+m)=0 \Rightarrow S^{\text{Full Match}}=Y-\frac{\alpha_i}{1+m}$$

Since $S^{\text{Full Match}}<S^{\text{cap}}$, the cap is non-binding, and the match structure is locally identical to full match. Therefore,

$$\hat{S}^{\text{Partial Match}}=S^{\text{Full Match}}$$

A.2.2 Case 2: Medium Saver

Assumptions: $S^{\text{Full Match}}>S^{\text{cap}}$, and $S^{\text{No Match}}<S^{\text{cap}}$
Claim: $\hat{S}^{\text{Partial Match}}=S^{\text{cap}}$
Proof:

- Under **no match**:

$$\max_S \alpha_i \ln(Y-S)+S \Rightarrow S^{\text{No Match}}=Y-\alpha_i$$

- Under **partial match**:

$$M(S)=\begin{cases} mS & \text{if } S\le S^{\text{cap}} \\ mS^{\text{cap}} & \text{if } S>S^{\text{cap}} \end{cases}$$

  For $S\le S^{\text{cap}}$, FOC:

$$-\frac{\alpha_i}{Y-S}+(1+m)=0 \Rightarrow S=Y-\frac{\alpha_i}{1+m}$$

  But this implies $S>S^{\text{cap}}$, violating the domain.

  For $S>S^{\text{cap}}$, FOC:

$$-\frac{\alpha_i}{Y-S}+1=0 \Rightarrow S=Y-\alpha_i=S^{\text{No Match}}<S^{\text{cap}}$$

  So neither interior solution is feasible.

- Therefore, optimal choice is at the **kink point**, $S=S^{\text{cap}}$.

A.2.3 Case 3: High Saver

Assumption: $S^{\text{No Match}}>S^{\text{cap}}$
Claim: $\hat{S}^{\text{Partial Match}}=S^{\text{No Match}}$
Proof:

- For $S>S^{\text{cap}}$, match is constant: $M(S)=mS^{\text{cap}}$

- Utility becomes:

$$\max_S \alpha_i \ln(Y - S) + S + mS^{\text{cap}} \Rightarrow \text{FOC: } -\frac{\alpha_i}{Y - S} + 1 = 0 \Rightarrow S = Y - \alpha_i = S^{\text{No Match}}$$

- Since $S^{\text{No Match}} > S^{\text{cap}}$, this solution is feasible under the cap.

A.3 Predicting Saving under General Match Formulas

The logic developed above generalizes beyond the one-tier match case. Suppose we can observe how an individual saves under each constant rate of return that appears in a given match formula. Then, using a piecewise model of the match and the structure of quasilinear utility, we can determine which piece binds and infer optimal saving under the full, non-linear formula.

Concave multi-tier match formulas. Consider a match formula that pays 100% on saving up to level S_1, and 25% on saving between S_1 and S_2:

$$M(S) = \begin{cases} m_1 S & \text{if } S \leq S_1 \\ m_1 S_1 + m_2(S - S_1) & \text{if } S_1 < S \leq S_2 \\ m_1 S_1 + m_2(S_2 - S_1) & \text{if } S > S_2 \end{cases}$$

Suppose we observe:

- $S^{\text{No Match}}$: the saving under no match (0% return),
- S^{m_2}: the saving under a uniform match at rate m_2,
- S^{m_1}: the saving under a uniform match at rate m_1.

Then we can apply the same three-case logic to each segment:

- If $S^{m_2} < S_1$, the m_2 tier never binds, so the agent saves S^{m_2}.
- If $S^{m_1} > S_1$ but $S^{m_2} < S_1$, then S_1 is a kink, and the individual will bunch there.
- If $S^{\text{No Match}} > S_2$, the cap is irrelevant and the agent saves as under no match.

Thus:

$$\hat{S}^{\text{Partial Match}} = \begin{cases} S^{m_2} & \text{if } S^{m_2} < S_1 \\ S_1 & \text{if } S^{m_1} > S_1 \text{ and } S^{m_2} < S_1 \\ S^{\text{No Match}} & \text{if } S^{\text{No Match}} > S_2 \end{cases}$$

Convex match formulas. Now consider a match formula in which the matching rate *increases* with saving. For example:

$$M(S) = \begin{cases} 0 & \text{if } S \leq S_1 \\ m_1(S - S_1) & \text{if } S_1 < S \leq S_2 \\ m_1(S_2 - S_1) + m_2(S - S_2) & \text{if } S > S_2 \end{cases}$$

Suppose we observe:

- $S^{\text{No Match}}$
- S^{m_1}
- S^{m_2}

Then we can again apply threshold comparisons:

- If $S^{m_2} < S_1$, the agent never reaches the match region: save as under no match.
- If $S^{\text{No Match}} > S_2$, the agent fully engages with the highest tier: save as under m_2.
- If $S^{m_2} > S_2$ but $S^{\text{No Match}} < S_1$, then prediction depends on how far $S^{\text{No Match}}$ lies from S_1, as in the one-tier convex medium saver case.

This logic holds because, with quasilinear utility, the marginal return to saving fully determines the agent's optimal choice. and at each segment, that marginal return is known. With observed behavior under each constant return that features in the match formula, we can recover the necessary first-order conditions to predict saving under the full formula.

B Additional Exhibits

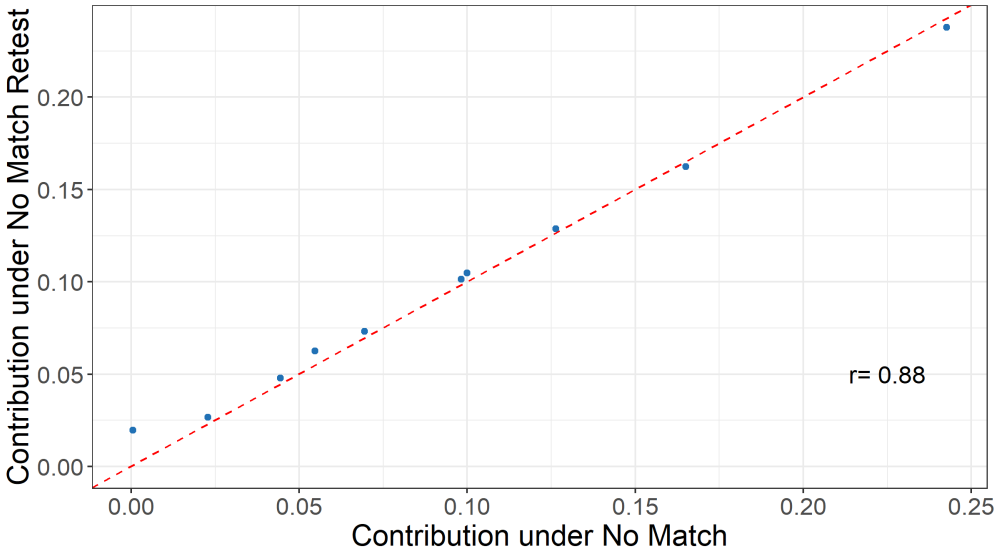


Figure 18: Test and Retest Responses to Saving Under No Match

Table 7: Moving to the frontier two tiers

Current Plan	CV Minimum			TE Maximum		
	Optimal Plan	Optimal NEC	Change CV	Optimal Plan	Optimal NEC	Change TE
100-5-0-0	50-1-25-6.75	3.1	-0.57	50-2-25-9.75	2.4	0.42
100-3-50-5	50-1-25-6.25	2.2	-0.46	50-1-25-8.75	1.8	0.37
100-6-0-0	50-2-25-7.75	3.5	-0.62	50-2-25-13.5	3	0.39
50-6-0-0	50-1-25-6.5	1.1	-0.37	50-1-25-8.5	0.8	0.29
100-1-50-6	50-1-25-7.5	1.4	-0.21	50-1-25-8.75	1.2	0.14
100-1.25-0-0	25-1.75-0-0	0.8	-0.58	50-1-25-2.25	0.5	0.94
100-4-0-0	50-1-25-4.75	2.5	-0.71	50-2-25-7.75	1.8	0.6
50-8-0-0	50-2-25-9.25	1.1	-0.32	50-2-25-12.5	0.8	0.2
50-4-0-0	50-1-25-3.75	0.8	-0.43	50-1-25-5.25	0.5	0.33
100-3-50-6	50-1-25-7.5	2.4	-0.41	50-2-25-9.75	1.8	0.26
100-3-0-0	50-1-25-3	2	-0.81	50-1-25-6	1.4	0.7
25-4-0-0	50-1-25-2.5	0.1	-0.37	50-1-25-2.75	0	0.11
100-4-50-6	50-1-25-7.5	2.8	-0.54	50-2-25-11.5	2.1	0.34
50-10-0-0	50-2-25-14.25	1.3	-0.21	25-18.5-0-0	1.5	0.2
100-8.5-0-0	25-17.5-0-0	5.2	-0.36	100-1-50-12.75	2.5	0.18
100-8-0-0	25-16.5-0-0	4.9	-0.32	25-20-0-0	4.7	0.2
100-2-50-6	50-1-25-7.5	1.9	-0.29	50-1-25-9.25	1.7	0.18
50-5-0-0	50-1-25-5	0.9	-0.42	50-1-25-6.75	0.6	0.34
100-5-50-6	50-1-25-7.75	3.3	-0.61	50-2-25-12.75	2.5	0.37
50-7-0-0	50-1-25-7.75	1.2	-0.38	50-2-25-9.75	0.7	0.24
25-6-0-0	50-1-25-4.5	0	-0.32	-	-	0
100-4-50-8	50-2-25-12.25	2.8	-0.24	25-16.25-0-0	3	0.19
100-3-50-7	50-2-25-9.25	2.2	-0.29	50-2-25-11.75	2	0.16
100-9-0-0	25-19.5-0-0	5.4	-0.36	100-1-50-13.75	2.7	0.16
25-12-0-0	-	-	0	-	-	0

Table 8: Moving to the frontier with single-tier plans

Current Plan	CV Minimum			TE Maximum		
	Optimal Plan	Optimal NEC	Change CV	Optimal Plan	Optimal NEC	Change TE
100-5-0-0	25-8-0-0	3.2	-0.37	25-10.75-0-0	2.8	0.31
100-3-50-5	25-7.75-0-0	2.2	-0.21	25-9-0-0	2	0.19
100-6-0-0	25-10-0-0	3.7	-0.35	25-14-0-0	3.4	0.35
50-6-0-0	25-8-0-0	1.1	-0.12	25-8.75-0-0	1	0.09
100-1-50-6	-	-	0	-	-	0
100-1.25-0-0	25-1.75-0-0	0.8	-0.58	25-2.75-0-0	0.6	0.58
100-4-0-0	25-6.5-0-0	2.5	-0.35	25-8-0-0	2.2	0.38
50-8-0-0	25-11.25-0-0	1.4	-0.15	25-12.75-0-0	1.3	0.13
50-4-0-0	-	-	0	-	-	0
100-3-50-6	25-9.5-0-0	2.4	-0.13	25-10.5-0-0	2.2	0.11
100-3-0-0	25-4.75-0-0	1.8	-0.38	25-6.25-0-0	1.6	0.39
25-4-0-0	-	-	0	-	-	0
100-4-50-6	25-9.5-0-0	2.8	-0.25	25-11.75-0-0	2.6	0.26
50-10-0-0	25-15-0-0	1.7	-0.19	25-18.5-0-0	1.5	0.2
100-8.5-0-0	25-17.5-0-0	5.2	-0.36	25-20-0-0	5.1	0.14
100-8-0-0	25-16.5-0-0	4.9	-0.32	25-20-0-0	4.7	0.2
100-2-50-6	25-9.5-0-0	1.9	0	-	-	0
50-5-0-0	25-6.75-0-0	0.9	-0.08	25-7-0-0	0.9	0.1
100-5-50-6	25-9.75-0-0	3.2	-0.33	25-13.25-0-0	2.9	0.32
50-7-0-0	25-9.5-0-0	1.2	-0.13	25-10.5-0-0	1.1	0.09
25-6-0-0	-	-	0	-	-	0
100-4-50-8	25-13.5-0-0	3.2	-0.18	25-16.25-0-0	3	0.19
100-3-50-7	25-11.25-0-0	2.5	-0.12	25-12.5-0-0	2.4	0.1
100-9-0-0	25-19.5-0-0	5.4	-0.36	50-14-0-0	3.2	0.09
25-12-0-0	-	-	0	-	-	0