Improving 401(k) Matches Using Hypothetical Choices*

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

May 2026

Abstract

How should employer 401(k) matching formulas, which allocate $250 billion annually, be designed to raise employee saving and reduce inequality in employer contributions? We use survey responses to hypothetical scenarios to predict how individuals would save under counterfactual policies. We then characterize the frontier of achievable saving-equity combinations. We find that survey responses accurately predict contribution choices in administrative 401(k) data, employee contributions are inelastic to the match rate, and non-elective contributions do not crowd out employee saving. Therefore, a lower match rate applied up to a higher cap paired with a non-elective contribution achieves higher savings and more equitable employer contributions. Many existing formulas, including those designated as safe harbors by regulation, are 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 James Choi, Janet Currie, Aaron Goodman, Kelly Hahn, Anna Madambda, Olivia Mitchell, Jonathan Parker, Alessandro Previtero, Ben Polak, and Dmitry Taubinsky 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. employers contribute over $250 billion to their employees' 401(k) plans, using more than 1,000 matching formulas that vary in their match rates, caps, and tier structure.[1] Employers choose these formulas as part of their compensation policy subject to strict regulatory constraints: federal non-discrimination rules limit the share of those contributions that can accrue to highly paid workers, and around half of firms adopt one of a small number of "safe harbor" matching formulas that automatically satisfy those rules. These formulas shape both how much workers save and how employer dollars are distributed across them, making their design a key feature of the U.S. retirement system. Yet they have never been evaluated against cost-equivalent alternatives. Two empirical patterns suggest there is substantial room for improvement. First, many employees do not take up matching incentives, especially low earners and those with lower-income parents (Choukhmane et al., 2024). Second, as we show, most match dollars accrue to inframarginal savers (employees contributing strictly above the maximum matched amount), and therefore do not affect the marginal returns to saving.

This paper asks how to improve the design of employer matching formulas. Answering this question requires solving two distinct problems: predicting behavior under counterfactual policies, and mapping that behavior into welfare-relevant criteria. We tackle these separately. For prediction, we combine administrative data from 1.8 million Vanguard participants with a novel survey eliciting responses to counterfactual matching designs, including formulas that are rare or non-existent in practice. We use the survey responses to build a predictive model of saving behavior that closely replicates contribution choices in the administrative data. This validates the use of hypothetical choices to predict contributions under a wide range of counterfactual policies, without needing to specify the distribution of heterogeneity in preferences or constraints.

For policy evaluation, we make progress without committing to a specific social welfare function. Two approaches point to similar welfare-relevant metrics. The first approach infers objectives from the revealed preferences of U.S. policymakers: since the Revenue Act of 1942, federal policy has aimed both to expand retirement saving and to prevent tax-preferred benefits from accruing disproportionately to highly compensated workers. The second approach derives objectives from a welfare framework: since Diamond (1977) and Feldstein (1985), the literature has emphasized that raising private saving can improve welfare when people undersave for retirement. Such undersaving might arise for rational reasons when people an-

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

ticipate relying on means-tested safety-net programs in retirement, or for behavioral reasons, such as present bias or inattention. Both perspectives point to similar evaluation criteria: how much a formula raises average employee saving, and how those gains are distributed across workers. We show that formulas with a more equal distribution of employer contribution rates are associated with better targeting: they raise total contributions more among lower-income workers (who rely more on means-tested safety nets in retirement) and among those who save little despite generous incentives (who are more likely to face behavioral frictions). These are the workers a corrective intervention is intended to reach, so targeting considerations alone favor lower inequality in employer contribution rates. A redistributive motive, if added, only sharpens the case. We then characterize the frontier of achievable equity-savings combinations, identifying formulas that are dominated in the sense that cost-equivalent alternatives improve on both dimensions simultaneously. We find that the most widely used matching formulas, including nearly all safe harbors, lie well inside this frontier.

Not specifying a social welfare function precludes identifying a single optimal policy. Two considerations motivate this choice. First, as Manski (2011, 2013) has cautioned, precise policy recommendations often rest on assumptions that are difficult to validate, what he calls "incredible certitude." Our frontier approach trades a single optimum for conclusions that are robust to a range of modeling choices. Second, optimal policies derived from economic models are rarely implemented in practice, whereas improvements over dominated formulas are within reach and the safe harbors offer a clear target for reform. The result is concrete guidance on a $250 billion policy without the strong assumptions such quantitative conclusions often require.

Employers ultimately choose their own formulas and will likely weigh objectives other than inequality and saving. But because safe harbors are widely adopted, reforming them offers a powerful lever to induce the adoption of new formulas even when employer objectives differ. A natural concern remains that as employers change their matching formulas to satisfy the new safe harbor, other components of compensation might adjust in equilibrium, undoing the intended effects of the reform. Whether this concern has force depends on the rationale for reform. If the aim is to correct undersaving, the reforms we identify remain attractive even if wages adjust one-for-one to offset the change in employer contributions: workers still end up with more allocated to retirement and less to current pay, and that reallocation helps most when it targets those more at risk of undersaving. If the aim is to redistribute, the concern carries more weight. But because the safe harbor formulas are used by around half of employers and over half of employees (shares that are growing), reforming them would shift retirement benefits at both workers' current jobs and their likely outside options, limiting the scope for offsetting compensating differentials.

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 matches, and yet only generate modestly more additional saving. Second, non-elective contributions (i.e., employer contributions made regardless of employee saving) do not appear to crowd out employee contributions. This holds both in survey responses (where hypothetical non-elective contributions do not reduce hypothetical employee contributions) and administrative data (where actual increases in non-elective contributions do not reduce actual employee contributions). Third, bringing together the survey evidence and the predictive model, we characterize the frontier of the equity-savings 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 increasing contributions requires more inequality. Formulas on the frontier tend to pair lower match rates with non-elective contributions. Our characterization is robust to using alternative measures of targeting: formulas that dominate on our equity metric better target the increase in contributions to workers who are lower income, would otherwise save little, and are less likely to take up financial incentives.

Fourth, and most importantly for current U.S. policy, many existing retirement formulas, including nearly all safe harbor formulas, are dominated: cost-neutral alternatives could increase saving, reduce inequality in employer contributions, or do both simultaneously. All 10 of the most widely used matching formulas in our sample, and 22 of the top 25, lie well inside the equity-savings frontier. Moving these plans to the saving-maximizing alternatives we identify would raise average employee contributions by about 0.4 pp of pay (a 7.0% increase in average employee contribution rates), with broad-based gains that extend to non-participants and financially constrained workers. Moving them instead to inequality-minimizing alternatives could reduce the coefficient of variation of employer contribution rates by nearly 40%, primarily by redistributing employer contributions from current participants to non-participants. Extrapolating to a broader sample of safe harbor plans from Form 5500 filings in 2023, we estimate that adopting the saving-maximizing alternatives would generate at least $6.7 billion in additional annual employee retirement contributions, at no additional cost to employers, and without exacerbating inequality.

We contribute to three strands of the literature. The first concerns the causal effects of saving incentives. This question has been studied with a variety of designs: using cross-sectional variation in match schedules (e.g., Engelhardt and Kumar (2007); Mitchell et al. (2007)), firm match rate changes (e.g., Papke (1995); Kusko et al. (1996); Choi et al. (2002)), tax incentive changes (e.g., Ramnath (2013); Chetty et al. (2014)) and experiments (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 contributions is ambiguous. We contribute to this literature by addressing a fundamental challenge: the difficulty of identifying responses to multidimensional and heterogeneous saving incentives that are often not observed in reality. Our approach estimates stated contribution preferences across match formulas that extend well beyond those observed empirically. By combining a model with 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 incentives.

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), the literature has characterized optimal paternalistic retirement saving policies. These models typically assume that individuals undersave either because of behavioral biases (e.g., Laibson (1997), Yu (2021), Beshears et al. (2025)) or because of the disincentives created by means-tested social safety nets (e.g., Hubbard et al. (1995), Sleet and Yeltekin (2006)). 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)). Recent work considers optimal savings policies under joint concerns for retirement undersaving and income inequality (Moser and Olea de Souza e Silva (2019); Choukhmane and Palmer (2025)). This literature provides important insights for the design of saving policies, but it often relies on stylized settings that abstract from the complexities of actual saving incentives. Real-world matching formulas involve multiple design parameters (match rates, caps, and often multiple tiers) and must account for 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 literature that combines administrative data, models of economic behavior, and custom surveys (for reviews, see Bachmann et al. (2022); Fuster and Zafar (2023); Stantcheva (2023); Caplin (2025)). These surveys bring rich variation to identify preferences, beliefs, and constraints that are generally missing from 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) show that survey responses correlate with behavior in administrative data, Allcott et al. (2022) combine survey-elicited beliefs with administrative payday loan records to calibrate a behavioral model of payday borrowing, while Ameriks et al. (2016, 2020) and Indarte et al.

(2025) use surveys to calibrate key parameters in lifecycle consumption models. Similar to our approach, Colarieti et al. (2024) show that survey responses to hypothetical scenarios (namely marginal propensities to save and borrow) align quantitatively with empirical estimates in the literature. We contribute to this literature by showing how survey responses to hypothetical policies not only predict actual behavior in administrative data but also enable quantitative policy evaluation.

The remainder of the paper proceeds as follows. Section 2 details the institutional setting and describes contribution behavior using the administrative data. Section 3 presents our survey design and results. Section 4 introduces a simple predictive model of retirement contributions. Section 5 brings together the model and survey data. Section 6 characterizes the savings-equity frontier and evaluates existing plans, and Section 7 concludes.

2 Matching Incentives: Institutional Context and Key Empirical Patterns

2.1 Employer Matches in the U.S.

Defined contribution (DC) plans are the dominant vehicle for private retirement saving in the United States. Two-thirds of U.S. workers have access to employer-sponsored accounts (U.S. Bureau of Labor Statistics (2024)), and Americans hold over $14 trillion of assets within them (Investment Company Institute (2026)).[2] 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, withdrawal options, and fund offerings.

[2]This dollar value underestimates the reach of employer-sponsored DC accounts; if IRAs are included, DC wealth is over $30 trillion, and most IRA balances represent rollovers from employer-sponsored accounts.

These plan design features may also influence whether the worker takes full advantage of the match and associated employer costs.[3]

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. Appendix A provides more details; here we provide a short overview.[4]

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 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) shows 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

Our primary administrative data come from records for Vanguard plan participants in plans for which Vanguard manages non-discrimination testing on behalf of the plan sponsor. These records include W-2 earnings reported directly by plan sponsors, plan formula details (match rates, caps, and tiers), and annual employee and employer contributions, with employer contributions separated into matching and non-elective components. We also observe whether each plan has auto-enrollment and, if so, the date of adoption. Plan sponsors additionally provide basic demographic information: age, gender, and tenure within the plan.

We complement these annual records with a monthly panel of elected contribution rates, capturing the sequence of rates participants actively choose over time.[5] We use the contribution rate in effect at the time of survey participation to validate our model predictions.

[3]For an analysis of vesting, see Carranza and Goodman (2024), of loans see Lu et al. (2014), and of plan quality and fees see Bhattacharya and Illanes (2025).

[4]For a historical account of the evolution of non-discrimination testing, safe harbor plans, and a thorough discussion of the current rules, see Israel Preninger and O'Dea (2026).

[5]We draw on a dataset of monthly contribution rates constructed by Aaron Goodman and used to measure income volatility among 401(k) participants in Goodman (2025).

2.4 Two Facts about DC Retirement Saving Behavior

Before analyzing counterfactual policies, we document two facts about defined contribution saving in our data. Together, they suggest current match formulas are poorly targeted. First, many workers do not capture the full match. Second, most of the employer's match budget goes to workers whose saving the match likely did not change.

1. Take-up is incomplete, especially among low-income workers. Figure 1(a) shows the fraction of workers who take up the full match across income groups. Only a small share of those with the very lowest income take up the full match offered to them, and just over half of those earning between \$30k and \$60k take up their full match. Furthermore, despite strong positive gradients in saving, many high earners do not contribute enough to receive the full employer match. Take-up remains incomplete even among very high earners: around 20% of workers earning above \$150,000 save below their match cap.

2. Most match dollars go to inframarginal savers. Figure 1(b) provides a breakdown of the fraction of employees and employer contribution dollars that accrue to different groups of savers. The final cluster of columns shows that close to half 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.

3 Survey Design and Empirical Patterns

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 use a survey that elicits hypothetical contributions from the same individuals under alternative match formulas. In this section we introduce the survey design and describe several patterns in the data. These patterns motivate our modeling assumptions, and the responses themselves serve as inputs to the predictive model introduced in Section 4.

3.1 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 as inputs to our model. In addition, we collect information on other respondent characteristics which are difficult to assess using administrative data and help to evaluate targeting of the savings responses induced by different formulas.

Figure 1: Match Take-Up and the Distribution of Employee Contributions

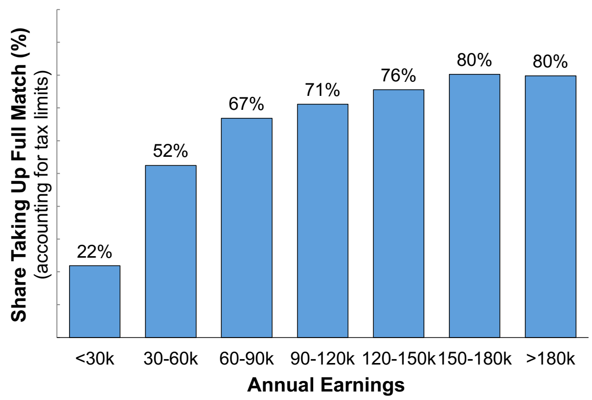

(a) Match Take Up by Earnings

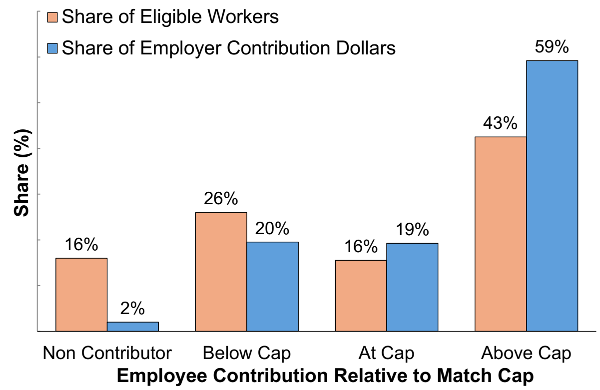

(b) Distribution of workers and match dollars

Notes: Panel (a) reports the share of workers fully taking up their employer match, by annual earnings. A worker is classified as fully taking up the match if their contribution is sufficient to receive the maximum match available to them after accounting for IRS annual contribution and compensation limits. Panel (b) partitions eligible workers by their employee contribution relative to the employer match cap: non-contributors, below, at, and above the cap. We report the share of workers and the share of total employer contribution dollars in each group. Sample: population in Vanguard administrative data in 2023.

Survey Design The core of the survey was a series of hypothetical contribution vignettes, in which respondents were asked to choose how much to contribute under several retirement plan designs. Before viewing the vignettes, they were shown a set of assumptions applying to all scenarios (see Appendix Figure C.2(a)): employer contributions would be immediately vested, employee contributions would receive the same tax treatment as in their current plan, and in scenarios where the match rate changed, all other aspects of their jobs would remain the same.

Following this, respondents were asked how much they would contribute under each hypothetical match schedule. Each vignette followed a standard format. Figure C.2(b) shows the interface for the "No Match" vignette. Respondents used a slider to choose their contribution rate as a percentage of income. As respondents moved the slider, 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%, 50%, and 100%.

3. One vignette with a match cap of 6%, randomly assigned a match rate of 25%, 50%, or 100%.

4. One vignette with no match but a universal employer contribution of 10% of pay.

The vignettes in items 1 and 2 provide the main inputs for predicting matching contributions under alternative plan designs, capturing $(s(0,.), s(0.25,1), s(0.5,1), s(1,1))$. The vignettes in item 3 are used to test the consistency of responses for respondents who actually have that match in reality. 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 financial literacy, liquidity constraints, and self-assessed savings adequacy. Appendix F reproduces the verbatim text of every question used in our analysis.

Survey Implementation The survey was fielded in February 2025. The sample consisted of eligible employees for whom we had administrative data on contributions and income in 2023. Table C.1 compares the respondent sample to the full administrative population across age, income, and tenure. The final sample matches the population in terms of income and tenure but includes a somewhat higher share of older participants. In our empirical analysis, we re-weight responses to match the joint distribution of age and income.

3.2 Survey Internal Validity

To test the internal validity of the survey responses, we compare survey responses to administrative data when the same measure is available in both sources. First, we compare self-reported contribution rates to employer-sponsored retirement accounts with the actual contribution rates recorded in administrative data. Appendix Figure C.3 shows the comparison. The horizontal axis groups respondents into deciles of the administrative contribution rate. For each decile, we plot the 25th, 50th, and 75th percentiles of the stated rate. Median stated responses track the administrative rates closely through the middle of the distribution, with a small divergence among the top two deciles.

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. Appendix Figure C.4 bins individuals by the administrative contribution rate, and compares that response to percentiles of the stated responses. In both panels, the median responses match those in the data, though with greater dispersion than in Appendix Figure C.3, partly reflecting smaller sample sizes. Together, these exercises provide evidence that respondents both report their actual behavior accurately and respond meaningfully to the hypothetical incentives presented in the survey.

As a final piece of internal validation, Appendix Figure C.5 examines internal consistency by comparing responses to a "no match" scenario elicited at the start of the survey with

Figure 2: Survey Treatment Effects and Cost Effectiveness, by Match Rate

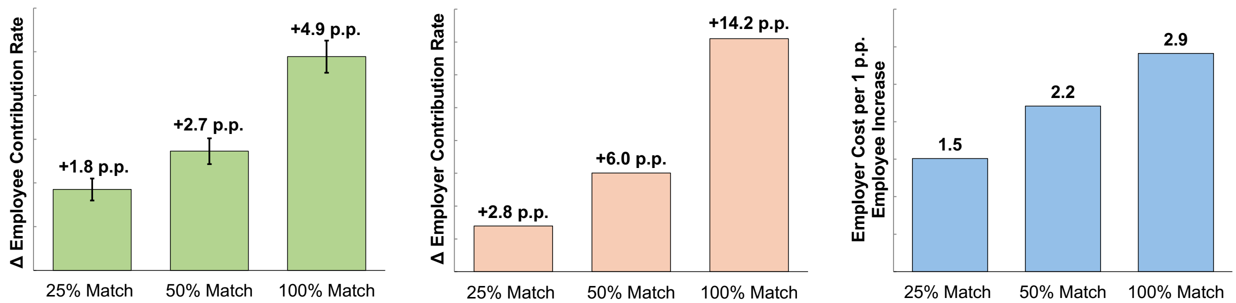

(a) Employee Contributions (b) Employer Contributions (c) Δ Employer / Δ Employee

Notes: This figure reports estimated treatment effects of introducing unlimited employer matching at different match rates, relative to a no-match baseline. Panel (a) shows the average change in employee contribution rates (as a share of earnings), calculated as the difference between saving under the match and saving under no match, controlling for participant fixed effects. Panel (b) shows the corresponding change in employer contribution rates. Panel (c) reports the ratio of the employer cost to the employee contribution response, i.e., the employer cost per percentage point increase in employee saving. Error bars in panel (a) represent 95% confidence intervals.

responses to the same scenario elicited at the end. The median responses at the start and end of the survey show a close match, indicating that answers are stable across the survey and not sensitive to fatigue.

3.3 Survey Evidence on Saving Responses

Before constructing model predictions from the survey responses, we present initial evidence on the economic mechanisms reflected in the survey data.

Employee contributions are inelastic to the match rate. Figure 2(a) illustrates how saving responds to different match rates offered without a cap. It shows the average increase in saving relative to the no-match baseline. Under unlimited matching, a 25% match rate increases saving by 1.8 pp; a 50% match rate yields 2.7 pp in extra saving; and a 100% match rate increases saving by approximately 4.9 pp.

Low match rates are more cost effective than high match rates. The patterns in Figure 2(a) suggest that while saving is responsive to the presence of a match, the elasticity with respect to the match rate is modest. Figure 2 panels (b) and (c) illustrate this point: panel (b) shows the employer cost and panel (c) shows the cost per percentage point increase in employee saving associated with the employee savings changes from panel (a). If the goal of the employer match is to encourage saving, lower match rates deliver substantially greater value per dollar spent. An unlimited 25% match generates one additional percentage point of employee saving at a cost of 1.5% of payroll. The same one percentage point increase in employee contributions costs 2.9% of payroll under a 100% match.

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 cause employees to reduce their own contributions. Our model assumes no such crowd-out. We present two pieces of evidence supporting this assumption.

First, we examine survey responses to a hypothetical scenario in which the employer contributes 10% of pay as a non-elective contribution. Appendix Figure C.6 plots stated contributions under this scenario against stated contributions under no match. Across the distribution, median stated contributions under the non-elective contribution are at least as high as those under no match, with no decile in which respondents reduce their saving in response. If anything, respondents in the lower deciles report somewhat higher contributions when offered the non-elective contribution, though these differences are small relative to the dispersion of responses. The pattern is inconsistent with crowd-out.

Second, we exploit discrete increases in non-elective contribution rates in our administrative data. We identify 20 plans that experienced sustained increases in their non-elective contribution rate during the period of observation: the 10 largest such plans with auto-enrollment and the 10 largest without. We pool across plans and estimate an event-study specification with plan fixed effects, year fixed effects, and demographic controls.[6] Figure 3 reports the results. The left panel confirms a sharp, persistent increase in NEC rates at the event date with no pre-trend. The right panel shows that employee contribution rates are very slightly increasing both before and after the NEC increase, with point estimates close to zero and confidence intervals that rule out economically meaningful crowd-out. This result has also been found by Card and Ransom (2011) and Chetty et al. (2014). The absence of crowd-out does not appear to be driven by inertia among employees whose plan changed mid-tenure. Appendix Figure C.7 shows similar results when we restrict the sample to new participants, defined as those observed only in their year of hire.

Across both the survey and administrative evidence, we find no indication that employer non-elective contributions crowd out participants' voluntary saving.

[6]Figure C.9 in the appendix plots the raw time series of non-elective contributions and employee contribution rates for each plan. Despite large jumps in NEC rates, employee contributions remain on stable trends with no visible breaks.

Figure 3: Change in Saving around Non-Elective Contribution Rate Increases

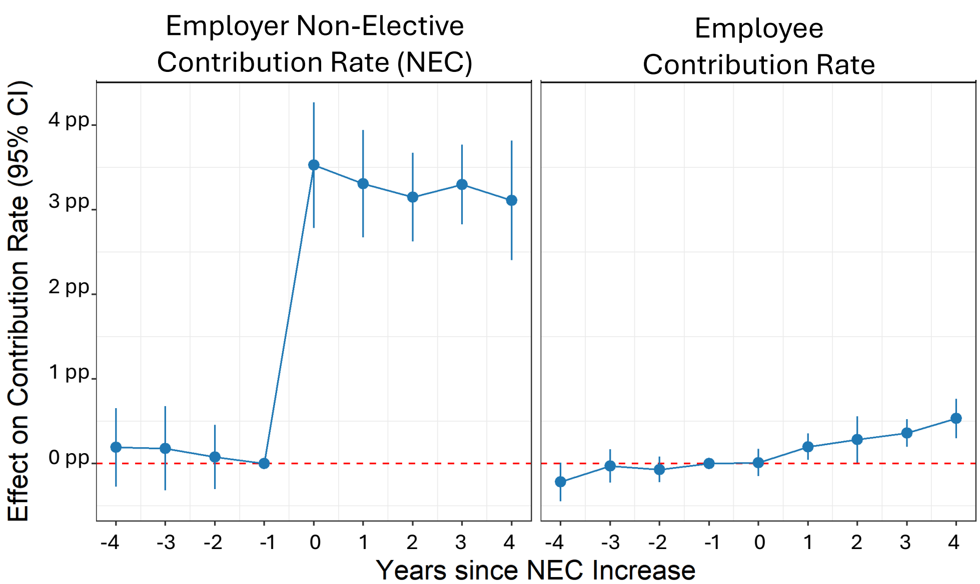

Notes: The figure plots event-study estimates from a staggered difference-in-differences design. The sample is a panel covering 2010 to 2023 and includes the 10 largest auto-enrollment (AE) and 10 largest non-auto-enrollment (non-AE) plans that experienced a non-elective contribution (NEC) rate increase. The specification controls for plan and year fixed effects and quadratics in age, income, and tenure. The left panel plots the estimated effect on the NEC rate and the right panel plots the estimated effect on the elective contribution rate. Standard errors are clustered at the plan level. See Figure C.7 for a version of this graph which uses only new participants.

4 From Survey Responses to Counterfactual Predictions

In this section, 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.

Let $s(m, c)$ denote the contribution rate an agent chooses under a match rate of m with a cap at $c \in [0, 1]$, where c is expressed as a proportion of salary. We show that under three testable assumptions, knowledge of two objects for an individual: $s(0, .)$ and $s(m, 1)$ (saving under no match, and saving under an unlimited match) is sufficient to predict $s(m, c)$ for any cap c.[7]

4.1 Assumptions

We impose three restrictions on saving behavior:

Assumption 1 (Incentives matter) *If $s(m, c) = c$ and $c' > c$ then*

$$s(m, c') \geq s(m, c)$$

Interpretation: When an employee is saving at the cap, raising the cap does not reduce contributions. This is a weak monotonicity condition: when a binding cap is relaxed, the marginal return to saving increases, and contributions must respond (weakly) positively.

Assumption 2 (Irrelevance of non-binding caps) *If $s(m, c) < c$ and $c' > c$ then*

$$s(m, c') = s(m, c)$$

Interpretation: When a non-binding cap is increased, savings do not change. Employees who contribute below the cap are unaffected by further increases in the cap. This assumes that agents respond only to the marginal incentive they face, not to incentives at higher contribution levels.

Assumption 3 (No wealth effects) *If $s(m, c) > c$ then*

$$s(m', c) = s(m, c) \quad \forall \quad m' \quad \textit{and} \quad s(m, c') = s(m, c) \quad \forall \quad c' < c$$

[7] Note that with no match, the cap is irrelevant, hence we do not give a second argument in $s(0, .)$.

Interpretation: When an employee contributes strictly above the cap, an inframarginal change in the cap or the match rate does not alter marginal incentives. Such changes (reductions in the cap that remain below the employee's contribution, or changes in the match rate) therefore generate pure income effects. This assumption requires that these income effects do not affect employee contributions. It is a strong but testable restriction, motivated by the evidence in Section 3.3 that non-elective contributions do not crowd out employee contributions, both in the survey responses and the administrative data.

These assumptions are mutually consistent and empirically testable.

4.2 Prediction Rule

Our aim is to predict saving under an arbitrary match rate $\hat{m}$ and cap c. The key insight is that the two observables $s(0,.)$ and $s(\hat{m},1)$ allow us to classify each individual into one of three mutually exclusive cases, each with a distinct prediction.

- **Case 1: High saver**: $s(0,.) > c$.

  Interpretation: This individual would save above c even with no match.

  - From Assumption 3, since $s(0,.) > c$, we obtain $s(\hat{m},c) = s(0,.)$.
  - Relative to the no match case, introducing a match provides an income effect but does not change the marginal return to saving above c.

- **Case 2: Medium saver**: $s(0,.) \leq c$ and $s(\hat{m},1) \geq c$.

  Interpretation: This individual would save no less than c with an unlimited match but no more than c with no match.

  - Suppose $s(\hat{m},c) < c$. By Assumption 2, $s(\hat{m},1) = s(\hat{m},c) < c$, which contradicts the medium-saver condition. Therefore $s(\hat{m},c) \geq c$.
  - Suppose $s(\hat{m},c) > c$. By Assumption 3, $s(0,.) = s(\hat{m},c) > c$, hence $s(0,.) > c$, which contradicts the medium-saver condition. Therefore $s(\hat{m},c) \leq c$.
  - Together, these yield $s(\hat{m},c) = c$. The individual saves exactly at the cap.

- **Case 3: Low saver**: $s(0,.) \leq c$ and $s(\hat{m},1) < c$.

  Interpretation: This individual would save below c even with an unlimited match. The cap is therefore irrelevant.

  - Suppose $s(\hat{m},c) \geq c$. We consider two subcases:

    * If $s(\hat{m}, c) = c$, then by Assumption 1, $s(\hat{m}, 1) \geq s(\hat{m}, c) = c$, which contradicts $s(\hat{m}, 1) < c$.
    * If $s(\hat{m}, c) > c$, then by Assumption 3, $s(0, .) = s(\hat{m}, c) > c$, hence $s(0, .) > c$, which contradicts $s(0, .) \leq c$.

  Both subcases yield a contradiction, so $s(\hat{m}, c) < c$.

  – Since $s(\hat{m}, c) < c$, the cap is non-binding. By Assumption 2, $s(\hat{m}, c) = s(\hat{m}, 1)$.

These three cases are exhaustive and mutually exclusive. Figure 4 summarizes the prediction rule. Panel (a) fixes the cap and classifies individuals. Those whose no-match saving exceeds the cap are the 'high savers' (by Assumption 3, the match is not relevant to their behavior as it represents a pure income effect). Those whose unlimited-match saving exceeds the cap but whose no-match saving does not are the 'medium savers' (they bunch at the cap). Those whose unlimited-match saving falls below the cap are the 'low savers' (the cap does not bind). Panel (b) offers a different way to visualize the prediction rule. It holds fixed individuals' responses to the questions and varies the cap, tracing out the three regimes as a function of c: flat at $s(0, .)$ when the cap is low, rising along the 45-degree line when the individual bunches at the cap, and flat at $s(\hat{m}, 1)$ when the cap is non-binding.

Appendix B.1 shows that the same logic extends to all concave multi-tier formulas, requiring only one additional observable (saving under an unlimited match) for each additional tier.

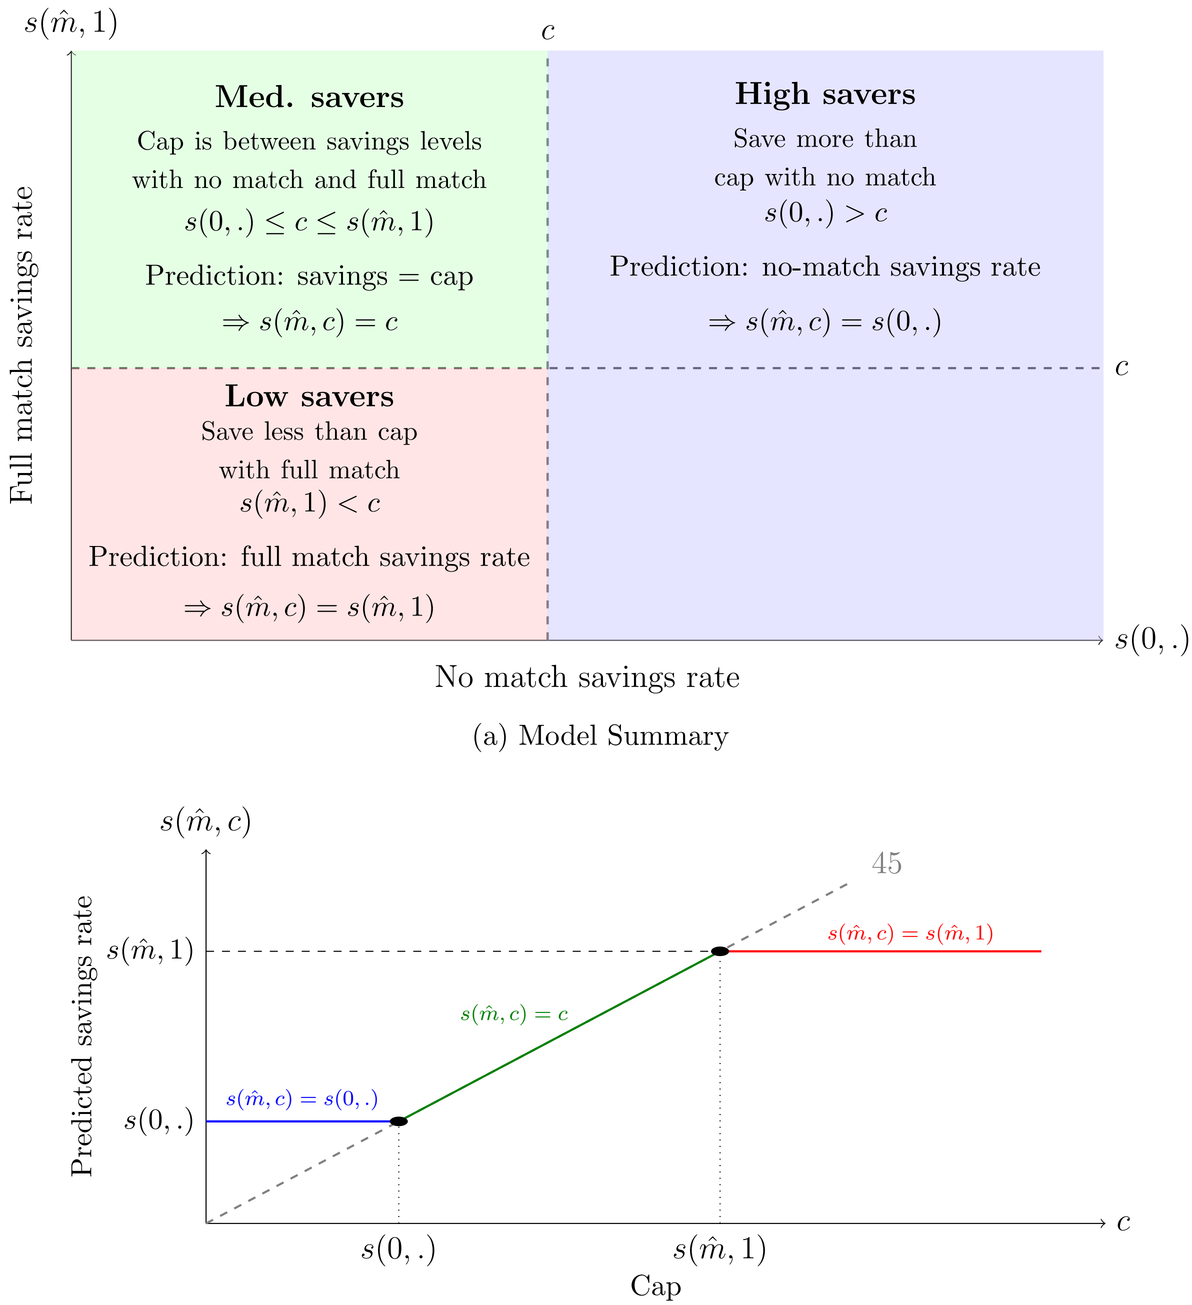

(a) Model Summary

(b) From data to prediction

Figure 4: Model illustration

Notes: This figure illustrates the prediction rule of saving under a match up to a cap. Panel (a) classifies individuals into three types based on the relationship between the cap and two observable quantities: saving under no match $s(0,.)$, and saving with an unlimited match rate of $\hat{m}$ denoted $s(\hat{m},1)$. High savers already contribute above the cap without a match, so the match does not affect their behavior. Medium savers would exceed the cap if the match applied to all contributions, so they bunch at the cap. Low savers contribute below the cap even with the match, so the cap does not bind. Panel (b) plots predicted saving $s(\hat{m},c)$ as a function of the cap c for a given individual. It shows three regimes: a flat region at $s(0,.)$ when the cap is below the individual's no-match saving, a 45-degree segment where the individual bunches at the cap, and a flat region at $s(\hat{m},1)$ when the cap is non-binding.

5 Prediction and Validation

The prediction rule from the previous section requires, for each respondent, saving with no match and saving with an unlimited match at the rates of interest. Our survey provides both. This section applies the rule, validates the predictions against administrative records of actual contributions, and assesses the robustness in an alternative sample.

5.1 Accounting for non-participants

A limitation of using hypothetical choices is that the survey is likely to overrepresent active decision-makers. Some eligible workers are passive: they do not make an active contribution choice, may remain at zero in opt-in plans, or may not pay attention to matching incentives. These passive workers are also plausibly less likely to respond to a retirement-plan survey. Consistent with this concern, participation is much higher in our survey than in the administrative data. In the survey, 90% of respondents report that they would contribute in the absence of a match, and over 97% report that they would contribute when a match is introduced. The implied participation response to matching is broadly consistent with estimates in the literature on the effect of matching on participation (Choi, 2015). The level of participation, however, is very different from that in the administrative data: participation is 64% among eligible workers in opt-in plans and 95% in plans with automatic enrollment.

We address this discrepancy by augmenting the survey with synthetic passive non-participants. These synthetic observations are intended to represent eligible workers who are underrepresented in the survey because of low engagement with the retirement plan. We assign them zero employee contributions under all match formulas, allow them to receive non-elective employer contributions when offered, and assume that they do not respond to matching incentives.

Let W denote the total weight of observed survey respondents and let $\hat{\pi}_0^{\text{survey}}$ denote the survey non-participation rate when a match is offered. For a target non-participation rate π_0 from the administrative data, we assign synthetic non-participants total (non-negative) weight

$$W_{\text{NP}}^{\text{syn}} = \frac{\pi_0 - \hat{\pi}_0^{\text{survey}}}{1 - \pi_0} W.$$

This choice ensures that non-participants account for a share π_0 of the augmented sample. We set $\pi_0 = 0.36$ for opt-in plans and $\pi_0 = 0.05$ for automatic-enrollment plans, corresponding to participation rates of 64% and 95%, respectively. We draw the age and income distribution of synthetic non-participants from that of non-participants in the administrative data.

5.2 Validation

We compare the model's predictions to administrative records of actual contribution rates. The analysis focuses on the more than 80% of the sample whose real-world formulas have one or two tiers with match rates included in our hypothetical scenarios (25%, 50%, or 100%), and can therefore be predicted using our data and approach.

Figure 5 compares the distribution of model predictions against the observed administrative contribution rates. The horizontal axis groups respondents into deciles of administrative contributions. The bars compare the administrative rate with the median predicted rate, with error bars showing the 25th-75th percentile range of predictions. Median predictions track the administrative rates closely through most of the distribution. The model overpredicts somewhat in the lowest deciles and underpredicts in the top decile, but the 25th and 75th prediction percentiles rise systematically with observed contribution rates throughout.

Figure 5: Predicted vs. Administrative Contribution Rates

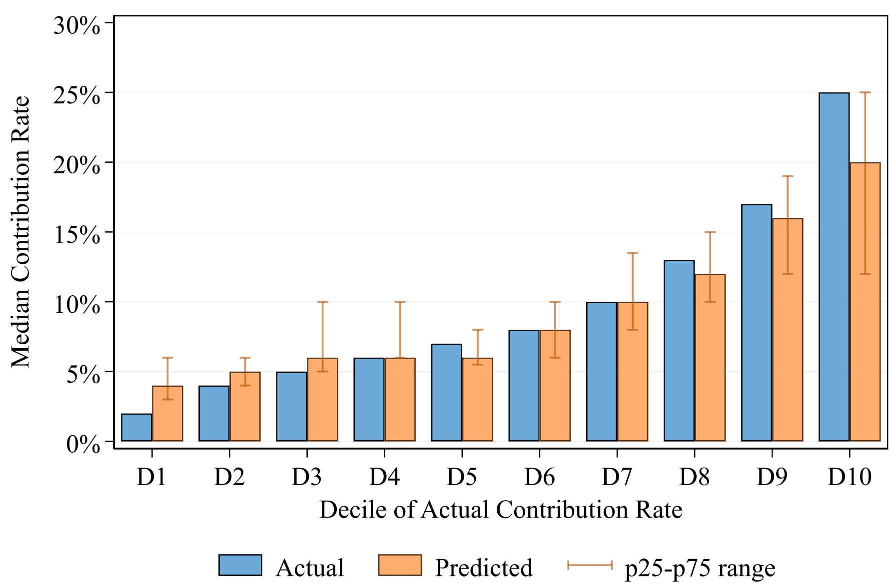

Notes: This figure compares respondents' saving as measured in the administrative data with the model-predicted contribution under their observed plan. The horizontal axis groups respondents into deciles of administrative contributions. Blue bars show the median administrative rate within each decile; orange bars show the median prediction, with error bars indicating the 25th-75th percentile range of predictions. Figure D.10(a) shows the robustness of this figure to our other chosen sample.

As a more formal test, Table 1 reports regressions of administrative contribution rates on our predicted rates, with alternative sets of controls. Column (1) includes only a constant and the predicted rate; column (2) adds two survey measures used in the prediction. Column (3) includes only demographics; column (4) includes the predicted rate, survey measures, and demographics. The results show that administrative rates rise nearly one-for-one with predicted rates, and R^2 values are between 0.41 and 0.44. Most of the explained variance comes from the prediction algorithm itself: demographics alone explain only 8% of the variance (column 3). The coefficient on the predicted rate remains large and statistically significant even after adding the survey inputs and demographic controls. We conclude that the model's predictions are strongly correlated with observed contribution rates, and that the predictive power is not simply due to the survey inputs or demographics alone, but reflects additional structure captured by the model.

Table 1: Model Predictive Power

	Dependent variable: Employee Contribution Rate			
	Model only (1)	Model + vignettes (2)	Demog. only (3)	All (4)
Predicted Rate	0.877^{***}	1.076^{***}		1.034^{***}
	(0.054)	(0.164)		(0.160)
Full Match		0.174^{***}		0.171^{***}
		(0.034)		(0.033)
No Match		-0.321^{**}		-0.315^{**}
		(0.135)		(0.130)
Demographic controls			✓	✓
Observations	1,326	1,326	1,326	1,326
R^2	0.405	0.429	0.075	0.443

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

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 a constant and the predicted rate. Column (2) includes the answers to the questions that are used to predict saving: saving under no match, and saving under an uncapped match (in the case of two-tier formulas we include the full uncapped match in the first tier of saving). Column (3) contains only demographics (age and income) and is shown to allow comparison of R^2 between specifications. Column (4) contains the model prediction, the survey questions, and includes demographics. The regression sample contains the participants with one and two-tier match formulas at rates 25%, 50% or 100%, for which we are able to predict contributions (approximately 85% of the survey sample). Standard errors are heteroskedasticity-robust.

Finally, we assess the extent to which our predictive algorithm used on our sample of plan participants is able to replicate plan-level distributions of contributions. Our ability to match real-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. Specifically, we divide both the survey sample and the administrative population into mutually exclusive age-income strata, and weight each survey observation by the ratio of the administrative share to the sample share in its stratum.

Figure 6: Actual vs Predicted Distributions of Employee Contribution Rates

(a) Aggregate Across 25 Most Common Formulas

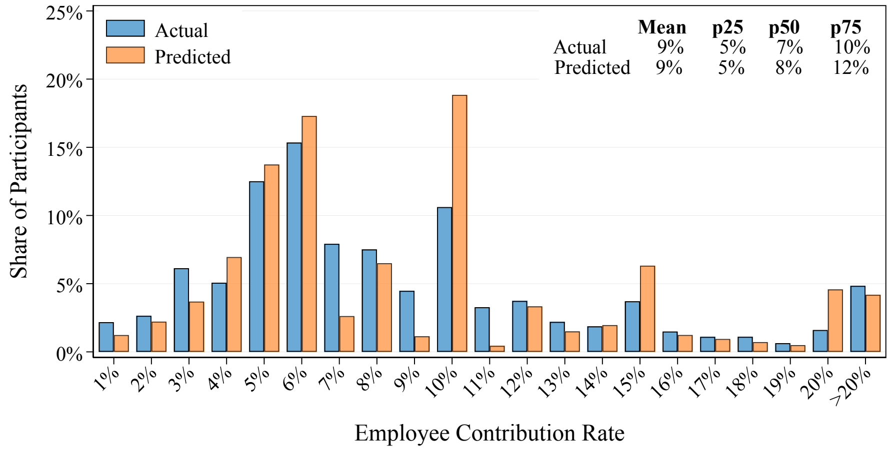

(b) Predictions Under Two Common Match Formulas

50% up to 6% Plan

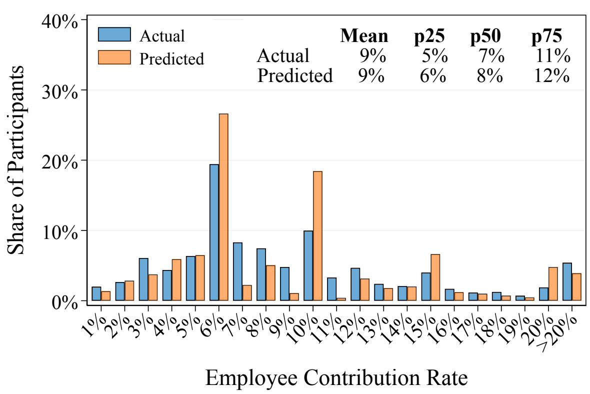

100% up to 6% Plan

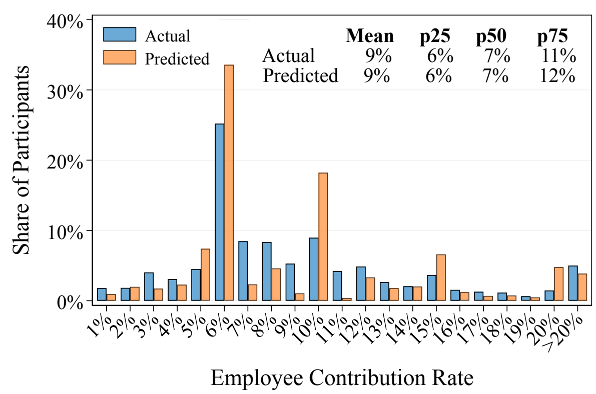

Notes: Panel (a) compares the distribution of saving in the administrative data with model-predicted distributions for the 25 most common formulas in our sample. Survey responses are reweighted to match administrative demographics of the plan with each formula using participant counts. Figure D.10(b) shows robustness to our other chosen sample. Panel (b) compares, for two common matching formulas, the distribution of employee contribution rates in the administrative data with model-predicted distributions; survey responses are reweighted to match the demographics of participants in plans using the corresponding formula.

Figure 6(a) shows the administrative distribution of contributions alongside our reweighted distribution of predicted rates. The means and medians match closely and the overall distributions exhibit the same broad shape, though the predicted distribution exhibits somewhat higher levels of bunching at round numbers (e.g., especially at 10%). This is a known feature of survey data (Manski and Molinari (2010); Gideon et al. (2017)), and we note that a large share of the excess mass at 10% appears to draw from a deficit in the mass at 9% and 11%.

Figure 6(b) shows the same comparison separately for respondents facing a 50% match up to 6% and a 100% match up to 6%, with similarly close alignment. We conclude that our prediction algorithm, combined with appropriate reweighting, can reproduce the distribution of real-world contributions under alternative matching formulas. Appendix Figure C.8 illustrates that our algorithm closely reproduces differences in average contributions by age and income in our sample.

5.3 Robustness

Our baseline analysis uses the full sample of 1,634 survey respondents. As a robustness check, we also repeat the analysis on a restricted sample that excludes respondents who fail a test–retest exercise. In the survey, we ask respondents the same hypothetical no-match contribution question twice: once near the beginning of the survey and once near the end. We exclude respondents whose stated contribution rate differs by more than one percentage point across the two responses. This restriction yields a sample of 1,259 respondents and focuses on individuals whose answers appear more stable, and therefore less likely to reflect measurement error. Our main findings are robust to this sample restriction (see Appendix D).

6 From Predictions to Policy Design: The Equity-Savings Frontier

Having validated our predictive model, we turn to evaluating alternative formula designs. Any such evaluation must first take a stand on what the policy is meant to achieve. For reasons we discuss in the next subsection, we choose to study the effect of formulas on average employee saving and on the inequality of employer contribution rates across workers. More broadly, our approach can be extended to any policy criterion that can be expressed as a function of workers' observable characteristics and their predicted behavioral responses.

6.1 Policy objectives

What should matching formulas achieve? We argue that two natural approaches to this question point to the same two objects: the effect of a formula on average employee saving (measured as the change in employee contribution rates) and the inequality in employer contributions across workers (measured as the coefficient of variation in employer contribution rates).

The first approach infers objectives from the revealed (and stated) preferences of U.S. policymakers. Since the Revenue Act of 1942, U.S. public policy has aimed both to expand retirement saving and to prevent tax-preferred retirement benefits from accruing disproportionately to highly compensated workers. This dual objective is reflected in the long history of non-discrimination rules. Analysis of legislative intent across successive reforms confirms this dual aim (Israel Preninger and O'Dea (2026)). Taking these aims as given, a natural policy evaluation asks whether alternative formulas can raise retirement saving while limiting inequality in the distribution of employer contributions.

The second approach derives the relevant outcomes from a welfare framework. A long tradition in public economics, going back to Diamond (1977) and Feldstein (1985), identifies the same two objectives. The standard welfare rationale for incentivizing retirement saving is that some households save too little privately, either because of savings distortions created by means-tested social insurance programs or because of behavioral frictions such as present bias or inattention. Under either rationale, the welfare value of a reform depends partly on its average effect on saving and partly on whether these increases are targeted at those who are most at risk of undersaving. Choukhmane and Palmer (2025) formalize this argument in a welfare framework with corrective retirement saving interventions.

This normative approach also calls for evaluating formulas along two dimensions. The first dimension, the average treatment effect on saving, comes directly from our predictive model. The second dimension, the quality of targeting, is harder to assess because it depends on covariances that are difficult to observe directly: how the treatment effect on total contributions covaries with worker-level fiscal externalities and with worker-level behavioral biases.[8] We use three observable proxies: worker income, which proxies for exposure to the savings disincentives created by means-tested programs in retirement[9]; and the worker's hypothetical contribution under no match and under an unlimited 100% match, which together proxy for behavioral biases such as present bias or inattention to financial incentives.[10] As we detail in Appendix E, the covariance between the formula's treatment effect and each of these three proxies is strongly correlated with the dispersion of employer contribution rates across formulas (with correlations of 0.94 to 0.98 across the three proxies). Specifically, low-dispersion formulas direct their saving gains more toward lower-income workers and toward

[8]Looking beyond average effects to consider targeting, specifically the covariance between treatment effects and the degree of behavioral bias, is a central theme in the theoretical behavioral public finance literature (Farhi and Gabaix (2020); Allcott et al. (2025)).

[9]Supplemental Security Income and Medicaid eligibility, for instance, require assets below a few thousand dollars, depressing private saving for those most likely to rely on these programs in retirement (Hubbard et al. (1995); Sleet and Yeltekin (2006)).

[10]Choukhmane and Palmer (2025) show that, in a model with heterogeneous present bias, the take-up of matching incentives is monotonically decreasing in the level of present bias.

workers who would contribute little under either no match or an unlimited 100% match. To the extent that the latter pattern reflects behavioral biases rather than genuinely low saving preferences, low-dispersion formulas better target the most behaviorally biased workers. These arguments hold absent any redistributive motive; adding one only strengthens the case.

Both the revealed-preferences of U.S. policymakers and the welfare-theoretic approaches thus point to similar objectives: raising average saving and shaping the distribution of these gains. We therefore evaluate formulas in a two-dimensional space. The vertical axis measures the increase in average employee contributions relative to a no-match benchmark. The horizontal axis measures inequality in employer contributions, using the coefficient of variation of employer contribution rates.[11] This approach does not require us to specify a social welfare function or to fix a particular weight on saving versus targeting. Instead, we characterize the frontier of feasible formulas and identify dominated formulas: those for which cost-equivalent alternatives generate more employee saving while producing a more equal distribution of employer contributions.

6.2 Characterizing the Frontier

Figure 7 shows the equity and savings outcomes for the set of formulas that meet three criteria: they cost the employer the same as a 50% match up to 6% of salary, include either one or two tiers, and have a match rate that does not increase with employee contributions.

Each point in the figure represents a feasible plan, and the upper-left envelope traces out the equity-savings frontier: the set of formulas that maximize employee saving for a given level of inequality, or equivalently, minimize inequality for a given level of saving. Two key patterns emerge. First, there is a clear equity-savings trade-off: along the frontier, holding employer cost constant, increasing employee saving requires accepting greater inequality in employer contributions. Second, many existing formulas are strictly dominated — alternative designs exist that cost the same, generate higher employee saving, and yield more equitable distributions of employer contributions.

Figure 7 shows that many formulas lie inside the frontier. Figure 8 turns to characterizing which plan features are associated with being on the frontier, and which are associated with being inside it. Each panel plots an equity-savings space, but now colors each region by a plan characteristic, averaged over all formulas in that region.[12]

[11] We additionally show results with three more direct proxies for fiscal externalities and behavioral biases.

[12] We show this for the set of plans that cost 3 pp of payroll, the results are very similar for other employer costs.

Figure 7: Equity-Savings Frontier

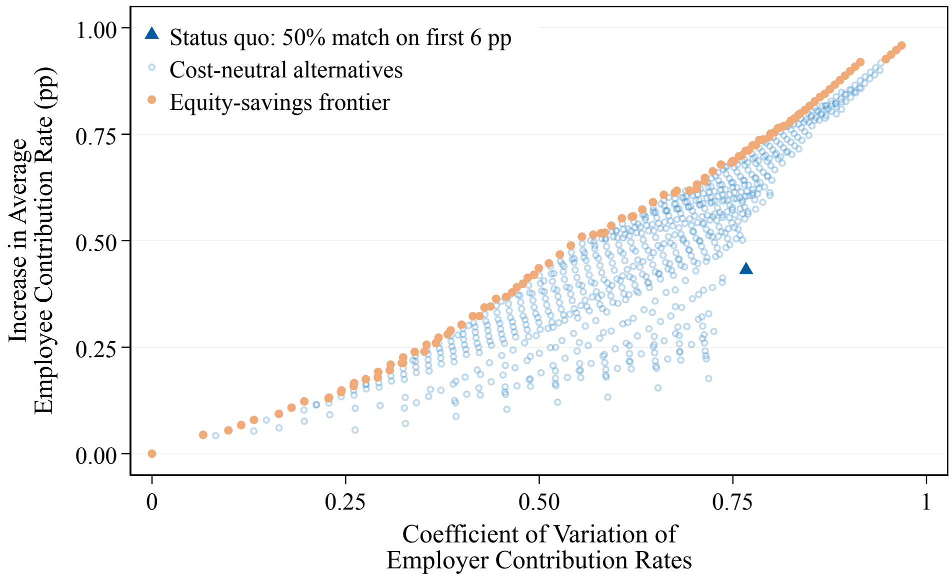

Notes: This figure plots the savings and equity properties of alternative employer matching formulas with the same expected cost as the status-quo formula (a 50% match on the first 6 pp of pay, without automatic enrollment). Each point is one candidate formula. The vertical axis is the change in the average employee contribution rate relative to their saving under no match; the horizontal axis is the cross-employee coefficient of variation of employer contributions under that formula. Orange points trace the equity–savings frontier: formulas for which no alternative achieves higher average contributions at equal or lower inequality. Candidate formulas are restricted to single- or two-tier schedules with a weakly decreasing match rate in employee contributions and without automatic enrollment.

Moving to the frontier. Panel (a) shows the average achievable match rate. Formulas on or near the frontier tend to have lower match rates. This reflects the fact that high match rates concentrate benefits toward saving that would have been done anyway. Panel (b) shows the match cap, and panel (c) shows the non-elective contribution. Together, these two panels reveal the key trade-off *along* the frontier: holding the budget constant, moving toward higher employee saving is associated with a higher cap and a lower non-elective contribution, while moving toward greater equity is associated with a lower cap and a higher non-elective contribution. This is intuitive: a non-elective contribution is distributed to all employees regardless of their saving decisions, which compresses the distribution of employer contributions, while a higher cap rewards those who save more, increasing both total saving and inequality. The frontier thus traces out how the planner allocates a fixed budget between an unconditional transfer (the non-elective contribution) that promotes equity and a conditional incentive (the match with a higher cap) that promotes saving.

Choosing a point on the frontier. The choice of a formula on the frontier ultimately depends on the relative weight placed on savings and equity. A policymaker who prioritizes

Figure 8: Characterizing the Equity-Savings Frontier

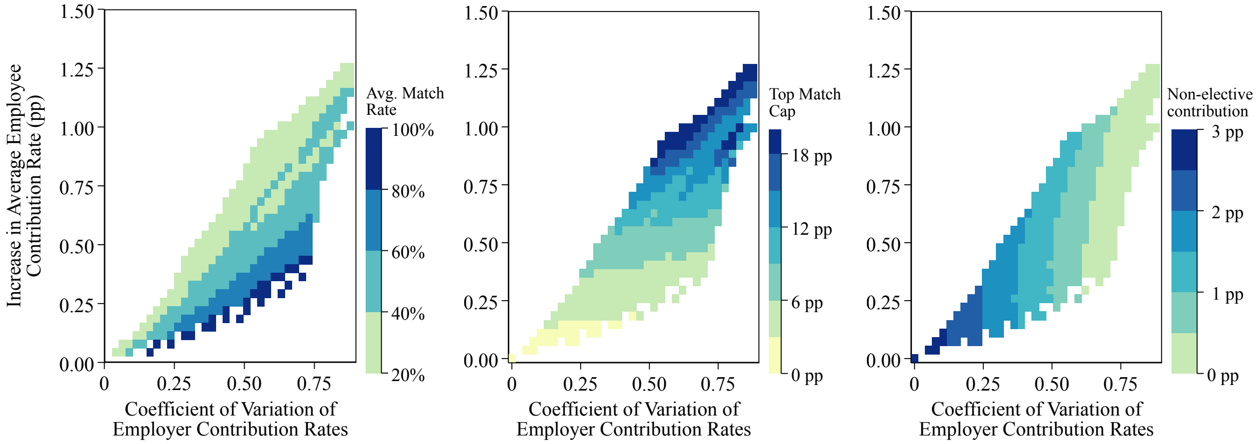

(a) Average Match Rate (b) Top Match Cap (c) Non-elective contribution

Notes: This figure characterizes how plan features map to positions in the equity–savings space. The set of formulas are those with a payroll cost of 3 pp of pay. The horizontal axis is the cross-employee coefficient of variation of employer contribution rates, and the vertical axis is the increase in the average employee contribution rate relative to a no-match baseline. Cells are colored by the average value of a given plan characteristic among formulas falling in that region of the space. Panel (a) shows the average match rate earned by an employee who contributes up to the highest match cap; formulas on or near the frontier tend to have lower average match rates. Panel (b) shows the match cap, or in the case of a two-tier formula, the highest cap. Moving rightward along the frontier (toward greater average contribution rates) is associated with a higher cap. Panel (c) shows the non-elective contribution rate. Moving leftward along the frontier (toward greater equity in benefits) is associated with a higher non-elective contribution.

aggregate saving will prefer formulas with higher caps and lower non-elective contributions, while one who places greater weight on equalizing employer contributions will prefer formulas with lower caps and higher non-elective contributions. Since our dominance results hold for objective functions that value both dimensions, we remain agnostic about where on the frontier a given plan should sit. Beyond these normative considerations, firms may also weigh practical factors when selecting among formulas on the frontier. One such consideration is simplicity: a single-tier match formula on the frontier may be easier to communicate to employees and administer than a two-tier formula, even if the two have broadly similar savings and equity properties.

The role of automatic enrollment. The analysis so far has focused on the equity-savings frontier among plans without automatic enrollment. Automatic enrollment is, however, an important additional lever for reducing inequality in employer contributions at a given matching cost. Participation in plans with automatic enrollment is around 95%, compared to roughly 65% in plans without. As a result, in plans without automatic enrollment, the 36% of non-participants can be brought into the plan through a non-elective contribu-

Figure 9: The Role of Automatic Enrollment

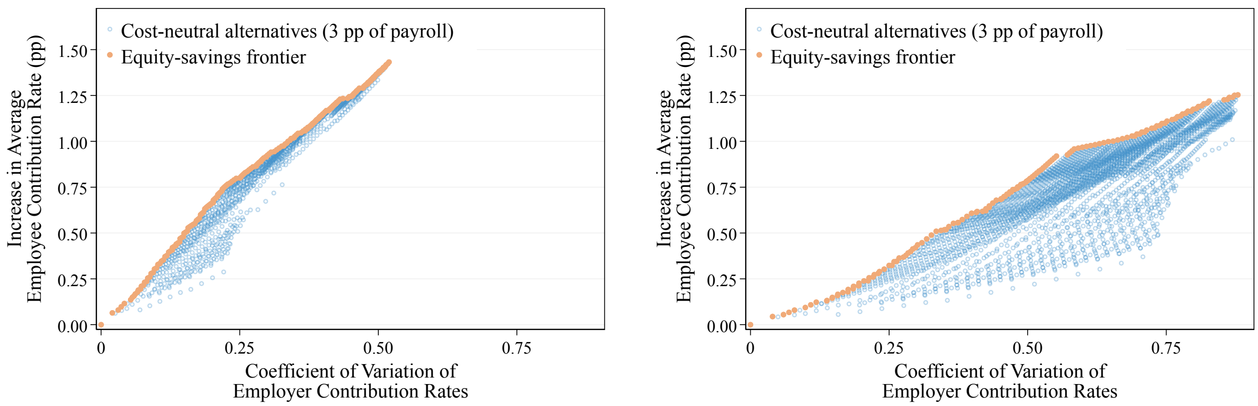

(a) Plans with auto-enrollment (b) Plans without auto-enrollment

Notes: These figures illustrate the equity and savings characteristics of the set of formulas which cost 3 pp of payroll, are either single tier or two tier, and have a match rate that is non-increasing in employee contribution. Panel (a) shows plans with auto-enrollment and panel (b) shows plans without auto-enrollment.

tion, generating substantial scope to compress the dispersion in employer contribution rates through formula design alone. In plans with automatic enrollment, by contrast, participation is already near-universal, so even absent a non-elective contribution few workers receive zero employer contribution and the dispersion in employer contribution rates is compressed.

Figure 9 illustrates this point for formulas with an average employer contribution cost of 3 pp of total payroll: the dispersion in our inequality measure across cost-neutral alternatives is substantially smaller under automatic enrollment. This is consistent with Choukhmane (2025), who shows that automatic enrollment has modest medium-run effects on average contributions but substantially reduces inequality in saving. The contrast across panels illustrates where the remaining design margins lie. In plans with automatic enrollment (panel a), cost-neutral formulas cluster tightly around the frontier: there is meaningful scope to raise saving incentives by moving vertically toward the frontier, but more limited additional scope to reduce inequality through formula redesign. In plans without automatic enrollment (panel b), the cloud of cost-neutral alternatives fans out widely, leaving substantial room to improve along both dimensions.

6.3 Robustness and Extensions

Our central result is that many existing matching formulas are dominated. This subsection addresses three concerns about this result. First, we show that formulas on the equity-savings frontier often also dominate under alternative measures of targeting quality. Second, we consider equilibrium wage adjustment, the concern that employers might offset the effects of any reform by adjusting other components of compensation. Finally, we ask how the predicted gains in retirement accounts interact with other margins: non-retirement saving, the tax advantage of retirement saving, and the rate of return on retirement assets relative to the alternative.

6.3.1 Is the Frontier Robust to Alternative Targeting Measures?

Our baseline analysis uses the coefficient of variation of employer contribution rates as a summary measure of targeting quality, on the grounds that it correlates strongly with three more direct proxies of targeting introduced in Section 6.1 and discussed further in Appendix E. A natural question is whether the equity-savings frontier we identify with this baseline measure would change materially if we used the underlying proxies instead. We show that it would not.

Figure 10 plots, in each of the three panels, the set of cost-neutral matching formulas (those with the same total expected employer cost of 3 pp of payroll) against a different targeting proxy. These are: the covariance between the formula's treatment effect on total contribution rates and worker income (panel a), the worker's hypothetical employee contribution rate under no match (panel b), and the worker's hypothetical employee contribution rate under an unlimited 100% match (panel c). The orange points mark the baseline equity-savings frontier from Figure 7, which uses the coefficient of variation of employer contributions as the equity measure. Across all three panels, formulas on the baseline equity-savings frontier lie on or near the frontier defined using each of the three covariance-based targeting measures. Thus, the same set of formulas is approximately dominating whether the targeting measure is the dispersion of employer contributions or any of the three direct proxies.

6.3.2 Are the Gains Robust to Compensation Adjustment?

A natural concern with reforms to employer matching formulas is that other components of compensation will adjust in equilibrium. Match formulas are part of a firm's compensation policy, and changing them may induce offsetting adjustments to wages or other benefits. The force of this concern depends on the rationale for reform. The corrective rationale, which motivates the targeting objective in our analysis, is largely robust to wage adjustment. The redistributive rationale, which a planner might hold alongside the corrective one, could be more sensitive.

Begin with the corrective case. Suppose, in the extreme, that wages adjust fully to any reform: every additional dollar of employer contribution is offset by a dollar of lower wage, leaving total individual compensation unchanged. Compensation still shifts from the present

Figure 10: Robustness of the Equity-Savings Frontier to Alternative Targeting Measures

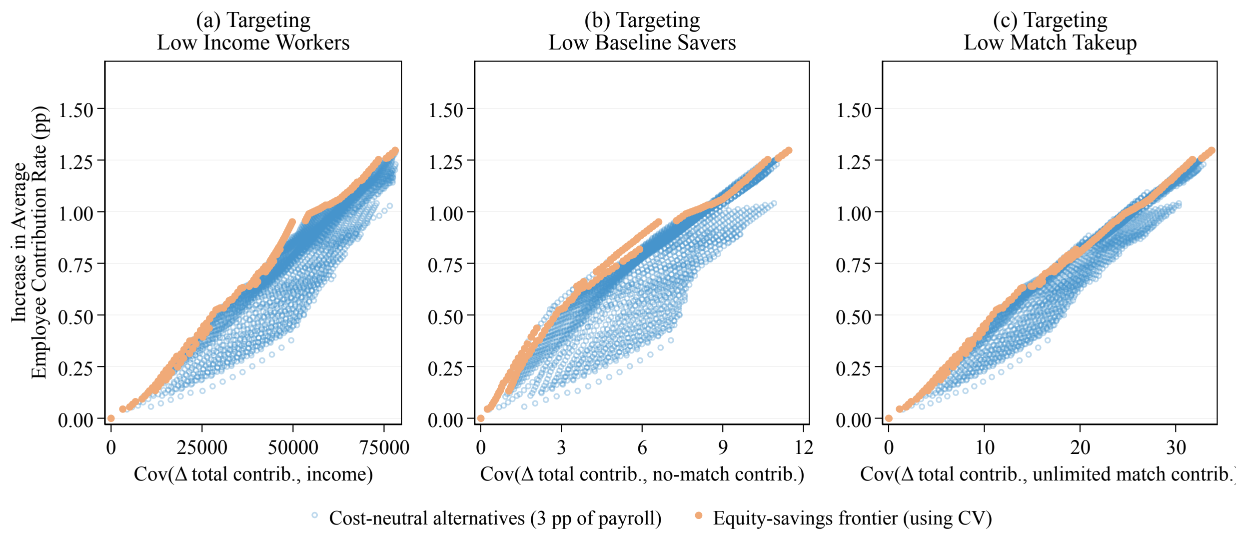

Notes: The figure plots, across cost-neutral matching formulas (total expected cost equal to 3 pp of payroll), the predicted increase in average employee contribution rate on the y-axis against three targeting measures on the x-axes. Each x-axis is the covariance between the formula's treatment effect on total contribution rates (defined as the predicted increase in combined employee and employer contribution rates relative to a no-match counterfactual) and, respectively: worker income in panel (a); the worker's hypothetical employee contribution rate under no match in panel (b); and the worker's hypothetical employee contribution rate under an unlimited 100% match in panel (c). Lower values of the covariance indicate better targeting of the relevant group, as discussed in Appendix E. Orange points mark formulas on the equity-savings frontier defined using the coefficient of variation of employer contribution rates as the targeting measure; blue points are the remaining cost-neutral alternatives. The frontier formulas identified on the coefficient of variation lie near the frontier in all three covariance spaces, indicating that formulas on our baseline equity-savings frontier are also approximately dominating when targeting is measured via direct covariances with income, baseline saving, and match responsiveness.

(wages) to the future (retirement resources), and welfare improves to the extent that this shift is better targeted toward workers with larger fiscal externalities or behavioral biases.

A redistributive motive could be more affected by equilibrium wage adjustments, since full adjustment at the individual level would leave the distribution of total compensation unchanged. Two features of our setting could attenuate this concern. First, much of the inequality reduction occurs within income groups, between high and low savers who presumably hold similar occupations and face similar labor markets. Second, these concerns are further attenuated when changing the safe harbor formulas. Since these formulas currently cover around half of employers and over half of employees (shares which are growing), reforming the safe harbors would shift the retirement benefit at both workers' current jobs and their likely outside options, limiting the scope for compensating differentials across employers.

6.3.3 Looking Beyond Retirement Accounts

Our analysis predicts how changes in matching formulas affect retirement contributions. A natural question is how these changes pass through to other aspects of household balance sheets and to government tax revenues. We discuss three channels that affect this pass-through.

Crowd-out of non-retirement saving. Increases in retirement contributions may partly come from reductions in saving outside the retirement account. The literature suggests that this crowd-out is incomplete: each additional dollar of retirement saving raises total saving by less than a dollar but more than zero. The contribution gains we report can therefore be read as upper bounds on the increase in total saving, and a reader who wants to incorporate a specific average crowd-out estimate can scale down our reported saving effects proportionally. Two findings in this literature are particularly relevant for our counterfactuals. First, Chetty et al. (2014) find that non-elective employer contributions generate very limited crowd-out of non-retirement savings, implying that the contribution gains from formulas with a larger NEC component are unlikely to be undone in total saving. Second, Choukhmane and Palmer (2025) document smaller crowd-out among financially constrained workers and larger crowd-out among those with ample liquid savings. Because lower inequality formulas raise contributions more for lower-income and lower-saving workers, the contribution gains from these formulas are likely concentrated in groups with lower crowd-out.

Tax advantage. At the individual level, retirement contributions enjoy tax-preferred treatment that magnifies their effect on terminal wealth relative to saving the same amount in a taxable account. Choukhmane et al. (2024) estimate that this tax advantage adds between a quarter (at the bottom of the lifetime earnings distribution) and a half (at the top) to the value of each dollar of employee contributions. The aggregate welfare implication of this tax subsidy, however, depends on how the foregone tax revenue is financed.

Returns on retirement assets. Retirement plans are subject to fiduciary regulation, and plan defaults are typically low-fee, well-diversified target date funds. To the extent that workers would otherwise hold non-retirement saving in cash or in higher-fee, less diversified investments, shifting contributions into the retirement account raises the rate of return on that saving and further amplifies the wealth gains from a given increase in retirement contributions.

A full welfare analysis would weight these three channels against each other and against the tax-financing assumption. Such an exercise is beyond the scope of our paper, but the qualitative point is straightforward: the first channel attenuates the gains we report, while the second and third amplify them.

6.4 Moving common and safe harbor formulas to the frontier

In this section, we use our predictive model to evaluate the most common and the safe harbor formulas against alternatives on the frontier. In our sample of Vanguard plans, we find that all 10 of the most common formulas, and 22 of the top 25, are dominated, in the sense that cost-neutral reforms could generate higher saving without increasing inequality in employer contributions. Extrapolating from this evidence to a broader set of plans using Form 5500 filings, we estimate that moving the most popular safe harbor formulas to their saving-maximizing alternatives on the frontier can raise employee retirement contributions by at least $6.7bn.

6.4.1 Saving-maximizing and inequality-minimizing alternatives

For each formula, we highlight two alternatives on the frontier: one that maximizes savings and one that minimizes inequality. We illustrate this approach using the second most common matching formula in our data: a 50% match on contributions up to 6% of pay. Using our model, we identify all single- and two-tier formulas that are equally costly to the firm (accounting for changes in employee contributions in response to a different formula), and compute, for each, the resulting increase in saving (relative to a no match benchmark) and the coefficient of variation in employer contribution rates. Figure 11 plots these outcomes. The blue triangle represents the 50% match up to 6% plan, which lies well inside the frontier.

Figure 11: Equity-Savings Frontier for Plans with Costs Similar to a 50% Match up to 6 pp

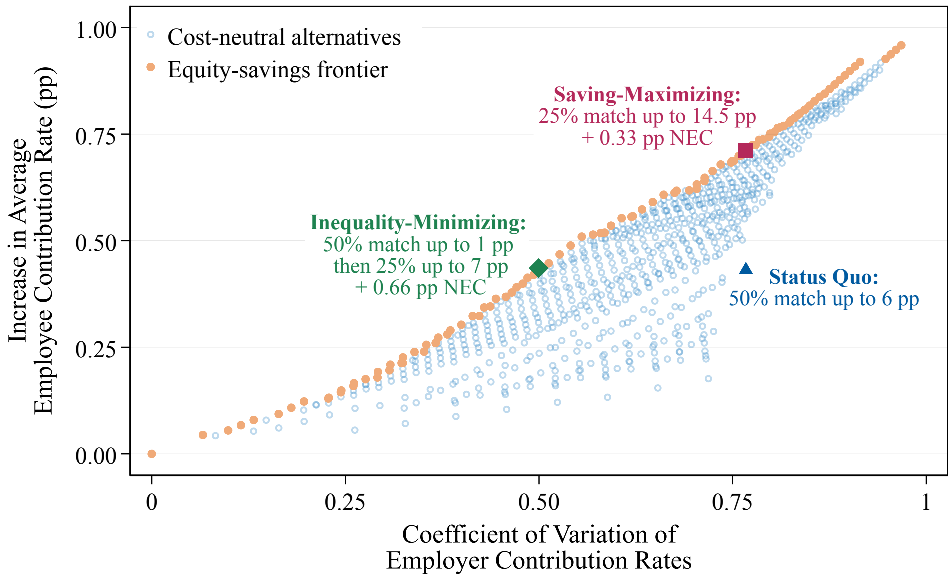

Notes: This figure plots the equity and savings characteristics of employer match formulas with the same expected cost as the status quo formula (50% match up to 6% of pay), shown as the blue triangle. Light blue squares denote the full set of cost-neutral alternatives, and orange circles trace the equity-savings frontier. We highlight two alternatives on the frontier. The red square (*Saving-Maximizing*) is the formula that maximizes the average employee contribution rate subject to no increase in the coefficient of variation of employer contribution rates relative to the status quo. The green diamond (*Inequality-Minimizing*) is the formula that minimizes the coefficient of variation of employer contribution rates subject to no decrease in the average employee contribution rate. NEC denotes a non-elective employer contribution. Absolute changes in average savings rates and employer-contribution inequality for these two formulas, relative to the status quo, are reported in Table 2.

We highlight two alternatives. The green diamond shows the formula that minimizes inequality, subject to achieving at least as much saving as the current formula. The red square shows the formula that maximizes saving, subject to no increase in the coefficient of variation in employer contributions. The frontier formulas have lower match rates and higher caps, and they balance the budget with non-elective contributions. The saving-maximizing formula increases employee contributions by almost 0.3 percentage points, an approximately 65% increase in the treatment effect (relative to no match) induced by the formula. The inequality-minimizing formula reduces the coefficient of variation in employer contributions from 0.77 to 0.50.

6.4.2 Moving the most common formulas in the Vanguard sample to the frontier

We can extend this analysis beyond this particular formula. All 10 of the most widely used matching formulas, and 22 of the top 25, lie well inside the equity-savings frontier. Table 2 reports the gains from moving the five most common formulas 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 Tables C.3 and C.4 present the corresponding results for the 25 most common formulas assuming no auto-enrollment, and Tables C.5 and C.6 show similar patterns for the case with auto-enrollment.

Table 2: The effect of moving to the frontier: five most common plans

Panel A: No Auto-Enrollment

	Minimize Inequality				Maximize Savings			
Current Plan	Plan	NEC	Baseline CV	Change CV	Plan	NEC	Baseline Savings	Change Savings
100-3-50-5	50-1-25-6.5	1.45	0.73	−0.40	50-2.25-25-20	0.54	6.16	0.59
50-6-0-0	50-1-25-7	0.66	0.77	−0.27	25-14.5-0-0	0.33	6.20	0.28
100-1-50-6	50-1-25-7.75	0.92	0.75	−0.30	25-19-0-0	0.55	6.24	0.40
100-5-0-0	50-1-25-7.25	2.02	0.74	−0.45	50-7-25-19.75	0.48	6.20	0.71
50-4-0-0	50-1-25-4.5	0.41	0.75	−0.24	50-1-25-7.25	0.09	6.01	0.20

Panel B: Auto-Enrollment

	Minimize Inequality				Maximize Savings			
Current Plan	Plan	NEC	Baseline CV	Change CV	Plan	NEC	Baseline Savings	Change Savings
100-3-50-5	50-1-25-6.5	2.10	0.24	−0.10	50-2.5-25-10	1.26	8.93	0.36
50-6-0-0	50-1-25-7	0.95	0.31	−0.09	50-2.5-25-9	0.32	8.99	0.22
100-1-50-6	50-1.75-25-7.5	1.19	0.28	−0.08	50-1.5-25-10	0.92	9.06	0.20
100-5-0-0	50-1.25-25-7	2.90	0.25	−0.13	50-6.5-25-10	1.32	9.00	0.50
50-4-0-0	50-1-25-4.5	0.60	0.27	−0.08	50-1.5-25-6	0.20	8.73	0.18

Notes: This table shows, for the five most common plans, the counterfactual plan which is equally costly and either minimizes inequality (subject to no fall in employee saving), or maximizes saving (subject to no increase in inequality). Panel A assumes no auto-enrollment; Panel B assumes auto-enrollment. See Appendix Tables C.3, C.4, C.5, and C.6 for further plans for the main sample. Table D.9 provides versions of this table for the sample that passes test-retest.

Average Effects. Averaging across all plans that use one of the 25 most common formulas, we estimate that replacing each formula with the dominating alternative we identify would increase average employee contributions by 0.4 pp of salary in plans without auto-enrollment and by 0.3 pp in plans with it. These increases occur without raising employer costs or increasing inequality. Alternatively, moving these plans to the frontier could reduce inequality in employer contributions, measured by the coefficient of variation, by 0.35 (or about 46%) in plans without auto-enrollment and by 0.10 (or about 37%) in plans with auto-enrollment, again without reducing voluntary saving or increasing costs. Both reforms achieve these gains in part by introducing a non-elective contribution, which mechanically raises the proportion of workers with positive retirement contributions by 5 pp in plans with auto-enrollment and 36 pp in plans without it.

Heterogeneous Effects of the Inequality-Minimizing Formulas. The reduction

Figure 12: Total contribution changes from moving to the saving-maximizing formula

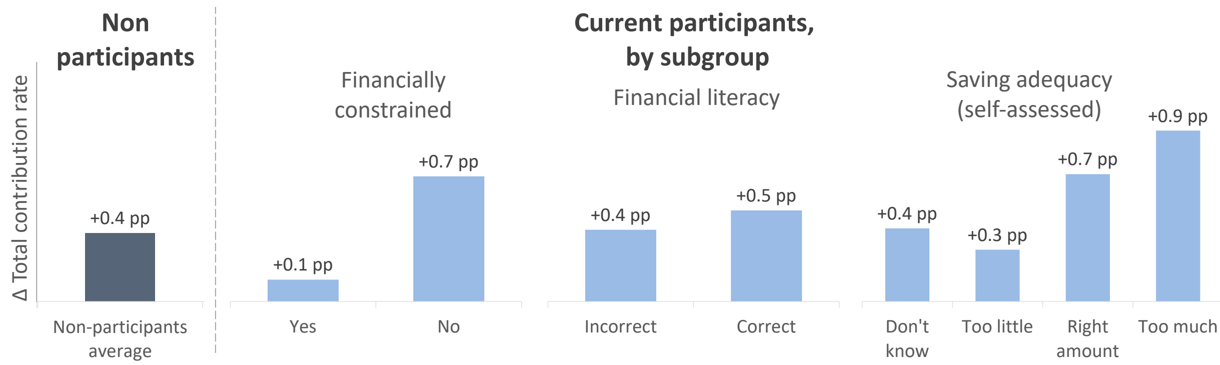

Notes: Each panel shows the predicted change in total contribution rates (in percentage points of salary) when each status-quo formula is replaced by its corresponding saving-maximizing formula. The left bar reports the average effect for current non-participants. The remaining bars report effects among current participants only, by self-assessed financial constraints, financial literacy, and self-assessed saving adequacy.

in inequality achieved by moving plans to the corresponding inequality-minimizing formulas operates primarily through a redistribution of employer contributions from participants to non-participating workers. Among the former, average employer contributions decline by 0.6 pp of salary; among the latter, they increase by 1.2 pp.

Heterogeneous Effects of the Saving-Maximizing Formulas. A natural concern is that the saving gains from moving all plans to the corresponding saving-maximizing formulas might be concentrated among a small number of very high savers, leaving most workers unaffected. In fact, we find the gains are broad-based (Figure 12). As expected, workers who report being financially unconstrained see some of the largest increases in total contributions, around 0.7 pp of salary, as they are most likely to have the capacity to save up to the higher cap. Yet even participants who report being financially constrained or saving too little are predicted to have total contributions increase by approximately 0.1 pp and 0.3 pp, respectively. Finally, non-participants, likely the most constrained group, see a substantial increase in total contributions of 0.4 pp, driven by the non-elective contribution component of most saving-maximizing formulas.

6.4.3 Moving safe harbor formulas in the Form 5500 sample to the frontier

The previous calculations are based on formulas commonly observed in the Vanguard sample. To assess the aggregate impact of reforming safe harbor regulations, we conduct a back-of-the-envelope exercise using a broader sample of plans with safe harbor matching formulas identified from 2023 Form 5500 filings by Choukhmane et al. (2026). Moving the most common safe harbor formulas in this sample to the saving-maximizing alternatives we identified in the Vanguard sample would increase employee contributions in those plans by 7.5%. This corresponds to an additional $6.7 billion in annual employee retirement contributions across the safe harbor plans in this sample. This estimate is a likely lower bound: the formulas in this sample cover around half of aggregate employee and employer contributions to DC plans in 2023, and the aggregate increase could be substantially larger (plausibly twice as large) if similar gains apply to safe harbor plans outside of this sample.

Appendix Table C.7 reports the calculation. We map each safe harbor formula in the Form 5500 sample to its corresponding saving-maximizing alternative from the Vanguard counterfactual analysis, separately for plans with and without automatic enrollment. We then apply the estimated percent increase in employee contributions from the Vanguard sample to aggregate employee contributions reported in the Form 5500 filings. This scaling exercise should be interpreted cautiously: it assumes that the predicted response in the Vanguard sample is informative about responses in the broader Form 5500 population, even though plans and participants may differ across the two samples. Subject to these caveats, the exercise suggests that safe harbor reform could generate billions of dollars in additional retirement contributions at no additional cost to employers and without increasing inequality in employer contributions.

7 Conclusion

This paper studies how 401(k) matching formulas should be designed to raise employee saving and reduce inequality in employer contributions. We use survey responses to hypothetical scenarios, together with a small number of testable assumptions on how people make their saving decisions, to predict saving behavior across a wide range of match schedules, including those not observed in reality. We trace out their implications for both saving levels and the distribution of employer saving subsidies.

Our findings highlight a central trade-off in plan design: raising saving often comes at the cost of greater inequality in employer contributions. Yet, many plans, including the most common ones and several safe harbor formulas, lie well inside the frontier. More saving and lower inequality in employer contribution rates are attainable within existing budgets.

These conclusions hold broadly for social planners who value both higher average saving and lower inequality: regardless of how the planner weights them, they should prefer the plans we identify as dominating. More broadly, our results point to the value of hypothetical surveys as a general tool for policy design. This approach is well suited to other settings, especially when the policy space is large relative to what is observed, behavioral responses are heterogeneous and hard to predict from observational data alone, and policy objectives can be expressed through a small number of measurable criteria.

This material is provided for informational purposes only and is not intended to be investment advice or a recommendation to take any particular investment action.

References

Allcott, H., D. Cohen, W. Morrison, and D. Taubinsky (2025). When do "nudges" increase welfare? *American Economic Review 115*(5), 1555–1596.

Allcott, H., J. Kim, D. Taubinsky, and J. Zinman (2022). Are high-interest loans predatory? theory and evidence from payday lending. *The Review of Economic Studies 89*(3), 1041–1084.

Ameriks, J., J. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2016). The long-term-care insurance puzzle: Modeling and measurement. Technical report, National Bureau of Economic Research.

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.

Bachmann, R., G. Topa, and W. van der Klaauw (2022). *Handbook of economic expectations.* Elsevier.

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

Bhattacharya, V. and G. Illanes (2025). The design of defined contribution plans. Technical report, National Bureau of Economic Research.

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

Card, D. and M. Ransom (2011). Pension plan characteristics and framing effects in employee savings behavior. *The Review of Economics and Statistics 93*(1), 228–243.

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. (2025). Default options and retirement saving dynamics. *American Economic Review 115*(11), 3749–3787.

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., J. Dedyo, C. O'Dea, and L. Schmidt (2026). Using large language models to measure U.S. retirement plan design at scale: Methods and evidence. Working paper.

Choukhmane, T. and C. Palmer (2025). The Effect of Increasing Retirement Saving on Consumption, Balance Sheets, and Welfare. Technical report.

Colarieti, R., P. Mei, and S. Stantcheva (2024). The how and why of household reactions to income shocks. Technical report, National Bureau of Economic Research.

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.

Farhi, E. and X. Gabaix (2020). Optimal taxation with behavioral agents. *American Economic Review 110*(1), 298–336.

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.

Gideon, M., B. Helppie-McFall, and J. W. Hsu (2017). Heaping at round numbers on financial questions: The role of satisficing. In *Survey research methods*, Volume 11, pp. 189.

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

Goodman, A. (2025). Impatience or illiquidity? the case of early 401(k) withdrawals. *Available at SSRN 5205483*.

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.

Investment Company Institute (2026, March). Release: Quarterly retirement market data, fourth quarter 2025. Technical report, Investment Company Institute, Washington, DC. Released March 26, 2026.

Israel Preninger, D. and C. O'Dea (2026). History and goals of nondiscrimination requirements and safe harbor 401(k) plans.

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.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Lu, T., O. Mitchell, S. Utkus, and J. Young (2014). Borrowing from the future. 401 (k) loans and loan defaults. Technical report, Technical report.

Manski, C. F. (2011). Policy analysis with incredible certitude. *The Economic Journal 121*(554), F261–F289.

Manski, C. F. (2013). *Public Policy in an Uncertain World: Analysis and Decisions*. Harvard University Press.

Manski, C. F. and F. Molinari (2010). Rounding probabilistic expectations in surveys. *Journal of Business & Economic Statistics 28*(2), 219–231.

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).

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.

Sleet, C. and Ş. Yeltekin (2006). Credibility and endogenous societal discounting. *Review of Economic Dynamics 9*(3), 410–437.

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.

Yu, P. C. (2021). Optimal retirement policies with present-biased agents. *Journal of the European Economic Association 19*(4), 2085–2130.

A Nondiscrimination Testing and Safe Harbor 401(k) Plans

This appendix provides a brief overview of nondiscrimination testing and the safe harbor 401(k) plan designs; Israel Preninger and O'Dea (2026) provides a more detailed treatment.

Since the Revenue Act of 1942, Congress has conditioned the tax-preferred status of employer-sponsored retirement plans on nondiscrimination requirements designed to prevent plan benefits from flowing disproportionately to highly compensated employees. For 401(k) plans, these requirements are enforced through two annual nondiscrimination tests: the Actual Deferral Percentage (ADP) test, which compares elective deferral rates of highly compensated and non-highly compensated employees, and the Actual Contribution Percentage (ACP) test, which performs the same comparison for employer matching contributions and after-tax employee contributions.[13] Plans that fail must take corrective action or risk losing their tax-qualified status.

Since 1996, employers have been able to satisfy the ADP test automatically by adopting one of a small set of statutorily specified plan designs known as *safe harbors*. There are two safe harbor regimes:

- **Traditional safe harbor** (Small Business Job Protection Act, 1996). The basic matching formula requires the employer to match 100% of employee contributions up to 3% of compensation, plus 50% on the next 2% of compensation, for a maximum employer cost of 4% of compensation. Safe harbor contributions must vest immediately, and automatic enrollment is optional.

- **Qualified Automatic Contribution Arrangement (QACA) safe harbor** (Pension Protection Act, 2006). The basic matching formula requires the employer to match 100% of the first 1% of compensation contributed, plus 50% on the next 5%, for a maximum employer cost of 3.5% of compensation. Plans must adopt automatic enrollment with an initial contribution rate of at least 3%, increasing by at least one percentage point per year until it reaches at least 6% (with the permissible default rate capped at 15%, raised from 10% by the SECURE Act of 2019), and may impose a two-year cliff vesting schedule on safe harbor contributions.

In addition to the two basic matching formulas, both regimes permit two further routes to satisfying the safe harbor:

- *Alternative matching formulas.* Either regime may be satisfied through any matching formula that (i) does not increase the match rate as the employee deferral rate rises, and (ii) is at least as generous as the corresponding basic formula at every deferral rate. The most common alternative under the traditional regime is a 100% match on the first 4% of compensation.

[13] Under the ADP test, the average deferral rate of highly compensated employees may not exceed the greater of (i) 1.25 times the average deferral rate of non-highly compensated employees, and (ii) the lesser of two times the non-highly compensated rate and that rate plus two percentage points. The ACP test applies the same thresholds to employer matching contributions and after-tax employee contributions.

- *Non-elective contribution path.* Either regime may also be satisfied without any match: the employer instead contributes a minimum of 3% of compensation to each eligible non-highly compensated employee, regardless of whether the employee defers. This satisfies the ADP test directly. ACP testing will still apply if the employer also makes matching contributions, unless those matches separately satisfy the ACP safe harbor.

Choukhmane et al. (2026) documents two facts about safe harbor adoption that motivate the focus of this paper. They show that the share of workers in safe harbor plans has grown substantially over the past two decades, and that today around half of plans operate under safe harbor formulas. They also show that transitions out of a safe harbor are rare: once a plan adopts a safe harbor design, it is unlikely to leave.

Figure A.1, reproduced from Choukhmane et al. (2026), shows that the share of workers in plans with safe-harbor matching formulas has grown substantially over the past two decades, and that today around half of plans offer such formulas.

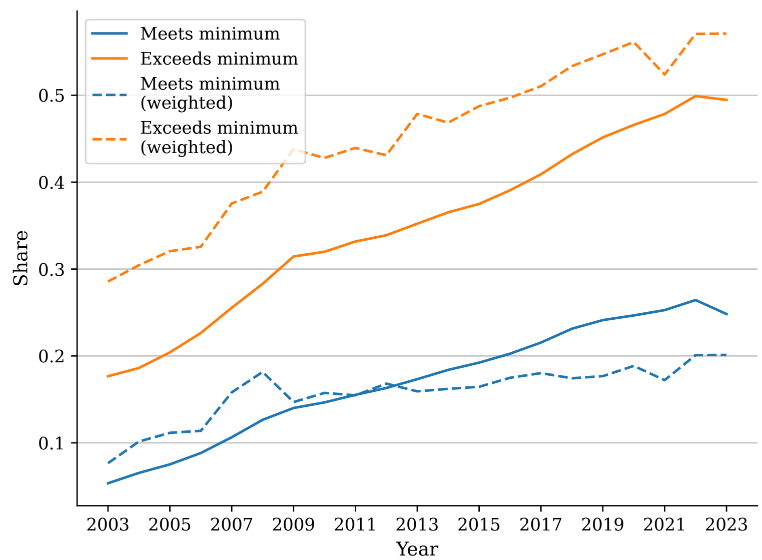

Figure A.1: Share of workers in safe harbor plans over time

Notes: The figure, reproduced from Choukhmane et al. (2026), shows the proportion of plans (weighted by participant count and unweighted) that have a matching formula satisfying the safe harbor matching requirements. The calculation does not account for other safe harbor requirements, such as vesting and auto-enrollment.

B Model Appendix

B.1 Extension to Multi-Tier Formulas

The prediction rule we laid out applies to single-tier match formulas. We now show that the same logic extends to formulas with multiple tiers as long as the match rate is non-increasing in employee saving. To derive this prediction rule, we naturally extend the economic logic of Assumptions 2 and 3 to multi-tier schedules: we assume that agents respond only to the local marginal incentives they face (a multi-tier analogue to Assumption 2) and that inframarginal changes to earlier tiers create no wealth effects that alter saving (a multi-tier analogue to Assumption 3).

Consider a two-tier plan that matches at rate m_1 on contributions up to c_1 and at rate m_2 on contributions between c_1 and c_2, where $m_1 > m_2$. We denote this $s(m_1, c_1; m_2, c_2)$. An individual facing this plan encounters three distinct marginal returns to saving: $1 + m_1$ below c_1, $1 + m_2$ between c_1 and c_2, and 1 above c_2.

To predict saving under this plan, we require three observable quantities: saving under no match, $s(0, .)$; saving under an unlimited match at the lower rate, $s(m_2, 1)$; and saving under an unlimited match at the higher rate, $s(m_1, 1)$.

The prediction proceeds in two steps, working from the top tier down.

Step 1: Predict saving under the second tier alone. Consider a hypothetical single-tier plan that matches at rate m_2 up to c_2. Applying the prediction rule from the single-tier case to the observables $s(0, .)$ and $s(m_2, 1)$ yields a prediction $s(m_2, c_2)$ that takes one of three values: $s(0, .)$ if the individual is a high saver relative to c_2, c_2 if a medium saver, and $s(m_2, 1)$ if a low saver.

Step 2: Predict saving under both tiers. We now introduce the first tier, which raises the match rate from m_2 to m_1 on contributions below c_1. We reapply the three-case analysis from the single-tier derivation, with $s(m_2, c_2)$ replacing $s(0, .)$ as the baseline:

- **High saver relative to** c_1: $s(m_2, c_2) > c_1$. The individual already saves above c_1 under the second tier alone, so introducing the first tier affects only inframarginal dollars and leaves the marginal return to saving unchanged. By Assumption 3, $s(m_1, c_1; m_2, c_2) = s(m_2, c_2)$.

- **Medium saver relative to** c_1: $s(m_2, c_2) \leq c_1$ and $s(m_1, 1) \geq c_1$. As in the single-tier medium-saver case, Assumption 2 rules out saving strictly below c_1 and Assumption 3 rules out saving strictly above c_1, yielding $s(m_1, c_1; m_2, c_2) = c_1$.

- **Low saver relative to** c_1: $s(m_2, c_2) \leq c_1$ and $s(m_1, 1) < c_1$. As in the single-tier low-saver case, the cap c_1 is non-binding and $s(m_1, c_1; m_2, c_2) = s(m_1, 1)$.

The prediction therefore takes one of three values, determined by the position of the two inputs, $s(m_2, c_2)$ and $s(m_1, 1)$, relative to c_1. The substantive change from the single-tier case is that the no-match benchmark $s(0, .)$ is replaced by $s(m_2, c_2)$: an individual who would already save above c_1 under the second tier alone is unaffected by the introduction of the first tier.

General K-tier formulas. The same logic applies recursively to plans with K tiers, where tier k matches at rate m_k on contributions between c_{k-1} and c_k (with $c_0 = 0$ and $m_1 > m_2 > \cdots > m_K$). The prediction proceeds from tier K down to tier 1, with each step applying the single-tier prediction rule using the prediction from the previous step as the baseline. This requires $K + 1$ observable quantities: saving under no match and saving under an unlimited match at each of the K match rates.

C Additional Exhibits

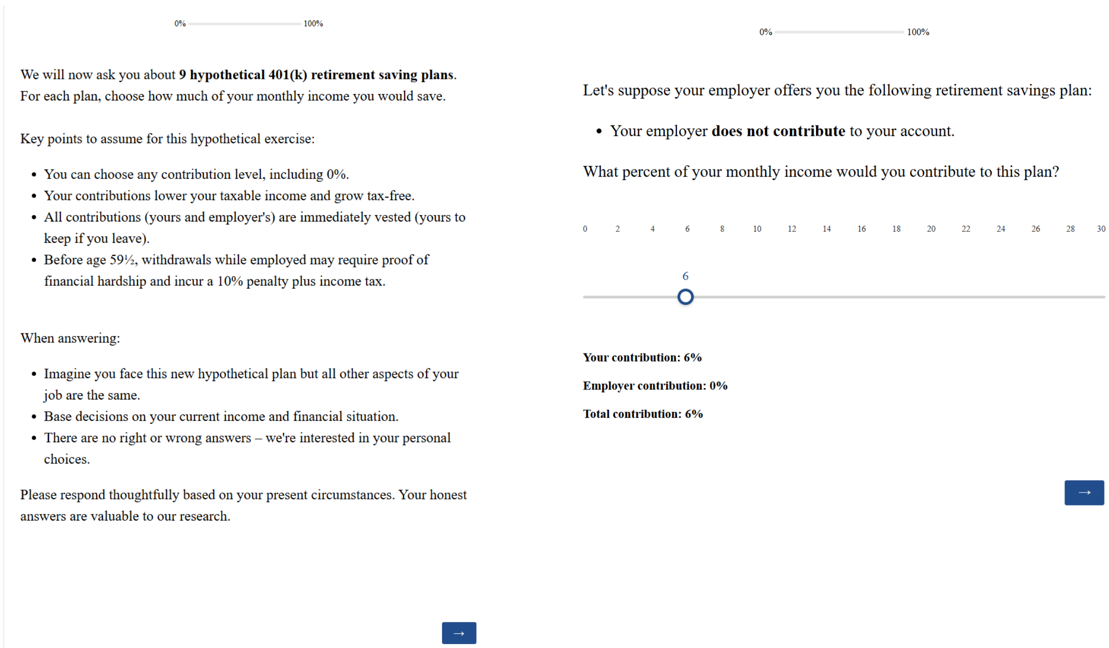

(a) Introductory Instructions (b) No Match Vignette

Figure C.2: 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 C.1: Demographic Characteristics of Survey Sample

Sample	**Administrative Population**	**Survey Population**	**Final Sample**
Count	1,859,488	340,653	1,634
Age			
18-35	26%	30%	7%
35-45	27%	27%	13%
45-55	24%	24%	29%
55-65	20%	15%	38%
65-75	4%	4%	13%
Income			
0-30k	10%	10%	9%
30-60k	24%	25%	22%
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	20%	22%	15%
6-9	14%	15%	11%
9-12	11%	11%	10%
12-15	7%	7%	8%
15+	27%	23%	40%

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. We weight our normative results in Section 6 to match the observable joint distribution of age and income in the administrative population.

Table C.2: Effect of NEC Rate on Contribution Rates

	Dependent variable: Administrative Contribution Rate					
	All Participants			New Participants		
	(1) All	(2) AE only	(3) Non-AE	(4) All	(5) AE only	(6) Non-AE
NEC Rate	0.007	0.063***	0.0005	0.076***	0.089***	0.031***
	(0.005)	(0.002)	(0.0004)	(0.009)	(0.004)	(0.011)
Observations	1,322,047	1,039,592	282,455	156,070	128,360	27,710
R^2	0.142	0.155	0.115	0.221	0.234	0.178
Plan FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses.

*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 the 10 largest auto-enrollment (AE) and 10 largest non-auto-enrollment (non-AE) plans that experienced a NEC rate increase from a 2010–2023 administrative panel. Columns (1) to (3) show all participants. Columns (4) to (6) restrict to new participants, defined as individuals hired and first eligible to contribute in a given year. Columns (1) and (4) report the baseline specification pooling plans with and without auto-enrollment. Columns (2) and (5) show results for auto-enrollment plans only. Columns (3) and (6) show results for plans without auto-enrollment only. All regressions include fixed effects for plan, year, and demographic controls.

Table C.3: Moving to the Savings-Maximizing Frontier: 25 Most Common Plans (No Auto-Enrollment)

			Employee Saving			Inequality (CV)					
Current Plan	Frontier Plan	NEC	Base	New	Δ	Base	New	Δ	Δ Part.	$\Delta S{\geq}6$	$\Delta S{\geq}12$
100-3-50-5	50-2.25-25-20	0.54	6.16	6.75	0.59	0.73	0.73	−0.00	34	−6	1
50-6-0-0	25-14.5-0-0	0.33	6.20	6.48	0.28	0.77	0.77	−0.00	35	−7	4
100-1-50-6	25-19-0-0	0.55	6.24	6.64	0.40	0.75	0.75	−0.00	34	−8	4
100-5-0-0	50-7-25-19.75	0.48	6.20	6.92	0.71	0.74	0.74	−0.00	34	−3	5
50-4-0-0	50-1-25-7.25	0.09	6.01	6.21	0.20	0.75	0.73	−0.01	35	−7	0
100-6-0-0	50-11-25-20	0.63	6.32	7.08	0.76	0.75	0.74	−0.00	34	−3	−13
50-7-0-0	25-18.5-0-0	0.45	6.29	6.62	0.33	0.78	0.78	−0.01	35	−7	5
100-4-0-0	50-2.5-25-20	0.53	6.07	6.75	0.68	0.73	0.73	−0.00	34	−6	1
100-1.25-0-0	50-1-25-4	0.03	5.84	5.99	0.16	0.72	0.72	−0.01	34	0	0
50-8-0-0	50-1.5-25-20	0.38	6.39	6.73	0.34	0.80	0.79	−0.00	35	−7	−9
50-10-0-0	50-5.5-25-19.75	0.14	6.58	6.84	0.26	0.83	0.83	−0.00	35	0	−9
25-4-0-0	—	—	—	—	0.00	—	—	—	—	—	—
100-3-0-0	50-1-25-13.25	0.35	5.96	6.49	0.53	0.73	0.72	−0.00	34	−11	4
100-3-50-6	50-4.25-25-20	0.49	6.27	6.79	0.53	0.74	0.73	−0.01	34	−4	1
100-2-50-6	50-2-25-20	0.52	6.25	6.74	0.49	0.74	0.74	−0.00	34	−2	1
100-4-50-6	50-6.5-25-20	0.49	6.29	6.90	0.61	0.74	0.74	−0.00	34	−4	5
100-1-25-4	50-1-25-6	0.07	6.00	6.13	0.13	0.73	0.72	−0.01	34	6	0
50-5-0-0	50-1-25-9.5	0.16	6.10	6.34	0.24	0.76	0.74	−0.02	35	−7	7
25-6-0-0	—	—	—	—	0.00	—	—	—	—	—	—
50-1-25-6	—	—	—	—	0.00	—	—	—	—	—	—
100-8-0-0	50-19-25-20	1.26	6.66	7.33	0.67	0.77	0.71	−0.05	34	−3	−13
100-3-50-7	50-6-25-20	0.46	6.38	6.87	0.50	0.75	0.74	−0.01	34	−4	−13
50-3-0-0	50-1-25-5	0.04	5.93	6.05	0.12	0.74	0.73	−0.01	35	7	0
100-8.5-0-0	50-19-25-20	1.49	6.71	7.33	0.61	0.77	0.68	−0.09	34	−3	−13
100-2-50-4	25-14.75-0-0	0.44	6.05	6.49	0.44	0.73	0.73	−0.00	34	−8	4

Notes: This table reports the effect of moving each of the 25 most common plans to its cost-equivalent, saving-maximizing frontier alternative, assuming No Auto-Enrollment. Column (1) gives the current plan. Columns (2) and (3) give the frontier plan formula and its non-elective contribution rate. The "Employee Saving" columns report the baseline and new employee saving rates (in percent of pay) and the change between them. The "Inequality" columns report the baseline and new coefficient of variation in employer contributions and the change between them. The final three columns report changes in participation and in the share of employees saving above specified thresholds. Δ Part. is the predicted change in the participation rate (in percentage points). $\Delta S{\geq}6$ and $\Delta S{\geq}12$ report the change in the fraction of employees with total contributions (inclusive of employer contribution) of at least 6% and 12% of pay, respectively (in percentage points). Plans marked with dashes are already on the frontier. Plans are ordered by the number of participants in the administrative data.

Table C.4: Moving to the Inequality-Minimizing Frontier: 25 Most Common Plans (No Auto-Enrollment)

			Employee Saving			Inequality (CV)					
Current Plan	Frontier Plan	NEC	Base	New	Δ	Base	New	Δ	Δ Part.	$\Delta S \geq 6$	$\Delta S \geq 12$
100-3-50-5	50-1-25-6.5	1.45	6.16	6.16	0.00	0.73	0.33	−0.40	34	−6	−3
50-6-0-0	50-1-25-7	0.66	6.20	6.20	0.00	0.77	0.50	−0.27	35	−6	−0
100-1-50-6	50-1-25-7.75	0.92	6.24	6.25	0.01	0.75	0.45	−0.30	34	−2	0
100-5-0-0	50-1-25-7.25	2.02	6.20	6.21	0.01	0.74	0.29	−0.45	34	−3	−2
50-4-0-0	50-1-25-4.5	0.41	6.01	6.02	0.01	0.75	0.51	−0.24	35	−6	−12
100-6-0-0	50-1-25-9	2.42	6.32	6.32	0.00	0.75	0.29	−0.45	34	−3	−15
50-7-0-0	50-1-25-8.75	0.72	6.29	6.30	0.01	0.78	0.53	−0.26	35	−7	1
100-4-0-0	50-1-25-5.5	1.59	6.07	6.09	0.02	0.73	0.29	−0.44	34	−5	−3
100-1.25-0-0	25-2-0-0	0.50	5.84	5.84	0.01	0.72	0.29	−0.43	34	−0	0
50-8-0-0	50-1-25-11.25	0.77	6.39	6.40	0.01	0.80	0.56	−0.23	35	−2	−8
50-10-0-0	25-17.75-0-0	1.08	6.58	6.59	0.01	0.83	0.58	−0.24	35	−2	−9
25-4-0-0	—	—	—	—	0.00	—	—	—	—	—	—
100-3-0-0	50-1-25-3.75	1.20	5.96	5.97	0.01	0.73	0.29	−0.44	34	−5	−1
100-3-50-6	50-1-25-8	1.55	6.27	6.27	0.01	0.74	0.36	−0.38	34	−6	−3
100-2-50-6	50-1-25-7.75	1.25	6.25	6.25	0.00	0.74	0.39	−0.35	34	−2	−3
100-4-50-6	50-1-25-8.75	1.82	6.29	6.30	0.01	0.74	0.34	−0.40	34	−6	6
100-1-25-4	50-1-25-4.25	0.30	6.00	6.00	0.01	0.73	0.55	−0.18	34	0	0
50-5-0-0	50-1-25-5.75	0.53	6.10	6.11	0.01	0.76	0.50	−0.25	35	−6	0
25-6-0-0	—	—	—	—	0.00	—	—	—	—	—	—
50-1-25-6	—	—	—	—	0.00	—	—	—	—	—	—
100-8-0-0	25-19.5-0-0	3.19	6.66	6.66	0.00	0.77	0.34	−0.43	34	−3	−16
100-3-50-7	25-12-0-0	1.68	6.38	6.38	0.00	0.75	0.40	−0.34	34	−6	−13
50-3-0-0	50-1-25-3.25	0.29	5.93	5.94	0.01	0.74	0.52	−0.22	35	0	0
100-8.5-0-0	50-1-25-19.75	3.25	6.71	6.71	0.00	0.77	0.34	−0.43	34	−3	−14
100-2-50-4	50-1-25-5	1.01	6.05	6.05	0.01	0.73	0.36	−0.37	34	−2	−1

Notes: This table reports the effect of moving each of the 25 most common plans to its cost-equivalent, inequality-minimizing frontier alternative, assuming No Auto-Enrollment. Column (1) gives the current plan. Columns (2) and (3) give the frontier plan formula and its non-elective contribution rate. The "Employee Saving" columns report the baseline and new employee saving rates (in percent of pay) and the change between them. The "Inequality" columns report the baseline and new coefficient of variation in employer contributions and the change between them. The final three columns report changes in participation and in the share of employees saving above specified thresholds. Δ Part. is the predicted change in the participation rate (in percentage points). $\Delta S \geq 6$ and $\Delta S \geq 12$ report the change in the fraction of employees with total contributions (inclusive of employer contribution) of at least 6% and 12% of pay, respectively (in percentage points). All changes are in percentage points. Plans marked with dashes are already on the frontier. Plans are ordered by the number of participants in the administrative data.

Table C.5: Moving to the Savings-Maximizing Frontier: 25 Most Common Plans (Auto-Enrollment)

			Employee Saving			Inequality (CV)					
Current Plan	Frontier Plan	NEC	Base	New	Δ	Base	New	Δ	Δ Part.	$\Delta S \geq 6$	$\Delta S \geq 12$
100-3-50-5	50-2.5-25-10	1.26	8.93	9.29	0.36	0.24	0.24	−0.00	4	−9	5
50-6-0-0	50-2.5-25-9	0.32	8.99	9.21	0.22	0.31	0.30	−0.00	5	−10	11
100-1-50-6	50-1.5-25-10	0.92	9.06	9.26	0.20	0.28	0.28	−0.00	4	−3	10
100-5-0-0	50-6.5-25-10	1.32	9.00	9.50	0.50	0.25	0.25	−0.00	5	−5	8
50-4-0-0	50-1.5-25-6	0.20	8.73	8.90	0.18	0.27	0.27	−0.00	5	−9	0
100-6-0-0	50-6-25-14.75	1.90	9.16	9.74	0.57	0.27	0.27	−0.00	5	−2	−21
50-7-0-0	50-4-25-10	0.18	9.13	9.34	0.21	0.33	0.33	−0.00	5	1	11
100-4-0-0	50-2.75-25-10	1.24	8.80	9.30	0.49	0.24	0.24	−0.00	5	−8	5
100-1.25-0-0	50-1-25-3.25	0.21	8.47	8.62	0.15	0.22	0.21	−0.01	5	0	0
50-8-0-0	50-2-25-13.75	0.69	9.27	9.48	0.21	0.35	0.35	−0.00	5	−3	−12
50-10-0-0	50-7-25-14.75	0.10	9.55	9.82	0.26	0.40	0.40	−0.00	5	0	−12
50-1-25-6	—	—	—	—	0.00	—	—	—	—	—	—
100-3-0-0	50-1.25-25-8	0.94	8.65	9.11	0.46	0.23	0.23	−0.00	5	−8	0
100-3-50-6	50-5-25-10	1.07	9.09	9.40	0.30	0.26	0.26	−0.00	4	−5	10
100-2-50-6	50-3.5-25-10	0.93	9.07	9.32	0.25	0.26	0.26	−0.00	4	0	5
100-4-50-6	50-6.25-25-10	1.29	9.13	9.48	0.35	0.25	0.25	−0.00	5	−5	8
100-1-25-4	50-1-25-5	0.28	8.70	8.79	0.09	0.24	0.24	−0.00	4	10	0
50-5-0-0	50-1.25-25-7.75	0.40	8.85	9.08	0.23	0.29	0.28	−0.00	5	−10	0
25-6-0-0	—	—	—	—	0.00	—	—	—	—	—	—
25-4-0-0	—	—	—	—	0.00	—	—	—	—	—	—
100-8-0-0	50-10-25-20	2.54	9.66	10.24	0.57	0.30	0.30	−0.00	5	−2	−14
100-3-50-7	50-6.75-25-10	1.06	9.26	9.52	0.26	0.27	0.27	−0.00	4	−5	−10
50-3-0-0	50-1-25-4.5	0.16	8.61	8.74	0.12	0.25	0.25	−0.01	5	10	0
100-8.5-0-0	50-11.5-25-20	2.72	9.74	10.31	0.57	0.31	0.31	−0.00	5	−2	−14
100-2-50-4	50-1-25-8	0.98	8.78	9.10	0.33	0.24	0.23	−0.00	4	−3	0

Notes: This table reports the effect of moving each of the 25 most common plans to its cost-equivalent, saving-maximizing frontier alternative, assuming Auto-Enrollment. Column (1) gives the current plan. Columns (2) and (3) give the frontier plan formula and its non-elective contribution rate. The "Employee Saving" columns report the baseline and new employee saving rates (in percent of pay) and the change between them. The "Inequality" columns report the baseline and new coefficient of variation in employer contributions and the change between them. The final three columns report changes in participation and in the share of employees saving above specified thresholds. Δ Part. is the predicted change in the participation rate (in percentage points). $\Delta S \geq 6$ and $\Delta S \geq 12$ report the change in the fraction of employees with total contributions (inclusive of employer contribution) of at least 6% and 12% of pay, respectively (in percentage points). Plans marked with dashes are already on the frontier. Plans are ordered by the number of participants in the administrative data.

Table C.6: Moving to the Inequality-Minimizing Frontier: 25 Most Common Plans (Auto-Enrollment)

			Employee Saving			Inequality (CV)					
Current Plan	Frontier Plan	NEC	Base	New	Δ	Base	New	Δ	Δ Part.	$\Delta S \geq 6$	$\Delta S \geq 12$
100-3-50-5	50-1-25-6.5	2.10	8.93	8.93	0.00	0.24	0.14	−0.10	4	−5	−5
50-6-0-0	50-1-25-7	0.95	8.99	9.00	0.01	0.31	0.21	−0.09	5	−3	−1
100-1-50-6	50-1.75-25-7.5	1.19	9.06	9.06	0.01	0.28	0.20	−0.08	4	−3	0
100-5-0-0	50-1.25-25-7	2.90	9.00	9.00	0.00	0.25	0.12	−0.13	5	−5	−2
50-4-0-0	50-1-25-4.5	0.60	8.73	8.74	0.01	0.27	0.19	−0.08	5	−8	−18
100-6-0-0	50-3.25-25-8	3.10	9.16	9.16	0.00	0.27	0.14	−0.13	5	1	−19
50-7-0-0	50-2-25-8	0.90	9.13	9.13	0.00	0.33	0.23	−0.10	5	−3	1
100-4-0-0	50-1.25-25-5.25	2.30	8.80	8.81	0.01	0.24	0.11	−0.13	5	−4	−5
100-1.25-0-0	25-2-0-0	0.73	8.47	8.48	0.01	0.22	0.11	−0.12	5	−0	0
50-8-0-0	50-2.25-25-10	0.90	9.27	9.28	0.01	0.35	0.27	−0.08	5	−3	−9
50-10-0-0	50-7.25-25-10	0.39	9.55	9.55	0.00	0.40	0.33	−0.08	5	0	−4
25-4-0-0	—	—	—	—	0.00	—	—	—	—	—	—
100-3-0-0	50-1-25-3.75	1.74	8.65	8.67	0.02	0.23	0.10	−0.13	5	−7	−1
100-3-50-6	50-1.75-25-7.75	2.11	9.09	9.09	0.00	0.26	0.16	−0.10	4	−5	12
100-2-50-6	50-1-25-7.75	1.82	9.07	9.07	0.00	0.26	0.18	−0.09	4	−3	0
100-4-50-6	50-2-25-8	2.49	9.13	9.13	0.00	0.25	0.15	−0.10	5	−5	10
100-1-25-4	50-1-25-4.25	0.43	8.70	8.71	0.01	0.24	0.20	−0.04	4	0	0
50-5-0-0	50-1-25-5.75	0.77	8.85	8.86	0.02	0.29	0.20	−0.09	5	−2	0
25-6-0-0	—	—	—	—	0.00	—	—	—	—	—	—
50-1-25-6	—	—	—	—	0.00	—	—	—	—	—	—
100-8-0-0	50-5-25-14.75	3.63	9.66	9.67	0.00	0.30	0.20	−0.11	5	1	−12
100-3-50-7	50-1.5-25-10	2.22	9.26	9.26	0.00	0.27	0.20	−0.08	4	−5	−10
50-3-0-0	50-1-25-3.25	0.42	8.61	8.62	0.01	0.25	0.18	−0.07	5	0	0
100-8.5-0-0	50-6.75-25-14	3.66	9.74	9.75	0.01	0.31	0.20	−0.11	5	1	−5
100-2-50-4	50-1-25-5	1.46	8.78	8.79	0.01	0.24	0.14	−0.10	4	−3	−1

Notes: This table reports the effect of moving each of the 25 most common plans to its cost-equivalent, inequality-minimizing frontier alternative, assuming Auto-Enrollment. Column (1) gives the current plan. Columns (2) and (3) give the frontier plan formula and its non-elective contribution rate. The "Employee Saving" columns report the baseline and new employee saving rates (in percent of pay) and the change between them. The "Inequality" columns report the baseline and new coefficient of variation in employer contributions and the change between them. The final three columns report changes in participation and in the share of employees saving above specified thresholds. Δ Part. is the predicted change in the participation rate (in percentage points). $\Delta S \geq 6$ and $\Delta S \geq 12$ report the change in the fraction of employees with total contributions (inclusive of employer contribution) of at least 6% and 12% of pay, respectively (in percentage points). Plans marked with dashes are already on the frontier. Plans are ordered by the number of participants in the administrative data.

Table C.7: Moving Safe Harbor Plans to Their Saving-Maximizing Alternatives

			Vanguard Sample			F5500 Sample		
Formula	AE	Target formula	Base EE pp of pay	Δ EE pp of pay	Δ / Base %	Workers millions	Agg. EE millions of $	Δ Agg. EE millions of $
Regular safe harbor								
100-3-50-5	No	50-2.25-25-20 + nec: 0.54	6.16	0.59	9.6%	4.75	13,836	1,330
100-3-50-5	Yes	50-2.5-25-10 + nec: 1.26	8.93	0.36	4.0%	1.41	8,736	352
100-6-0-0	No	50-11-25-20 + nec: 0.63	6.32	0.76	12.0%	2.55	7,485	900
100-6-0-0	Yes	50-6-25-14.75 + nec: 1.90	9.16	0.57	6.2%	1.61	13,147	818
100-5-0-0	No	50-7-25-19.75 + nec: 0.48	6.20	0.71	11.5%	2.08	6,933	794
100-5-0-0	Yes	50-6.5-25-10 + nec: 1.32	9.00	0.50	5.6%	1.29	9,805	545
100-4-0-0	No	50-2.5-25-20 + nec: 0.53	6.07	0.68	11.2%	1.81	8,098	907
100-4-0-0	Yes	50-2.75-25-10 + nec: 1.24	8.80	0.49	5.6%	1.11	8,228	458
100-3-50-6	No	50-4.25-25-20 + nec: 0.49	6.27	0.53	8.5%	0.48	2,617	221
100-3-50-6	Yes	50-5-25-10 + nec: 1.07	9.09	0.30	3.3%	0.36	3,120	103
100-4-50-6	No	50-6.5-25-20 + nec: 0.49	6.29	0.61	9.7%	0.26	1,191	115
100-4-50-6	Yes	50-6.25-25-10 + nec: 1.29	9.13	0.35	3.8%	0.29	1,927	74
Auto-enrollment safe harbor (QACA)								
100-1-50-6	Yes	50-1.5-25-10 + nec: 0.92	9.06	0.20	2.2%	0.65	3,483	77
100-2-50-6	Yes	50-3.5-25-10 + nec: 0.93	9.07	0.25	2.8%	0.18	1,341	37
Total for safe harbor formulas in F5500 sample						**18.75m**	**89.0bn**	**+ 6.7bn**

Notes: This table reports the predicted effects of moving safe harbor formulas to their corresponding saving-maximizing alternatives. For the Vanguard sample, *Base EE* and Δ EE are taken from Table C.3 for plans without auto-enrollment and Table C.5 for plans with auto-enrollment. *Base EE* is the average employee contribution rate under the status-quo formula, and Δ EE is the predicted change in the employee contribution rate from moving to the target formula. Δ/Base is the ratio of these two quantities and therefore reports the percent change in employee contribution rates. The F5500 sample corresponds to the sample of safe harbor plans collected from plans' form 5500 fillings in 2023 using LLMs by Choukhmane et al. (2026). In the F5500 columns, *Workers* reports the number of eligible workers covered by each formula and *Agg. EE* reports aggregate employee contributions in these plans in 2023 from Form 5500 schedule H filings. Δ Agg. EE is calculated by multiplying aggregate employee contributions by the corresponding Δ/Base value estimated in the Vanguard sample, thereby assuming that the same percentage increase in employee contribution rates applies to aggregate employee contributions in the F5500 sample. Dollar amounts are reported in millions.

Figure C.3: Comparison of Administrative and Stated Contribution Rates

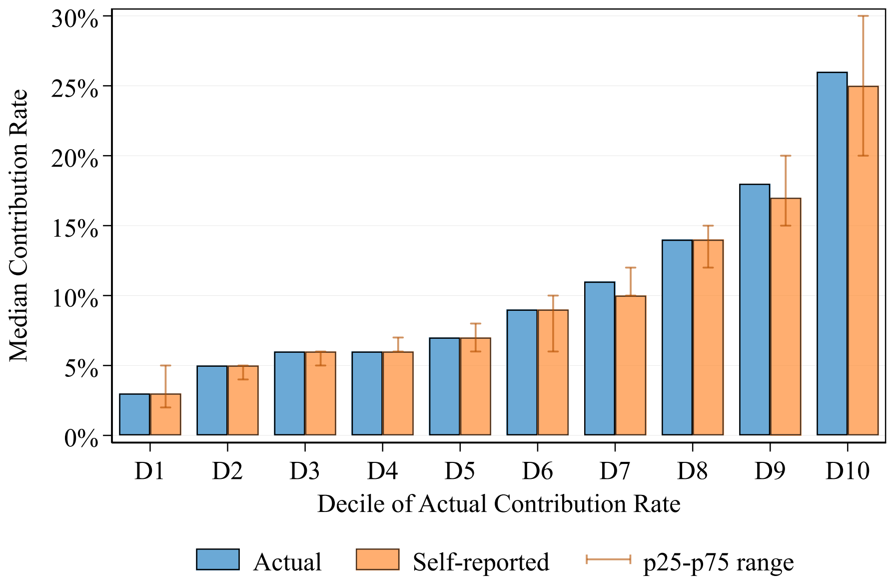

Notes: This figure compares self-reported contribution rates from the survey with actual contribution rates from administrative data. Each point on the horizontal axis represents a bin of the administrative contribution rate distribution, with a comparison provided to the median and 25th-75th percentile range of the stated rate.

Figure C.4: Comparison of Hypothetical and Actual Contribution Rates for Aligned Plans

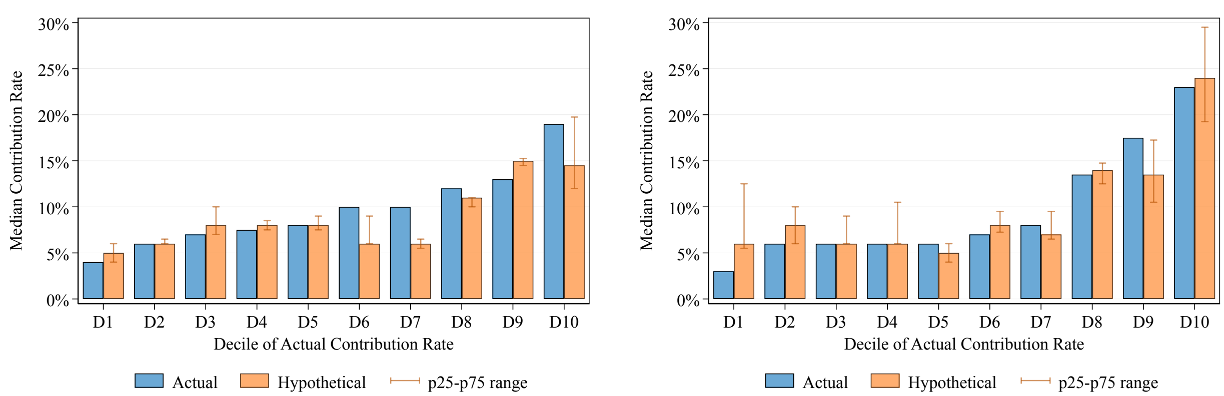

(a) 50% up to 6% plan

(b) 100% up to 6% plan

Notes: This figure compares hypothetical contribution rates from the survey with actual contribution rates from administrative data, for the subset of respondents whose actual plan matches the hypothetical scenario. Panel (a) shows respondents facing a 50% match up to 6% of pay, and panel (b) shows respondents facing a 100% match up to 6% of pay. Each point on the horizontal axis represents a bin of the administrative contribution rate distribution, with a comparison provided to the median and 25th-75th percentile range of the stated rate.

Figure C.5: Test-Retest Responses to Saving Under No Match

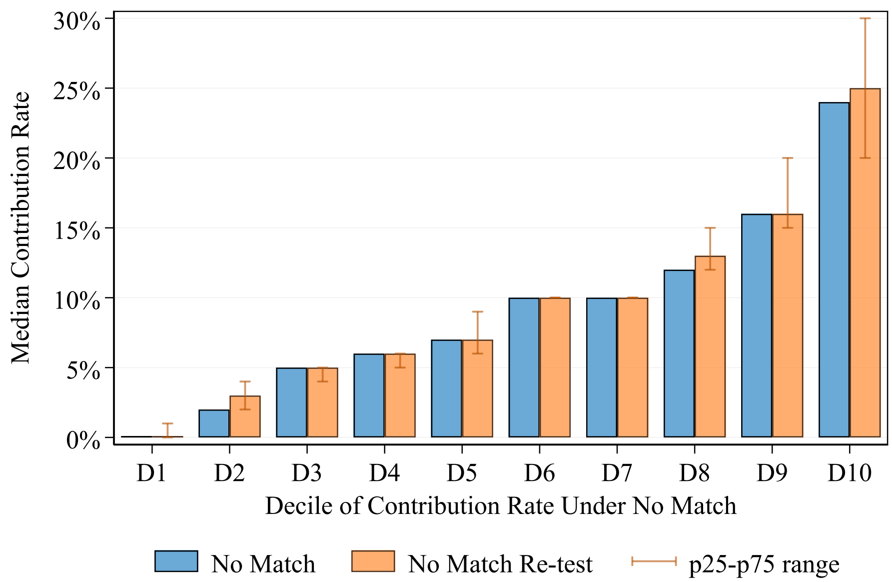

Notes: This figure compares responses to a "no match" hypothetical scenario asked at the beginning of the survey with responses to the same scenario asked at the end. Each point represents a bin of the initial response, with the 25th, 50th, and 75th percentiles of the retest response plotted for each bin.

Figure C.6: Wealth Effects Survey

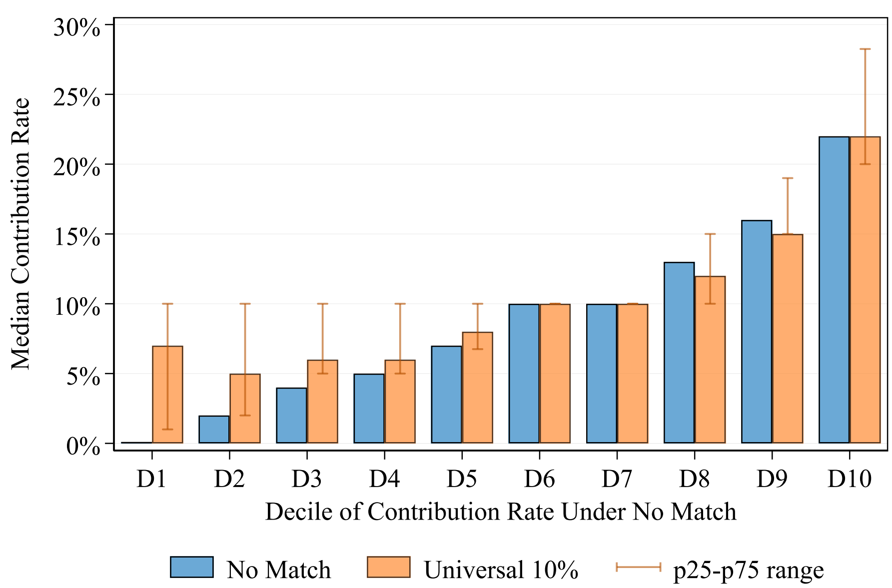

Notes: This figure compares responses to a "no match" hypothetical scenario with responses to a "universal contribution" scenario, in which the employer contributes a fixed amount regardless of employee contributions. Each point represents a bin of the initial response, with the 25th, 50th, and 75th percentiles of the universal contribution response plotted for each bin.

Figure C.7: Change in Saving when Non-Elective is increased: New Participants Only

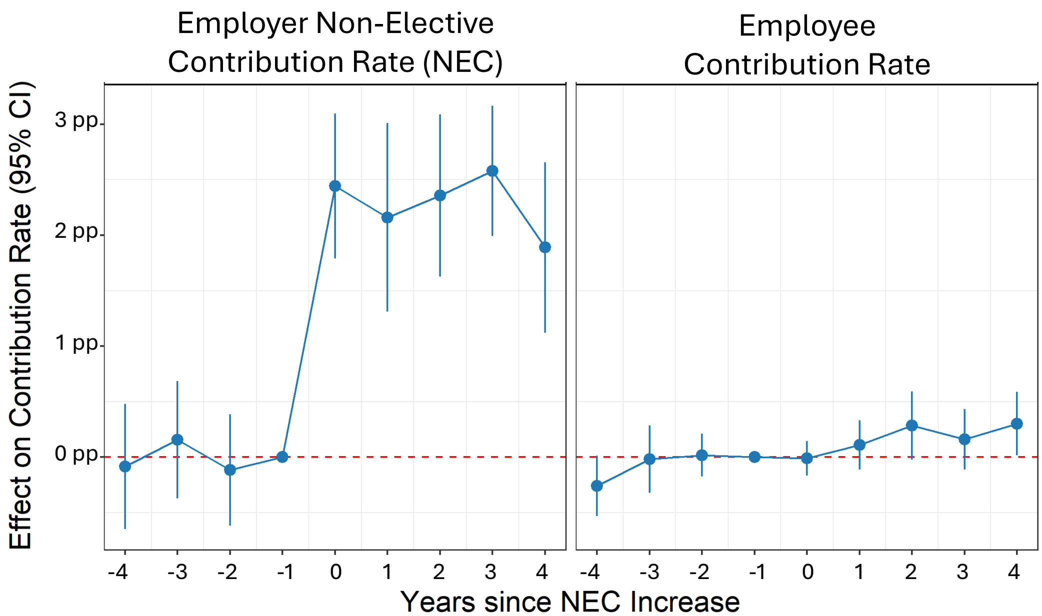

Notes: This figure reports event-study estimates from a staggered difference-in-differences design using a panel spanning 2010 to 2023, restricted to new participants, defined as individuals hired and first eligible to contribute in a given year. The sample comprises the 10 largest auto-enrollment (AE) and 10 largest non-auto-enrollment (non-AE) plans that experienced a non-elective contribution (NEC) rate increase. All specifications include plan and year fixed effects and control for quadratics in age, income, and tenure. The left panel reports the estimated effect on the NEC rate; the right panel reports the estimated effect on the elective contribution rate. Standard errors are clustered at the plan level. See Figure 3 for the analogous estimates using all participants.

Figure C.8: Predicted vs. Administrative Distributions by Demographics

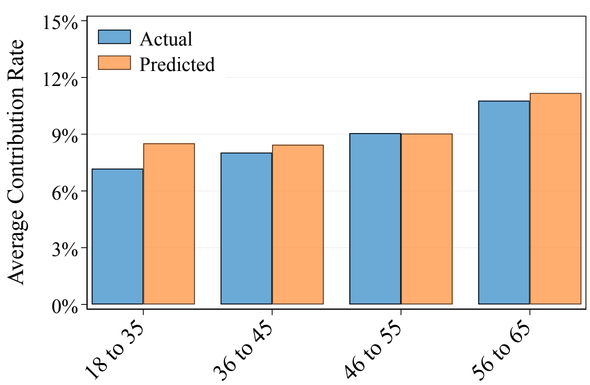

(a) By age

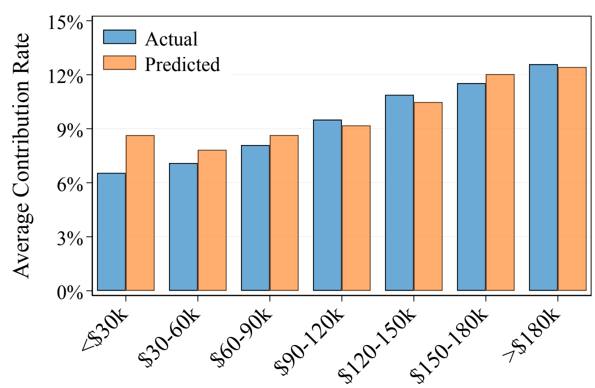

(b) By income

Notes: This figure compares average contribution rates in the administrative data with model-predicted average contribution rates across the 25 largest plans in our sample, disaggregated by age group (Panel a) and annual earnings group (Panel b). Survey responses are reweighted to match administrative demographics within each plan and aggregated across plans using participant counts. See Figure 6 for the aggregate distributional fit.

Figure C.9: Evolution of non-elective and contribution rates for 20 large plans

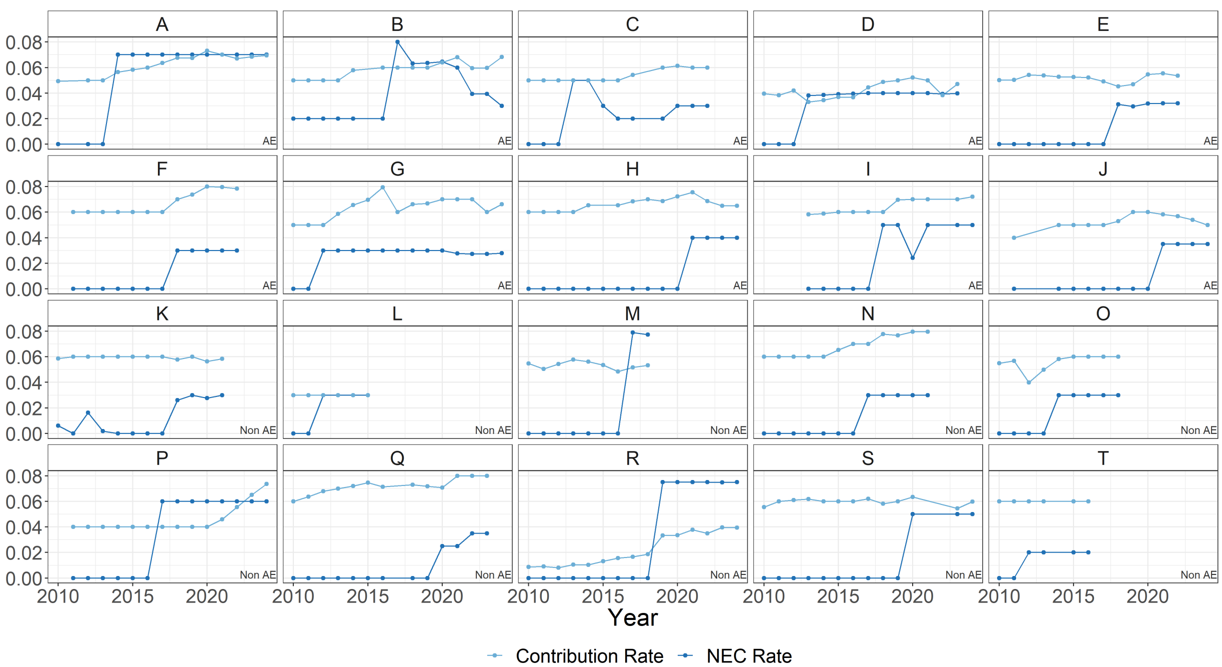

Notes: The figure displays the time series of non-elective contribution (NEC) rates and employee contribution rates for the 20 plans used in the event-study regression: the 10 largest auto-enrollment (AE) and 10 largest non-auto-enrollment (non-AE) plans that experienced a NEC rate increase. Each panel corresponds to a single plan, labeled by an anonymized identifier, with the plan type (AE or Non-AE) indicated in the bottom corner.

D Robustness to Dropping Those Who Fail Test-Retest

We evaluate how our results change if we exclude respondents who fail the test-retest check, defined as those whose responses to the no-match vignette at the beginning and end of the survey differ by more than one percentage point. This yields a sample of 1,259 respondents. This appendix reports our main results under this sample.

Figure D.10: Model Predictions: Passes Test-Retest

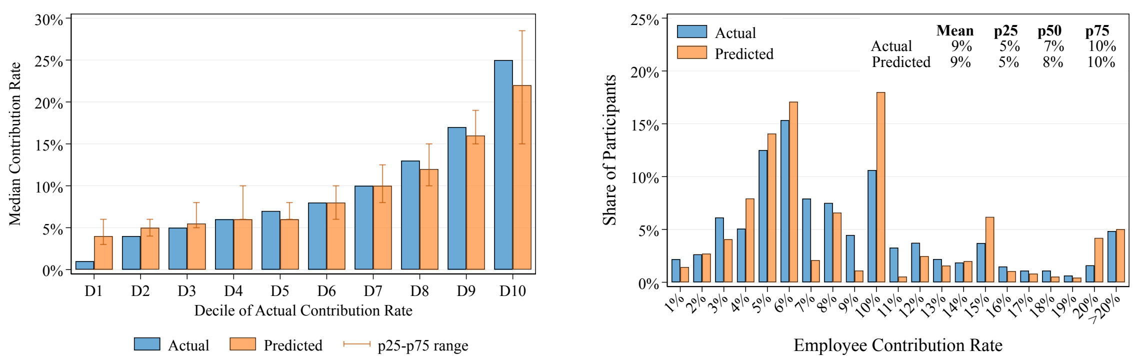

(a) Current Administrative Plan Saving (b) Predicted Distributions

Notes: Panel (a) compares respondents' saving as measured in the administrative data with the model-predicted contribution under their observed plan. The horizontal axis groups respondents into deciles of administrative contributions. Blue bars show the median administrative rate within each decile; orange bars show the median prediction, with error bars indicating the 25th-75th percentile range of predictions. Panel (b) compares the distribution of saving in the administrative data with the model-predicted distribution under each respondent's plan. We weight survey responses to match demographics in the administrative data. Sample restricted to respondents who pass the test-retest consistency check.

Figure D.11: Characterizing the frontier - Test-Retest Passing Sample

Passes Test Retest

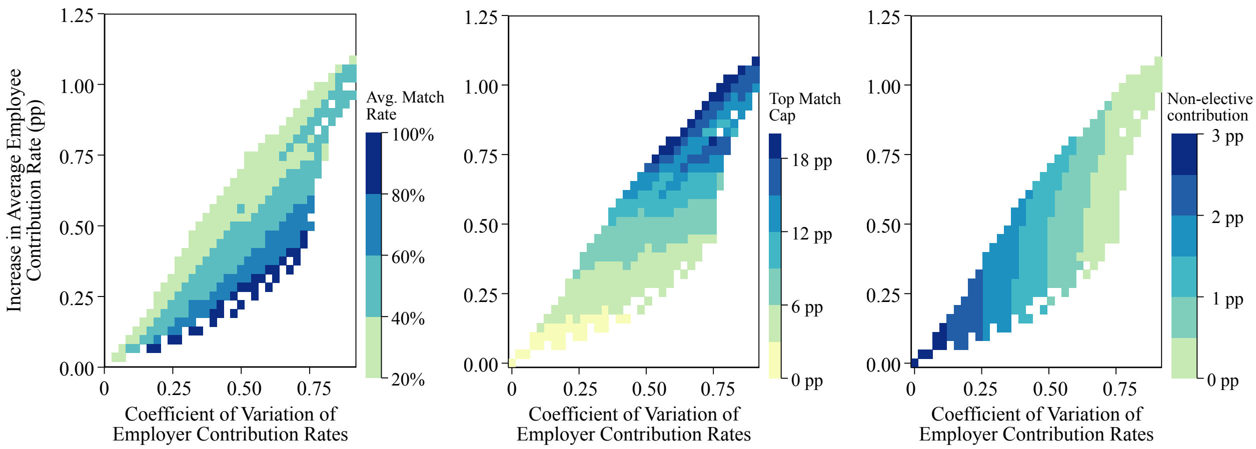

(a) Average Match Rate (b) Top Match Cap (c) Non-elective contribution

Notes: This figure characterizes how plan features map to positions in the equity–savings space. The horizontal axis is the cross-employee coefficient of variation of employer contribution rates, and the vertical axis is the increase in the average employee contribution rate relative to a no-match baseline. Cells are colored by the average value of a given plan characteristic among plans falling in that region of the space. Panel (a) shows the average match rate earned by an employee who contributes up to the highest match cap; plans on or near the frontier tend to have lower average match rates. Panel (b) shows the match cap, or in the case of a two-tier formula, the highest cap. Moving rightward along the frontier (toward greater average contribution rates) is associated with a higher cap. Panel (c) shows the non-elective contribution rate. Moving leftward along the frontier (toward greater equity in benefits) is associated with a higher non-elective contribution.

Figure D.12: Moving to the frontier for plans with a 50% match up to 6%: Test-Retest Passing Sample

Passes Test Retest

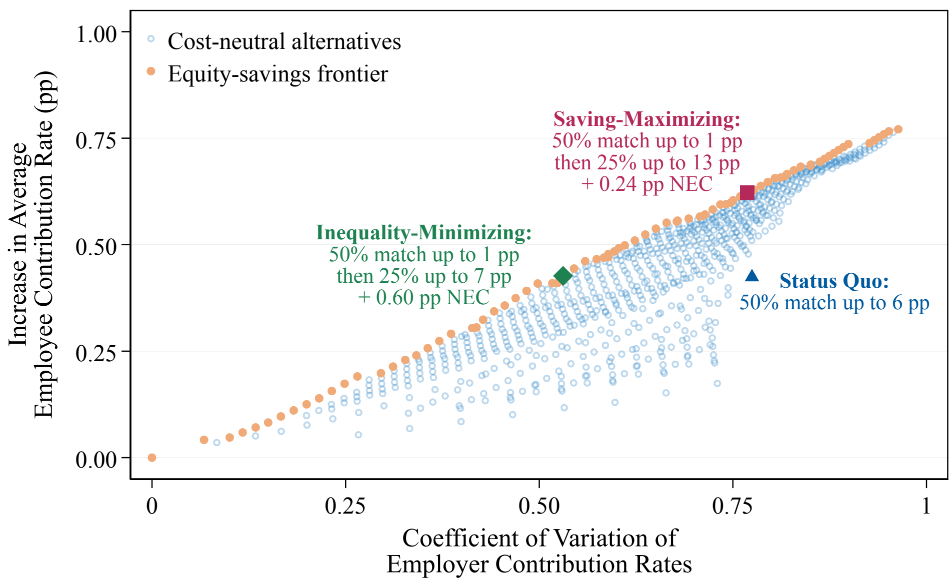

Notes: This figure plots the equity and savings characteristics of employer match formulas with the same expected cost as the status quo formula (50% match up to 6% of pay), shown as the blue triangle. Light blue squares denote the full set of cost-neutral alternatives, and orange circles trace the equity-savings frontier. We highlight two alternatives on the frontier. The red square (*Saving-Maximizing*) is the formula that maximizes the average employee contribution rate subject to no increase in the coefficient of variation of employer contribution rates relative to the status quo. The green diamond (*Inequality-Minimizing*) is the formula that minimizes the coefficient of variation of employer contribution rates subject to no decrease in the average employee contribution rate. NEC denotes a non-elective employer contribution. Absolute changes in average savings rates and employer-contribution inequality are reported in the next two tables.

Table D.8: Model Predictive Power

	Dependent variable: Contribution Rate							
	Full Sample				Passes Test Retest			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Predicted Rate	0.877***	1.076***		1.034***	0.938***	0.993***		0.982***
	(0.054)	(0.164)		(0.160)	(0.061)	(0.169)		(0.170)
Full Match		0.174***		0.171***		0.154***		0.151***
		(0.034)		(0.033)		(0.037)		(0.037)
No Match		−0.321**		−0.315**		−0.173		−0.187
		(0.135)		(0.130)		(0.131)		(0.133)
Log Income			1.957***	0.981***			2.149***	0.824**
			(0.342)	(0.316)			(0.411)	(0.378)
Age			−0.468***	−0.338**			−0.213	−0.070
			(0.159)	(0.136)			(0.172)	(0.148)
Exp(Age,2)			0.006***	0.004***			0.003**	0.001
			(0.002)	(0.001)			(0.002)	(0.001)
Tenure			0.008	−0.017			−0.031	−0.089
			(0.081)	(0.065)			(0.097)	(0.076)
Exp(Tenure,2)			0.001	0.001			0.001	0.002
			(0.002)	(0.002)			(0.003)	(0.002)
Observations	1,326	1,326	1,326	1,326	1,014	1,014	1,014	1,014
R^2	0.405	0.429	0.075	0.443	0.480	0.494	0.070	0.503

*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 only a constant and the predicted rate. Column (2) adds the two survey responses used to construct the prediction. Column (3) includes only demographic controls: log income, age, age squared, tenure, and tenure squared. Column (4) includes the predicted rate, survey responses, and demographic controls. Columns (5)–(8) repeat the same specifications for the sample that passes the test-retest check.

Table D.9: The effect of moving to the frontier: five most common plans, sample passing test retest

Panel A: No Auto-Enrollment

	Minimize Inequality				Maximize Savings			
Current Plan	Plan	NEC	Baseline CV	Change CV	Plan	NEC	Baseline Savings	Change Savings
100-3-50-5	25-7.75-0-0	1.49	0.74	−0.40	50-2.75-25-20	0.51	6.07	0.42
50-6-0-0	50-1-25-7.5	0.60	0.77	−0.24	50-1-25-13	0.24	6.10	0.20
100-1-50-6	50-1-25-8	0.89	0.76	−0.29	25-20-0-0	0.57	6.14	0.28
100-5-0-0	50-1-25-8	1.94	0.74	−0.43	50-7.25-25-20	0.52	6.12	0.52
50-4-0-0	25-5.75-0-0	0.43	0.75	−0.22	25-8.5-0-0	0.16	5.94	0.15

Panel B: Auto-Enrollment

	Minimize Inequality				Maximize Savings			
Current Plan	Plan	NEC	Baseline CV	Change CV	Plan	NEC	Baseline Savings	Change Savings
100-3-50-5	50-1.5-25-6.75	1.94	0.25	−0.10	50-2.75-25-10	1.22	8.80	0.29
50-6-0-0	50-1-25-7.5	0.88	0.32	−0.08	50-2-25-9.25	0.41	8.86	0.15
100-1-50-6	50-1-25-8	1.30	0.29	−0.07	50-1.75-25-10	0.87	8.91	0.15
100-5-0-0	50-1.25-25-7.75	2.79	0.27	−0.12	50-6.5-25-10	1.36	8.89	0.40
50-4-0-0	50-1-25-4.75	0.55	0.28	−0.07	50-1.5-25-6	0.21	8.62	0.13

Notes: This table shows, for the five most common plans, the counterfactual plan which is equally costly and either minimizes inequality (subject to no fall in employee saving), or maximizes saving (subject to no increase in inequality). Panel A assumes no auto-enrollment; Panel B assumes auto-enrollment.

E Targeting Effectiveness and the Dispersion of Employer Contributions

In Section 6.1, we argued that the coefficient of variation of employer contribution rates is a useful summary measure for targeting quality, on the grounds that it correlates strongly with three more direct proxies for targeting: the covariance between the formula's treatment effect on total contributions and (i) worker income, (ii) the worker's hypothetical contribution under no match, and (iii) the worker's hypothetical contribution under an unlimited 100% match. This appendix shows the underlying correlations and justifies the use of each proxy.

We compute the three covariances across the set of formulas that form the basis for analysis in Figure 8 (the set of formulas that cost 3 pp of payroll). Specifically, for each formula, we apply the predictive model from Section 4 to forecast each survey respondent's contribution under that formula and the implied employer contribution. We then compute, across respondents, the coefficient of variation of employer contribution rates and the three covariances.

Figure E.13 plots, in each of three panels, the coefficient of variation on the horizontal axis against one of the three targeting covariances on the vertical axis.

(a) Income. Panel (a) shows that formulas with less dispersion in employer contribution rates generate a lower covariance between income and the treatment effect of the formula on total contributions, i.e. the saving gains from lower dispersion plans accrue disproportionately to lower-income workers. Lower-income workers are more likely to rely on means-tested social insurance programs in retirement (Medicaid, Supplemental Security Income), and the asset tests for these programs depress their private saving below the social optimum (Hubbard et al. (1995); Sleet and Yeltekin (2006)). Formulas that direct contributions toward lower-income workers therefore better target the workers whose saving is most affected by such programs, and most likely to generate fiscal externalities.

(b) Hypothetical contributions under no match. Panel (b) shows the analogous pattern using the worker's hypothetical contribution rate under no match as the proxy. Formulas with less dispersion in employer contributions generate a smaller covariance between this proxy and the treatment effect on total contributions. The saving gains from low dispersion plans thus accrue disproportionately to workers who would contribute little absent any intervention, the group for whom a corrective intervention might be most valuable.

(c) Hypothetical contributions under an unlimited 100% match. Panel (c) uses the worker's hypothetical contribution rate under an unlimited 100% match. The pattern is similar: low-dispersion formulas direct their saving gains more toward workers who would contribute little even when offered a generous match. Choukhmane and Palmer (2025) show that in a model with heterogeneous present bias, the take-up of matching incentives is monotonically decreasing in the level of present bias.

If low baseline saving (in panel (b)) or low take-up of generous financial incentives (in panel (c)) reflect behavioral biases such as present bias or a lack of attention to match incentives rather than genuinely low saving preferences, then lower covariance signals better targeting of behavioral undersaving.

Across all three proxies, the same set of formulas (i.e., those with low dispersion in employer contribution rates) is identified as better-targeted. This justifies our use of the coefficient of variation as a summary statistic for targeting quality in the main text.

Figure E.13: Targeting Quality and Inequality in Employer Contributions

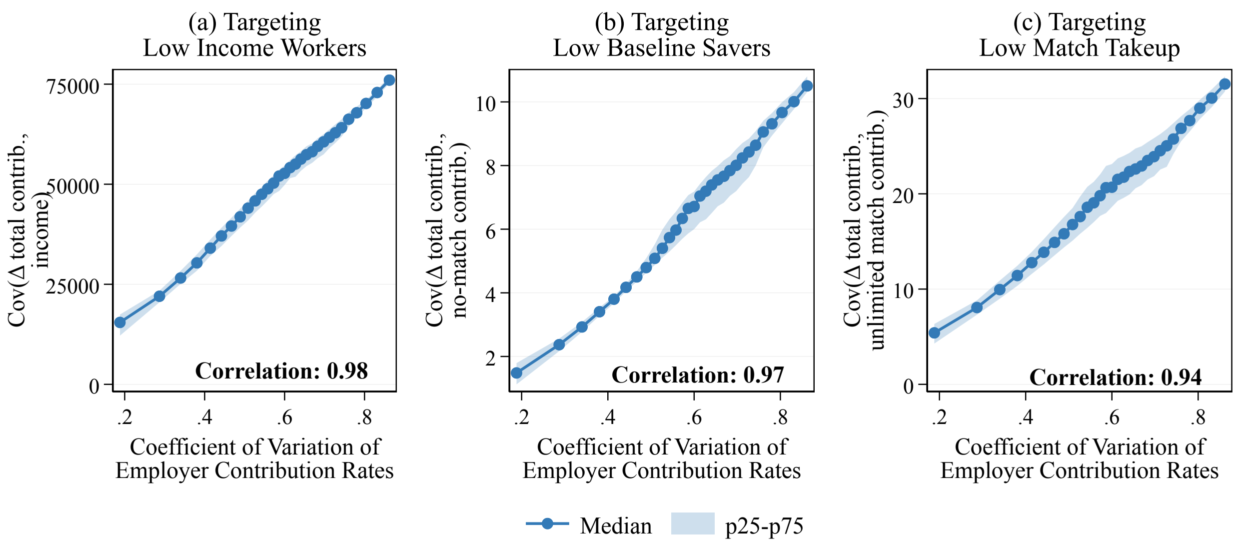

Notes: Each panel plots, across matching formulas, the coefficient of variation of employer contribution rates on the x-axis against a proxy for targeting quality on the y-axis. In all panels, lower covariance on the vertical axes indicates better targeting. Panel (a) uses the covariance between income and the formula's treatment effect on total contribution rates, defined as the predicted increase in combined employee and employer contribution rates relative to a no-match counterfactual. Because lower-income workers rely more on means-tested programs in retirement, lower covariance with income indicates better targeting of workers who face stronger saving disincentives from means-tested old-age safety nets. Panel (b) uses the covariance between the treatment effect on total contribution rates and the worker's hypothetical employee contribution rate under no match, capturing whether the formula targets workers who would contribute little absent any intervention. Panel (c) uses the analogous covariance with employee contribution rates under an unlimited 100% match. If low baseline saving (in panel (b)) or low take-up of generous financial incentives (in panel (c)) reflect behavioral biases such as present bias or a lack of attention to match incentives rather than genuinely low saving preferences, then lower covariance signals better targeting of behavioral undersaving.

F Online Appendix: Survey Instrument

This appendix reproduces the text of the survey questions used in our empirical analysis. The questions that follow are presented in the order in which respondents encountered them. Slider questions show the slider widget, the scale, and (where given) the endpoint labels exactly as they appeared in the survey. Answer choices are reproduced verbatim, with the internal Qualtrics response code shown in parentheses. In the hypothetical matching-design scenarios in Section F.2, the bold label preceding each scenario (e.g., "No Match", "100% Match") is our internal name for the plan and was not shown to respondents.

F.1 401(k) Background

Thank you for participating in our study. We will start by asking you a few questions about your current employer-sponsored retirement savings plan with Vanguard. Please answer these questions to the best of your knowledge.

You currently participate in a Vanguard employer-sponsored 401(k)/403(b) retirement savings plan. In 401(k) or 403(b) plans, you and your employer can contribute each month to your retirement savings account. Your employer contributions may be in the form of direct contributions or contributions that match your own savings.

1. **We would like to understand how your employer makes matching contributions into your current Vanguard retirement plan. Please select which of the following is true.**
   - My employer does not make matching contributions. (1)
   - My employer contributes $1 for every $1 I contribute, up to a certain percentage of my salary (also called the match cap). Please enter the match cap as a number below. (2) ____________
   - My employer contributes 50¢ for every $1 I contribute, up to a certain percentage of my salary (also called the match cap). Please enter the match cap as a number below. (3) ____________
   - My employer makes contributions with multiple match rates or has some other type of matching contributions. (4)
   - I do not know / I am not sure (5)

2. **How much do you currently save in your Vanguard retirement plan each month? Please include only contributions which are deducted from your paycheck. Do not include any employer contributions (whether matching or some other type).** You can choose to enter your monthly contributions as a percent share of your salary or a dollar amount.
   - I contribute ___% of my monthly paycheck. Please enter the percentage below: (1) ____________
   - I contribute $___ a month. Please enter the dollar amount below. (2) ____________
   - I do not currently contribute to this plan. (3)
   - I do not know / I am not sure (4)

3. **We are interested in understanding how you decide which percentage of your income to contribute into your Vanguard retirement plan.**
   Which of the following reasons, if any, are important factors in driving your contribution decision? Select all that apply.
   - I limit my contributions because otherwise I will be unable to cover my financial needs right now. (1)
   - I limit my contributions because otherwise I may be unable to cover my financial needs in the near-future. (2)
   - I contribute enough to maximize my employer's matching contribution. (3)
   - I contribute enough to ensure a comfortable retirement. (4)
   - I am following the retirement savings advice of a coworker, family member or financial planner. (8)

- I am following a common rule-of-thumb around retirement savings. (9)
- I set my contributions to the amount that maximizes the tax advantages. (11)
- Other. (14)

F.2 Hypothetical 401(k) Plans

We will now ask you about 9 hypothetical 401(k) retirement saving plans.[14] For each plan, choose how much of your monthly income you would save.
Key points to assume for this hypothetical exercise:

- You can choose any contribution level, including 0%.
- Your contributions lower your taxable income and grow tax-free.
- All contributions (yours and employer's) are immediately vested (yours to keep if you leave).
- Before age $59\frac{1}{2}$, withdrawals while employed may require proof of financial hardship and incur a 10% penalty plus income tax.

When answering:

- Imagine you face this new hypothetical plan but all other aspects of your job are the same.
- Base decisions on your current income and financial situation.
- There are no right or wrong answers – we're interested in your personal choices.

Please respond thoughtfully based on your present circumstances. Your honest answers are valuable to our research.

1. **No Match.**[15] Let's suppose your employer offers you the following retirement savings plan:
   - Your employer does not contribute to your account.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

   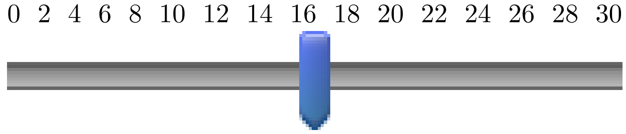

2. **100% Match.** Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 100% of your contributions.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

   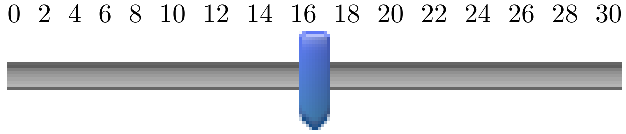

[14] Note that this number differs from the number of questions detailed below because we asked about several additional formulas which were not used in the study.

[15] As noted above, these bold labels summarizing match formulas are not observed by survey respondents.

3. **50% Match.** Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 50% (half) of your contributions.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

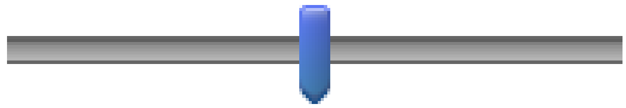

4. **25% Match.** Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 25% (a fourth) of your contributions.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

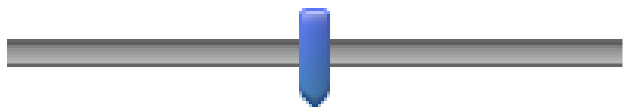

5. **100% on 6.** *[Randomized matching slope: each respondent saw only one of the three "X% on 6" scenarios in this section (100%, 50%, or 25%); the slope was assigned at random across respondents.]*
   Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 100% of your contributions up to 6% of your paycheck.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

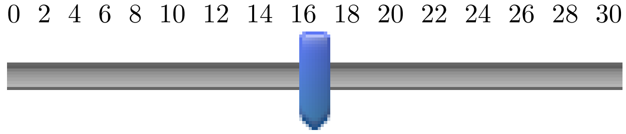

6. **50% on 6.** *[Randomized variant of "X% on 6"; see note above.]*
   Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 50% of your contributions up to 6% of your paycheck.

   What percent of your monthly income would you contribute to this plan?

   0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

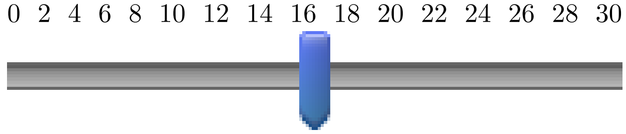

7. **25% on 6.** *[Randomized variant of "X% on 6"; see note above.]*
   Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer matches 25% of your contributions up to 6% of your paycheck.

What percent of your monthly income would you contribute to this plan?

0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

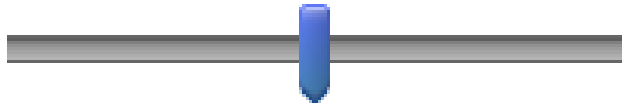

8. **Universal 10%.** Now, let's suppose your employer offers you the following retirement savings plan:
   - Your employer makes a fixed contribution of 10% of your paycheck, regardless of how much you contribute.

   What percent of your monthly income would you contribute to this plan?

0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30

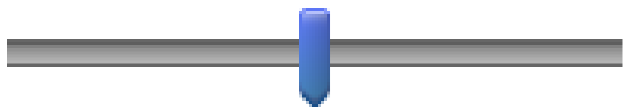

F.3 Self-Assessed Retirement Savings Adequacy

The next questions explore how you value present versus future financial rewards. We are interested in your personal preferences – there are no right or wrong answers.

1. **It can be hard to exercise self-control, and some people feel that there are things they do too much or too little – for example, exercise, save money, or eat junk food. Do you feel like you save for retirement too little, too much, or the right amount?**
   - Far too little (10)
   - Slightly too little (12)
   - The right amount (13)
   - Slightly too much (14)
   - Far too much (17)
   - I don't know / I am not sure (18)

F.4 Financial Literacy Question

Now, let's explore how 401(k) taxes work for retirement planning.

1. **When someone makes withdrawals from a traditional 401(k) after age 60, how are the taxes handled?**
   - You pay income taxes on what you save now but do not pay income taxes on withdrawals (1)
   - You pay no income taxes on what you save now but pay income taxes on withdrawals (2)
   - You pay income taxes both on what you save now and on withdrawals (3)
   - You pay no income taxes on what you save now and no income tax on withdrawals (4)
   - I do not know (5)